



Optimal and Automated Deployment for Microservices

Mario Bravetti, Saverio Giallorenzo, Jacopo Mauro, Iacopo Talevi, Gianluigi
Zavattaro

► **To cite this version:**

Mario Bravetti, Saverio Giallorenzo, Jacopo Mauro, Iacopo Talevi, Gianluigi Zavattaro. Optimal and Automated Deployment for Microservices. FASE 2019 - 22nd International Conference on Held as Part of the European, Apr 2019, Prague, Czech Republic. 10.1007/978-3-030-16722-6_21 . hal-02387483

HAL Id: hal-02387483

<https://hal.inria.fr/hal-02387483>

Submitted on 29 Nov 2019

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.



Optimal and Automated Deployment for Microservices

Mario Bravetti¹, Saverio Giallorenzo^{2(✉)}, Jacopo Mauro², Iacopo Talevi¹,
and Gianluigi Zavattaro¹

¹ FOCUS Research Team, University of Bologna/Inria, Bologna, Italy

² University of Southern Denmark, Odense, Denmark

saverio@imada.sdu.dk

Abstract. Microservices are highly modular and scalable Service Oriented Architectures. They underpin automated deployment practices like Continuous Deployment and Autoscaling. In this paper we formalize these practices and show that automated deployment — proven undecidable in the general case — is algorithmically treatable for microservices. Our key assumption is that the configuration life-cycle of a microservice is split into two phases: (i) creation, which entails establishing initial connections with already available microservices, and (ii) subsequent binding/unbinding with other microservices. To illustrate the applicability of our approach, we implement an automatic optimal deployment tool and compute deployment plans for a realistic microservice architecture, modeled in the Abstract Behavioral Specification (ABS) language.

1 Introduction

Inspired by service-oriented computing, Microservices structure software applications as highly modular and scalable compositions of fine-grained and loosely-coupled services [18]. These features support modern software engineering practices, like continuous delivery/deployment [30] and application autoscaling [3]. Currently, these practices focus on single microservices and do not take advantage of the information on the interdependencies within an architecture. On the contrary, architecture-level deployment supports the global optimization of resource usage and avoids “domino” effects due to unstructured scaling actions that may cause cascading slowdowns or outages [27, 35, 39].

In this paper, we formalize the problem of automatic deployment and reconfiguration (at the architectural level) of microservice systems, proving formal properties and presenting an implemented solution.

In our work, we follow the approach taken by the *Aeolus component model* [13–15], which was used to formally define the problem of deploying component-based software systems and to prove that, in the general case, such problem is undecidable [15]. The basic idea of Aeolus is to enrich the specification of components with a finite state automaton that describes their deployment life cycle. Previous work identified decidable fragments of the Aeolus model: e.g.,

removing from Aeolus replication constraints (e.g., used to specify a minimal amount of services connected to a load balancer) makes the deployment problem decidable, but non-primitive recursive [14]; removing also conflicts (e.g., used to express the impossibility to deploy in the same system two types of components) makes the problem PSpace-complete [34] or even poly-time [15], but under the assumption that every required component can be (re)deployed from scratch.

Our intuition is that the Aeolus model can be adapted to formally reason on the deployment of microservices. To achieve our goal, we significantly revisit the formalization of the deployment problem, replacing Aeolus components with a model of *microservices*. The main difference between our model of microservices and Aeolus components lies in the specification of their deployment life cycle. Here, instead of using the full power of finite state automata (like in Aeolus and other TOSCA-compliant deployment models [10]), we assume microservices to have two states: (i) creation and (ii) binding/unbinding. Concerning creation, we use *strong* dependencies to express which microservices must be immediately connected to newly created ones. After creation, we use *weak* dependencies to indicate additional microservices that can be bound/unbound. The principle that guided this modification comes from state-of-the-art microservice deployment technologies like Docker [36] and Kubernetes [29]. In particular, the weak and strong dependencies have been inspired by Docker Compose [16] (a language for defining multi-container Docker applications) where it is possible to specify different relationships among microservices using, e.g., the `depends_on` (resp. `external_links`) modalities that force (resp. do not force) a specific startup order similarly to our strong (resp. weak) dependencies. Weak dependencies are also useful to model horizontal scaling, e.g., a load balancer that is bound to/unbound from many microservice instances during its life cycle.

In addition, w.r.t. the Aeolus model, we also consider resource/cost-aware deployments, taking inspiration from the `memory` and `CPU` resources found in Kubernetes. Microservice specifications are enriched with the amount of resources they need to run. In a deployment, a system of microservices runs within a set of computation *nodes*. Nodes represent computational units (e.g., virtual machines in an Infrastructure-as-a-Service Cloud deployment). Each node has a cost and a set of resources available to the microservices it hosts.

On the model above, we define the *optimal deployment problem* as follows: given an initial microservice system, a set of available nodes, and a new target microservice to be deployed, find a sequence of reconfiguration actions that, once applied to the initial system, leads to a new deployment that includes the target microservice. Such a deployment is expected to be *optimal*, meaning that the total cost (i.e., the sum of the costs) of the nodes used is minimal. We show that this problem is decidable by presenting an algorithm working in three phases: (1) generate a set of constraints whose solution indicates the microservices to be deployed and their distribution over the nodes; (2) generate another set of constraints whose solution indicates the connections to be established; (3) synthesize the corresponding deployment plan. The set of constraints includes optimization metrics that minimize the overall cost of the computed deployment.




Fig. 1. Example of microservice deployment (blue boxes: nodes; green boxes: microservices; continuous lines: the initial configuration; dashed lines: full configuration). (Color figure online)

The algorithm has NEXPTIME complexity because, in the worst-case, the length of the deployment plan could be exponential in the size of the input. However, we consider this worst-case unfeasible in practice, as the number of microservices deployable on one node is limited by the available resources. Under the assumption that each node can host at most a polynomial amount of microservices, the deployment problem is NP-complete and the problem of deploying a system minimizing its total cost is an NP-optimization problem. Moreover, having reduced the deployment problem in terms of constraints, we can exploit state-of-the-art constraint solvers [12, 23, 24] that are frequently used in practice to cope with NP-hard problems.

To concretely evaluate our approach, we consider a real-world microservice architecture, inspired by the reference email processing pipeline from Iron.io [22]. We model that architecture in the Abstract Behavioral Specification (ABS) language, a high-level object-oriented language that supports deployment modeling [31]. We use our technique to compute two types of deployments: an initial one, with one instance for each microservice, and a set of deployments to horizontally scale the system depending on small, medium or large increments in the number of emails to be processed. The experimental results are encouraging in that we were able to compute deployment plans that add more than 30 new microservice instances, assuming availability of hundreds of machines of three different types, and guaranteeing optimality.

