

This is a repository copy of *A Multi-agent Framework for Dependable Adaptation of Evolving System Architectures*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/93151/>

Version: Accepted Version

Proceedings Paper:

Johnson, Kenneth Harold Anthony, Sinha, Roopak, Calinescu, Radu orcid.org/0000-0002-2678-9260 et al. (1 more author) (2015) A Multi-agent Framework for Dependable Adaptation of Evolving System Architectures. In: 41st Euromicro Conference on Software Engineering and Advanced Applications (SEAA). IEEE , pp. 159-166.

<https://doi.org/10.1109/SEAA.2015.49>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A Multi-Agent Framework for Dependable Adaptation of Evolving System Architectures

Kenneth Johnson* Roopak Sinha* Radu Calinescu† Ji Ruan*

* School of Computer and Mathematical Sciences, Auckland University of Technology, Auckland, New Zealand

† Department of Computer Science, University of York, York, United Kingdom

Abstract—We present a multi-agent framework for the formal verification of component-based systems after changes such as addition, removal and modification of components. The core of our approach is an Agent Verification Engine (AVE) that constructs evolvable Belief-Desire-Intention (BDI) agents to coordinate and plan the re-verification of component models after system changes. The engine provides BDI-agents with existing techniques for the compositional verification of component-based systems. We illustrate this integration for Satisfiability Modulo Theories (SMT) constraint analysis and demonstrate our framework on requirements arising from industrial control systems.

I. INTRODUCTION

Component-based software systems are increasingly common, and include business- and safety-critical systems from domains as diverse as healthcare, transportation and finance [25]. Critical systems are expected to be reliable since downtime often results in a decrease of revenue or unavailability of essential services. Formal verification techniques such as model checking or automated theorem proving can be used to provide irrefutable proof of a system’s compliance to requirements by analysing mathematical models of the system against properties derived from the requirements.

Compositional verification techniques [3], [2] decompose the verification of a large, monolithic system model into a sequence of small verification steps that are applied to the components of the system, thus enabling formal verification of much larger systems. However, these approaches do not take into account the system’s runtime environment, where components are frequently added and updated or may unexpectedly fail. Agent-based modelling is widely used in safety or business critical systems as a means for intelligent adaptation in response to planned and unplanned runtime changes. This approach is well suited to adapt large scale cloud deployed systems [12], [26] and industrial manufacturing applications [20], [17] whose compliance to quality-of-service (QoS) requirements must be maintained during runtime.

In this paper, we introduce the *Agent Verification Engine* (AVE) which constructs agents to perceive, react, and adapt to runtime changes of a component-based system. These agents are based on the *Belief-Desire-Intention* (BDI) architecture [21], in which agents operate in terms of motivation and beliefs. BDI-agents have a formal basis in logic to formalize an agent’s decision making in an attempt to achieve its goals [9]. BDI-agents are constructed by AVE according to the system’s architecture and observe components to detect changes. When changes happen, agents perform *verification tasks* to verify the model of components against properties to determine compliance with system requirements.

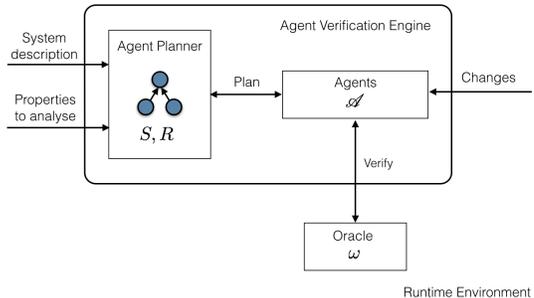


Fig. 1. Agent Verification Engine (AVE) workflow

The main contributions of this paper are (a) an algebraic specification of an agent verification engine based on BDI-agents, (b) the integration of a satisfiability modulo theories (SMT) solver to provide agents with verification capabilities, (c) an application of the engine to a real-world industrial control software case study, and (d) a prototype implementation of the agent verification engine using Jason AgentSpeak [4].

The rest of this paper is organised as follows. Section II provides key mathematical preliminaries. Section III formulates the agent verification engine algebraically and describes evolvable BDI-agent plans. Section IV integrates an existing SMT-based compositional verification approach into the engine. Section V applies the engine to analyse an industrial item sorting system. Section VI shows the effectiveness and scalability of our approach. Section VII discusses related work and Section VIII presents concluding remarks.

II. PRELIMINARIES

Assuming readers to have basic understanding of universal algebra [19] and BDI-agent architectures [21], [4], we define the following key tools and notations.

Signatures and Terms: A signature Σ consists of a sort s and a finite number of operation symbols $f : s^n \rightarrow s$ for $n \geq 0$. Symbols of the form $c \rightarrow S$ are signature constants. A signature Σ contains the sort *Bool* with constants *true*, *false* \rightarrow *Bool* and names the standard logical connectives \wedge and \neg . Operation symbols of the form $b : s^n \rightarrow \text{Bool}$ are Σ -predicates. Σ -terms are defined by the rules

$$t ::= c_1 \mid \dots \mid c_m \mid z \mid f_1(t_1, \dots, t_{m_1}) \mid \dots \mid f_n(t_1, \dots, t_{m_n})$$

for constants c_1, \dots, c_m , variable z from the set Z , operation symbols f_1, \dots, f_n and Σ -terms t_{i_j} .

The AgentSpeak Language: AgentSpeak [21] is an agent-oriented programming based on logic programming and the BDI architecture for autonomous agents. BDI-agents operate within an environment and receive continuous input *perceptions*. Agents respond by performing actions that affect the environment. The *beliefs, desires and intentions* of an agent correspond with the agent’s informational, motivational and decision making components that comprise its mental state.

Beliefs: An agent’s belief is constructed from predicate terms as follows. If $\mathbf{t} = t_1, \dots, t_n$ is a tuple of Σ -terms then $b(\mathbf{t})$ is a belief atom for the predicate symbol b . If $b(\mathbf{t})$ and $c(\mathbf{t})$ are belief atoms then $\neg b(\mathbf{t})$, $b(\mathbf{t}) \wedge c(\mathbf{t})$ are beliefs. Beliefs represent the information available to the agent and are stored in its *belief base* \mathcal{B} .

Goals: If g is a predicate symbol and $\mathbf{t} = t_1, \dots, t_n$ a tuple of Σ -terms then $!g(\mathbf{t})$ is an *achievement goal* an agent aims to achieve and $?g(\mathbf{t})$ is a *test goal* that tests if $g(\mathbf{t})$ is a true belief. Acquiring a new belief or goal creates a *trigger event* within an agent. If $b(\mathbf{t})$ is a belief atom then $+b(\mathbf{t})$, $-b(\mathbf{t})$ are triggering events corresponding to the addition and removal of beliefs. Similarly, $!g(\mathbf{t})$, $-!g(\mathbf{t})$ and $?g(\mathbf{t})$, $-?g(\mathbf{t})$ correspond to the addition and removal of achievement and test goals, respectively.

