On the Complexity of Reconfiguration in Systems with Legacy Components

Jacopo Mauro, Gianluigi Zavattaro

► To cite this version:

Jacopo Mauro, Gianluigi Zavattaro. On the Complexity of Reconfiguration in Systems with Legacy Components. MFCS, Aug 2015, Milan, Italy. pp.382–393, 10.1007/978-3-662-48057-1_30 . hal-01233482

HAL Id: hal-01233482

<https://hal.inria.fr/hal-01233482>

Submitted on 25 Nov 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Complexity of Reconfiguration in Systems with Legacy Components

Jacopo Mauro and Gianluigi Zavattaro

Department of Computer Science and Engineering - Univ. of Bologna / INRIA

Abstract. In previous works we have proved that component reconfiguration in the presence of conflicts among components is non-primitive recursive, while it becomes poly-time if there are no conflicts and under the assumption that there are no components in the initial configuration. The case with non-empty initial configurations was left as an open problem, that we close in this paper by showing that, if there are legacy components that cannot be generated from scratch, the problem turns out to be PSpace-complete.

1 Introduction

Modern software systems are obtained as combination of software artefacts having complex interdependencies. Their composition, configuration and management is a difficult task, traditionally performed manually or by writing low level configuration scripts. Recently, many high level languages and tools like, for instance, TOSCA [17] or Engage [8] have been proposed to support the application manager in this difficult task. By adopting these tools, it is possible to describe the software components required to realise the system, define their interdependencies and specify the configuration actions to be executed to actually deploy an instance of the desired system. In some limited cases and under some specific assumptions (no circular component dependencies), such tools automatically synthesise the configuration actions to be executed. Automatic deployment is becoming more and more important for these tools especially due to the advent of virtualization technologies, like in Cloud Computing, that makes it possible to quickly acquire and release computing resources in order to deploy new software systems or reconfigure running applications on-demand.

In previous works [5, 7, 14] we have performed a rigorous and systematic analysis of the automatic deployment problem. We have proved that in general the problem is undecidable, it is non-primitive recursive if component interdependencies do not include numerical constraints, and it is poly-time if also conflicts among components are not considered. This last result was proved by restricting our attention on the deployment of an application from scratch, that is, by assuming that the initial configuration is empty. This result is of particular interest because it underpins the recent industrial trend of using the so called “immutable servers” [16]. The application is divided in stateless components/services that are deployed on virtual machines. When a new version of the component is developed or the virtual machine needs updates (e.g., new security patches have to

be installed), instead of upgrading in-place the virtual machine, a new one is created and the old one destroyed. According to this approach, since all the needed components can be freshly generated, the result proven in [14] shows that a new deployment can be efficiently computed simply generating a new configuration from scratch, without considering or reusing existing components.

Unfortunately the “immutable servers” approach has also some disadvantages. First of all, it requires that every application is carefully designed to ensure that important data is stored and not lost when the old servers are destroyed. System upgrades are usually slower because creating new virtual servers takes more time than performing an upgrade in-place. But, most importantly, this approach cannot be adopted in presence of legacy components, a scenario that often happens in practice due to software applications that for several reasons, like incompatibility with novel computing architectures or cost purposes, cannot be replaced and must be kept in-place.

Given these premises, the following question arises. How complex is the *reconfiguration* problem of deciding if a final configuration can be reached in the presence of components that cannot be switched off and re-deployed from scratch? The goal of this paper is to address this last question, proving that *reconfiguration* is no longer polynomial, but it turns out to be PSpace-complete.

More precisely, we first report the formalisation of the *reconfiguration* problem using the Aeolus component model adopted in [14] (Sect. 2). Then we show that the problem can be solved by performing a symbolic forward search of the new configurations that can be reached from a given initial one (Sect. 3). The symbolic approach allows for a finite representation of all the (possibly infinite) reachable configurations. Unfortunately, the number of possible symbolic configurations is exponential; we mitigate this blow up by adopting a nondeterministic polynomial-space visit of the (symbolic) search space. Finally, we show that it is not possible to significantly improve our algorithm as we prove that the reconfiguration problem is indeed PSpace-hard (Sect. 4). The proof is by reduction from the reachability problem in 1-safe Petri nets [2].

For space reasons, proofs are reported in [15].

2 Formalising the Reconfiguration Problem

In this section we recapitulate the fragment of the Aeolus model used to formally define the reconfiguration problem. This fragment of Aeolus [14] is exactly the one used by the planner Metis [13], a tool for finding deployment plans starting from an empty initial configuration integrated in an industrial deployment platform [6].¹ In the Aeolus model, a component is a grey-box showing relevant internal states and the actions that can be acted on the component to change its state during (re)configuration. Each state activates provide-ports and require-ports representing functionalities that the component provides and needs. Active require-ports must be bound to active provide-ports of other components.