2 The Microservice Optimal Deployment Problem

We model microservice systems as aggregations of components with ports. Each port exposes provided and required interfaces. Interfaces describe offered and required functionalities. Microservices are connected by means of bindings indicating which port provides the functionality required by another port. As discussed in the Introduction, we consider two kinds of requirements: strong required interfaces, that need to be already fulfilled when the microservice is created, and weak required interfaces, that must be fulfilled at the end of a

deployment (or reconfiguration) plan. Microservices are enriched with the specification of the resources they need to properly run; such resources are provided to the microservices by nodes. Nodes can be seen as the unit of computation executing the tasks associated to each microservice.

As an example, in Fig. 1 we have reported the representation of the deployment of a microservice system inspired by the email processing pipeline that we will discuss in Sect. 3. Here, we consider a simplified pipeline. A **Message Receiver** microservice handles inbound requests, passing them to a **Message Analyzer** that checks the email content and sends the attachments for inspection to an **Attachment Analyzer**. The **Message Receiver** has a port with a *weak* required interface that can be fulfilled by **Message Analyzer** instances. This requirement is weak, meaning that the **Message Receiver** can be initially deployed without any connection to instances of **Message Analyzer**. These connections can be established afterwards and reflect the possibility to horizontally scale the application by adding/removing instances of **Message Analyzer**. This last microservice has instead a port with a *strong* required interface that can be fulfilled by **Attachment Analyzer** instances. This requirement is strong to reflect the need to immediately connect a **Message Analyzer** to its **Attachment Analyzer**.

Figure 1 presents a reconfiguration that, starting from the initial deployment depicted in continuous lines, adds the elements depicted with dashed lines. Namely, a couple of new instances of **Message Analyzer** and a new instance of **Attachment Analyzer** are deployed. This is done in order to satisfy numerical constraints associated to both required and provided interfaces. For required interfaces, the numerical constraints indicate lower bounds to the outgoing bindings, while for provided interfaces they specify upper bounds to the incoming connections. Notice that the constraint ≥ 3 associated to the weak required interface of **Message Receiver** is not initially satisfied; this is not problematic because constraints on weak interfaces are relevant only at the end of a reconfiguration. In the final deployment, such a constraint is satisfied thanks to the two new instances of **Message Analyzer**. These two instances need to be immediately connected to an **Attachment Analyzer**: only one of them can use the initially available **Attachment Analyzer**, because of the constraint ≤ 2 associated to the corresponding provided interface. Hence, a new instance of **Attachment Analyzer** is added.

We also model resources: each microservice has associated resources that it consumes (see the CPU and RAM quantities associated to the microservices in Fig. 1). Resources are provided by nodes, that we represent as containers for the microservice instances, providing them the resources they require. Notice that nodes have also costs: the total cost of a deployment is the sum of the costs of the used nodes (e.g., in the example the total cost is 598 cents per hour, corresponding to the cost of 4 nodes: 2 C4 large and 2 C4 xlarge virtual machine instances of the Amazon public Cloud).

We now move to the formal definitions. We assume the following disjoint sets: \mathcal{I} for interfaces, \mathcal{Z} for microservices, and a finite set \mathcal{R} for kinds of resources. We use \mathbb{N} to denote natural numbers, \mathbb{N}^+ for $\mathbb{N} \setminus \{0\}$, and \mathbb{N}_∞^+ for $\mathbb{N}^+ \cup \{\infty\}$.

Definition 1 (Microservice type). *The set Γ of microservice types, ranged over by $\mathcal{T}_1, \mathcal{T}_2, \dots$, contains 5-tuples $\langle P, D_s, D_w, C, R \rangle$ where:*

- $P = (\mathcal{I} \mapsto \mathbb{N}_\infty^+)$ are the provided interfaces, defined as a partial function from interfaces to corresponding numerical constraints (indicating the maximum number of connected microservices);
- $D_s = (\mathcal{I} \mapsto \mathbb{N}^+)$ are the strong required interfaces, defined as a partial function from interfaces to corresponding numerical constraints (indicating the minimum number of connected microservices);
- $D_w = (\mathcal{I} \mapsto \mathbb{N})$ are the weak required interfaces (defined as the strong ones, with the difference that also the constraint 0 can be used indicating that it is not strictly necessary to connect microservices);
- $C \subseteq \mathcal{I}$ are the conflicting interfaces;
- $R = (\mathcal{R} \rightarrow \mathbb{N})$ specifies resource consumption, defined as a total function from resources to corresponding quantities indicating the amount of required resources.

We assume sets $\text{dom}(D_s)$, $\text{dom}(D_w)$ and C to be pairwise disjoint.¹

Notation: given a microservice type $\mathcal{T} = \langle P, D_s, D_w, C, R \rangle$, we use the following postfix projections `.prov`, `.reqs`, `.reqw`, `.conf` and `.res` to decompose it; e.g., $\mathcal{T}.\text{reqw}$ returns the partial function associating interfaces to weak required interfaces. In our example, for instance, the Message Receiver microservice type is such that $\text{Message Receiver}.\text{reqw}(\text{MA}) = 3$ and $\text{Message Receiver}.\text{res}(\text{RAM}) = 4$. When the numerical constraints are not explicitly indicated, we assume as default value ∞ for provided interfaces (i.e., they can satisfy an unlimited amount of ports requiring the same interface) and 1 for required interfaces (i.e., one connection with a port providing the same interface is sufficient).

Inspired by [14], we allow a microservice to specify a conflicting interface that, intuitively, forbids the deployment of other microservices providing the same interface. Conflicting interfaces can be used to express conflicts among microservices, preventing both of them to be present at the same time, or cases in which only one microservice instance can be deployed (e.g., a consistent and available microservice that can not be replicated).

Since the requirements associated with strong interfaces must be immediately satisfied, it is possible to deploy a configuration with circular dependencies only if at least one weak required interface is involved in the cycle. In fact, having a cycle with only strong required interfaces would mean to deploy all the microservices involved in the cycle simultaneously. We now formalize a well-formedness condition on microservice types to guarantee the absence of such configurations.