III. AGENT VERIFICATION ENGINE

This section algebraically specifies the *Agent Verification Engine (AVE)*. The engine’s workflow is depicted in Figure 1 and comprises the agent planner, that reads an algebraic system specification S and requirements R , a set \mathcal{A} of BDI-agents constructed by the agent planner to verify components of S and an oracle ω . The workflow has two phases. In the *agent construction phase*, the agent planner constructs a set $\mathcal{A} = \{\alpha_{t_1}, \dots, \alpha_{t_n}\}$ of BDI-agents for components t_1, \dots, t_n of the system S . Each agent is constructed with a plan library to determine compliance of its component with the requirements given by R . The agent invokes an external *oracle* ω which represents a compositional verification *service* to determine the component’s satisfaction of requirements. In the *agent runtime phase*, an agent perceives localised component changes t' in the runtime environment and notifies the agent planner. The agent planner transforms S by performing the corresponding term substitution $S' = S[t'/t]$ of the agent’s component term t with the new component term t' to obtain S' . The agent planner uses t' to update the agent’s beliefs and plan library.

A. The Oracle

Let \mathcal{M} be a set of component models, \mathcal{P} a set of properties and \mathcal{V} a set of values. The *oracle* is an automated process that accepts as input a verification task $w = (m, g, a)$ from the set $W = \mathcal{M} \times \mathcal{P} \times [\mathcal{P} \rightarrow \mathcal{V}]$ and is modelled mathematically by the total function $\omega : W \rightarrow \mathcal{V}$ such that

$$\omega(m, g, a) = \text{the value } v \text{ obtained from the process} \quad (1)$$

of verifying property g on model m
under assumptions of a

for model $m \in \mathcal{M}$, $g \in \mathcal{P}$ and $v \in \mathcal{V}$ and the map $a : \mathcal{P} \rightarrow \mathcal{V}$ of properties to verification values. The oracle is programmed to always terminate, potentially due to failure. If a failure occurs, then a special unverified value \mathbf{u} is returned

by the service. We extend the set of values to include this element by setting $\mathcal{V} = \mathcal{V} \cup \{\mathbf{u}\}$.

Our choice of notation in (1) generalises *assume-guarantee* [11], [14] model checking approaches to include decomposition techniques used in formal analysis of component-based systems. In assume-guarantee reasoning, a monolithic model m is decomposed into smaller, more manageable models m_1, \dots, m_q . Each model m_i is associated with a property $a_i \in \mathcal{P}$, formalising requirements of the component. Component-wise verification is performed in a sequence of steps whereby the verification of property a_i on model m_i obtains a verification value $v_i \in \mathcal{V}$. In symbols, we write $v_i \models a_i$, for the model m_i , $1 \leq i \leq q$. We denote the list of assumptions and verification results obtained by component-wise verification as \mathbf{a} . Using the assume-guarantee approach, the system requirements formalised as the property $g \in \mathcal{P}$ for the monolithic model m has verification result v under the assumptions \mathbf{a} . In symbols, we write $v \models g$.

B. Agent Construction Phase

We give an algebraic specification of component-based systems and use their inductive properties to derive the agent planner’s construction of the agent set \mathcal{A} .

Specifying Component-Based Systems: A component signature Σ is a finite set C of component sorts that represent parts of a system and a set $\{f_1, \dots, f_n\}$ of operations to be performed on components. By choosing some basic components c_1, \dots, c_m , e.g. constants in Σ and applying a sequence of operations, we form a component-based system as a high-level syntactic *component term* S . The term defines the hierarchical structure of a component-based system by expressing the sequence of operations carried out in the construction of the system and specifies the order in which they are applied. Let the set of all component terms over the signature Σ be denoted as $C(\Sigma)$. For example, the component term $comp(c_1, comp(c_2, c_3))$ in $C(\Sigma)$ is formed from the composition of basic component c_1 and the composition of basic components c_2 and c_3 , where $comp : C(\Sigma) \times C(\Sigma) \rightarrow C(\Sigma)$ is an operation in Σ . The inductive properties of component terms are a natural data structure for specifying BDI-agent verification plans.

Each component $c \in C(\Sigma)$ is associated with a model $m \in \mathcal{M}$ and a property $g \in \mathcal{P}$ formulated from R .

Agent Plans for Basic Components: Let c be a basic component of the component signature Σ , $m \in \mathcal{M}$ a component model and $g \in \mathcal{P}$ a property. The agent planner constructs a BDI-agent α_c with initial beliefs

$$\mathcal{B}_c = \{model(m), property(g)\} \quad (2)$$

where the predicates *model* and *property* specify agent α_c ’s belief of the component model m and property g . Initially, the agent has the achievement goal *!verify* to determine the verification result $v \in \mathcal{V}$ of property g for model m . We define the context expression

$$E_c = model(M) \wedge property(G) \quad (3)$$

such that the agent holds sufficient information about its component’s model and property for the verification plan

constructed by the agent planner. This plan has the form

$$+verify : E_c \leftarrow ?model(M) \wedge ?property(G) \wedge \omega(M, G, V) \wedge +result(V). \quad (4)$$

Since the initial belief base (2) of α_c satisfies the plan's context E_c , the agent's goal $!verify$ triggers plan (4). The test goals $?model(M)$ and $?property(G)$ unify the model m and property g with variables M and G respectively, according to the agent's beliefs (2). The agent α_c performs the action $\omega(M, G, V)$ and invokes the oracle to compute $v = \omega(m, g, \eta)$, with the null function $\eta : \emptyset \rightarrow \emptyset$ specifying that no assumptions are required for the computation. The result v is unified with variable V and $+result(V)$ adds the value to the belief base with the predicate symbol *result*.

Agent Plans for Composite Components: We give an inductive definition of the plans for agent α_t assigned to the component term $t = f(t_1, \dots, t_m)$ with sub-components t_1, \dots, t_m and the operation symbol $f \in \Sigma$. The agent has beliefs of its component model m and property g , and thus the initial belief base of α_t is set as $\mathcal{B}_t = \{model(m), property(g)\}$, similar to (2).