¹ W.r.t. the Aeolus model [5], the fragment used by Metis does not allow the use of capacity constraints, conflicts, and multiple state changes.

The problem that we address in this paper is verifying the existence of a plan (i.e., a correct sequence of configuration actions like component instantiation, binding, or internal state changes) that, given a universe of available components and an initial component configuration, leads to a configuration where a target component is in a given state.

As an example, consider the task of reconfiguring a system setting up a *MySQL master-slave replication* avoiding the downtime of an existing legacy MySQL database. Reconfigurations of this kind are frequent in practice, and are nowadays performed by system administrators who execute reconfiguration receipts that are part of their know-how. According to the Aeolus model, the problem can be formalised as follows. The involved components are two distinct database instances, one in master mode and one in slave mode. We assume to start from a configuration with only one legacy running instance, that will become the master in the new configuration. To activate the slave, a *dump* of the data stored in the master is needed. Moreover, the master has to authorise the slave. This is a circular dependency that is resolved by forcing a precise order in which the reconfiguration actions can be performed: the master first requires authentication of the slave that, subsequently, requires the dump from the master.

Fig. 1: MySQL master-slave instances (in black the initial configuration, in grey the parts added by the reconfiguration and the new state of the master).

In Fig. 1, following the Aeolus model, we depict how to configure MySQL components as master or as slave. We assume the master component to be a legacy one, meaning that it can not be created from scratch but has to be used as deployed in the initial configuration. This is technically obtained setting a dummy state with no outgoing transitions as the initial one. In this way, no newly legacy component could be generated and moved in a state that is different from the dummy one. Apart from the initial dummy state, the master component has 5 more states. The `uninst` state is followed by `inst` and `serving`. In `serving`, the master activates the provide-port `mysql` used by the clients to access the database service. When replication is needed, in order to enter the final `master`

serving state, it first traverses the state `auth` that requires the IP address from the slave, and the state `dump` to provide the dump to the slave. The slave has instead 4 states, an initial `uninst` state and 3 states which complement those of the master during the replication process.

The formal definition of the Aeolus model is based on the notion of *component type*, used to specify the behaviour of a particular kind of component. In the following, $\mathcal{I}$ denotes the set of port names and $\mathcal{Z}$ the set of components.

Definition 1 (Component type). *The set Γ of component types ranged over by $\mathcal{T}, \mathcal{T}_1, \mathcal{T}_2, \dots$ contains 4-tuples $\langle Q, q_0, T, D \rangle$ where:*

- Q is a finite set of states containing the initial state q_0 ;
- $T \subseteq Q \times Q$ is the set of transitions;
- D is a function from Q to a pair $\langle \mathbf{P}, \mathbf{R} \rangle$ of port names (i.e., $\mathbf{P}, \mathbf{R} \subseteq \mathcal{I}$) indicating the provide-ports and require-ports that each state activates. We assume that the initial state q_0 has no requirements (i.e., $D(q_0) = \langle \mathbf{P}, \emptyset \rangle$).

Configurations describe systems composed by components and their bindings. A binding connects a component providing a functionality with a component requiring it. Each component has a unique identifier, taken from the set $\mathcal{Z}$. A configuration, ranged over by $\mathcal{C}_1, \mathcal{C}_2, \dots$, is given by a set of available component types, a set of component instances in some state, and a set of bindings.

Definition 2 (Configuration). *A configuration $\mathcal{C}$ is a quadruple $\langle U, Z, S, B \rangle$ where:*

- $U \subseteq \Gamma$ is the finite universe of the available component types;
- $Z \subseteq \mathcal{Z}$ is the set of the currently deployed components;
- S is the component state description, i.e., a function that associates to components in Z a pair $\langle \mathcal{T}, q \rangle$ where $\mathcal{T} \in U$ is a component type $\langle Q, q_0, T, D \rangle$, and $q \in Q$ is the current component state;
- $B \subseteq \mathcal{I} \times Z \times Z$ is the set of bindings, namely 3-tuples composed by a port, the component that provides that port, and the component that requires it; we assume that the two components are distinct.

