






Universal Invariant Checking of Parametric Systems with Quantifier-free SMT Reasoning

Alessandro Cimatti , Alberto Griggio , and Gianluca Redondi 

Fondazione Bruno Kessler, Trento, Italy
{cimatti, griggio, gredondi}@fbk.eu

Abstract. The problem of invariant checking in parametric systems – which are required to operate correctly regardless of the number and connections of their components – is gaining increasing importance in various sectors, such as communication protocols and control software. Such systems are typically modeled using quantified formulae, describing the behaviour of an unbounded number of (identical) components, and their automatic verification often relies on the use of decidable fragments of first-order logic in order to effectively deal with the challenges of quantified reasoning.

In this paper, we propose a fully automatic technique for invariant checking of parametric systems which does not rely on quantified reasoning. Parametric systems are modeled with array-based transition systems, and our method iteratively constructs a quantifier-free abstraction by analyzing, with SMT-based invariant checking algorithms for non-parametric systems, increasingly-larger finite instances of the parametric system. Depending on the verification result in the concrete instance, the abstraction is automatically refined by leveraging candidate lemmas from inductive invariants, or by discarding previously computed lemmas.

We implemented the method using a quantifier-free SMT-based IC3 as underlying verification engine. Our experimental evaluation demonstrates that the approach is competitive with the state of the art, solving several benchmarks that are out of reach for other tools.

Keywords: Parametric Systems · Array-based transitions systems · Abstraction-refinement · SMT

1 Introduction

Parametric systems consist of a finite but unbounded number of components. Examples include communication protocols (e.g. leader election), feature systems, or control algorithms in various application domains (e.g. railways interlocking logics). The key challenge is to prove the correctness of the parametric system for all possible configurations corresponding to instantiations of the parameters.

Parametric systems can be described as symbolic array-based transition systems [10], where the dependence on the configuration is expressed with first-order quantifiers in the initial condition and the transition relation of the model.

In this paper, we propose a fully automated approach for solving the universal invariant problem of array-based systems. The distinguishing feature is that the approach, grounded in SMT, does not require dealing with quantified theories, with obvious computational advantages. The algorithm implements an abstraction-refinement loop, where the abstract space is a quantifier-free transition system over some SMT theories. Our inspiration and starting point is the Parameter Abstraction of [3,15], which we extend in two directions. First, we modify the definition of the abstraction, by introducing a set of different *environment variables*, which intuitively overapproximate the behaviour of all the instances not precisely tracked by the abstraction, and by introducing a special *stuttering transition* in which the environment is allowed to change non-deterministically. Second, we combine the abstraction with a method for *automatically* inferring candidate universal lemmas, which are used to strengthen the abstraction in case of spurious counterexamples. The candidate lemmas are obtained by generalization from the spuriousness proof carried out in a finite-domain instantiation of the concrete system. However, we do not require quantified reasoning to prove that they universally hold; rather, the algorithm takes into account the fact that candidate lemmas may turn out not to be universally valid. In such cases, the method is able to automatically discover such bad lemmas and discard them, by examining increasingly-higher-dimension bounded instances of the parametric system.

We implemented the method in a tool called LAMBDA. At its core, LAMBDA leverages modern model checking approaches for quantifier-free infinite-state systems, i.e. the SMT-based approach of IC3 with implicit abstraction [4], in contrast to other approaches [19] where the abstract space is Boolean. In our experimental evaluation, we compared LAMBDA with the state-of-the-art tools MCMT [11] and CUBICLE [7]. The results show the advantage of the approach, that is able to solve multiple benchmarks that are out of reach for its competitors.

The rest of the paper is structured as follows. In Section 2 we present some logical background, and in Section 3 we describe array-based systems. We give an informal overview of the algorithm in Section 4. In Section 5 we define the abstraction and state its formal properties. In Section 6 we discuss the approach to concretization and refinement, and we present the techniques for inferring candidate lemmas. We discuss the related work in Section 7, and we present our experimental evaluation in Section 8. Finally, in Section 9 we draw some conclusions and present directions for future work. For lack of space, the proofs of our theoretical results, as well as further details on our experiments, are reported in an extended technical report [5].

2 Preliminaries

Our setting is standard first order logic. A theory \mathcal{T} in the SMT sense is a pair $\mathcal{T} = (\Sigma, \mathcal{C})$, where Σ is a first order signature and \mathcal{C} is a class of models over Σ . A theory \mathcal{T} is closed under substructure if its class \mathcal{C} of structures is such

that whenever $\mathcal{M} \in \mathcal{C}$ and \mathcal{N} is a substructure of \mathcal{M} , then $\mathcal{N} \in \mathcal{C}$. We use the standard notions of Tarskian interpretation (assignment, model, satisfiability, validity, logical consequence). We refer to 0-arity predicates as Boolean variables, and to 0-arity uninterpreted functions as (theory) variables. A literal is an atom or its negation. A clause is a disjunction of literals. A formula is in conjunctive normal form (CNF) iff it is a conjunction of clauses. If x_1, \dots, x_n are variables and ϕ is a formula, we might write $\phi(x_1, \dots, x_n)$ to indicate that all the variables occurring free in ϕ are in x_1, \dots, x_n .

If ϕ is a formula, t is a term and v is a variable which occurs free in ϕ , we write $\phi[v/t]$ for the substitution of every occurrence of v with t . If \underline{t} and \underline{v} are vectors of the same length, we write $\phi[\underline{v}/\underline{t}]$ for the simultaneous substitution of each v_i with the corresponding term t_i . We use an if-then-else notation for formulae. We write **if** ϕ_1 **then** ψ_1 **elif** ϕ_2 **then** ψ_2 **elif** $\dots \psi_{n-1}$ **else** ψ_n to denote the formula $(\phi_1 \rightarrow \psi_1) \wedge ((\neg\phi_1 \wedge \phi_2) \rightarrow \psi_2) \wedge \dots \wedge ((\neg\phi_1 \dots \neg\phi_{n-1} \wedge \neg\phi_n) \rightarrow \psi_n)$.