In general, the verification result $v \models g$ for model m of agent α_t is dependent upon the verification values obtained by agents $\mathcal{A}_t = \{\alpha_{t_1}, \dots, \alpha_{t_m}\}$ for the sub-terms of t . Agent α_t therefore must communicate with each agent in \mathcal{A}_t and construct assumptions $\mathbf{a} : \mathcal{P} \rightarrow \mathcal{V}$ to be supplied to the oracle. To coordinate with agents $\alpha_{t_1}, \dots, \alpha_{t_m}$, agent α_t is constructed with the plan

$$+!queryagents : true \leftarrow \begin{aligned} &send(t_1, givep, A_1) \wedge \dots \wedge send(t_m, givep, A_m) \\ &\wedge send(t_1, givev, V_1) \wedge \dots \wedge send(t_m, givev, V_m) \end{aligned} \quad (5)$$

which asks each sub-component agent α_{t_i} to retrieve the verification task information stored within its belief base. The properties a_1, \dots, a_m in \mathcal{P} and values v_1, \dots, v_m in \mathcal{V} are unified with variables A_1, \dots, A_m and V_1, \dots, V_m respectively. They are used to construct the assumption function $\mathbf{a} : \{a_1, \dots, a_m\} \rightarrow \{v_1, \dots, v_m\}$ such that

$$\mathbf{a}(a_i) = v_i. \quad (6)$$

To facilitate communication with α_t , each agent α_{t_i} is constructed with the plan

$$+?givep : true \leftarrow ?property(A) \wedge send(t, property_{t_i}(A)). \quad (7)$$

by the agent planner, which is triggered whenever its mental state is changed to the test goal $?givep$, for $1 \leq i \leq m$.

The plan body of (7) contains a special *send* action which enables transmission of messages between agents using the Knowledge Query and Manipulation Language (KQML) [16] which expresses an agent's intention for sending a message. Agents send messages for two purposes: i) to change the receiver's beliefs, or ii) to change the receiver's goals. In the plan body of (7), agent α_{t_i} 's property belief is unified with the variable P and sent to agent α_t , who receives the belief and adds it to \mathcal{B}_t . Similarly, the verification value obtained from the oracle ω by α_{t_i} is sent to α_t by the plan

$$+?givev : true \leftarrow ?result(V) \wedge send(t, result_{t_i}(V)). \quad (8)$$

The goal $?result(V)$ is complete only when α_{t_i} completes its verification task and $result_{t_i}(V)$ is added to its belief base.

For the composite component represented by the component term t with sub-terms t_1, \dots, t_m , the context expression (3) is extended to

$$E'_t = E_t \wedge property_{t_1}(A_1) \wedge \dots \wedge property_{t_m}(A_m) \wedge result_{t_1}(V_1) \wedge \dots \wedge result_{t_m}(V_m) \quad (9)$$

to ensure the values from agents $\alpha_{t_1}, \dots, \alpha_{t_m}$ have been received. The agent planner formulates two plans so that agent α_t can achieve its verification goal $!verify$. The first plan is applicable whenever α_t has not obtained values from agents $\alpha_{t_1}, \dots, \alpha_{t_m}$. In this context, the plan

$$+!verify : \neg E'_t \leftarrow !queryagents \wedge !verify$$

comprises the achievement subgoal $!queryagents$ which triggers plan (5) and continues to attempt the achievement $!verify$. In the context of receiving values from all sub-component agents, the plan invoked to achieve $!verify$ is

$$+verify : E'_t \leftarrow ?model(M) \wedge ?property(G) \wedge ?property_{t_1}(A_1) \wedge \dots \wedge ?property_{t_m}(A_m) \wedge ?result_{t_1}(V_1) \wedge \dots \wedge ?result_{t_m}(V_m) \wedge \omega(M, G, V, A_1, \dots, A_m, V_1, \dots, V_m) \wedge +result(V). \quad (10)$$

comprising additional test goals unifying A_1, \dots, A_m and V_1, \dots, V_m with properties and values from the verification tasks of agents $\alpha_{t_1}, \dots, \alpha_{t_m}$. The properties and values construct the assumption function \mathbf{a} defined in (6) used as input for action $\omega(M, G, V, P_1, \dots, P_m, V_1, \dots, V_m)$, invoking the external oracle to compute $v = \omega(m, g, \mathbf{a})$. The variable V is unified with the value v and stored in the agent's belief base with predicate symbol *result*.

C. Agent Runtime Phase

Runtime environments are dynamic and changes such as the addition, removal and modification of components occur in rapid succession. In this section, we specify plans enabling an agent to sense localised changes of the components modelled by their verification task. When a component change occurs, the agent is required to adapt and evolve its beliefs and behaviour during system operation in order to maintain a verification task that correlates to the actual state of the system. Formally, agent α_t is constructed with the plan

$$+!detect : true \leftarrow sense \wedge !detect. \quad (11)$$

which executes the action *sense* to detect changes in the component $t \in C(\Sigma)$. Component change detection only occurs after α_t achieves its initial verification goal. We extend the agent's verification plan by appending the achievement goal $!detect$ to the plan bodies listed in (4) and (10).

Agent Change Perceptions: The agent action *sense* notifies the engine that the system represented algebraically by the component term S is updated according to a component change sensed in the *runtime environment*. Mathematically, a component change is modelled as a component term $t' \in C(\Sigma)$ with an associated model $m' \in \mathcal{M}$ and property $g' \in \mathcal{P}$. When a component change occurs, the engine first updates

S and associated model and property. It then notifies agent α_t of the new model m' and property g' , and finally adds a *change perception* to the belief base of agent α_t corresponding to addition, removal or modification changes.

Modification Component Change: Change perceptions are handled by agents in the following cases. When component t_i of the component $t \equiv f(t_1, \dots, t_m)$ for $1 \leq i \leq m$ is modified, a change perception *modified* is issued by the engine to agent α_{t_i} . In this case, the engine first updates the component model and property to m'_{t_i} and g'_{t_i} respectively. The engine updates the belief base \mathcal{B}_{t_i} of the agent, replacing its model and property beliefs with information of the new verification task *model*(m'_{t_i}) and *property*(g'_{t_i}), respectively. The engine's change perception triggers the belief addition event $+modified$ in the agent and is handled by the plan

$$+modified : true \leftarrow !verify \wedge !notify.$$

supplied by the agent planner. The agent proceeds to call the oracle to verify its modified task, providing the agent with a new belief *result*(V). Next, the agent achieves the goal *!notify* to notify agent α_t that re-verification is complete. The plan

$$+!notify : true \leftarrow send(t, !reverify_{t_i}). \quad (12)$$

communicates the *!reverify_{t_i}* achievement goal to agent α_t . Agent α_t attempts to achieve this goal using the plan

$$\begin{aligned} +!reverify_{t_i} : true \leftarrow & -result_{t_i} \wedge -property_{t_i} \quad (13) \\ & \wedge send(t_i, givep, A_i) \\ & \wedge send(t_i, givev, V_i) \wedge !verify \wedge !notify. \end{aligned}$$

which obtains the modified property a'_i and updated verification result v'_i from α_{t_i} . The *!verify* goal uses the updated assumptions function $\mathbf{a}' : \mathcal{P} \rightarrow \mathcal{V}$ to invoke the oracle where

$$\mathbf{a}'(a_j) = \begin{cases} v'_i & \text{if } a_j = a_i, \\ \mathbf{a}(a_j) & \text{otherwise,} \end{cases}$$

and notifies agents that re-verification has occurred.