Definition 2 (Well-formed Universe). *Given a finite set of microservice types U (that we also call universe), the strong dependency graph of U is as follows: $G(U) = (U, V)$ with $V = \{(\mathcal{T}, \mathcal{T}') \mid \mathcal{T}, \mathcal{T}' \in U \wedge \exists p \in \mathcal{I}. p \in \text{dom}(\mathcal{T}.\text{reqs}) \cap \text{dom}(\mathcal{T}'.\text{prov})\}$. The universe U is well-formed if $G(U)$ is acyclic.*

¹ Given a partial function f , we use $\text{dom}(f)$ to denote the domain of f , i.e., the set $\{e \mid \exists e' : (e, e') \in f\}$.

In the following, we always assume universes to be well-formed. Well-formedness does not prevent the specification of microservice systems with circular dependencies, which are captured by cycles with at least one weak required interface.

Definition 3 (Nodes). *The set \mathcal{N} of nodes is ranged over by o_1, o_2, \dots . We assume the following information to be associated to each node o in \mathcal{N} .*

- A function $R = (\mathcal{R} \rightarrow \mathbb{N})$ that specifies node resource availability: we use $o.\text{res}$ to denote such a function.
- A value in \mathbb{N} that specifies node cost: we use $o.\text{cost}$ to denote such a value.

As example, in Fig. 1, the node `Node1_large` is such that `Node1_large.res(RAM) = 4` and `Node1_large.cost = 100`.

We now define configurations that describe systems composed of microservice instances and bindings that interconnect them. A configuration, ranged over by $\mathcal{C}_1, \mathcal{C}_2, \dots$, is given by a set of microservice types, a set of deployed microserves (with their associated type), and a set of bindings. Formally:

Definition 4 (Configuration). *A configuration \mathcal{C} is a 4-ple $\langle Z, T, N, B \rangle$ where:*

- $Z \subseteq \mathcal{Z}$ is the set of the currently deployed microserves;
- $T = (Z \rightarrow \mathcal{T})$ are the microservice types, defined as a function from deployed microserves to microservice types;
- $N = (Z \rightarrow \mathcal{N})$ are the microservice nodes, defined as a function from deployed microserves to nodes that host them;
- $B \subseteq \mathcal{I} \times Z \times Z$ is the set of bindings, namely 3-ples composed of an interface, the microservice that requires that interface, and the microservice that provides it; we assume that, for $(p, z_1, z_2) \in B$, the two microserves z_1 and z_2 are distinct and $p \in (\text{dom}(T(z_1).\text{req}_s) \cup \text{dom}(T(z_1).\text{req}_w)) \cap \text{dom}(T(z_2).\text{prov})$.

In our example, if we use `mr` to refer to the instance of `Message Receiver`, and `ma` for the initially available `Message Analyzer`, we will have the binding $(\text{MA}, \text{mr}, \text{ma})$. Moreover, concerning the microservice placement function N , we have $N(\text{mr}) = \text{Node1_large}$ and $N(\text{ma}) = \text{Node2_xlarge}$.

We are now ready to formalize the notion of correctness of configuration. We first define a *provisional correctness*, considering only constraints on strong required and provided interfaces, and then we define a general notion of configuration correctness, considering also weak required interfaces and conflicts. The former is intended for transient configurations traversed during the execution of a reconfiguration, while the latter for the final configuration.

Definition 5 (Provisionally correct configuration). *A configuration $\mathcal{C} = \langle Z, T, N, B \rangle$ is provisionally correct if, for each node $o \in \text{ran}(N)$, it holds²*

$$\forall r \in \mathcal{R}. \quad o.\text{res}(r) \geq \sum_{z \in Z, N(z)=o} T(z).\text{res}(r)$$

and, for each microservice $z \in Z$, both following conditions hold:

² Given a (partial) function f , we use $\text{ran}(f)$ to denote the range of f , i.e., the function image set $\{f(e) \mid e \in \text{dom}(f)\}$.

- $(p \mapsto n) \in T(z).\mathbf{reqs}$ implies that there exist n distinct microservices $z_1, \dots, z_n \in Z \setminus \{z\}$ such that, for every $1 \leq i \leq n$, we have $\langle p, z, z_i \rangle \in B$;
- $(p \mapsto n) \in T(z).\mathbf{prov}$ implies that there exist no m distinct microservices $z_1, \dots, z_m \in Z \setminus \{z\}$, with $m > n$, such that, for every $1 \leq i \leq m$, we have $\langle p, z_i, z \rangle \in B$.

Definition 6 (Correct configuration). A configuration $\mathcal{C} = \langle Z, T, N, B \rangle$ is correct if \mathcal{C} is provisionally correct and, for each microservice $z \in Z$, both following conditions hold:

- $(p \mapsto n) \in T(z).\mathbf{req}_w$ implies that there exist n distinct microservices $z_1, \dots, z_n \in Z \setminus \{z\}$ such that, for every $1 \leq i \leq n$, we have $\langle p, z, z_i \rangle \in B$;
- $p \in T(z).\mathbf{conf}$ implies that, for each $z' \in Z \setminus \{z\}$, we have $p \notin \text{dom}(T(z').\mathbf{prov})$.

Notice that, in the example in Fig. 1, the initial configuration (in continuous lines) is only provisionally correct in that the weak required interface MA (with arity 3) of the Message Receiver is not satisfied (because there is only one outgoing binding). The full configuration — including also the elements in dotted lines — is instead correct: all the constraints associated to the interfaces are satisfied.

We now formalize how configurations evolve by means of atomic actions.

Definition 7 (Actions). The set \mathcal{A} contains the following actions:

- $\mathit{bind}(p, z_1, z_2)$ where $z_1, z_2 \in \mathcal{Z}$, with $z_1 \neq z_2$, and $p \in \mathcal{I}$: add a binding between z_1 and z_2 on port p (which is supposed to be a weak-require port of z_1 and a provide port of z_2);
- $\mathit{unbind}(p, z_1, z_2)$ where $z_1, z_2 \in \mathcal{Z}$, with $z_1 \neq z_2$, and $p \in \mathcal{I}$: remove the specified binding on p (which is supposed to be a weak required interface of z_1 and a provide port of z_2);
- $\mathit{new}(z, T, o, B_s)$ where $z \in \mathcal{Z}$, $T \in \Gamma$, $o \in \mathcal{N}$ and $B_s = (\text{dom}(T.\mathbf{reqs}) \rightarrow 2^{Z \setminus \{z\}})$; with B_s (representing bindings from strong required interfaces in T to sets of microservices) being such that, for each $p \in \text{dom}(T.\mathbf{reqs})$, it holds $|B_s(p)| \geq T.\mathbf{reqs}(p)$: add a new microservice z of type T hosted in o and bind each of its strong required interfaces to a set of microservices as described by B_s ;³
- $\mathit{del}(z)$ where $z \in \mathcal{Z}$: remove the microservice z from the configuration and all bindings involving it.