Given a set of variables \underline{v} , we denote with \underline{v}' the set $\{v' | v \in \underline{v}\}$. A symbolic transition system is a triple $(\underline{v}, I(\underline{v}), T(\underline{v}, \underline{v}'))$, where \underline{v} is a set of variables, and $I(\underline{v})$, $T(\underline{v}, \underline{v}')$ are first order formulae over some signature. An assignment to the variables in \underline{v} is a state. A state s is initial iff it is a model of $I(\underline{v})$, i.e. $s \models I(\underline{v})$. The states s, s' denote a transition iff $s \cup s' \models T(\underline{v}, \underline{v}')$, also written $T(s, s')$. A path is a sequence of states s_0, s_1, \dots such that s_0 is initial and $T(s_i, s'_{i+1})$ for all i . We denote paths with π , and with $\pi[j]$ the j -th element of π . A state s is reachable iff there exists a path π such that $\pi[i] = s$ for some i . A variable v is frozen iff for all π, i it holds that $\pi[i](v) = \pi[0](v)$. In the following, when we define a frozen variable v , we assume that this is done by having a constraint $v' = v$ as a top-level conjunct of the transition formula. A formula $\phi(\underline{v})$ is an invariant of the transition system $C = (\underline{v}, I(\underline{v}), T(\underline{v}, \underline{v}'))$ iff it holds in all the reachable states. Following the standard model checking notation, we denote this with $C \models \phi(\underline{v})$.¹ A formula $\phi(\underline{v})$ is an inductive invariant for C iff $I(\underline{v}) \models \phi(\underline{v})$ and $\phi(\underline{v}) \wedge T(\underline{v}, \underline{v}') \models \phi(\underline{v}')$.

3 Modeling Parametric Systems as Array-based Transition Systems

In order to describe parametric systems, we adapt from [10] the notion of array-based systems. In the following, we fix a theory of indexes $\mathcal{T}_I = (\Sigma_I, \mathcal{C}_I)$ and a theory of elements $\mathcal{T}_E = (\Sigma_E, \mathcal{C}_E)$. In order to model the parameters, we require that the class \mathcal{C}_I is closed under substructure. Then with A_I^E we denote the theory whose signature is $\Sigma = \Sigma_I \cup \Sigma_E \cup \{[-]\}$, and a model for it is given by a set of total functions from a model of \mathcal{T}_I to a model of \mathcal{T}_E . In general, we can have several array theories with multiple sorts for indexes and elements.

¹ Note that we use the symbol \models with three different denotations: if ϕ, ψ are formulae, $\phi \models \psi$ denotes that ψ is a logical consequence of ϕ ; if μ is an interpretation, and ψ is a formula, $\mu \models \psi$ denotes that μ is a model of ψ ; if C is a transition system, $C \models \psi$ denotes that ψ is an invariant of C .

For simplicity, we fix only an *index* sort and an *elem* sort. In the following, an array-based transition system

$$C = (a, \iota(a), \tau(a, a'))$$

is a symbolic transition system, with the additional constraints that:

- a is a variable of sort $index \mapsto elem$. We use a single variable for the sake of simplicity: additional variables of arbitrary type (also of index or element type) can be added without loss of generality.
- $\iota(a)$ is a first-order formula of the form $\forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$, where \underline{i} is of index sort and ϕ is a quantifier-free formula.
- $\tau(a, a')$ is a finite disjunction of formulae, $\bigvee_{k=1}^n \tau_k$, such that every τ_k is a formula of the following type (with $\underline{i}, \underline{j}$ of index sort):

$$\exists \underline{i} \forall \underline{j}. \psi(\underline{i}, \underline{j}, a[\underline{i}], a'[\underline{j}], a'[\underline{i}], a'[\underline{j}])$$

with ψ a quantifier-free formula.

This syntactic requirement subsumes the common guard and update formalism used for the description of parametric systems, used e.g in [10, 12, 15].

In the following, we shall refer to the disjuncts τ_k of τ as *transition rules* (or simply *rules* when clear from the context).

An array-based transition system can be seen as a family of transition systems, one for each cardinality of the finite models \mathcal{M}_I of \mathcal{T}_I . In the following, given d an integer, we denote with C^d the *finite instance of C of size d* obtained by instantiating the quantifiers of C over a set of fresh index variables of cardinality d (considered implicitly different from each other). Note that this C^d is a *symmetric presentation* [15]: if $\underline{c} = \{c_1, \dots, c_d\}$ are the fresh index variables, and σ is a permutation of \underline{c} , we have that, for every formula $\phi(\underline{c}, a[\underline{c}])$, $C^d \models \phi(\underline{c}, a[\underline{c}]) \Leftrightarrow C^d \models \phi(\sigma(\underline{c}), a[\sigma(\underline{c})])$.

Example 1 (Mutex Protocol for Ring Topology). Here we describe a simple protocol for accessing a shared resource, with processes in a ring-shaped topology. As an index theory, we use the finite sets of integers. As an element theory, we use both the Booleans and an enumerated data type of two elements, namely $\{idle, critical\}$. The array variable t , with sort $index \mapsto boolean$, is true in an index variable x if x holds the token. The variable s , with sort $index \mapsto \{idle, critical\}$ holds the current state of the process. In addition, we have an integer frozen variable $length$, which represents the length of the ring. The transition system is described by the following formulae:

Initial states. Initially, only one process holds the token, and every process is idle. We model this initial process with an additional constant *init_token* of sort *index*. Moreover, each index is bounded by the value of *length*. The initial formula is:

$$\begin{aligned} \forall j. p[j] = idle \wedge j \geq 1 \wedge j \leq length \wedge length > 0 \\ \wedge \left\{ \begin{array}{l} \text{if } j = \text{init_token} \text{ then } t[j] = \text{true} \\ \text{else } t[j] = \text{false} \end{array} \right. \end{aligned}$$

Transition rule 1. A process which holds the token can enter the critical section:

$$\begin{aligned} \exists i. s[i] = \text{idle} \wedge t[i] = \text{true} \wedge s'[i] = \text{critical} \wedge t'[i] = t[i] \wedge \\ \forall j, j \neq i. (s'[j] = s[j] \wedge t'[j] = t[j]) \end{aligned}$$

Transition rule 2. A process exits from the critical section and passes the token to the process at its right:

$$\begin{aligned} \exists i. \wedge s[i] = \text{critical} \wedge s'[i] = \text{idle} \wedge t'[i] = \text{false} \wedge \\ \forall j, j \neq i. \begin{cases} \text{if } j = 1 \wedge i = \text{length} \text{ then } s'[j] = s[j] \wedge t'[j] = \text{true} \\ \text{elif } j = i + 1 \wedge i < \text{length} \text{ then } s'[j] = s[j] \wedge t'[j] = \text{true} \\ \text{else } s'[j] = s[j] \wedge t'[j] = t[j] \end{cases} \end{aligned}$$

3.1 Universal invariant problem for array-based systems

In the following, given an array-based transition system

$$C = (a, \iota(a), \tau(a, a')),$$

the *universal invariant problem* is the problem of proving (or disproving) that a formula of the form $\Phi \stackrel{\text{def}}{=} \forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$ is an invariant for C .

