Automatic Application Deployment in the Cloud: from Practice to Theory and Back

Gianluigi Zavattaro

University of Bologna - Italy FoCUS research team INRIA - France

Based on joint work with: T.A. Lascu, J. Mauro R. Di Cosmo, S. Zacchiroli, J. Zwolakowski A. Eiche

Univ. of Bologna - Italy

Univ. Paris Diderot - France

Mandriva SA - France

Cloud applications

 Cloud computing offers the possibility to build **sophisticated** software systems on virtualized infrastructures at a fraction of the time / cost necessary just few years ago We can give a look at Juju to have an idea of how cloud applications can be easily deployed nowadays



Cloud application management is any way a complex task

Even if tool-supported, the cloud application operator must decide:
 which software components to select
 the overall application architecture
 the order of the configuration actions

The challenge

Understand how much of the operator's activities can be automatised:
 selection of the software components (selected from appropriate repositories)
 synthesis of the overall architecture
 planning of the configuration actions to be executed to realize the expected architecture

Definition of a language for describing component's repositories

{	{	{
"states": ["provide": {},	"provide": {
{	"require": {	"@Wordpress/ActiveWithNfs/get_website": 1000
"provide": {},	"@Haproxy/Active/add_database": 1	},
"require": {},	},	"require": {
"initial": true,	"successors": ["@Haproxy/Active/add_database": 1,
"name": "Installed",	"Active"	"@Httpd/Active/start": 1,
"successors": [1,	"@Httpd/Configured/get_document_root": 1,
"Template"	"name": "Configured"	"@Nfs_client/Active/mount": 1
ĺ	},	},
},		"name": "ActiveWithNfs"
{	<pre>"provide": {},</pre>	}
"provide": {},	"require": {],
"require": {},	"@Haproxy/Active/add_database": 1,	"name": "Wordpress"
"successors": ["@Httpd/Active/start": 1,	····}···
"Configured"	"@Httpd/Configured/get_document_root": 1	·
],	}	
"name": "Template"	"successors": [
} /	"ActiveWithNfs"	
],	
	"name": "Active"	

Automatic Application Deployment in the Cloud

Definition of a language for describing component's repositories



Automatic Application Deployment in the Cloud

Realization of a tool for component's selection and architecture synthesis



Realization of a tool for component's selection and architecture synthesis



Automatic Application Deployment in the Cloud

Realization of a tool for planning the configuration actions to be executed



Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

Component types

A component has provide and require ports
A component has an internal state machine
Ports are active or inactive according to the current internal state

Provide ports



Require ports

Automatic Application Deployment in the Cloud

Example: the Wordpress component type



Automatic Application Deployment in the Cloud

Conflicts

Conflicts are expressed as special ports The apache web server is in conflict with the lighttpd web server



Automatic Application Deployment in the Cloud

Capacity constraints

 Provide (resp. require) ports could have an associated upper (resp. lower)
 bound to the number of connections



Automatic Application Deployment in the Cloud

Configurations

- Component instances, with a current state, and complementary provide/require ports connected by bindings
 - Example: Kerberos with Idap support in Debian (example of circular dependency)



Automatic Application Deployment in the Cloud

Configurations

 Component instances, with a current state, and complementary provide/require ports connected by bindings



Formalizing the "deployment" problem

▶ Definition 1 (Component type). The set Γ of *component types* of the Aeolus model, ranged over by $\mathcal{T}_1, \mathcal{T}_2, \ldots$ contains 5-ple $\langle Q, q_0, T, P, D \rangle$ where:

- \blacksquare Q is a finite set of states;
- $q_0 \in Q$ is the initial state and $T \subseteq Q \times Q$ is the set of *transitions*;
- $P = \langle \mathbf{P}, \mathbf{R} \rangle$, with $\mathbf{P}, \mathbf{R} \subseteq \mathcal{I}$, is a pair composed of the set of *provide* and the set of *require*-ports, respectively;
 - D is a function from Q to 2-ple in $(\mathbf{P} \to \mathbb{N}^+_{\infty}) \times (\mathbf{R} \to \mathbb{N}).$
- ▶ Definition 2 (Configuration). A configuration C is a quadruple (U, Z, S, B) where:
 - $U \subseteq \Gamma$ is the finite *universe* of all 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 $\langle Q, q_0, T, P, D \rangle$, and $q \in Q$ is the current component state;
- $B \subseteq \mathcal{I} \times Z \times Z$ is the set of *bindings*, namely 3-ples composed by an interface, the component that requires that interface, and the component that provides it; we assume that the two components are distinct.