Addition Component Change: In the general case, the component term $t = f(t_1, \dots, t_m)$ has m sub-components. The addition of the new component t_{m+1} to t is specified by the component term $t' = f(t_1, \dots, t_m, t_{m+1})$. Mathematically, we define the substitution operation

$$S' = S[f(t_1, \dots, t_m)/f(t_1, \dots, t_m, t_{m+1})] \quad (14)$$

to replace instances of t with t' in the component term representation S to obtain S' . When the system representation is updated, the engine adds the change perception *add*(t_{m+1}) to agent α_t 's belief base, triggering a belief addition event that is handled accordingly by the plan

$$\begin{aligned} +add(T) : true \leftarrow & create(T) \wedge replan \\ & \wedge !verify \wedge !notify. \end{aligned}$$

The action *create*(T) invokes the agent planner to construct a new agent $\alpha_{t_{m+1}}$ and assign it the verification task w associated with the new component t_{m+1} . Once constructed, the agent completes its initial achievement goal *!verify*. The action *replan* described in the next section is used by α_t to take into account the added component. Once α_t obtains new plans it re-verifies to take into account the new component. The goal *!notify* ensures all relevant agents are notified.

Removal Component Change: A component t_i is removed from $t = f(t_1, \dots, t_i, \dots, t_m)$, resulting in a new component term $t' = f(t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_m)$. The term substitution operation (14) is applied to S such that $S' = S[f(t_1, \dots, t_i, \dots, t_m)/f(t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_m)]$ replaces instances of t with t' in S , forming the component term S' . When the system representation is updated, the engine adds the change perception *remove_{t_i}* to agent α_t 's belief base, triggering a belief addition event handled by the plan

$$\begin{aligned} +remove_{t_i} : true \leftarrow & -property_{t_i} \wedge -result_{t_i} \wedge stop_{t_i} \\ & \wedge replan \wedge !verify \wedge !notify. \end{aligned}$$

which removes agent α_t 's beliefs regarding the verification result and property obtained from agent α_{t_i} . The action *stop_{t_i}* terminates and removes agent α_{t_i} from the AVE. The agent plan *replan* requests new plans that enable α_t to reverify their verification task and notify all relevant agents.

Agent Evolution: Agent action *replan* requests advice from the planner after a change. The planner modifies the agent's plan library as follows. Given the component term t' obtained by adding or removing the sub-component term t_i from $t = f(t_1, \dots, t_i, \dots, t_m)$, the agent planner constructs

- a new *!queryagents* plan body (5), adding or removing actions *send*($t_i, givep, A_i$) and *send*($t_i, givev, V_i$) requesting agent α_{t_i} 's property and verification value. This corresponds to constructing the assumption function in (6).
- a context rule (9), adding or removing *property_{t_i}*(A_i) and *result_{t_i}*(V_i) as belief requirements for verification,
- a verification plan body (10), adding or removing test goals *?property_{t_i}*(A_i) and *?test_{t_i}*(V_i). The parameter listing in the oracle invocation is accordingly updated to add or remove parameters P_i and V_i .
- a re-verification plan (13), adding and removing plans corresponding to sub-components dependencies.

When the agent receives the new plans, the agent reinitialises and attempts to achieve their initial verification goal.

Re-verification Policies: Plans (12) and (13) perform re-verification that essentially follow recently introduced incremental verification techniques [13], [14] based on component dependencies defined by the system's architecture. By using a BDI-agent approach, these techniques are extendable to generalised *verification policies* comprising programmable agent behaviours to perform remedial actions taken in response to component change. Agent behaviour can express policies to send conditional notifications only when their component becomes non-compliant to requirements after a change, perform verification steps involving other oracles, or notify other agents specified by the application domain requirements.

IV. A COMPOSITIONAL APPROACH TO SMT

We instantiate AVE with an existing SMT-based technique. The following concepts are defined in [13] and reproduced for completeness.

Component Models: To integrate an SMT verification service with the engine we define $\mathcal{M} = \mathbb{P}(Z)^+$, such that each component term t is modelled as a finite subset Z_t of variables from the set Z inductively on the structure of component terms:

- $Z_{c_i} \in \mathbb{P}(Z)$, for each basic component term c_i such that models are pairwise disjoint $Z_{c_i} \cap Z_{c_j} = \emptyset$, for $i \neq j$.
- $Z_{f(t_1, \dots, t_m)} = \cup_{i=1}^m Z_{t_i}$ for each operation $f \in \Sigma$ and component terms t_1, \dots, t_m .

Component Properties: Let $\mathcal{P} = F(\Gamma, Z)$ such that component requirements are quantifier-free first order formulae over the signature Γ , defined inductively by the rules

$$\begin{aligned} \phi ::= & t_1 = t_2 \mid r_1(t_1, \dots, t_{n_1}) \mid \dots \mid r_m(t_1, \dots, t_{n_m}) \mid \\ & \neg \phi_1 \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \end{aligned}$$

where t_1 and t_2 are terms, and ϕ_1 and ϕ_2 are formulae.

We often display the variables that appear in a formula by writing $\phi(\mathbf{z})$ to mean that the formula ϕ has at least one instance of the variables in the tuple $\mathbf{z} = (z_1, \dots, z_p)$.

SMT Resolution: Satisfiability modulo theories represent a class of theoretical techniques and practical tools that determine the satisfiability of a formula expressed in terms of tests and operations in a signature Γ . An SMT technique $smt : F(\Gamma, Z) \rightarrow [Z \rightarrow \mathcal{V}]$ is applied to the formula $\phi(\mathbf{z})$ in order to compute and return an assignment $\mathbf{v} : Z \rightarrow \mathcal{V}$ of values from the set \mathcal{V} to the variables in \mathbf{z} , such that the formula is satisfied. In symbols, we write $\mathbf{v} \models \phi(\mathbf{z}) \iff \llbracket \phi(\mathbf{z}) \rrbracket(\mathbf{v}) = true$ where $\mathbf{v} = smt(\phi(\mathbf{z}))$. If no such \mathbf{v} exists then the unsatisfiable assignment $\mathbf{u} : Z \rightarrow \mathcal{V}$ is returned. We note the empty formula ϵ is valid (i.e. satisfiable by all assignment mappings).