Guard Strengthening In order to prove that $\forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$ is an invariant of a system $C = (a, \iota(a), \tau(a, a'))$, we can first strengthen the rules of C by adding the candidate invariant in conjunction with the transition relation, and then prove that the formula is an invariant of the newly-restricted system. This induction principle is justified by the following proposition:

Proposition 1 (Guard strengthening [15]) *Let $C = (a, \iota(a), \tau(a, a'))$ be a transition system and let Φ be $\forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$. Let $C_\Phi = (a, \iota(a), \tau(a, a') \wedge \Phi)$ be the guard-strengthening of C with respect to Φ . Then, if Φ is an invariant of C_Φ , it is also an invariant of C .*

Prophecy variables The universal quantifiers in the candidate invariant can be replaced with fresh frozen variables, called *prophecy variables*, that intuitively contain the indexes of the processes witnessing the violation of the property.

Proposition 2 (Removing quantifiers [19]) *Let $C = (a, \iota(a), \tau(a, a'))$ be an array-based system. The formula $\forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$ is an invariant for C iff the formula $\phi(\underline{p}, a[\underline{p}])$ is an invariant for $C_{+\underline{p}} = (a \cup \underline{p}, \iota(a), \tau(a, a'))$, where \underline{p} is a set of fresh frozen variables of index sort.*

For better readability, in the following we will omit the subscript $+\underline{p}$. Moreover, we assume that the index variables universally quantified in the candidate invariant are considered to be different. This does not limit expressiveness, and simplifies our discourse. Therefore, the prophecy variables induced by a candidate invariant are considered to be *implicitly different*.

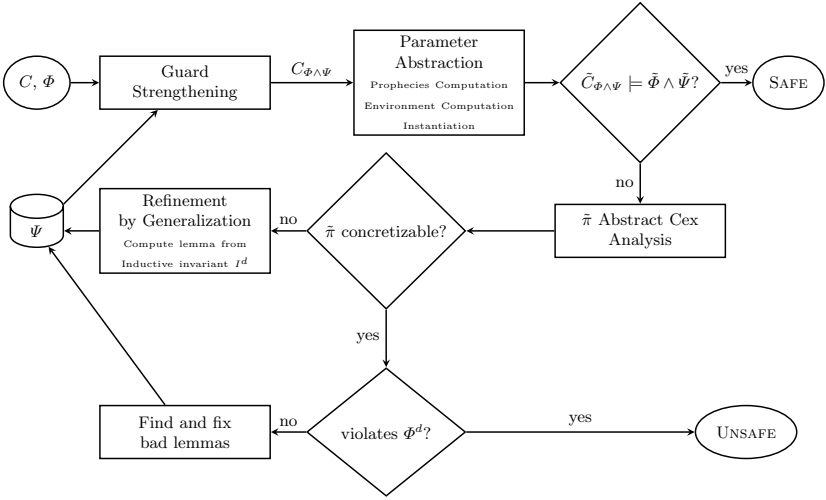


Fig. 1. An overview of the algorithm. C is an array-based transition system; Φ is a quantified candidate invariant; $\Psi \stackrel{\text{def}}{=} \{\psi_1, \dots, \psi_n\}$ is the set of candidate lemmas; $C_{\Phi \wedge \Psi}$ is a quantified transition system resulting from the strengthening of C ; $\tilde{C}_{\Phi \wedge \Psi}$ is a quantifier-free transition system.

4 Overview of the Method

In the following, let an array-based transition system $C \stackrel{\text{def}}{=} (a, \iota(a), \tau(a, a'))$, and a candidate universal invariant $\Phi \stackrel{\text{def}}{=} \forall \underline{i}. \phi(\underline{i}, a[\underline{i}])$ for C be given.

We now summarize the algorithm that attempts to solve the universal invariant problem for C and Φ . The algorithm, depicted in Figure 1, iterates trying either to construct an abstraction sufficiently precise to prove the property (exit with SAFE), or to find a finite instantiation of the problem exhibiting a concrete counterexample (exit with UNSAFE). The abstract space is quantifier-free, and obtained by instantiating the universally quantified formulae over two sets of index variables: the prophecy variables, which arise from the candidate invariant (as explained in Proposition 2), and are denoted with \underline{p} ; and the *environmental* variables, denoted with \underline{x} , which arise from the transition formula and are intended to represent the environment surrounding the \underline{p} indexes, interacting with them in the behaviour leading to the violation. While prophecy variables are frozen, thus representing the same indexes for the whole run, environmental variables are free to change at each time step, hence producing possibly spurious behaviours. The algorithm maintains a set of *candidate lemmas* $\Psi \stackrel{\text{def}}{=} \{\psi_i\}_i$, composed of universally quantified formulae, that are used to strengthen the property and to tighten the abstraction. Initially, Ψ is empty. In the following, if C^d is a finite instance of C and Φ is a candidate universal invariant, with Φ^d we denote the formula obtained from Φ by instantiating the quantifiers in variables used for the domain of cardinality d .