In our example, assuming that the initially available Attachment Analyzer is named aa , we have that the action to create the initial instance of Message Analyzer is $\mathit{new}(\mathit{ma}, \text{MessageAnalyzer}, \text{Node2_xlarge}, (\mathit{AA} \mapsto \{\mathit{aa}\}))$. Notice that it is necessary to establish the binding with the Attachment Analyzer because of the corresponding strong required interface.

The execution of actions can now be formalized using a labeled transition system on configurations, which uses actions as labels.

³ Given sets S and S' we use: 2^S to denote the power set of S , i.e., the set $\{S' \mid S' \subseteq S\}$; $S - S'$ to denote set difference; and $|S|$ to denote the cardinality of S .

Definition 8 (Reconfigurations). *Reconfigurations are denoted by transitions $C \xrightarrow{\alpha} C'$ meaning that the execution of $\alpha \in \mathcal{A}$ on the configuration C produces a new configuration C' . The transitions from a configuration $C = \langle Z, T, N, B \rangle$ are defined as follows:*

$$\begin{array}{ll}
C \xrightarrow{\text{bind}(p, z_1, z_2)} \langle Z, T, N, B \cup \langle p, z_1, z_2 \rangle \rangle & C \xrightarrow{\text{unbind}(p, z_1, z_2)} \langle Z, T, N, B \setminus \langle p, z_1, z_2 \rangle \rangle \\
\text{if } \langle p, z_1, z_2 \rangle \notin B \text{ and} & \text{if } \langle p, z_1, z_2 \rangle \in B \text{ and} \\
p \in \text{dom}(T(z_1).\text{req}_w) \cap \text{dom}(T(z_2).\text{prov}) & p \in \text{dom}(T(z_1).\text{req}_w) \cap \text{dom}(T(z_2).\text{prov}) \\
\\
C \xrightarrow{\text{new}(z, T, o, B_s)} \langle Z \cup \{z\}, T', N', B' \rangle & C \xrightarrow{\text{del}(z)} \langle Z \setminus \{z\}, T', N', B' \rangle \\
\text{if } z \notin Z \text{ and} & \text{if } T' = \{(z' \mapsto T) \in T \mid z \neq z'\} \text{ and} \\
\forall p \in \text{dom}(T.\text{req}_s). \forall z' \in B_s(p). & N' = \{(z' \mapsto o) \in N \mid z \neq z'\} \text{ and} \\
p \in \text{dom}(T(z').\text{prov}) \text{ and} & B' = \{\langle p, z_1, z_2 \rangle \in B \mid z \notin \{z_1, z_2\}\} \\
T' = T \cup \{(z \mapsto T)\} \text{ and} & \\
N' = N \cup \{(z \mapsto o)\} \text{ and} & \\
B' = B \cup \{\langle p, z, z' \rangle \mid z' \in B_s(p)\} &
\end{array}$$

A *deployment plan* is simply a sequence of actions that transform a provisionally correct configuration (without violating provisional correctness along the way) and, finally, reach a correct configuration.

Definition 9 (Deployment plan). *A deployment plan P from a provisionally correct configuration C_0 is a sequence of actions $\alpha_1, \dots, \alpha_m$ such that:*

- there exist C_1, \dots, C_m provisionally correct configurations, with $C_{i-1} \xrightarrow{\alpha_i} C_i$ for $1 \leq i \leq m$, and
- C_m is a correct configuration.

Deployment plans are also denoted with $C_0 \xrightarrow{\alpha_1} C_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_m} C_m$.

In our example, a deployment plan that reconfigures the initial provisionally correct configuration into the final correct one is as follows: a *new* action to create the new instance of **Attachment Analyzer**, followed by two *new* actions for the new **Message Analyzers** (as commented above, the connection with the **Attachment Analyzer** is part of these *new* actions), and finally two *bind* actions to connect the **Message Receiver** to the two new instances of **Message Analyzer**.

We now have all the ingredients to define the *optimal deployment problem*, that is our main concern: given a universe of microservice types, a set of available nodes and an initial configuration, we want to know whether and how it is possible to deploy at least one microservice of a given microservice type T by optimizing the overall cost of nodes hosting the deployed microservices.

Definition 10 (Optimal deployment problem). *The optimal deployment problem has, as input, a finite well-formed universe U of microservice types, a finite set of available nodes O , an initial provisionally correct configuration C_0 and a microservice type $T_t \in U$. The output is:*

- A **deployment plan** $P = C_0 \xrightarrow{\alpha_1} C_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_m} C_m$ such that
 - for all $C_i = \langle Z_i, T_i, N_i, B_i \rangle$, with $1 \leq i \leq m$, it holds $\forall z \in Z_i. T_i(z) \in U \wedge N_i(z) \in O$, and
 - $C_m = \langle Z_m, T_m, N_m, B_m \rangle$ satisfies $\exists z \in Z_m : T_i(z) = T_t$;
 if there exists one. In particular, among all deployment plans satisfying the constraints above, one that minimizes $\sum_{o \in O. (\exists z. N_m(z)=o)} o.\text{cost}$ (i.e., the overall cost of nodes in the last configuration C_m), is outputted.
- **no** (stating that no such plan exists); otherwise.

We are finally ready to state our main result on the decidability of the optimal deployment problem. To prove the result we describe an approach that splits the problem in three incremental phases: (1) the first phase checks if there is a possible solution and assigns microservices to deployment nodes, (2) the intermediate phase computes how the microservices need to be connected to each other, and (3) the final phase synthesizes the corresponding deployment plan.

Theorem 1. *The optimal deployment problem is decidable.*

Proof. The proof is in the form of an algorithm that solves the optimal deployment problem. We assume that the input to the problem to be solved is given by U (the microservice types), O (the set of available nodes), C_0 (the initial provisionally correct configuration), and $T_t \in U$ (the target microservice type). We use $\mathcal{I}(U)$ to denote the set of interfaces used in the considered microservice types, namely $\mathcal{I}(U) = \bigcup_{T \in U} \text{dom}(T.\text{reqs}) \cup \text{dom}(T.\text{reqw}) \cup \text{dom}(T.\text{prov}) \cup T.\text{conf}$. The algorithm is based on three phases.