The requirements function $R : C(\Sigma) \rightarrow \mathcal{P}$ supplied to the agent verification engine associates each component term t with a logical formula ϕ_t such that $var(\phi_t) \subseteq Z_t$ where $var(\phi_t)$ is the set of variables in ϕ_t and Z_t is the component model of t . We use the inductive properties of component terms to compute the conjunction $\phi = \phi_{t_1} \wedge \dots \wedge \phi_{t_m}$ of formulae associated to the components t_1, \dots, t_m of S . By resolving ϕ using smt , we get an assignment $\mathbf{a} : \mathcal{P} \rightarrow \mathcal{V}$ of values from \mathcal{V} to variables in ϕ such that system requirements R are satisfied. We also define the function $mono : C(\Sigma) \rightarrow F(\Gamma, Z)$ such that $mono(t)$ is the monolithic formula of t , by induction over the structure of component terms.

Let t be a component term in $C(\Sigma)$.

Base case: for $t \equiv c_i$ the basic component c_i , we have $mono(c_i) = \phi_{c_i}$.

Inductive step: for $t \equiv f(t_1, \dots, t_m)$ with sub-components t_1, \dots, t_m and operation symbol $f \in \Sigma$, we have $mono(t) = \phi_{f(t_1, \dots, t_m)} \wedge (\bigwedge_{i=1}^m mono(\phi_{t_i}))$.

Compositional SMT Resolution: As the component-based system S typically contains a large number of components and requirements, resolving a monolithic formula $mono(S)$ involves a huge constraint problem and is generally unfeasible for a single agent to complete during runtime. Instead, we describe a compositional approach based on the results of [13], [24] which constructs smaller resolution steps to be carried out independently. The solutions obtained in each step are combined to form a solution for the monolithic formula.

The resolution of the conjunction $\phi_1(\mathbf{y}) \wedge \phi_2(\mathbf{z})$ of independent formulae (no shared variables) may be decomposed into two smaller steps $a_1 = smt(\phi_1(\mathbf{y}))$ and $a_2 = smt(\phi_2(\mathbf{z}))$,

combined using the operation $\oplus : [\mathcal{P} \rightarrow \mathcal{V}]^2 \rightarrow [\mathcal{P} \rightarrow \mathcal{V}]$ such that $(a_1 \oplus a_2)(z) = a_1(z)$ if $z \in \mathbf{z}$ and $a_2(z)$ otherwise. We can prove that $a_1 \oplus a_2 \models \phi_1(\mathbf{y}) \wedge \phi_2(\mathbf{z})$ (cf. Lemma 4.1 [13]). We extend this observation to $n > 1$ independent logical formulae ϕ_1, \dots, ϕ_n and construct an assignment $v \models \bigwedge_{i=1}^n \phi_i$ through a sequence of independent resolution steps $v_1 \models \phi_1, \dots, v_n \models \phi_n$. (cf. Theorem 4.1 [13]).

The SMT Oracle: Compositional SMT resolution serves as the basis for instantiating the assume-guarantee oracle ω specified by (1) to perform compositional SMT resolution. The oracle accepts as input a *verification task* comprising

- mapping $\mathbf{a} : \{\phi_1, \dots, \phi_n\} \rightarrow \{v_1, \dots, v_n\}$ of independent logical formulae in $F(\Gamma, Z)$ to values in \mathcal{V} representing assumptions such $v_i \models \phi_i$ for $\phi_i \in dom(\mathbf{a})$ and $v_i \in range(\mathbf{a})$,
- $g \in F(\Gamma, Z)$, independent of each formula in $dom(\mathbf{a})$,
- model $m = var(g) \cup (\cup_{i \in \{1, \dots, n\}} var(a_i))$, the union of the variable sets for g and a_1, \dots, a_n .

We model the oracle as the total function $\omega : W \rightarrow \mathcal{V}$ as

$$\omega(m, g, \mathbf{a}) = smt(g) \oplus (\oplus_{i=1}^n v_i) \quad (15)$$

that performs smt resolution to compute an assignment of values in \mathcal{V} to the model variables in m that satisfies the conjunction $g \wedge a_1 \wedge \dots \wedge a_n$ of logical formulae. Thus, we conclude $\omega(m, g, \mathbf{a}) \models g \wedge (\bigwedge_{i=1}^n a_i)$.

V. CASE STUDY

We introduce an industrial control system to illustrate our agent-based verification engine. Figure 2 depicts an item sorting system (ISS), typically used to sort high volume item streams within systems such as airport baggage handling and food packaging. Items are placed in an initial position by a robotic arm. Horizontal and vertical pushers use sensors to detect items and then direct them towards user-specified end-points. Additional robotic arms remove items from the end-points. The system software is written using the IEC 61499 [18] language, which enables component-based design where components contain finitely nested networks of (sub) components. Figure 3 shows the architecture of the ISS. The system \mathbf{S} comprises components **vert**, **horiz** and **robot**, representing the control programs for the vertical and horizontal pushers, and the robotic arms, respectively. Each component has three sub-components labelled from **A** to **I**.

Components of IEC 61499 programs execute on *devices* such as programmable logic controllers (PLCs). Components deployed onto the same PLC communicate using faster local mechanisms such as shared buffers. Components on different PLCs communicate using networks like Ethernet. For the ISS, which handles high item volumes, system speed can be significantly reduced if critical path sub-components like communicating pushers execute on different PLCs. Also, the ISS can be made more dependable by deploying multiple instances of sub-components **A** – **I** across available PLCs, numbered 1 to p ($p \geq 2$), according to the following high-level requirements R :

- R1** each component has 1 to 10 instances deployed
- R2** **A** – **C** have at least 4 instances deployed

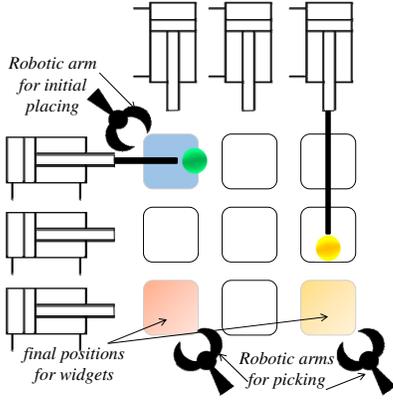


Fig. 2. An item sorting system

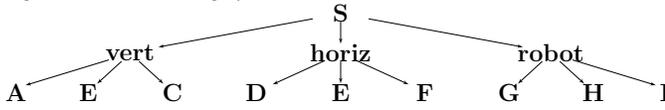


Fig. 3. Architecture of the item sorting system

R3 **G** has exactly 3 instances on PLC 2

R4 **H** must not be deployed on PLC 1

R5 all instances of **D**, **E** and **F** are identically distributed

The *deployment configuration problem* determines a distribution of component instances of the ISS **S** across p PLCs to satisfy the high-level requirements R .