Formalizing the "deployment" problem

- **Definition 5** (Actions). The set \mathcal{A} contains the following actions:
- stateChange(z, q₁, q₂) where z ∈ Z: change the state of the component z from q₁ to q₂;
 bind(r, z₁, z₂) where z₁, z₂ ∈ Z and r ∈ I: add a binding between z₁ and z₂ on port r;
 unbind(r, z₁, z₂) where z₁, z₂ ∈ Z and r ∈ I: remove the specified binding;
 new(z : T) where z ∈ Z and T is a component type: add a new component z of type T;
 del(z) where z ∈ Z: remove the component z from the configuration.

Formalizing the "deployment" problem

▶ **Definition 6** (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:

$$\mathcal{C} \xrightarrow{\text{stateChange}(z,q_1,q_2)} \langle U, Z, S', B \rangle$$

if $\mathcal{C}[z].\texttt{state} = q_1$
and $(q_1,q_2) \in \mathcal{C}[z].\texttt{trans}$
and $S'(z') = \begin{cases} \langle \mathcal{C}[z].\texttt{type}, q_2 \rangle & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases}$

$$\begin{array}{c} \xrightarrow{new(z:\mathcal{T})} \langle U, Z \cup \{z\}, S', B \rangle \\ \text{if } z \notin Z, \ \mathcal{T} \in U \\ \text{and } S'(z') = \begin{cases} \langle \mathcal{T}, \mathcal{T}.\texttt{init} \rangle & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases}$$

$$\mathcal{C} \xrightarrow{bind(r,z_1,z_2)} \langle U, Z, S, B \cup \langle r, z_1, z_2 \rangle \rangle$$

if $\langle r, z_1, z_2 \rangle \notin B$
and $r \in \mathcal{C}[z_1]$.req $\cap \mathcal{C}[z_2]$.prov

$$\mathcal{C} \xrightarrow{del(z)} \langle U, Z \setminus \{z\}, S', B' \rangle$$

if $S'(z') = \begin{cases} \bot & \text{if } z' = z \\ \mathcal{C}[z'] & \text{otherwise} \end{cases}$
and $B' = \{ \langle r, z_1, z_2 \rangle \in B \mid z \notin \{z_1, z_2\} \}$

$$\mathcal{C} \xrightarrow{unbind(r,z_1,z_2)} \langle U, Z, S, B \setminus \langle r, z_1, z_2 \rangle \rangle \quad \text{if } \langle r, z_1, z_2 \rangle \in B$$

Automatic Application Deployment in the Cloud

"Deployment" problem

Input:

- A set of component types (called Universe)
- One target component type-state pair
- Output:
 - Yes, if there exists a deployment plan
 - No, otherwise

Deployment plan:

a sequence of actions leading to a final configuration containing at least one component of the given target type, in the given target state

mysql

>3

>2

target

target

Deployment problem: example

 Consider the problem of installing kerberos with Idap support in Debian
 Universe: packages krb5 and openIdap
 Target: krb5 in normal state



Automatic Application Deployment in the Cloud

Deployment problem: example

Deployment plan:

new(k:krb5),new(o:openldap),
stateChange(k,uninst,stage1),
bind(libkrb5-dev,o,k),stateChange(o,uninst,normal),
bind(libldap2-dev,k,o),
stateChange(k,stage1,normal)



Automatic Application Deployment in the Cloud

Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

Summary of decidability/complexity results

Component model Deployment is

Full component model Undecidable [SEFM12]

No capacity constraints Ackermann-hard [ICALP13,I&C14]

No capacity constraints, Quadratic No conflicts [SEFM12]

Automatic Application Deployment in the Cloud

Quadratic algorithm without constraints and conflicts [SEFM12]

Forward reachability algorithm

all reachable states computed by saturation

Algorithm 1 Checking achievability in the Aeolus⁻ model function ACHIEVABILITY (U, \mathcal{T}, q) $absConf := \{ \langle \mathcal{T}', \mathcal{T}'.\texttt{init} \rangle \mid \mathcal{T}' \in U \}$ $provPort := \bigcup_{\langle \mathcal{T}', q' \rangle \in absConf} \{ dom(\mathcal{T}'.\mathbf{P}(q')) \}$ repeat $new := \{ \langle \mathcal{T}', q' \rangle \ | \ \langle \mathcal{T}', q'' \rangle \in absConf, (q'', q') \in \mathcal{T}'.\texttt{trans} \} \setminus absConf$ $newPort := \bigcup_{\langle \mathcal{T}', q' \rangle \in new} \{ dom(\mathcal{T}'.\mathbf{P}(q')) \}$ while $\exists \langle \mathcal{T}', q' \rangle \in new$. $dom(\mathcal{T}'.\mathbf{R}(q')) \not\subseteq provPort \cup newPort$ do $new := new \setminus \{ \langle \mathcal{T}', q' \rangle \}$ $newPort := \bigcup_{\langle \mathcal{T}', q' \rangle \in new} \{ dom(\mathcal{T}'.\mathbf{P}(q')) \}$ end while $absConf := absConf \cup new$ $provPort := provPort \cup newPort$ until $new = \emptyset$ if $\langle \mathcal{T}, q \rangle \in absConf$ then return true else return false end if end function