At each iteration, we carry out the following high-level steps (described in detail in the next sections):

- the property Φ to be proved is conjoined with the candidate lemmas in Ψ , and its quantifiers are moved in prenex form;²
- we construct the guard-strengthening $C_{\Phi \wedge \Psi}$ (cfr. Proposition 1), conjoining $\Phi \wedge \Psi$ to the transition rules of C ;
- we compute our modified Parameter Abstraction of $C_{\Phi \wedge \Psi}$ (defined in §5.1). First, we define the necessary prophecy variables \underline{p} and environmental variables \underline{x} . Then, we instantiate the quantifiers obtaining the quantifier-free array transition system $\tilde{C}_{\Phi \wedge \Psi}$.
- we (try to) solve the invariant checking problem $\tilde{C}_{\Phi \wedge \Psi} \models \tilde{\Phi} \wedge \tilde{\Psi}$ by calling a model checker for quantifier-free transition systems. $\tilde{\Phi} \wedge \tilde{\Psi}$ is obtained from $\Phi \wedge \Psi$ by removing quantifiers with prophecy variables, as in Proposition 2
- if the model checker concludes that there is no violation, then Φ holds in C (for the properties of the Parameter Abstraction), and we exit with SAFE.
- otherwise, we try to check whether the property violation in the abstract space corresponds to a real counterexample. We do so by checking whether the current property $\Phi \wedge \Psi$ is falsified in C^d , a suitable finite instance of C . That is, we check whether $C^d \models (\Phi \wedge \Psi)^d$.
- if $C^d \models (\Psi \wedge \Phi)^d$, then the abstraction must be tightened. When the verification of the finite instance succeeds, an inductive invariant I^d is produced, which is used to compute (candidate) lemmas by generalization from d to the universal case.
- if $C^d \not\models (\Psi \wedge \Phi)^d$, two cases are possible. First, we check if the (instantiation of the) property Φ is indeed violated. If so, we exit with UNSAFE, and we produce a concrete counterexample to the original problem, finitely witnessed in C^d .
- However, it is also possible that C^d does not violate Φ^d , but it falsifies some lemmas. In fact, the candidate lemmas obtained at previous iterations, by generalization on C^{d^-} with $d^- \neq d$, may not hold universally in C . In that case, the bad lemmas must be fixed, and the iteration is restarted.

When the algorithm terminates with UNSAFE, we are able to exhibit a finite counterexample trace in a finite instance of C violating the property. When the algorithm terminates with safe, then the property holds in C . The result is obtained by the following chain of implications: from Theorem 3, stated in the next section, we have that $\tilde{C}_{\Phi \wedge \Psi} \models \tilde{\Phi} \wedge \tilde{\Psi}$ implies $C_{\Phi \wedge \Psi} \models \tilde{\Phi} \wedge \tilde{\Psi}$. From Proposition 2, we have that $C_{\Phi \wedge \Psi} \models \Phi \wedge \Psi$. Therefore, from Proposition 1, we have $C \models \Phi \wedge \Psi$. In particular, we have $C \models \Phi$.

5 Modified Parameter Abstraction

We describe here our Parameter Abstraction. The first version of this approach was introduced in [3], and later formalized in [15]. In the following, we describe

² In the following, with $\Phi \wedge \Psi$ we denote the prenex form $\Phi \wedge \bigwedge_i \Psi_i$

a novel version of the abstraction, and how it can be applied to array-based transition systems. The main novelty is that, instead of using a special abstract index “*” that overapproximates the behaviour of the system in the array locations that are not explicitly tracked, we use n *environmental (index) variables* which are not abstracted, but are allowed to change nondeterministically in some transitions. This can be achieved by the usage of an additional **stuttering transition**: this rule allows the environmental variables to change value arbitrarily, while not changing the values of the array in the prophecies.

5.1 Abstraction Computation

Let an array-based transition system C and a universal invariant Φ be given³. By conjoining Φ to the transition rules in C , we obtain C_Φ , the guard strengthening of C with respect to Φ . Then, we define two sets of variables: the prophecy variables \underline{p} , in number determined by Proposition 2, and the environmental variables \underline{x} , in number determined by the greatest existential quantification depth in the transition rules of C_Φ . While the prophecies are frozen variables, the interpretation of the environmental variables is not fixed. Moreover, we assume that the values taken by \underline{p} and \underline{x} are different. We now define \tilde{C} , the parameter abstraction of C .

Initial formula Let $\iota(a)$ be $\forall \underline{i}.\phi(\underline{i}, a[\underline{i}])$, the initial formula of C in prenex form, with $\phi(\underline{i}, a[\underline{i}])$ quantifier-free. The initial formula of the abstract system is a quantifier-free first order formula, denoted $\tilde{\iota}(\underline{p}, a[\underline{p}])$ obtained by instantiating all the universal quantifiers in ι over the set of prophecy variables \underline{p} .

Transition formula The transition formula of C_Φ is still represented by a disjunction of formulae of the form⁴

$$\tau(a, a') \stackrel{\text{def}}{=} \exists \underline{i} \forall \underline{j}.\psi(\underline{i}, \underline{j}, a[\underline{i}], a[\underline{j}], a'[\underline{i}], a'[\underline{j}]).$$

For simplicity, we can assume that we have only one rule $\tau(a, a')$. First, we compute the set of all substitutions of the \underline{i} over $\underline{p} \cup \underline{x}$, and we consider the set of formulae $\{\tilde{\tau}_j(\underline{p}, \underline{x}, a, a')\}$, where j ranges over the substitutions, and $\tilde{\tau}_j$ is the result of applying the substitution to τ .

Then, for each formula in the set $\{\tilde{\tau}_j\}$, we instantiate the universal quantifiers over the set $\underline{p} \cup \underline{x}$, obtaining a quantifier-free formula over prophecy and environmental variables.

Moreover, we consider an additional transition formula, called the **stuttering transition**, defined by:

$$\tilde{\tau}_S \stackrel{\text{def}}{=} \bigwedge_{p \in \underline{p}} a'[p] = a[p] \wedge p' = p$$

³ These represent the system and the property in input to each iteration of the loop.

⁴ Possibly by performing trivial logical manipulations to distribute the guard strengthening inside the rules.

The disjunction of all the abstracted transition formulae is the transition formula $\tilde{\tau}$. So, we can now define the transition system

$$\tilde{C} \stackrel{\text{def}}{=} (\{a, \underline{p}, \underline{x}\}, \tilde{l}(\underline{p}, a[\underline{p}]), \tilde{\tau}(\underline{p}, \underline{x}, a[\underline{p} \cup \underline{x}], a'[\underline{p} \cup \underline{x}])).$$

Example 2. We apply the abstraction procedure to the transition rule 2 of the token in the ring protocol of Example 1.