System Specification in AVE: We define a component signature Σ and express an algebraic specification of the system architecture depicted in Figure 3. Let the set of component terms be denoted as C . We set the basic components of the system **A**, ..., **I** and define an operation symbol $c : C \times C \rightarrow C$. For the ISS, we have the component term

$$\mathbf{S} \equiv c(\mathbf{vert}, c(\mathbf{horiz}, \mathbf{robot})) \quad (16)$$

where, $\mathbf{vert} \equiv c(\mathbf{A}, c(\mathbf{B}, \mathbf{C}))$, $\mathbf{horiz} \equiv c(\mathbf{D}, c(\mathbf{E}, \mathbf{F}))$ and $\mathbf{robot} \equiv c(\mathbf{G}, c(\mathbf{H}, \mathbf{I}))$.

Component models: To apply the compositional SMT resolution technique described in Section IV to the deployment configuration problem, we model sub-components **A**, ..., **I** by finite sets $Z_{\mathbf{A}}, \dots, Z_{\mathbf{I}}$ of variables. E.g., **A** is modelled by the set $Z_{\mathbf{A}} = \{\mathbf{Az}_1, \dots, \mathbf{Az}_p\}$ of integer-typed variables from the set Z such that the assignment $\mathbf{Az}_i = n$ models the deployment of n instances of component **A** onto the i^{th} PLC. By definition, we have the model $Z_{\mathbf{S}} = Z_{\mathbf{vert}} \cup Z_{\mathbf{horiz}} \cup Z_{\mathbf{robot}}$ associated with the component term **S**, where $Z_{\mathbf{vert}} = Z_{\mathbf{A}} \cup Z_{\mathbf{B}} \cup Z_{\mathbf{C}}$, $Z_{\mathbf{horiz}} = Z_{\mathbf{D}} \cup Z_{\mathbf{E}} \cup Z_{\mathbf{F}}$ and $Z_{\mathbf{robot}} = Z_{\mathbf{G}} \cup Z_{\mathbf{H}} \cup Z_{\mathbf{I}}$.

Requirements Specification We define a mapping $\phi : C(\Sigma) \rightarrow F(\Gamma, Z)$ by formalising requirements **R1** - **R5** of the ISS architecture as logical formulae over integer-type variables to be associated with components of **S**. Let Γ be a signature comprising sorts for integers and Boolean, symbols for the standard arithmetic operations, equality and inequality predicates. Initially, we set $\phi_t = \epsilon$, that signifies the case where component t is not constrained. Requirement **R1** sets the range of deployed instances for each sub-component of **S** between 1 and 10, with non-negative number of instances over any PLC. For sub-component **A**, we express this requirement as

the formula

$$\phi_{\mathbf{A}}^1 := \bigwedge_{i=1}^p \mathbf{Az}_i \geq 0 \wedge \sum_{i=1}^p \mathbf{Az}_i \geq 1 \wedge \sum_{i=1}^p \mathbf{Az}_i \leq 10. \quad (17)$$

Similarly, all other sub-components **B** - **I** are assigned formulae $\phi_{\mathbf{B}}^1 - \phi_{\mathbf{I}}^1$ respectively. Requirement **R2** further constrains the deployment, and is expressed as $\phi_{\mathbf{A}}^2 := \sum_{i=1}^p \mathbf{Az}_i \geq 4$. We add this formula to the existing formula for **R1** by assigning $\phi_{\mathbf{A}} := \phi_{\mathbf{A}}^1 \wedge \phi_{\mathbf{A}}^2$. We assign similar formulae $\phi_{\mathbf{B}} := \phi_{\mathbf{B}}^1 \wedge \phi_{\mathbf{B}}^2$ and $\phi_{\mathbf{C}} := \phi_{\mathbf{C}}^1 \wedge \phi_{\mathbf{C}}^2$ for sub-components **B** and **C**. **R3** is formalised as $\phi_{\mathbf{G}}^2 := (\mathbf{Gz}_2 = 4)$ and we construct $\phi_{\mathbf{G}} := \phi_{\mathbf{G}}^1 \wedge \phi_{\mathbf{G}}^2$, using the formula $\phi_{\mathbf{G}}^1$ formalising **R1** for **G**. For **R4** we specify $\phi_{\mathbf{H}}^2 := (\mathbf{Hz}_1 = 0)$ and add this formula to the existing requirements by assigning $\phi_{\mathbf{H}} := \phi_{\mathbf{H}}^1 \wedge \phi_{\mathbf{H}}^2$. Lastly, we formalise Requirement **R5** which states that the sub-components of **D**, **E** and **F** must have identical distributions of instances across the PLCs. We define the formula

$$\phi_{\mathbf{vert}} := \bigwedge_{i=1}^p (\mathbf{Dz}_i = \mathbf{Ez}_i) \wedge (\mathbf{Fz}_i = \mathbf{Ez}_i) \quad (18)$$

where the model for **vert** comprises variables from its sub-component models $Z_{\mathbf{D}}$, $Z_{\mathbf{E}}$, and $Z_{\mathbf{F}}$.

VI. IMPLEMENTATION AND SIMULATION

We developed a generic multi-agent simulator of the agent verification engine as an open-source Java application. The engine's implementation is based on Jason AgentSpeak [4]. Jason has a theoretical basis which is amenable to our formal approach for system adaptation and has built-in, extensible support for multi-agent system distribution over networks. We used Jason's standard architecture for agent perceptions, inter-communication and actions. The AVE's implementation comprises the following core classes

- `RunTimeEnvironment` $\langle M, P, V \rangle$ simulates perceptible component changes within the system's runtime environment and executes agent actions (M for models, P for properties and V for verification results)
- `Engine` $\langle M, P, V \rangle$ maintains the system representation and compositional verification technique
- `AgentPlanner` $\langle M, P, V \rangle$ constructs agents and provides new plans after changes occur.
- `CompVerify`, a compositional verification paradigm providing abstract methods `add`, `modify` and `remove` for updating the system's algebraic specification.

Agent Verification Engine for the ISS: We instantiated AVE for the deployment configuration problem for the ISS as follows. First, `Z3CompVerify`, a concrete class utilising the compositional SMT approach described in Section IV was developed. The oracle ω defined by (15) is implemented using the Microsoft satisfiability modulo theories solver Z3 [8]. AVE was also provided with **S**, the architecture (16) of the ISS, as well as a mapping $\phi : C(\Sigma) \rightarrow F(\Gamma, Z)$ that assigns logical formulae to components of the system.