Automatic Application Deployment in the Cloud

Example: the kerberos case-study



Automatic Application Deployment in the Cloud

Lesson learned from the foundational study

Deployment can be reasonably **fully automatised** if we do not consider capacity constraints and conflicts

Automatic Application Deployment in the Cloud

Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

Use the graph of the reachability algorithm
 bottom-up from the target state
 select the bindings (red arrows)





Automatic Application Deployment in the Cloud



 Generate an
 abstract plan
 (one component for each maximal path)

Arrows represent a **precedence** relation: **blue**: start requirement **red**: end requirement

Automatic Application Deployment in the Cloud

Plan as a topological visit until target:

new(k:krb5),new(o:openldap),
stateChange(k,uninst,stage1),
bind(libkrb5-dev,o,k),stateChange(o,uninst,normal),
bind(libldap2-dev,k,o),
stateChange(k,stage1,normal)



Arrows represent a precedence relation: blue: start requirement red: end requirement

Automatic Application Deployment in the Cloud

 Problem: cycles could forbid the topological visit
 Example: krb5 in normal requires an openIdap in uninst state



Automatic Application Deployment in the Cloud





Automatic Application Deployment in the Cloud

Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

Capacity constraints and conflicts strike back [ASE14,CONCUR15]

 We have investigated the problem of synthesising the **final** configuration considering capacity constraints and conflicts but... ...abstracting away from the internal configuration automata Idea for computing the final configuration: first perform component selection... ...and then establish the **bindings**

Automatic Application Deployment in the Cloud

Component selection [CONCUR15] Component selection is NP-complete but we can use Contraint Solving technology $\bigwedge \quad \bigwedge \ \mathcal{T}.\mathbf{R}(q)(p) \times \operatorname{comp}(\langle \mathcal{T},q\rangle) \leq \ \sum \ \operatorname{bind}(p,\langle \mathcal{T}',q'\rangle,\langle \mathcal{T},q\rangle)$ $p \in \mathcal{I} \quad \langle \mathcal{T}, q \rangle$ $\langle \mathcal{T}', q' \rangle$ $\bigwedge \qquad \qquad \mathcal{T}.\mathbf{P}(q)(p)\times \operatorname{comp}(\langle \mathcal{T},q\rangle) \geq \sum \ \operatorname{bind}(p,\langle \mathcal{T},q\rangle,\langle \mathcal{T}',q'\rangle)$ $p \in \mathcal{I}$ $\langle \mathcal{T}, q \rangle$. $\mathcal{T}.\mathbf{P}(q)(p) < \infty$ $\bigwedge \qquad \operatorname{comp}(\langle \mathcal{T},q\rangle)=0 \ \Rightarrow \ \sum \ \operatorname{bind}(p,\langle \mathcal{T},q\rangle,\langle \mathcal{T}',q'\rangle)=0$ $\langle \mathcal{T}, q \rangle$. $\mathcal{T}.\mathbf{P}(q)(p) = \infty$ $p \in \mathcal{I}$ $\langle \mathcal{T}', q' \rangle$ $\operatorname{comp}(\langle \mathcal{T}, q \rangle) \leq 1$ $p \in \mathcal{I} \quad \langle \mathcal{T}, q \rangle \ . \ \mathcal{T}. \mathbf{R}(q)(p) = 0 \land$ $\mathcal{T}.\mathbf{P}(q)(p) > 0$ $\bigwedge \quad \operatorname{comp}(\langle \mathcal{T}, q \rangle) > 0 \quad \Rightarrow \quad \operatorname{comp}(\langle \mathcal{T}', q' \rangle) = 0$ $\begin{array}{ccc} \langle \mathcal{T}, q \rangle. & \langle \mathcal{T}', q' \rangle \neq \langle \mathcal{T}, q \rangle \\ \mathcal{T}. \mathbf{R}(q)(p) = 0 & \mathcal{T}'. \mathbf{P}(q')(p) > 0 \end{array}$ $p \in \mathcal{I}$ $\mathtt{bind}(p, \langle \mathcal{T}, q \rangle, \langle \mathcal{T}', q' \rangle) \leq \mathtt{comp}(\langle \mathcal{T}, q \rangle) \times \mathtt{comp}(\langle \mathcal{T}', q' \rangle)$ $\langle \mathcal{T},q \rangle = \langle \mathcal{T}',q' \rangle \neq \langle \mathcal{T},q \rangle$ $p \in \mathcal{I}$ $\bigwedge \ \operatorname{bind}(p, \langle \mathcal{T}, q \rangle, \langle \mathcal{T}, q \rangle) \leq \operatorname{comp}(\langle \mathcal{T}, q \rangle) \times (\operatorname{comp}(\langle \mathcal{T}, q \rangle) - 1)$ $p \in \mathcal{I} \quad \langle \mathcal{T}, q \rangle$ Automatic Application Deployment in the Cloud CONCUR'15 - 1.9.2015