Notation. We write $\mathcal{C}[z]$ as a lookup operation that retrieves the pair $\langle \mathcal{T}, q \rangle = S(z)$, where $\mathcal{C} = \langle U, Z, S, B \rangle$. On such a pair we then use the postfix projection operators `.type` and `.state` to retrieve $\mathcal{T}$ and q , respectively. Similarly, given a component type $\langle Q, q_0, T, D \rangle$, we use projections to decompose it: `.states`, `.init`, and `.trans` return the first three elements; `.P(q)` and `.R(q)` return the two elements of the $D(q)$ tuple. Moreover, we use `.prov` (resp. `.req`) to denote the union of all the provide-ports (resp. require-ports) of the states in Q . When there is no ambiguity we take the liberty to apply the component type projections to $\langle \mathcal{T}, q \rangle$ pairs. *Example:* $\mathcal{C}[z].\mathbf{R}(q)$ stands for the require-ports of component z in configuration $\mathcal{C}$ when it is in state q .

As formalised below, a configuration is correct if all the active require-ports are bound to active provide-ports.

Definition 3 (Correctness). Let us consider the configuration $\mathcal{C} = \langle U, Z, S, B \rangle$.

We write $\mathcal{C} \models_{req} (z, r)$ to indicate that the require-port of component z , with port r , is bound to an active port providing r , i.e., there exists a component $z' \in Z \setminus \{z\}$ such that $\langle r, z', z \rangle \in B$, $\mathcal{C}[z'] = \langle \mathcal{T}', q' \rangle$ and r is in $\mathcal{T}'.\mathbf{P}(q')$.

The configuration $\mathcal{C}$ is correct if for every component $z \in Z$ with $S(z) = \langle \mathcal{T}, q \rangle$ we have that $\mathcal{C} \models_{req} (z, r)$ for every $r \in \mathcal{T}.\mathbf{R}(q)$.

In Aeolus configurations evolve by means of (deployment) actions.

Definition 4 (Actions). The set $\mathcal{A}$ contains the following actions:

- *stateChange*(z, q, q') changes the state of the component $z \in \mathcal{Z}$ from q to q' ;
- *bind*(r, z_1, z_2) creates a binding between the provide-port $r \in \mathcal{I}$ of the component z_1 and the require-port r of z_2 ($z_1, z_2 \in \mathcal{Z}$);
- *unbind*(r, z_1, z_2) deletes the binding between the provide-port $r \in \mathcal{I}$ of the component z_1 and the require-port r of z_2 ($z_1, z_2 \in \mathcal{Z}$);
- *new*($z : \mathcal{T}$) creates a new component of type $\mathcal{T}$ in its initial state. The new component is identified by a unique and fresh identifier $z \in \mathcal{Z}$;
- *del*(z) deletes the component $z \in \mathcal{Z}$.

The execution of actions is formalised by means of a labelled transition system on configurations, which uses actions as labels.

Definition 5 (Reconfigurations). Reconfigurations are denoted by transitions $\mathcal{C} \xrightarrow{\alpha} \mathcal{C}'$ meaning that the execution of $\alpha \in \mathcal{A}$ on the configuration $\mathcal{C}$ produces a new configuration $\mathcal{C}'$. The transitions from a configuration $\mathcal{C} = \langle U, Z, S, B \rangle$ are defined as follows:

$$\begin{array}{ll}
\mathcal{C} \xrightarrow{\text{stateChange}(z, q, q')} \langle U, Z, S', B \rangle & \mathcal{C} \xrightarrow{\text{bind}(r, z_1, z_2)} \langle U, Z, S, B \cup \langle r, z_1, z_2 \rangle \rangle \\
\text{if } \mathcal{C}[z].\text{state} = q \text{ and} & \text{if } \langle r, z_1, z_2 \rangle \notin B \\
(q, q') \in \mathcal{C}[z].\text{trans} \text{ and} & \text{and } r \in \mathcal{C}[z_1].\text{prov} \cap \mathcal{C}[z_2].\text{req} \\
S'(z') = \begin{cases} \langle \mathcal{C}[z].\text{type}, q' \rangle & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases} & \mathcal{C} \xrightarrow{\text{unbind}(r, z_1, z_2)} \langle U, Z, S, B \setminus \langle r, z_1, z_2 \rangle \rangle \\
& \text{if } \langle r, z_1, z_2 \rangle \in B \\
\\
\mathcal{C} \xrightarrow{\text{new}(z : \mathcal{T})} \langle U, Z \cup \{z\}, S', B \rangle & \mathcal{C} \xrightarrow{\text{del}(z)} \langle U, Z \setminus \{z\}, S', B' \rangle \\
\text{if } z \notin Z, \mathcal{T} \in U \text{ and} & \text{if } S'(z') = \begin{cases} \perp & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases} \text{ and} \\
S'(z') = \begin{cases} \langle \mathcal{T}, \mathcal{T}.\text{init} \rangle & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases} & B' = \{ \langle r, z_1, z_2 \rangle \in B \mid z \notin \{z_1, z_2\} \}
\end{array}$$

A deployment plan is simply a sequence of actions that transform a correct configuration without violating correctness along the way.