Since the invariant is the formula $\forall i, j. \neg(s[i] = \text{critical} \wedge s[j] = \text{critical})$ it follows that we have two prophecy variables p_1, p_2 . Recall that the invariant itself is added to the transition as an additional conjunct. Since the existential quantification depth is one, we have only one environment variable x_1 . In the abstraction system we obtain three transition formulae from the original transition; we report the one indexed by the substitution mapping i into x_1 ; such a formula is equivalent to the following:

$$\begin{aligned} & s[x_1] = \text{crit} \wedge t[x_1] = \text{true} \wedge s'[x_1] = \text{idle} \wedge t'[x_1] = \text{false} \wedge \\ & \bigwedge_{j \in \{p_1, p_2\}} \left\{ \begin{array}{l} \text{if } j = 1 \wedge x_1 = \text{length} \text{ then } s'[j] = s[j] \wedge t'[j] = \text{false} \\ \text{elif } j = x_1 + 1 \wedge x_1 < \text{length} \text{ then } s'[j] = s[j] \wedge t'[j] = \text{false} \\ \text{else } s'[j] = s[j] \wedge t'[j] = t[j] \end{array} \right. \\ & \bigwedge_{\substack{i, j \in \{p_1, p_2, x_1\} \\ i \neq j}} \neg(s[i] = \text{critical} \wedge s[j] = \text{critical}) \end{aligned}$$

5.2 Stuttering Simulation

We define here the stuttering simulation induced by our version of the Parameter Abstraction. The proof of the main theorem can be found in the appendix. The stuttering is induced by $\tilde{\tau}_S$: this is a weaker version than the simulation induced by [15], yet it is sufficient for preserving invariants.

Definition 1 (Stuttering simulation) *Given two symbolic transition systems $C_1 = (\underline{x}_1, \iota_1, \tau_1)$ and $C_2 = (\underline{x}_2, \iota_2, \tau_2)$, with sets of states S_1 and S_2 , a stuttering simulation \mathcal{S} is a relation $\mathcal{S} \subset S_1 \times S_2$, such that:*

- for every $s_1 \in S_1$ such that $s_1 \models \iota_1$, there exists some $s_2 \in S_2$ such that $(s_1, s_2) \in \mathcal{S}$ and $s_2 \models \iota_2$;
- for every $(s_1, s_2) \in \mathcal{S}$, and for every $s'_1 \in S_1$ such that $s_1 \cup s'_1 \models \tau_1$, there exists either some $s'_2 \in S_2$ such that $(s'_1, s'_2) \in \mathcal{S}$ and $s_2 \cup s'_2 \models \tau_2$, or some $(s'_2, s''_2) \in S_2 \times S_2$ such that $(s'_1, s''_2) \in \mathcal{S}$, and $s_2 \cup s'_2 \models \tau_2$, $s'_2 \cup s''_2 \models \tau_2$.

If such a relation exists, we say that C_2 stutter simulates C_1 .

We write $\mathcal{S}(s_1)$ for $\{s_2 \mid (s_1, s_2) \in \mathcal{S}\}$. We recall that stutter simulation preserves reachability, i.e. if C_2 stutter simulates C_1 , then if s_1 is reachable in C_1 then the set $\mathcal{S}(s_1)$ is reachable in C_2 . Formally, the stuttering simulation induced by the Parameter Abstraction is defined as follows.

Definition 2 (Simulation) *Let C be the original transition system and let \tilde{C} be its Parameter Abstraction. Let s and \tilde{s} denote states of C and \tilde{C} , respectively. We define \mathcal{S} as follows:*

$$\mathcal{S}(s, \tilde{s}) \text{ iff } s(a)[i] = \tilde{s}(a)[i] \text{ for all } i \in \bigcup_{p \in \underline{p}} \tilde{s}(p).$$

Intuitively, we require that in the concrete state s and the abstract state \tilde{s} , the array is interpreted in the same way for all the locations referred by the prophecy variables. We then have the following:

Theorem 3. *The relation \mathcal{S} is a stuttering simulation between C and \tilde{C} . Moreover, if $\tilde{C} \models \Phi(\underline{p}, a[\underline{p}])$, then $C \models \Phi(\underline{p}, a[\underline{p}])$.*

6 Refinement

If $\Phi(\underline{p}, a[\underline{p}])$ does not hold in \tilde{C} , in general we cannot conclude anything, since the abstraction could be too coarse. So, if an abstract counterexample is encountered, we try to explore a small instance of the system to see if this counterexample occurs in it. To choose the appropriate size, our algorithm keeps a counter d , whose value is equal to the size to explore. Initially, d is equal to the number of (universally-quantified) index variables in the property Φ .⁵ When an abstract counterexample is encountered, we check whether $C^d \models (\Phi \wedge \Psi)^d$. For this check, we use a model checker able to return, in case of success, an inductive invariant I^d . From the inductive invariant we compute some first order formulae J which will be a new set of candidate lemmas. We will see later how to obtain this generalization. After computing the new lemmas, we set $d = d + 1$. If a concrete counterexample is found, then there are two cases: (i) the counterexample falsifies the original property, and we exit from the algorithm with a concrete counterexample; (ii) the counterexample falsifies some lemmas; in this case we remove the lemma and restart the loop (without changing d).