When the engine receives these inputs, the agent planner constructs the necessary agents that co-operate to solve the deployment configuration problem. Figure 4 depicts a tree comprising ten agents in the set \mathcal{A} constructed by the agent planner as described in Section III. Each node is labeled with

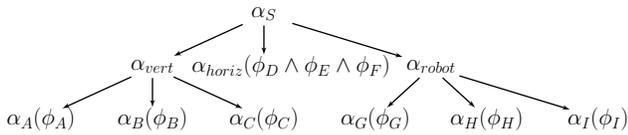


Fig. 4. ISS Agent Tree
TABLE I. AN ASSIGNMENT MAPPING SATISFYING ALL ISS REQUIREMENT IN R FOR 2 PLCs.

v_S	α_A A	α_B B	α_C C	α_{vert} D E F	α_G G	α_H H	α_I I
z_1	1	4	2	1 1 1	0	0	1
z_2	3	0	2	1 1 1	3	1	0

the name of the agent corresponding to a component in the ISS. The logical formulae to be resolved by the agent using the oracle are written in parenthesis. The formula ϕ_{vert} defined in (18) shares variables with its sub-component formulae ϕ_D , ϕ_E and ϕ_F . None of these formulae can be resolved independently thus the engine assigns the task of resolving $\phi_{vert} \wedge \phi_D \wedge \phi_E \wedge \phi_F$ to agent α_{vert} . The edges of the tree denote communication links between agents in which queries, properties and values are transmitted according to plans (5), (7) and (8) respectively.

Once the engine has been initialised, agents in \mathcal{A} attempt to achieve their verification goal. We consider this process for agent α_A as follows. The agent has initial beliefs $model(Z_A)$ and $property(\phi_A)$. It then attempts to achieve goal $!verify$ using plan (10) by invoking $\omega(Z_A, \phi_A, \eta)$ to obtain assignment $v_A = \{AZ_1 \rightarrow 1, AZ_1 \rightarrow 3\}$, storing the belief $result(v_A)$. It also attempts the achievement goal $!detect$ defined in (11) to sense changes in component **A** during runtime.

The configuration problem is resolved when the result belief is stored by α_S . α_S carries out the following steps for this purpose. α_S first queries the agents α_{vert} , α_{horiz} and α_{robot} and forms the assumption function $a: \mathcal{P} \rightarrow \mathcal{V}$ such that $a(\phi_{vert}) = v_{vert}$, $a(\phi_{horiz}) = v_{horiz}$, and $a(\phi_{robot}) = v_{robot}$. α_S then invokes the oracle $v_S = \omega(Z_S, \epsilon, a)$, where Z_S is the component model of **S**, ϵ is the empty formula. Finally, α_S stores $+result(v_S)$ in its belief base.

Table I depicts the value v_S obtained by the SMT process carried out by the oracle. The agent's belief represents one possible assignment mapping to satisfy all requirements in R .

Experimental Results

We evaluated the performance of AVE by running a series of experiments starting with the ISS architecture shown in Figure 3 and subsequently increasing the system size by five components at each step, corresponding to an increase of 5 agents and 50 variables in the Z_3 specification of the system requirements (assuming 10 available PLCs). The experiments hence considered system sizes containing between 10 and 60 agents. A standard PC running Windows 7 Enterprise 64-bits with an Intel i7 2.1GHz Processor and 16 GB RAM was used for the tests. We measured the time for α_S to acquire the belief $result(v_S)$ and compared it to the cumulative time across all oracle invocations, corresponding to the time needed to complete minimal number of steps to resolve the deployment configuration problem using standard compositional SMT resolution. The results of this experiment are presented in Figure 5. All runs completed under 7 seconds. For small system sizes, the difference in speed of the two approaches is negligible. As

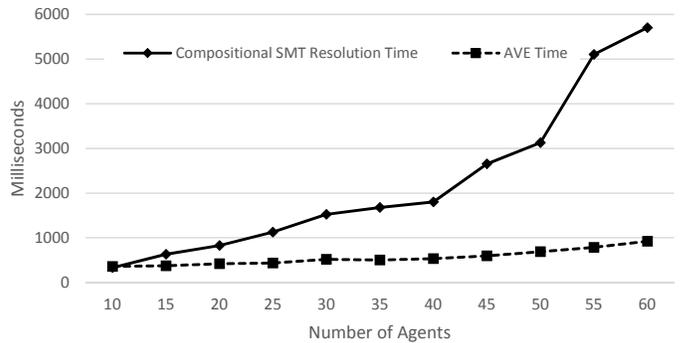


Fig. 5. Configuration using Compositional SMT resolution and AVE

the system grows with additional components, AVE benefits from concurrent oracle invocations by the agents.

We also carried out the following experiments to evaluate the performance of our implementation for *re-configuration* after a system change is perceived by an agent.

1. Adding a Robotic Arm: The ISS requires an additional robotic arm to pick up items from a new end-point and the runtime environment simulates the addition of a new component **R** to the system. The component is modelled by the variable set Z_R and its requirements are formalised by the standard logical formula $\phi_R = \phi_R^1$ as defined in (17). The agent verification engine performs the following steps. First, a substitution on the component term **S** is performed, obtaining $S' \equiv S[\mathbf{robot}/c(c(\mathbf{G}, c(\mathbf{H}, \mathbf{I})), \mathbf{R})]$ and $add(\mathbf{R})$ is added to α_{robot} belief base. Next, agent α_{robot} issues the action to $create(\mathbf{R})$. The agent planner then submits new plans to α_{robot} , who invokes the oracle with assumptions from α_G to α_I and Z_R . Agent α_R obtains a new belief of the result. Finally, α_{robot} notifies α_S , resulting in a new result belief for α_S that includes the new verification result.

Figure 6 shows AVE consistently took about 37 milliseconds to re-configure the system, ranging in size from 10 to 60 agents. In contrast, global reconfiguration of the entire system took increasingly longer as system size increased.

2. Modifying Requirements: We considered updated requirements to change the deployment of the pushers over available PLCs. Such updates may happen often for the ISS, in order to react to changing item volumes and types. As Figure 6 shows, the re-configuration effort grew from about 67 milliseconds to 110 milliseconds for systems sized from 10 to 60 agents. In contrast, a complete re-configuration can take as much as nine times longer to reverify, requiring 1000 milliseconds to completely reverify the largest system containing 60 agents. Experiments to compare re-verification time after an agent removal due to component failure showed that on average it takes 58 ms for agent α_S to obtain an updated result belief, whereas global reconfiguration takes 897 ms on average, for systems containing 10 to 60 agents.