Definition 6 (Deployment plan). A deployment plan P from a correct configuration $\mathcal{C}_0$ is a sequence of actions $\alpha_1, \dots, \alpha_m$ s.t. there exists $\mathcal{C}_1, \dots, \mathcal{C}_m$ correct configurations s.t. $\mathcal{C}_{i-1} \xrightarrow{\alpha_i} \mathcal{C}_i$.

In the following, exploiting the fact that reconfigurations are deterministic, we denote the deployment plan $\alpha_1, \dots, \alpha_m$ from $\mathcal{C}_0$ also with the sequence of reconfigurations steps $\mathcal{C}_0 \xrightarrow{\alpha_1} \mathcal{C}_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_m} \mathcal{C}_m$.

We now have all the ingredients to define the *reconfiguration problem*, that is our main concern: given a universe of component types and an initial configuration, we want to know whether and how it is possible to deploy at least one component of a given component type $\mathcal{T}$ in a given state q .

Definition 7 (Reconfiguration problem). *The reconfiguration problem has as input a universe U of component types, an initial correct configuration C_0 , a component type $\mathcal{T}_t$, and a target state q_t . The output is **yes** if there exists a deployment plan $P = C_0 \xrightarrow{\alpha_1} C_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_m} C_m$ s.t. $C_m[z] = \langle \mathcal{T}_t, q_t \rangle$, for some component z in C_m . Otherwise, it returns **no**, stating that no such plan exists.*

As an example, considering Fig. 1, we can see that there are deployment plans that lead from the initial configuration (in black) to the final MySQL master-slave replication configuration. For instance, such a plan could start with the creation of the slave instance, followed by a state change to the `inst` state and the creation of a binding between the ports `slave_ip` of the two components. At this point, the master component can perform two state changes, reaching the `dump` state. Then, after another binding is established between the `dump` ports, the slave can be moved to its `serving` state by performing two state changes. Finally, the master can enter in the `master serving` state by performing a state change. Note that every action in the deployment plan will correspond to one or more concrete instructions. For instance, the state change from the `serving` to the `auth` state in the master corresponds to issue the command `grant replication slave on *.* to user@'slave_ip'`.

The addition of a dummy initial state to define the master component captures its legacy nature. Indeed, since no other state of the master component is reachable from the initial one, no component created from scratch can provide the same functionalities of the deployed master. For this reason, only the master component present in the initial configuration can be used to reach the target.

Notice that the restriction to consider one target state only in the definition of the reconfiguration problem is not limiting: one can require several target pairs $\langle \mathcal{T}_t, q_t \rangle$ by adding dummy provide-ports enabled only by the components of type $\mathcal{T}_t$ in state q_t and a dummy target component that requires all such provides. For instance, Fig. 2 depicts the dummy target component that in `inst` state requires both an active master and an active slave as needed in the MySQL master-slave reconfiguration discussed above.

Fig. 2: Target

3 Solving the Reconfiguration Problem

In this section we present a nondeterministic polynomial space algorithm that resolves the reconfiguration problem, thus the problem is proved to be in PSpace (as a consequence of the Savitch's theorem [18] stating the equivalence between

NPSpace and PSpace). The idea is to perform a nondeterministic forward exploration of the reachable configurations. This visit could be in principle arbitrarily long because infinitely many different configurations could be potentially reached. The main result that we prove in this section is that it is sufficient to consider a bounded amount of possibly reachable *abstract* configurations. In abstract configurations the bindings are not considered, but only the component type and state of the components are taken into account. Moreover, in abstract configurations, only the components present in the initial configuration are precisely represented, while for all the other components that are dynamically created, it is only considered the presence or absence of instances of components of type $\mathcal{T}$ in state q , thus abstracting away from their precise number.

In order to abstract away from the bindings and consider only the component types and states, we define the following equivalence among configurations.

Definition 8 (Configuration equivalence). *Two configurations $\langle U, Z, S, B \rangle$ and $\langle U, Z', S', B' \rangle$ are equivalent ($\langle U, Z, S, B \rangle \equiv \langle U, Z', S', B' \rangle$) iff there exists a bijective function ρ from Z to Z' s.t. $S(z) = S'(\rho(z))$ for every $z \in Z$.*

The research of the existence of the deployment plan is done on abstract configurations where bindings are not considered. We now show that this is not restrictive because every plan has a corresponding *normalised* plan where unbinding actions are absent and binding actions are generated as soon as possible.