6.1 From Invariants to Universal Lemmas

Definition 3 *Let d be an integer, and let I^d be a set of clauses containing d variables. A generalization of I^d is a first-order formula J such that, when evaluating the quantifiers in J in a domain with precisely d elements, we obtain a formula equivalent to I^d .*

We use the following technique for generalization. Suppose that I^d is in CNF, and that we used c_1, \dots, c_d as variables for an instance with d elements. Then, $I^d = \mathcal{C}_1 \wedge \dots \wedge \mathcal{C}_n$ is a conjunction of clauses. From each of those clauses we

⁵ Recall that we assume that quantified index variables are required to be different. Therefore, the property holds vacuously on instances of size smaller than the number of index variables in Φ .

will obtain a new candidate lemma. Let $AllDiff(\underline{i})$ be the formula which states that all variables in \underline{i} are different from each other. Since every C^d is given by a symmetric presentation [15], we have that, for every $i \in \{1, \dots, n\}$, $C^d \models \forall i_1, \dots, i_h. AllDiff(i_1, \dots, i_h) \rightarrow C_i(i_1, \dots, i_h)$, where the quantifiers range over c_1, \dots, c_d and $h \leq d$ is the number of variables which occur in C_i . This means that $J \stackrel{\text{def}}{=} \bigwedge_i \forall \underline{i}. AllDiff(\underline{i}) \rightarrow C_i(\underline{i})$ is a generalization of I^d . In our algorithm, we add the set $\{\forall \underline{i}. C_i(\underline{i})\}_{i=1}^n$ of new candidate lemmas to Ψ . Note that we omitted the formula $AllDiff$ for our assumption on the different values of index variables.

Fixing Unsound Lemmas Unfortunately, we know a priori that a lemma holds only for the instance from which it was generalized. In general, its universal generalization obtained as outlined above might not hold in the system.

Suppose that the formula ψ_1 is a candidate lemma, obtained by generalization after the successful verification of an instance of size d . Suppose that later, a counterexample for ψ_1 is found by exploring a different instance $C^{d'}$ (with $d' > d$). This means that the lemma ψ_1 does not hold universally, but only for some finite instances of the system (including C^d), and not in general. In this case, we simply remove ψ_1 from the set of candidate lemmas Ψ , thus effectively weakening our working property (from $\Phi \wedge \Psi$ to $\Phi \wedge (\Psi \setminus \{\psi_1\})$). While this may cause a particular (abstract) counterexample to be encountered more than once during the main loop of the algorithm, since the finite instances are explored monotonically and their size d is increased after every successful verification of a bounded instance, the overall procedure still makes progress by exploring increasingly-large instances of the system. The hope is that eventually the algorithm will discover enough good lemmas that block the abstract counterexample. This notion of (weak) progress is justified by the following:

Proposition 4 *Let $\tilde{\pi}$ be an abstract counterexample, Ψ be the current set of universally quantified lemmas, and d be the size of the bounded instance to explore. During every execution of the algorithm, the same triple $(\tilde{\pi}, \Psi, d)$ never occurs twice.*

7 Related Work

Parametric verification is a challenging problem, and there is a large body of work in the literature devoted to this problem. Here, we (necessarily) focus on the approaches that are most related to ours.

Several methods are based on quantifier elimination using decidable fragments of first order logic, with notable examples in [7, 10, 22]. These methods guarantee a high degree of automation, but typically impose strong syntactic requirements in the input problem, and may suffer from scalability issues. A second popular approach is based on abstraction and abstraction refinement. Within this family of abstractions, earlier versions of the Paramater Abstraction [3, 15] have been used successfully also for industrial protocols [24]. The

main drawback is that the degree of automation is limited, and substantial expertise is required to obtain the desired results. The first steps of our abstraction algorithm are inspired by the ones in [19] and [15]. The key difference from [19] is that in that work the abstract transition system \tilde{C} is given by an eager propositional abstraction, with the axioms of the background theories recovered by the usage of some schemata. Here we retain the theory of arrays in the abstract space \tilde{C} . Moreover, differently from both [15] and [19], our procedure includes an automatic refinement of the abstraction in a counterexample-driven manner.

Ivy [20, 22] implements both semi-automatic invariant checking with decidable logics (namely, Effectively Propositional Logic – EPR) and compositional abstraction with eager axioms [19]. MYPYVY [13, 14] is a model checker inspired by the language of Ivy. It implements a version of IC3 capable of dealing with universal formulas [13]; the algorithm is completely automatic, but it is still based on quantifier elimination via reduction to decidable logics. In a more recent work, MYPYVY has gained the capability of inferring invariants with quantifier alternations, using a procedure that combines separators and first-order logic [14]. At the moment, our framework is capable of handling only universally quantified invariants. On the other hand, our approach is not limited to EPR, but it can in principle handle formulae with arbitrary SMT theories.

Exploring small instances of a parameterized system for candidate lemmas is a popular approach for parametric verification. In [8], this idea is used to over-approximate backward reachable states inside an algorithm which combines backward search and quantifier elimination. In [16], a finite-instance exploration is used together with a theorem prover to check the validity of candidate lemmas. In [17], candidate invariants are obtained from the set of reachable states of small instances. Similarly to our approach, these lemmas are used to strengthen an earlier version of the parameter abstraction. However, human intervention is still needed for the refinement.

A similar approach is presented in [23], where lemmas are obtained from a generalization of the proof of the property in a small instance of the protocol. The main difference with our technique, besides the methods used to extract such invariants, is the following: in [23], the authors show that to prove that a property (conjoined with lemmas) is inductive for all N , it is enough to prove that it is inductive for a particular N_0 , which is computable from the number of variables in the description of the system. This result is obtained from the imposed syntactic structure of the system. On the other hand, we impose less structure, and we rely on proving the property in an abstract version (and not a concrete instance) of the system. Moreover, our approach is integrated in an abstraction/refinement loop, which is missing from [23].

Another SMT-based approach for parametric verification is in [12]. The method is based on a reduction of invariant checking to the satisfiability of non-linear Constrained Horn Clauses (CHCs). Besides differing substantially in the overall approach, the method is more restrictive in the input language, and handles invariants only with a specific syntactic structure.

The use of prophecy variables for inferring universally quantified invariants has been explored also in non-parametric contexts, such as [18]. The main difference with our work is that [18] focuses on finding quantified invariants for quantifier-free transition systems with arrays, rather than array-based systems with quantifiers. The overall abstraction-refinement approach is also substantially different.