VII. RELATED WORK

There have been many approaches to solve the system reconfiguration problem that we used to demonstrate the agent verification engine. Various formal notations [1], [27], [28], [15], [6] specify and verify component-based systems whose architectures can evolve at run-time with the addition or removal of components. We address a similar problem in this

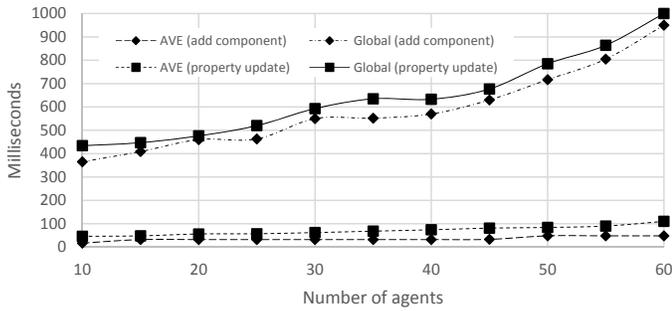


Fig. 6. Running time of AVE versus full compositional verification paper but by contrast, our work delegates verification tasks to agents rather than handling them in a centralised way. In [23], two configuration protocols for deploying a cloud application over multiple virtual machines are proposed and verified using formal methods. In contrast to our work, the protocols proposed in this article do not take into account component failures and subsequent re-configuration actions. In [5], a reconfiguration protocol is proposed that can handle any number of failures during a reconfiguration. This protocol expects reconfiguration decisions to be made before-hand, unlike our approach which uses agents to make reconfiguration decisions based on changes in the system architecture, requirements, or device availability. Another approach that incorporates agent technology is Lira [7], a light-weight, agent-based reconfiguration engine. Each component in a given distributed software system has a unique agent associated with it. An agent handles reconfiguration requests for its associated component. RECoMa [10] is a configuration manager can help find appropriate computer platforms to deploy software agents of a multi-agent system. All available devices on which available components can be deployed are treated as equal, each capable of deploying the components allocated to it. However, our framework can be extended to heterogeneous platforms by the use of requirements that specify constraints on the device's deployment.

VIII. CONCLUDING REMARKS

This paper presents a generic Agent Verification Engine in which agents observe components and determine their compliance to system requirements using a supplied compositional verification technique. Agents communicate verification results such that validation of localised components infer validation of the entire system. A key benefit of the approach is the inherently distributed nature of the agents constructed by the engine, providing a decentralised method to apply compositional verification to large-scale component-based systems.

Future work includes extending the theoretical underpinnings of our approach. We defined the engine as a translation from an algebraic specification of the system architecture to agent behaviours. BDI-agents were chosen to express re-verification behaviours since they have a rich logical framework [22] that is required to prove the correctness of the agents constructed by the engine. Secondly, the engine can be extended to include agent interpretation of verification results, forming instructions to be provided as input for actuators to affect the environment in a way that component compliance is restored after a change. Lastly, there is no end to the kinds of verification techniques to instantiate the engine and solve domain specific problems from areas such as industrial control

systems or cloud computing technologies.

REFERENCES

- [1] R. Allen, R. Douence, and D. Garlan. Specifying and analyzing dynamic software architectures. In *FASE*, pages 21–37, 1998.
- [2] S. Bensalem et al. Compositional verification for component-based systems and application. *IET Software*, 4(3):181–193, 2010.
- [3] S. Berezin, S. V. A. Campos, and E. M. Clarke. Compositional reasoning in model checking. In *COMPOS*, pages 81–102, 1997.
- [4] R. H. Bordini, J. F. Hübner, and M. Wooldridge. *Programming multi-agent systems in AgentSpeak using Jason*. John Wiley & Sons, 2007.
- [5] F. Boyer, O. Gruber, and D. Pous. Robust reconfigurations of component assemblies. In *ICSE*, pages 13–22. IEEE Press, 2013.
- [6] A. Cansado et al. A formal framework for structural reconfiguration of components under behavioural adaptation. *Electr. Notes Theor. Comput. Sci.*, 263:95–110, 2010.
- [7] M. Castaldi et al. A lightweight infrastructure for reconfiguring applications. In *ICSE*, pages 231–244, 2003.
- [8] L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340. Springer, 2008.
- [9] M. Fisher, L. Dennis, and M. Webster. Verifying autonomous systems. *Communications of the ACM*, 56(9):84–93, Sept. 2013.
- [10] J. A. Giampapa, O. H. Juarez-Espinosa, and K. P. Sycara. Configuration management for multi-agent systems. In *AAMAS*, pages 230–231. ACM, 2001.
- [11] O. Grumberg and D. E. Long. Model checking and modular verification. *ACM Trans. Program. Lang. Syst.*, 16(3):843–871, 1994.
- [12] F.-S. Hsieh and J.-B. Lin. A self-adaptation scheme for workflow management in multi-agent systems. *J Intell Manuf*, pages 1–18, 2013.
- [13] K. Johnson and R. Calinescu. Efficient re-resolution of SMT specifications for evolving software architectures. In *QoSA*, pages 93–102. ACM, 2014.
- [14] K. Johnson, R. Calinescu, and S. Kikuchi. An incremental verification framework for component-based software systems. In *CBSE*, pages 33–42. ACM, 2013.
- [15] S. S. Kulkarni and K. N. Biyani. Correctness of component-based adaptation. In *CBSE*, volume 3054, pages 48–58. Springer, 2004.
- [16] Y. Labrou and T. Finin. Semantics and conversations for an agent communication language. In *Readings in Agents*, pages 235–242. Morgan Kaufmann, 1998.
- [17] W. Lopuschitz et al. Toward self-reconfiguration of manufacturing systems using automation agents. *IEEE Trans. Syst., Man, Cybern., Syst.*, 41(1):52–69, 2011.
- [18] R. Lewis. *Modelling distributed control systems using IEC 61499: Applying function blocks to distributed systems*. Number 59. IET, 2001.
- [19] K. Meinke and J. V. Tucker. Universal algebra. In *Handbook of Logic in Computer Science*, pages 189–368. Oxford University Press, 1992.
- [20] M. Metzger and G. Polakow. A survey on applications of agent technology in industrial process control. *IEEE Trans. Ind. Inform.*, 7(4):570–581, 2011.
- [21] A. S. Rao. Agentspeak (I): BDI agents speak out in a logical computable language. In *Agents Breaking Away*, pages 42–55. Springer, 1996.
- [22] A. S. Rao and M. P. Georgeff. Decision procedures for BDI logics. *J Logic Comput*, 8(3):293–343, 1998.
- [23] G. Salaün et al. An experience report on the verification of autonomic protocols in the cloud. *Innov Syst Softw Eng*, 9(2):105–117, 2013.
- [24] R. Sinha, K. Johnson, and R. Calinescu. A scalable approach for re-configuring evolving industrial control systems. In *ETFA*, 2014.
- [25] I. Sommerville et al. Large-scale complex IT systems. *COMMUN ACM*, 55(7):71–77, 2012.
- [26] D. Talia. Clouds meet agents: Toward intelligent cloud services. *IEEE Internet Computing*, 16(2):78–81, 2012.
- [27] M. Wermelinger, A. Lopes, and J. L. Fiadeiro. A graph based architectural (re)configuration language. In *ACM SIGSOFT SEN*, pages 21–32, 2001.
- [28] J. Zhang and B. H. C. Cheng. Model-based development of dynamically adaptive software. In *ICSE*, pages 371–380, 2006.