Phase 1 The first phase consists of the generation of a set of constraints that, once solved, indicates how many instances should be created for each microservice type T (denoted with $\text{inst}(T)$), how many of them should be deployed on node o (denoted with $\text{inst}(T, o)$), and how many bindings should be established for each interface p from instances of type T — considering both weak and strong required interfaces — and instances of type T' (denoted with $\text{bind}(p, T, T')$). We also generate an optimization function that guarantees that the generated configuration is minimal w.r.t. its total cost.

We now incrementally report the generated constraints. The first group of constraints deals with the number of bindings:

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{T \in U, p \in \text{dom}(T.\text{reqs})} T.\text{reqs}(p) \cdot \text{inst}(T) \leq \sum_{T' \in U} \text{bind}(p, T, T') \quad (1a)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{T \in U, p \in \text{dom}(T.\text{reqw})} T.\text{reqw}(p) \cdot \text{inst}(T) \leq \sum_{T' \in U} \text{bind}(p, T, T') \quad (1b)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{T \in U, T.\text{prov}(p) < \infty} T.\text{prov}(p) \cdot \text{inst}(T) \geq \sum_{T' \in U} \text{bind}(p, T', T) \quad (1c)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{T \in U, T.\text{prov}(p) = \infty} \text{inst}(T) = 0 \Rightarrow \sum_{T' \in U} \text{bind}(p, T', T) = 0 \quad (1d)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{T \in U, p \notin \text{dom}(T.\text{prov})} \sum_{T' \in U} \text{bind}(p, T', T) = 0 \quad (1e)$$

Constraint **1a** and **1b** guarantee that there are enough bindings to satisfy all the required interfaces, considering both strong and weak requirements. Symmetrically, constraint **1c** guarantees that the number of bindings is not greater than the total available capacity, computed as the sum of the single capacities of each provided interface. In case the capacity is unbounded (i.e., ∞), it is sufficient to have at least one instance that activates such port to support any possible requirement (see constraint **1d**). Finally, constraint **1e** guarantees that no binding is established connected to provided interfaces of microservice types that are not deployed.

The second group of constraints deals with the number of instances of microservices to be deployed.

$$\text{inst}(\mathcal{T}_i) \geq 1 \quad (2a)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{\substack{\mathcal{T} \in U, \\ p \in \mathcal{T}.\text{conf}}} \bigwedge_{\substack{\mathcal{T}' \in U - \{\mathcal{T}\}, \\ p \in \text{dom}(\mathcal{T}'.\text{prov})}} \text{inst}(\mathcal{T}) > 0 \Rightarrow \text{inst}(\mathcal{T}') = 0 \quad (2b)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{\substack{\mathcal{T} \in U, p \in \mathcal{T}.\text{conf} \wedge \\ p \in \text{dom}(\mathcal{T}.\text{prov})}} \text{inst}(\mathcal{T}) \leq 1 \quad (2c)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{\mathcal{T} \in U} \bigwedge_{\mathcal{T}' \in U - \{\mathcal{T}\}} \text{bind}(p, \mathcal{T}, \mathcal{T}') \leq \text{inst}(\mathcal{T}) \cdot \text{inst}(\mathcal{T}') \quad (2d)$$

$$\bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{\mathcal{T} \in U} \text{bind}(p, \mathcal{T}, \mathcal{T}) \leq \text{inst}(\mathcal{T}) \cdot (\text{inst}(\mathcal{T}) - 1) \quad (2e)$$

The first constraint **2a** guarantees the presence of at least one instance of the target microservice. Constraint **2b** guarantees that no two instances of different types will be created if one activates a conflict on an interface provided by the other one. Constraint **2c**, consider the other case in which a type activates the same interface both in conflicting and provided modality: in this case, at most one instance of such type can be created. Finally, the constraints **2d** and **2e** guarantee that there are enough pairs of distinct instances to establish all the necessary bindings. Two distinct constraints are used: the first one deals with bindings between microservices of two different types, the second one with bindings between microservices of the same type.

The last group of constraints deals with the distribution of microservice instances over the available nodes O .

$$\text{inst}(\mathcal{T}) = \sum_{o \in O} \text{inst}(\mathcal{T}, o) \quad (3a)$$

$$\bigwedge_{r \in \mathcal{R}} \bigwedge_{o \in O} \sum_{\mathcal{T} \in U} \text{inst}(\mathcal{T}, o) \cdot \mathcal{T}.\text{res}(r) \leq o.\text{res}(r) \quad (3b)$$

$$\bigwedge_{o \in O} \left(\sum_{\mathcal{T} \in U} \text{inst}(\mathcal{T}, o) > 0 \right) \Leftrightarrow \text{used}(o) \quad (3c)$$

$$\min \sum_{o \in O, \text{used}(o)} o.\text{cost} \quad (3d)$$

Constraint 3a simply formalizes the relationship among the variables $\mathbf{inst}(\mathcal{T})$ and $\mathbf{inst}(\mathcal{T}, o)$ (the total amount of all instances of a microservice type, should correspond to the sum of the instances locally deployed on each node). Constraint 3b checks that each node has enough resources to satisfy the requirements of all the hosted microservices. The last two constraints define the optimization function used to minimize the total cost: constraint 3c introduces the boolean variable $\mathbf{used}(o)$ which is true if and only if node o contains at least one microservice instance; constraint 3d is the function to be minimized, i.e., the sum of the costs of the used nodes.

These constraints, and the optimization function, are expected to be given in input to a constraint/optimization solver. If a solution is not found it is not possible to deploy the required microservice system; otherwise, the next phases of the algorithm are executed to synthesize the optimal deployment plan.

Phase 2 The second phase consists of the generation of another set of constraints that, once solved, indicates the bindings to be established between any pair of microservices to be deployed. More precisely, for each type \mathcal{T} such that $\mathbf{inst}(\mathcal{T}) > 0$, we use $s_i^{\mathcal{T}}$, with $1 \leq i \leq \mathbf{inst}(\mathcal{T})$, to identify the microservices of type \mathcal{T} to be deployed. We also assume a function N that associates microservices to available nodes O , which is compliant with the values $\mathbf{inst}(\mathcal{T}, o)$ already computed in Phase 1, i.e., given a type \mathcal{T} and a node o , the number of $s_i^{\mathcal{T}}$, with $1 \leq i \leq \mathbf{inst}(\mathcal{T})$, such that $N(s_i^{\mathcal{T}}) = o$ coincides with $\mathbf{inst}(\mathcal{T}, o)$.