8 Experimental Evaluation

We have implemented our algorithm in a tool called LAMBDA (for **L**earning **A**bstractions **f**ro**M** **B**ounded **A**nalysis). LAMBDA is written in Python, and uses the SMT-based IC3 with implicit predicate abstraction of [4] as underlying quantifier-free verification engine.⁶ LAMBDA accepts as input array-based systems specified either in the language of MCMT [11] or in VMT format (a light-weight extension of SMT-LIB to model transition systems [25]). In case of successful termination, LAMBDA generates either a counterexample trace (for violated properties) in a concrete instance of the parametric system, or a quantified inductive invariant that proves the property for any instance of the system. In the latter case, LAMBDA can also generate proof obligations that can be independently checked with an SMT solver supporting quantifiers, such as Z3 [21] or CVC4 [2]. More specifically, the quantified inductive invariant can be generated by LAMBDA by simply universally quantifying all the (index) variables in the inductive invariant generated for \tilde{C} , and conjoining it with the lemmas Ψ discovered during the main loop iterations. Computing such an invariant is immediate after the termination of the algorithm, and does not require additional reasoning.

In order to evaluate the effectiveness of our method, we have compared LAMBDA with two state-of-the-art tools for the verification of array-based systems, namely CUBICLE [7] and MCMT. We could not include MYPYVY in the comparison, due to the many differences in input languages and modeling formalisms, which make an automatic translation of the benchmarks very difficult. We would also have liked to compare with the technique of [12], however the prototype tool mentioned in the paper doesn't seem to be available.

For our evaluation, we have collected a total of 116 benchmarks, divided in three different groups:

Protocols consists of 42 instances taken from the MCMT or the CUBICLE distributions, and used in previous works on verification of array-based systems. We have used all the instances which were available in both input formats, and we have split benchmarks containing multiple properties into different files.

DynArch consists of 57 instances of verification problems of dynamic architectures, taken from [6]. These benchmarks make use of arithmetic constraints on

⁶ In our implementation, we use the theory of integers as an index theory. At first, this may seem odd, since we should consider all finite subsets of the integers. However, this is not a problem, since the satisfiability of a quantifier-free UFLIA formula is equivalent to its satisfiability in a finite index model.

Table 1. Summary of experimental results.

Benchmark family	# of instances	Lambda		MCMT		Cubicle	
		Solved	Unique	Solved	Unique	Solved	Unique
Protocols	42	34	3	24	0	30	1
DynArch	57	48	5	48	5	–	–
Trains	17	17	–	–	–	–	–

index terms, which are not supported by CUBICLE. Therefore, we could only compare LAMBDA with MCMT on them.

Trains consists of 17 instances derived by (a simplified version of) verification problems on railway interlocking logics [1]. These benchmarks make use of several features that are not fully supported by CUBICLE and MCMT (such as non-functional updates in the transition relation, transition rules with more than one universally-quantified variable, real-valued variables). None of such restrictions applies to LAMBDA, which in general accepts models with significantly fewer syntactic constraints than CUBICLE and MCMT. Since these instances are inspired by relevant real-world verification problems, we believe that it is interesting to include them in the evaluation even though we could only run LAMBDA on them.

Our implementation, all the benchmarks, and the scripts for reproducing the results are available at <http://es.fbk.eu/people/griggio/papers/cade21-lambda.tar.gz>. We have run our experiments on a cluster of machines with a 2.90GHz Intel Xeon Gold 6226R CPU running Ubuntu Linux 20.04.1, using a time limit of 1 hour and a memory limit of 4GB for each instance. We have used the default settings for MCMT, whereas for CUBICLE we have also enabled the BRAB algorithm.⁷ A summary of the results of our evaluation are presented in Table 1. More details are provided in our extended version [5].

Overall, LAMBDA is very competitive with the state of the art, and in fact it solves the largest number of instances (even when disregarding the Trains group, which cannot be handled by the other tools). When considering the Protocols group, CUBICLE is often significantly faster than LAMBDA, especially on easier problems, thanks to its explicit-state exploration component (part of the BRAB algorithm). However, the symbolic techniques used by LAMBDA allow it to generally scale better to larger, more challenging problems: in the end, LAMBDA solves 4 more instances than CUBICLE, and 10 more than MCMT. The situation is different for the DynArch group, in which LAMBDA and MCMT solve the same number of instances. However, it is interesting to observe that both tools can solve 5 instances that the other tool cannot solve; more in general, it seems that the two approaches have somewhat complementary strengths. Moreover, as already stated above, the fact that LAMBDA imposes significantly less syntactic restrictions than the other two tools considered allowed it to handle all the instances of the Trains group, which cannot be easily modeled in the languages of MCMT or CUBICLE.

⁷ The results reported were obtained using `-brab 2`; we have however experimented also with other (small) values for `-brab`, without noticing any significant difference.

Finally, we wish to remark that we have generated SMT proof obligations for checking the correctness of all the (universally quantified) inductive invariants produced by LAMBDA, and checked them with both CVC4 and Z3. None of the solvers reported any error, and overall the combination of the two solvers was able to successfully verify all the proof obligations for 65 of the 67 instances reported as safe.⁸ We believe that the fact that we can easily produce proof obligations that can be independently checked is another strength of our approach. This is in contrast to the approach of CUBICLE, where generating proof obligations is nontrivial [9].

9 Conclusions

In this paper we tackled the problem of universal invariant checking for parametric systems. We proposed a fully-automated abstraction-refinement approach, based on quantifier-free reasoning. The abstract model, that stutter simulates the concrete model, is a quantifier-free symbolic transition system refined by (the instantiation of) candidate universal lemmas. These are obtained by analyzing the proofs of validity of the property in a finite instance of the parametric system. We experimentally evaluated an implementation on standard benchmarks from the literature. The results show the effectiveness of the method, also in comparison with state-of-the-art tools (CUBICLE, MCMT). We are able to prove, in a fully automated manner and without manual intervention, several benchmarks that are considered challenging. In the future, we plan to work on generalization, to improve the ability of inferring the right lemmas from a small instance, and to find more effective ways to filter out bad candidates. On the theoretical side, we will investigate the relation between the termination of the algorithm and decidable classes of parametric systems (e.g. those that enjoy a cut-off property). Finally, we will work on the verification of temporally extended properties which are also preserved by stuttering simulations (such as fragments of Linear Temporal Logic).