Definition 9 (Normalised deployment plan). *A deployment plan $P = C_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_m} C_m$ is normalised iff:*

- it does not contain unbind actions,
- if C_i for $i \in [1, m-1]$ can be extended with a bind action then $\xrightarrow{\alpha_{i+1}}$ is a bind action,
- C_m cannot be extended with a bind action.

Lemma 1. *Given a deployment plan $P = C_0 \xrightarrow{\alpha_1} C_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_m} C_m$ there exists a normalised deployment plan $P' = C_0 \xrightarrow{\alpha'_1} C'_1 \xrightarrow{\alpha'_2} \dots \xrightarrow{\alpha'_n} C_n$ such that $C_n \equiv C_m$.*

In the remainder of the section, we assume a given universe U of component types; so we can consider that the set of distinct component type and state pairs $\langle \mathcal{T}, q \rangle$ is finite. Let k be its cardinality. Moreover, we assume a given initial configuration C_0 having the initial set of components Z_0 .

We are now ready to define our abstractions $\mathcal{B}$ consisting of pairs of functions $\langle \mathcal{B}_i, \mathcal{B}_c \rangle$. Components are divided into two groups, those that were present in the *initial* configuration and those that were dynamically *created*: the first ones are precisely counted by the function $\mathcal{B}_i$, while for the second ones only the presence of a component type and state pair $\langle \mathcal{T}, q \rangle$ is checked by the function $\mathcal{B}_c$.

Definition 10 (Abstract configuration). *An abstract configuration $\mathcal{B}$ is a pair of functions $\langle \mathcal{B}_i, \mathcal{B}_c \rangle$ that associate to every pair $\langle \mathcal{T}, q \rangle$ respectively a natural number and a boolean value.*

It is immediate to see that (given a universe U of component types and an initial set Z_0 of components) the set of possible abstract configurations is finite: both functions have a domain bound by k , $\mathcal{B}_c$ is a boolean function, and the sum of the values in the codomain of $\mathcal{B}_i$ is bound by $|Z_0|$, i.e., the number of initial components, because such components can only be destroyed and not created.

A concretisation of an abstract configuration $\langle \mathcal{B}_i, \mathcal{B}_c \rangle$ is defined w.r.t. a set of initial components Z . These components occur according to the component type/state pairs counted by $\mathcal{B}_i$, while the other components satisfy the presence/absence indication of the boolean function $\mathcal{B}_c$. In the definition of concretisation we use the following notations: $\mathcal{C}_{\langle \mathcal{T}, q \rangle}^\#(Z)$ is the number of components in Z of type $\mathcal{T}$ in state q in the configuration $\mathcal{C}$, while $Z - Z'$ is the set difference between two sets of components Z and Z' .

Definition 11 (Concretisation). *Given an abstract configuration $\mathcal{B} = \langle \mathcal{B}_i, \mathcal{B}_c \rangle$ and a set of components Z we say that a correct configuration $\mathcal{C} = \langle U, Z', S, B \rangle$ is one concretisation of $\mathcal{B}$ w.r.t. Z if the following hold:*

- $\mathcal{B}_i(\langle \mathcal{T}, q \rangle) = \mathcal{C}_{\langle \mathcal{T}, q \rangle}^\#(Z)$;
- if $\neg \mathcal{B}_c(\langle \mathcal{T}, q \rangle)$ then $\mathcal{C}_{\langle \mathcal{T}, q \rangle}^\#(Z' - Z) = 0$;
- if $\mathcal{B}_c(\langle \mathcal{T}, q \rangle)$ then $\mathcal{C}_{\langle \mathcal{T}, q \rangle}^\#(Z' - Z) > 0$.

We denote with $\gamma(\mathcal{B}, Z)$ the set of concretisations of $\mathcal{B}$ w.r.t. Z . We say that an abstract configuration $\mathcal{B}$ is correct w.r.t. Z if it has at least one concretisation (formally $\gamma(\mathcal{B}, Z) \neq \emptyset$).

In the following, we usually consider concretisations w.r.t. the initial set of components Z_0 , and we simply use $\gamma(\mathcal{B})$ to denote $\gamma(\mathcal{B}, Z_0)$.

We now define the notion of deployment plan on abstract configurations and formalise its correspondence with *concrete* normalised plans.

Definition 12 (Abstract deployment plan). *We write $\mathcal{B} \rightarrow \mathcal{B}'$ with $\mathcal{B} \neq \mathcal{B}'$ if there exists $\mathcal{C} \xrightarrow{\alpha} \mathcal{C}'$ for some $\mathcal{C} \in \gamma(\mathcal{B})$ and $\mathcal{C}' \in \gamma(\mathcal{B}')$.*

A first lemma proves that each normalised deployment plan has a corresponding abstract version.