In the constraints below we use the variables $\mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'})$ (with $i \neq j$, if $\mathcal{T} = \mathcal{T}'$): its value is 1 if there is a connection between the required interface p of $s_i^{\mathcal{T}}$ and the provided interface p of $s_j^{\mathcal{T}'}$, 0 otherwise. We use n and m to denote $\mathbf{inst}(\mathcal{T})$ and $\mathbf{inst}(\mathcal{T}')$, respectively, and an auxiliary total function $\mathit{limProv}(\mathcal{T}', p)$ that extends $\mathcal{T}'.\mathit{prov}$ associating 0 to interfaces outside its domain.

$$\bigwedge_{\mathcal{T} \in U} \bigwedge_{p \in \mathcal{I}(U)} \bigwedge_{i \in 1 \dots n} \bigwedge_{j \in (1 \dots m) \setminus \{i | \mathcal{T} = \mathcal{T}'\}} \sum \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'}) \leq \mathit{limProv}(\mathcal{T}', p) \quad (4a)$$

$$\bigwedge_{\mathcal{T} \in U} \bigwedge_{p \in \mathbf{dom}(\mathcal{T}.\mathit{reqs})} \bigwedge_{i \in 1 \dots n} \bigwedge_{j \in (1 \dots m) \setminus \{i | \mathcal{T} = \mathcal{T}'\}} \sum \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'}) \geq \mathcal{T}.\mathit{reqs}(p) \quad (4b)$$

$$\bigwedge_{\mathcal{T} \in U} \bigwedge_{p \in \mathbf{dom}(\mathcal{T}.\mathit{reqw})} \bigwedge_{i \in 1 \dots n} \bigwedge_{j \in (1 \dots m) \setminus \{i | \mathcal{T} = \mathcal{T}'\}} \sum \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'}) \geq \mathcal{T}.\mathit{reqw}(p) \quad (4c)$$

$$\bigwedge_{\mathcal{T} \in U} \bigwedge_{p \notin \mathbf{dom}(\mathcal{T}.\mathit{reqs}) \cup \mathbf{dom}(\mathcal{T}.\mathit{reqw})} \bigwedge_{i \in 1 \dots n} \bigwedge_{j \in (1 \dots m) \setminus \{i | \mathcal{T} = \mathcal{T}'\}} \sum \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'}) = 0 \quad (4d)$$

Constraint 4a considers the provided interface capacities to fix upper bounds to the bindings to be established, while constraints 4b and 4c fix lower bounds based on the required interface capacities, considering both the weak (see 4b) and the strong (see 4c) ones. Finally, constraint 4d indicates that it is not possible to establish connections on interfaces that are not required.

A solution for these constraints exists because, as also shown in [13], the constraints 1a ... 2e (already solved during Phase 1) guarantee that the config-

uration to be synthesized contains enough capacity on the provided interfaces to satisfy all the required interfaces.

Phase 3 In this last phase we synthesize the deployment plan that, when applied to the initial configuration \mathcal{C}_0 , reaches a new configuration \mathcal{C}_t with nodes, microservices and bindings as computed in the first two phases of the algorithm. Without loss of generality, in this decidability proof we show the existence of a simple plan that first removes the elements in the initial configuration and then deploys the target configuration from scratch. However, as also discussed in Sect. 3, in practice it is possible to define more complex planning mechanisms that re-use microservices already deployed.

Reaching an empty configuration is a trivial task since it is always possible to perform in the initial configuration unbind actions for all the bindings connected to weak required interfaces. Then, the microservices can be safely deleted. Thanks to the well-formedness assumption (Definition 2) and using a topological sort, it is possible to order the microservices to be removed without violating any strong required interface (e.g., first remove the microservice not requiring anything and repeat until all the microservices have been deleted).

The deployment of the target configuration follows a similar pattern. Given the distribution of microservices over nodes (computed in the first phase) and the corresponding bindings (computed in the second phase), the microservices can be created by following a topological sort considering the microservices dependencies following from the strong required interfaces. When all the microservices are deployed on the corresponding nodes, the remaining bindings (on weak required ports) may be added in any possible order. \square

Remark 1. The constraints generated during Phase 2 of the algorithm, in order to establish the microservice bindings, are expected to be given in input to a constraint/optimization solver. One can enrich such constraints with metrics to optimize, e.g., the number of local bindings (i.e., give a preference to the connections among microservices hosted in the same node):

$$\min \sum_{\mathcal{T}, \mathcal{T}' \in U, i \in 1 \dots \text{inst}(\mathcal{T}), j \in 1 \dots \text{inst}(\mathcal{T}'), p \in \mathcal{I}(U), N(s_i^{\mathcal{T}}) \neq N(s_j^{\mathcal{T}'})} \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'})$$

Another example, used in the case study discussed in Sect. 3, is the following metric that maximizes the number of bindings⁴:

$$\max \sum_{s_i^{\mathcal{T}}, s_j^{\mathcal{T}'}, p \in \mathcal{I}(U)} \mathbf{b}(p, s_i^{\mathcal{T}}, s_j^{\mathcal{T}'})$$

From the complexity point of view, it is possible to show that the decision versions of the optimization problem solved in Phase 1 is NP-complete, in Phase

⁴ We model a load balancer as a microservice having a weak required interface, with arity 0, that can be provided by its back-end service. By adopting the above maximization metric, the synthesized configuration connects all possible services to such required interface, thus allowing the load balancer to forward requests to all of them.




Fig. 2. Microservice architecture for email processing.

2 is in NP, while the planning in Phase 3 is synthesized in polynomial time. Unfortunately, due to the fact that numeric constraints can be represented in log space, the output of Phase 2 requiring the enumeration of all the microservices to deploy can be exponential in the size of the output of Phase 1 (indicating only the total number of instances for each type). For this reason, the optimal deployment problem is in NEXPTIME. However, we consider unfeasible in practice the deployment of an exponential number of microservices on one node having limited resources. If at most a polynomial number of microservices can be deployed on each node, we have that the optimal deployment problem becomes an NP-optimization problem and its decision version is NP-complete. See the companion technical report [8] for the formal proofs of complexity.