References

1. Amendola, A., Becchi, A., Cavada, R., Cimatti, A., Griggio, A., Scaglione, G., Susi, A., Tacchella, A., Tesse, M.: A model-based approach to the design, verification and deployment of railway interlocking system. In: Margaria, T., Steffen, B. (eds.) *Leveraging Applications of Formal Methods, Verification and Validation: Applications - 9th International Symposium on Leveraging Applications of Formal Methods, ISoLA 2020, Rhodes, Greece, October 20-30, 2020, Proceedings, Part III*. Lecture Notes in Computer Science, vol. 12478, pp. 240–254. Springer (2020)
2. Barrett, C.W., Conway, C.L., Deters, M., Hadarean, L., Jovanovic, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: *CAV*. Lecture Notes in Computer Science, vol. 6806, pp. 171–177. Springer (2011)

⁸ In the remaining two cases, both solvers returned `unknown` when trying to prove the validity of some of the proof obligations.

3. Chou, C.T., Mannava, P.K., Park, S.: A simple method for parameterized verification of cache coherence protocols. In: Hu, A.J., Martin, A.K. (eds.) *Formal Methods in Computer-Aided Design*. pp. 382–398. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
4. Cimatti, A., Griggio, A., Mover, S., Tonetta, S.: Infinite-state invariant checking with IC3 and predicate abstraction. *Formal Methods Syst. Des.* **49**(3), 190–218 (2016)
5. Cimatti, A., Griggio, A., Redondi, G.: *Universal Invariant Checking of Parametric Systems with Quantifier-Free SMT Reasoning (extended version)*. Tech. rep., Fondazione Bruno Kessler (2021), <https://es-static.fbk.eu/people/griggio/papers/cade21extended.pdf>
6. Cimatti, A., Stojic, I., Tonetta, S.: Formal specification and verification of dynamic parametrized architectures. In: Havelund, K., Peleska, J., Roscoe, B., de Vink, E.P. (eds.) *Formal Methods - 22nd International Symposium, FM 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 15-17, 2018, Proceedings*. Lecture Notes in Computer Science, vol. 10951, pp. 625–644. Springer (2018)
7. Conchon, S., Goel, A., Krstic, S., Mebsout, A., Zaïdi, F.: Cubicle: A Parallel SMT-based Model Checker for Parameterized Systems. In: Parthasarathy, M., Seshia, S.A. (eds.) *CAV 2012: Proceedings of the 24th International Conference on Computer Aided Verification*. Lecture Notes in Computer Science, Springer Verlag, Berkeley, California, USA (July 2012)
8. Conchon, S., Goel, A., Krstic, S., Mebsout, A., Zaïdi, F.: Invariants for finite instances and beyond. In: *Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR, USA, October 20-23, 2013*. pp. 61–68. IEEE (2013)
9. Conchon, S., Mebsout, A., Zaïdi, F.: Certificates for parameterized model checking. In: *FM. Lecture Notes in Computer Science*, vol. 9109, pp. 126–142. Springer (2015)
10. Ghilardi, S., Nicolini, E., Ranise, S., Zucchelli, D.: Towards smt model checking of array-based systems. In: Armando, A., Baumgartner, P., Dowek, G. (eds.) *Automated Reasoning*. pp. 67–82. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
11. Ghilardi, S., Ranise, S.: Backward reachability of array-based systems by SMT solving: Termination and invariant synthesis. *Log. Methods Comput. Sci.* **6**(4) (2010)
12. Gurfinkel, A., Shoham, S., Meshman, Y.: Smt-based verification of parameterized systems. In: *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. p. 338–348. FSE 2016, Association for Computing Machinery, New York, NY, USA (2016)
13. Karbyshev, A., Bjørner, N., Itzhaky, S., Rinetzky, N., Shoham, S.: Property-directed inference of universal invariants or proving their absence. In: Kroening, D., Păsăreanu, C.S. (eds.) *Computer Aided Verification*. pp. 583–602. Springer International Publishing, Cham (2015)
14. Koenig, J.R., Padon, O., Immerman, N., Aiken, A.: First-order quantified separators. In: Donaldson, A.F., Torlak, E. (eds.) *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*. pp. 703–717. ACM (2020)
15. Krstic, S.: *Parametrized system verification with guard strengthening and parameter abstraction* (2005)
16. Li, Y., Duan, K., Jansen, D.N., Pang, J., Zhang, L., Lv, Y., Cai, S.: An automatic proving approach to parameterized verification. *ACM Trans. Comput. Logic* **19**(4) (Nov 2018)

17. Lv, Y., Lin, H., Pan, H.: Computing invariants for parameter abstraction. In: 2007 5th IEEE/ACM International Conference on Formal Methods and Models for Codesign (MEMOCODE 2007). pp. 29–38 (2007)
18. Mann, M., Irfan, A., Griggio, A., Padon, O., Barrett, C.W.: Counterexample-guided prophecy for model checking modulo the theory of arrays. CoRR **abs/2101.06825** (2021)
19. McMillan, K.L.: Eager abstraction for symbolic model checking. In: Chockler, H., Weissenbacher, G. (eds.) Computer Aided Verification. pp. 191–208. Springer International Publishing, Cham (2018)
20. McMillan, K.L., Padon, O.: Ivy: A multi-modal verification tool for distributed algorithms. In: Lahiri, S.K., Wang, C. (eds.) Computer Aided Verification. pp. 190–202. Springer International Publishing, Cham (2020)
21. de Moura, L.M., Bjørner, N.: Z3: an efficient SMT solver. In: TACAS. Lecture Notes in Computer Science, vol. 4963, pp. 337–340. Springer (2008)
22. Padon, O., McMillan, K.L., Panda, A., Sagiv, M., Shoham, S.: Ivy: Safety verification by interactive generalization. SIGPLAN Not. **51**(6), 614–630 (Jun 2016)
23. Pnueli, A., Ruah, S., Zuck, L.D.: Automatic deductive verification with invisible invariants. In: Margaria, T., Yi, W. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001 Genova, Italy, April 2-6, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2031, pp. 82–97. Springer (2001)
24. Talupur, M., Tuttle, M.R.: Going with the flow: Parameterized verification using message flows. In: 2008 Formal Methods in Computer-Aided Design. pp. 1–8 (2008)
25. VMT-LIB. <http://www.vmt-lib.org>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