Lemma 2. *Given a normalised deployment plan $P = \mathcal{C}_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_m} \mathcal{C}_m$ there is an abstract deployment plan $\mathcal{B}_0 \rightarrow \dots \rightarrow \mathcal{B}_n$ s.t. $\mathcal{C}_0 \in \gamma(\mathcal{B}_0)$ and $\mathcal{C}_m \in \gamma(\mathcal{B}_n)$.*

The opposite correspondence (each abstract plan has at least one corresponding normalised *concrete* plan) is more complex to be formalised and proved. The intuition is that, given an abstract configuration $\mathcal{B} = \langle \mathcal{B}_i, \mathcal{B}_c \rangle$ that can be reached by an abstract plan, there exist normalised deployment plans able to reconfigure exactly the initial components as indicated by $\mathcal{B}_i$, and deploy an arbitrary number of instances of other components in the type and state indicated by the boolean function $\mathcal{B}_c$.

Lemma 3. *Given a correct configuration $\mathcal{C}_0$ that cannot be extended with bind actions and an abstract deployment plan $\mathcal{B}_0 \rightarrow \dots \rightarrow \mathcal{B}_n = \langle \mathcal{B}_i, \mathcal{B}_c \rangle$ such that $\mathcal{C}_0 \in \gamma(\mathcal{B}_0)$ then there is a normalised deployment plan $\mathcal{C}_0 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{m-1}} \mathcal{C}_m$ s.t.:*

- $\mathcal{C}_m \in \gamma(\mathcal{B}_n)$;
- for all natural numbers $j_{\langle \mathcal{T}, q \rangle} > 0$, for every component type $\mathcal{T}$ and state q such that $\mathcal{B}_c(\langle \mathcal{T}, q \rangle)$, then $\mathcal{C}_m^{\#}_{\langle \mathcal{T}, q \rangle}(Z_m - Z_0) = j_{\langle \mathcal{T}, q \rangle}$ where Z_m and Z_0 are the components of $\mathcal{C}_m$ and $\mathcal{C}_0$ respectively.

Algorithm 1 Nondeterministic check for $\mathcal{C}_0 = \langle U, Z_0, S, B \rangle$ and target $\mathcal{T}_t, q_t$

```

for all  $\langle \mathcal{T}, q \rangle$  pairs in the universe  $U$  do
     $\mathcal{B}_i(\langle \mathcal{T}, q \rangle) = \mathcal{C}^{\#}_{\langle \mathcal{T}, q \rangle}(Z_0)$ 
     $\mathcal{B}_c(\langle \mathcal{T}, q \rangle) = \text{False}$ 
    counter = 0
    while counter  $\leq |Z_0|^k * 2^k$  do                                 $\triangleright k$  is the number of  $\langle \mathcal{T}, q \rangle$  pairs in  $U$ 
        guess  $\mathcal{B}'_i, \mathcal{B}'_c$ 
        if  $\langle \mathcal{B}_i, \mathcal{B}_c \rangle \not\rightarrow \langle \mathcal{B}'_i, \mathcal{B}'_c \rangle$  then return Failure
        if  $\mathcal{B}'_i(\mathcal{T}_t, q_t) > 0$  or  $\mathcal{B}'_c(\mathcal{T}_t, q_t)$  then return Success
        counter = counter + 1;  $\mathcal{B}_i = \mathcal{B}'_i$ ;  $\mathcal{B}_c = \mathcal{B}'_c$ 
    return Failure

```

In order to check if a solution to the reconfiguration problem exists, it is possible to consider all the possible abstract plans. This can be done using the nondeterministic Algorithm 1. Starting from the abstract representation $\langle \mathcal{B}_i, \mathcal{B}_c \rangle$ of the initial configuration $\mathcal{C}_0$, it performs a nondeterministic exploration of the reachable abstract configurations until either a configuration containing the target $\langle \mathcal{T}_t, q_t \rangle$ is reached or at least $K = |Z_0|^k * 2^k$ abstract steps have been considered, where $|Z_0|$ is the quantity of components of the initial configuration and k is the number of different $\langle \mathcal{T}, q \rangle$ pairs in the universe U . K is an upper bound to the number of different abstract configurations: $|Z_0|^k$ is an upper bound to the different combinations of states for the initially available components, while 2^k is the number of possible sets of $\langle \mathcal{T}, q \rangle$ pairs.