3 Application of the Technique to the Case-Study

Given the asymptotic complexity of our solution (NP under the assumption of polynomial size of the target configuration) we have decided to evaluate its applicability in practice by considering a real-world microservice architecture, namely the email processing pipeline described in [22]. The considered architecture separates and routes the components found in an email (headers, links, text, attachments) into distinct, parallel sub-pipelines with specific tasks (e.g., remove malicious attachments, tag the content of the mail). We report in Fig. 2 a depiction of the architecture. When an email reaches the Message Receiver it is forwarded to the Message Parser, which sends each component into a specific sub-pipeline. In the sub-pipelines, some microservices — e.g., Text Analyzer and Attachment Analyzer — coordinate with other microservices — e.g., Sentiment Analyzer and Virus Scanner — to process their inputs. Each microservice in the architecture has a given resource consumption (expressed in terms of CPU and memory). As expected, the processing of each email component entails a specific load. Some microservices can handle large inputs, e.g., in the range of 40K simultaneous requests (e.g., Header Analyzer that processes short and uniform inputs). Other microservices sustain heavier computations (e.g., Image Recognizer) and can handle smaller simultaneous inputs, e.g., in the range of 10K requests.

To model the system above, we use the Abstract Behavioral Specification (ABS) language, a high-level object-oriented language that supports deployment modeling [31]. ABS is agnostic w.r.t. deployment platforms (Amazon AWS, Microsoft Azure) and technologies (e.g., Docker or Kubernetes) and it offers high-level deployment primitives for the creation of new *deployment components* and the instantiation of objects inside them. Here, we use ABS deployment components as computation nodes, ABS objects as microservice instances, and ABS object references as bindings. Finally, to describe the requirements in our model, we use ABS with SmartDepl [25], an extension that supports deployment annotations. Strong required interfaces are modeled as class annotations indicating mandatory parameters for the class constructor: such parameters contain the references to the objects corresponding to the microservices providing the strongly required interfaces. Weak required interfaces are expressed as annotations concerning specific methods used to pass, to an already instantiated object, the references to the objects providing the weakly required interfaces. We define a class for each microservice type, plus one *load balancer* class for each microservice type. A load balancer distributes requests over a set of instances that can scale horizontally. Finally, we model nodes corresponding to Amazon EC2 instances: `c4.large`, `c4.xlarge`, and `c4.2xlarge` (with the corresponding provided resources and costs).

Microservice (max computational load)	Initial (10K)	+20K	+50K	+80K
MessageReceiver (∞)	1	-	-	-
MessageParser (40K)	1	-	+1	-
HeaderAnalyzer (40K)	1	-	+1	-
LinkAnalyzer (40K)	1	-	+1	-
TextAnalyzer (15K)	1	+1	+2	+2
SentimentAnalyzer (15K)	1	+3	+4	+6
AttachmentsManager (30K)	1	+1	+2	+2
VirusScanner (13K)	1	+3	+4	+6
ImageAnalyzer (30K)	1	+1	+2	+2
NSFWDetector (13K)	1	+3	+4	+6
ImageRecognizer (13K)	1	+3	+4	+6
MessageAnalyzer (70K)	1	+1	+2	+2

In the table above, we report the result of our algorithm w.r.t. four incremental deployments: the initial in column 2 and under incremental loads in 3–5. We also consider an availability of 40 nodes for each of the three node types. In the first column of the Table, next to a microservice type, we report its corresponding maximum computational load, i.e., the maximal number of simultaneous requests that it can manage. As visible in columns 2–5, different maximal computational loads imply different scaling factors w.r.t. a given

number of simultaneous requests. In the initial configuration we consider 10K simultaneous requests and we have one instance of each microservice type (and of the corresponding load balancer). The other deployment configurations deal with three scenarios of horizontal scaling, assuming three increasing increments of inbound messages (20K, 50K, and 80K). In the three scaling scenarios, we do not implement the planning algorithm described in Phase 3 of the proof of Theorem 1. Contrarily, we take advantage of the presence of the load balancers and, as described in Remark 1, we achieve a similar result with an optimization function that maximizes the number of bindings of the load balancers. For every scenario, we use SmartDepl [33] to generate the ABS code for the plan that deploys an optimal configuration, setting a timeout of 30 min for the computation of every deployment scenario.⁵ The ABS code modeling the system and the generated code are publicly available at [7]. A graphical representation of the initial configuration is available in the companion technical report [8].

4 Related Work and Conclusion

In this work, we consider a fundamental building block of modern Cloud systems, microservices, and prove that the generation of a deployment plan for an architecture of microservices is decidable and fully automatable; spanning from the synthesis of the optimal configuration to the generation of the deployment actions. To illustrate our technique, we model a real-world microservice architecture in the ABS [31] language and we compute a set of deployment plans.

The context of our work regards automating Cloud application deployment, for which there exist many specification languages [5, 11], reconfiguration protocols [6, 19], and system management tools [26, 32, 37, 38]. Those tools support the specification of deployment plans but they do not support the automatic distribution of software instances over the available machines. The proposals closest to ours are those by Feinerer [20] and by Fischer et al. [21]. Both proposals rely on a solver to plan deployments. The first is based on the UML component model, which includes conflicts and dependencies, but lacks the modeling of nodes. The second does not support conflicts in the specification language. Neither proposals support the computation of optimal deployments.

Three projects inspire our proposal: Aeolus [13, 14], Zephyrus [1], and ConfSolve [28]. The Aeolus model paved the way to reason on deployment and reconfiguration, proving some decidability results. Zephyrus is a configuration tool based on Aeolus and it constitutes the first phase of our approach. ConfSolve is a tool for the optimal allocation of virtual machines to servers and of applications to virtual machines. Both tools do not synthesize deployment plans.

⁵ Here, 30 min are a reasonable timeout since we predict different system loads and we compute in advance a different deployment plan for each of them. An interesting future work would aim at shortening the computation to a few minutes (e.g., around the average start-up time of a virtual machine in a public Cloud) to obtain on-the-fly deployment plans tailored to unpredictable system loads.

Regarding autoscaling, existing solutions [2, 4, 17, 29] support the automatic increase or decrease of the number of instances of a service/container, when some conditions (e.g., CPU average load greater than 80%) are met. Our work is an example of how we can go beyond single-component horizontal scaling policies (as analyzed, e.g., in [9]).

As future work, we want to investigate local search approaches to speed-up the solution of the optimization problems behind the computation of a deployment plan. Shorter computation times would open our approach to contexts where it is unfeasible to compute plans ahead of time, e.g., due to unpredictable loads.