Bindings establishment [CONCUR15]

Bindings decided as solution of a max-flow problem



Automatic Application Deployment in the Cloud

Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

Putting everything together: Aeolus Blender [ICSOC15] Armonic components User desiderata Armonic: library Builder of components Zephyrus: Universe + Config file Universe specification synthesis of the Optimisation: Merge Config data final architecture Zephyrus Metis: plan the Unmerge configuration Metis Filler Final conf actions Plan Config file Launcher

Automatic Application Deployment in the Cloud

Reconfiguration vs. Deployment

- Reconfiguration problem:
 - same as deployment, but with **non empty** initial configuration
 - reconfiguration is already **PSpace**-complete without capacity constraints and conflicts
 - [MFCS15]

Open issue:

 Find further restrictions to the model that make reconfiguration tractable (seems very useful in practice)

Automatic Application Deployment in the Cloud

Other open issues

 In real systems there is a flow of configuration data among components: Room for name-passing models? Hierarchical modeling (administrative domains, cloud providers, geographical areas, ...): Room for ambient-like models? • QoS and resource consumption: • Room for **quantitative** models?

Automatic Application Deployment in the Cloud

Structure of the talk

 Formalizing the deployment problem Decidability/complexity results Fully automatic deployment (without capacity constraints and conflicts) Constraints and conflicts strike back Conclusion and Open issues Related work

TOSCA

[OASIS standard 2013]

Language for topologies and deployment plans



Automatic Application Deployment in the Cloud

CloudMF

[J.Ferry et al. - NordiCloud13]

CONCUR'15 - 1.9.2015

Similar language for component description



Automatic Application Deployment in the Cloud

- ConfSolve [J.A.Hewson, P.Anderson, A.D.Gordon LISA12]
 Object-oriented language for services and machines
 - Type system for checking configuration correctness
 class DatabaseServer
 - Constraint solver for automatic placement of services on machines

class DatabaseServer extends Role {
 var role as DatabaseRole;

```
// slave or master
var peer as ref DatabaseServer;
```

// the peer cannot be itself
peer != this;

// a master's peer must be a slave, // and a slave's peer must be a master role != peer.role;

class Machine {
 var os as OperatingSystem;
 var cpus as 1..4;
 var memory as int;

- Engage [J.Fischer, R.Majumdar, S.Esmaeilsabzali PLDI12]
 - Architectural specification in terms of inside / peer / environment relationships
 - Automata with resource lifecycle and transient dependencies
 - Assumption on acyclic relationships (to always guarantee topological visit)



Mentioned publications

- [SEFM12] R. Di Cosmo, S. Zacchiroli, G. Zavattaro. *Towards a Formal Component Model for the Cloud.* Proc. of SEFM'12: 156-171. LNCS 7504, Springer.
- [ICALP13] R. Di Cosmo, J. Mauro, S. Zacchiroli, G. Zavattaro. *Component Reconfiguration in the Presence of Conflicts.* Proc. of ICALP'13: 187-198. LNCS 7966, Springer.
- [ICTAI13] T. A. Lascu, J. Mauro, G. Zavattaro.
 A Planning Tool Supporting the Deployment of Cloud Applications.
 Proc. of ICTAI'13: 213-220. IEEE Press.
- [I&C14] R. Di Cosmo, J. Mauro, S. Zacchiroli, G. Zavattaro.
 Aeolus: A component model for the cloud.
 Information and Computation, 239: 100-121 (2014).
- [MFCS15] J. Mauro, G. Zavattaro.
 On the Complexity of Reconfiguration in Systems with Legacy Components.
 Proc. of MFCS'15: 382-393. LNCS 9234, Springer.
- [ASE14] R. Di Cosmo, M. Lienhardt, R. Treinen, S. Zacchiroli, J. Zwolakowski, A. Eiche, A. Agahi. Automated synthesis and deployment of cloud applications.
 Proc. of ASE'14: 211-222. ACM Press.
- [ICSOC15] R. Di Cosmo, A. Eiche, J. Mauro, S. Zacchiroli, G. Zavattaro, J. Zwolakowski. *Automatic Deployment of Services in the Cloud with Aeolus Blender.* Proc. of ICSOC'15, to appear, Springer.

Automatic Application Deployment in the Cloud