Assuming n the size of the input we have that $|Z_0| \leq n, k \leq n$ and therefore all the variables of the nondeterministic Algorithm 1 can be encoded in $O(n \log(n))$ space. For this reason (and for Savitch's theorem [18]) we can conclude that the reconfiguration problem is in PSpace.

Theorem 1. *The reconfiguration problem is PSpace.*

4 The Reconfiguration Problem is PSpace-hard

PSpace-hardness of the reconfiguration problem is proved by reduction from the reachability problem in 1-safe Petri nets, which is indeed known to be a PSpace-hard problem [2]. We start with some background on Petri nets.

Fig. 3: 1-safe Petri net encoding

A *Petri net* is a tuple $N = (P, T, \mathbf{m}_0)$, where P and T are finite sets of *places* and *transitions*, respectively. A finite multiset over the set P of places is called a *marking*, and $\mathbf{m}_0$ is the initial marking. Given a marking $\mathbf{m}$ and a place p , we say that the place p contains a number of *tokens* equal to the number of instances of p in $\mathbf{m}$. A transition $t \in T$ is a pair of markings denoted with $\bullet t$ and $t^\bullet$. A transition t can fire in the marking $\mathbf{m}$ if $\bullet t \subseteq \mathbf{m}$ (where $\subseteq$ is multiset inclusion); upon transition firing the new marking of the net becomes $\mathbf{n} = (\mathbf{m} \setminus \bullet t) \uplus t^\bullet$ (where $\setminus$ and $\uplus$ are the difference and union operators for multisets, respectively). This is written as $\mathbf{m} \mapsto \mathbf{n}$. We use $\mapsto^*$ to denote the reflexive and transitive closure of $\mapsto$. We say that $\mathbf{m}'$ is *reachable from* $\mathbf{m}$ if $\mathbf{m} \mapsto^* \mathbf{m}'$. A Petri net P is 1-safe if in every reachable marking every place has at most one token. Reachability of a specific marking $\mathbf{m}_t$ from the initial marking $\mathbf{m}_0$ is PSpace-complete for 1-safe nets [2].

We now consider a given 1-safe Petri net $N = \langle P, T, \mathbf{m}_0 \rangle$ and discuss how to encode it in Aeolus component types. We will use two types of legacy components: one modelling the places and one for the transitions. The simplest component type, denoted with $\mathcal{T}_p$ and depicted in Fig. 3a, is the one used to model a place $p \in P$. Namely, a place p is encoded as one instance of $\mathcal{T}_p$. A token is present in p if the component of type $\mathcal{T}_p$ is in the *on* state. There could be just one of these components deployed simultaneously. This can be obtained simply adding this component to the initial configuration in the *on* or *off* state, according to the initial marking, and make these two states non reachable from the initial state q_0 . The token could be created starting from the *off* state following a protocol consisting of providing the port a_p and then requiring the port on_p . Symmetrically, a token can be removed by providing the port b_p and then requiring the port off_p . The component provides the port on_p when it is in the *on* state, the port off_p when it is in the *off* state.

The transitions in T can be represented with a single component of type $\mathcal{T}_T$ depicted in Fig. 3b. The uniqueness of this component is guaranteed, as done for

$\mathcal{T}_p$, by adding it to the initial configuration and forbidding outgoing transitions from the initial state q_0 . This component is assumed to be present in the initial configuration in state q . From this state it can nondeterministically select one transition t to fire, by entering a corresponding q_t state. The subsequent state changes can be divided into two phases: consumption and production. These phases respectively model the consumption of tokens from the places in the preset of t and the production of tokens in the places in the postset of t . The consumption and production of tokens have been already discussed above: consumption (see Fig. 3c) is obtained by providing and requiring the ports b_p and off_p , production (see Fig. 3d) by providing and requiring the ports a_p and on_p .

We now consider a marking $\mathbf{m}_t$ of the 1-safe Petri net N . We can check whether $\mathbf{m}_t$ is reachable in N by considering the following Aeolus reconfiguration problem. The initial configuration consists of an instance of $\mathcal{T}_T$, in state q , plus a component of type $\mathcal{T}_p$ for every place p , in *on* or *off* depending on the initial marking $\mathbf{m}_0$. The target to be considered consists of a configuration in which a port on_p is active for all places $p \in \mathbf{m}_t$, a port off_p is active for all places $p \notin \mathbf{m}_t$ and the port t is active indicating that no transition is currently in execution. Checking these requirements can be easily done, as explained in Section 2, by adding a dummy component having a target state requiring all the ports as explained above. Hence, we have the following.

Theorem 2. *The reconfiguration problem is PSpace-hard.*

5 Related Work and Conclusions

To the best of our knowledge, there is no work that formally studies the complexity of automatic reconfiguration of component systems. A significant part of the related literature focuses on the problem of dynamic re-allocation of resources, e.g., [3, 10]. Other works focus on the nature of the reconfiguration problem, like in [19] where a classification of the reconfiguration problems is made based on its causes, namely failures, system updates, and user requests. This work, however, does not consider the complexity of establishing the reconfiguration steps.

Different tools to compute the (optimal) final configuration exist, e.g., [4, 12]. However, all these approaches just focus on the target configuration to reach without computing the deployment steps. AI Planning Technologies [9] have been used to generate automatically the actions to reconfigure a system [1, 11]. However, these techniques have scalability issues. Conversely, tools like Metis [13, 14] or Engage [8] are able to compute the deployment steps needed to reach a target configuration but in simplified contexts: Metis imposes empty initial configurations while Engage forbids circular dependencies.

In this work we proved that extending these tools to deal also with reconfigurations may be too computationally expensive. Indeed, PSpace-completeness means that there are at least some cases where solving a reconfiguration problem requires a huge computational effort. For instance, our hardness proof shows that this can happen in the presence of legacy components that can not be recreated from scratch and may be required to perform cycles of deployment actions.

As a future work we plan to investigate limitations to be imposed to the Aeolus model (e.g., limiting the shape of the automata describing the components lifecycle) in order to have more efficient solutions for the reconfiguration problem. Another approach could be to relax completeness, by designing algorithms that could give negative answers even if a solution exists.

References

1. M. Chen, P. Poizat, and Y. Yan. Adaptive Composition and QoS Optimization of Conversational Services Through Graph Planning Encoding. In *Web Services Foundations*. 2014.
2. A. Cheng, J. Esparza, and J. Palsberg. Complexity results for 1-safe nets. *Theor. Comput. Sci.*, 1995.
3. H. W. Choi, H. Kwak, A. Sohn, and K. Chung. Autonomous learning for efficient resource utilization of dynamic VM migration. In *ICS*, 2008.
4. R. D. Cosmo, M. Lienhardt, R. Treinen, S. Zacchiroli, J. Zwolakowski, A. Eiche, and A. Agahi. Automated synthesis and deployment of cloud applications. In *ASE*, 2014.
5. R. D. Cosmo, J. Mauro, S. Zacchiroli, and G. Zavattaro. Aeolus: A component model for the cloud. *Inf. Comput.*, 2014.
6. R. Di Cosmo, A. Eiche, J. Mauro, G. Zavattaro, S. Zacchiroli, and J. Zwolakowski. Automatic Deployment of Software Components in the Cloud with the Aeolus Blender. Technical report, Inria Sophia Antipolis, 2015.
7. R. Di Cosmo, J. Mauro, S. Zacchiroli, and G. Zavattaro. Component Reconfiguration in the Presence of Conflicts. In *ICALP*, 2013.
8. J. Fischer, R. Majumdar, and S. Esmailsabzali. Engage: a deployment management system. In *PLDI*, 2012.
9. M. Ghallab, D. S. Nau, and P. Traverso. *Automated planning - Theory and Practice*. Elsevier, 2004.
10. D. Gmach, J. Rolia, L. Cherkasova, G. Belrose, T. Turicchi, and A. Kemper. An integrated approach to resource pool management: Policies, efficiency and quality metrics. In *DSN*, 2008.
11. H. Herry, P. Anderson, and G. Wickler. Automated Planning for Configuration Changes. In *LISA*, 2011.
12. J. A. Hewson, P. Anderson, and A. D. Gordon. A Declarative Approach to Automated Configuration. In *LISA*, 2012.
13. T. A. Lascu, J. Mauro, and G. Zavattaro. A Planning Tool Supporting the Deployment of Cloud Applications. In *ICTAI*, 2013.
14. T. A. Lascu, J. Mauro, and G. Zavattaro. Automatic Component Deployment in the Presence of Circular Dependencies. In *FACS*, 2013.
15. J. Mauro and G. Zavattaro. On the Complexity of Reconfiguration in Systems with Legacy Components. Technical report, INRIA Sophia Antipolis, 2015.
16. K. Morris. Immutableserv. <http://martinfowler.com/bliki/ImmutableServer.html>, 2013.
17. OASIS. Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0.
18. W. J. Savitch. Relationships Between Nondeterministic and Deterministic Tape Complexities. *J. Comput. Syst. Sci.*, 1970.
19. S. Wang, F. Du, X. Li, Y. Li, and X. Han. Research on dynamic reconfiguration technology of cloud computing virtual services. In *CCIS*, 2011.