References

1. Ábrahám, E., Corzilius, F., Johnsen, E.B., Kremer, G., Mauro, J.: Zephyrus2: on the fly deployment optimization using SMT and CP technologies. In: Fränzle, M., Kapur, D., Zhan, N. (eds.) SETTA 2016. LNCS, vol. 9984, pp. 229–245. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47677-3_15
2. Amazon: Amazon CloudWatch. <https://aws.amazon.com/cloudwatch/>. Accessed Jan 2019
3. Amazon: AWS auto scaling. <https://aws.amazon.com/autoscaling/>. Accessed Jan 2019
4. Apache: Apache Mesos. <http://mesos.apache.org/>. Accessed Jan 2019
5. Bergmayr, A., et al.: A systematic review of cloud modeling languages. *ACM Comput. Surv.* **51**(1), 22:1–22:38 (2018)
6. Boyer, F., Gruber, O., Pous, D.: Robust reconfigurations of component assemblies. In: ICSE, pp. 13–22. IEEE Computer Society (2013)
7. Bravetti, M., Giallorenzo, S., Mauro, J., Talevi, I., Zavattaro, G.: Code repository for the email processing example. <https://github.com/IacopoTalevi/SmartDeploy-ABS-ExampleCode>. Accessed Jan 2019
8. Bravetti, M., Giallorenzo, S., Mauro, J., Talevi, I., Zavattaro, G.: Optimal and automated deployment for microservices. Technical Report (2019). <https://arxiv.org/abs/1901.09782>
9. Bravetti, M., Gilmore, S., Guidi, C., Tribastone, M.: Replicating web services for scalability. In: Barthe, G., Fournet, C. (eds.) TGC 2007. LNCS, vol. 4912, pp. 204–221. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78663-4_15
10. Brogi, A., Canciani, A., Soldani, J.: Modelling and analysing cloud application management. In: Dustdar, S., Leymann, F., Villari, M. (eds.) ESOC 2015. LNCS, vol. 9306, pp. 19–33. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24072-5_2
11. Chardet, M., Coullon, H., Pertin, D., Pérez, C.: Madeus: a formal deployment model. In: HPCS, pp. 724–731. IEEE (2018)
12. Chuffed Team: The CP solver. <https://github.com/geoffchu/chuffed>. Accessed Jan 2019
13. Di Cosmo, R., Lienhardt, M., Mauro, J., Zacchiroli, S., Zavattaro, G., Zwolakowski, J.: Automatic application deployment in the cloud: from practice to theory and back (invited paper). In: CONCUR. LIPIcs, vol. 42, pp. 1–16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2015)

14. Di Cosmo, R., Mauro, J., Zacchiroli, S., Zavattaro, G.: Aeolus: a component model for the cloud. *Inf. Comput.* **239**, 100–121 (2014)
15. Di Cosmo, R., Zacchiroli, S., Zavattaro, G.: Towards a formal component model for the cloud. In: Eleftherakis, G., Hinchey, M., Holcombe, M. (eds.) SEFM 2012. LNCS, vol. 7504, pp. 156–171. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33826-7_11
16. Docker: Docker compose documentation. <https://docs.docker.com/compose/>. Accessed Jan 2019
17. Docker: Docker swarm. <https://docs.docker.com/engine/swarm/>. Accessed Jan 2019
18. Dragoni, N., et al.: Microservices: yesterday, today, and tomorrow. In: Mazzara, M., Meyer, B. (eds.) Present and Ulterior Software Engineering, pp. 195–216. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67425-4_12
19. Durán, F., Salaün, G.: Robust and reliable reconfiguration of cloud applications. *J. Syst. Softw.* **122**, 524–537 (2016)
20. Feinerer, I.: Efficient large-scale configuration via integer linear programming. *AI EDAM* **27**(1), 37–49 (2013)
21. Fischer, J., Majumdar, R., Esmailsabzali, S.: Engage: a deployment management system. In: PLDI (2012)
22. Fromm, K.: Thinking Serverless! How New Approaches Address Modern Data Processing Needs. <https://read.acloud.guru/thinking-serverless-how-new-approaches-address-modern-data-processing-needs-part-1-af6a158a3af1>. Accessed Jan 2019
23. GECODE: an open, free, efficient constraint solving toolkit. <http://www.gecode.org>. Accessed Jan 2019
24. Google: Optimization tools. <https://developers.google.com/optimization/>. Accessed Jan 2019
25. de Gouw, S., Mauro, J., Nobakht, B., Zavattaro, G.: Declarative elasticity in ABS. In: Aiello, M., Johnsen, E.B., Dustdar, S., Georgievski, I. (eds.) ESOCC 2016. LNCS, vol. 9846, pp. 118–134. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44482-6_8
26. Red Hat Ansible. <https://www.ansible.com/>. Accessed Jan 2019
27. Hellerstein, J.M., et al.: Serverless computing: one step forward, two steps back. arXiv preprint [arXiv:1812.03651](https://arxiv.org/abs/1812.03651) (2018)
28. Hewson, J.A., Anderson, P., Gordon, A.D.: A declarative approach to automated configuration. In: LISA (2012)
29. Hightower, K., Burns, B., Beda, J.: Kubernetes: Up and Running Dive into the Future of Infrastructure, 1st edn. O’Reilly Media, Inc., Sebastopol (2017)
30. Humble, J., Farley, D.: Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation. Addison-Wesley Professional, Boston (2010)
31. Johnsen, E.B., Hähnle, R., Schäfer, J., Schlatter, R., Steffen, M.: ABS: a core language for abstract behavioral specification. In: Aichernig, B.K., de Boer, F.S., Bonsangue, M.M. (eds.) FMCO 2010. LNCS, vol. 6957, pp. 142–164. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25271-6_8
32. Kanies, L.: Puppet: next-generation configuration management. *login: USENIX Mag.* **31**(1), 19–25 (2006)
33. Mauro, J.: Smartdepl. https://github.com/jacopoMauro/abs_deployer. Accessed Jan 2019

34. Mauro, J., Zavattaro, G.: On the complexity of reconfiguration in systems with legacy components. In: Italiano, G.F., Pighizzini, G., Sannella, D.T. (eds.) MFCS 2015. LNCS, vol. 9234, pp. 382–393. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48057-1_30
35. Mccombs, S.: Outages? Downtime? <https://sethmccombs.github.io/work/2018/12/03/Outages.html>. Accessed Jan 2019
36. Merkel, D.: Docker: lightweight Linux containers for consistent development and deployment. *Linux J.* **2014**(239), 2 (2014)
37. Opscode: Chef. <https://www.chef.io/chef/>. Accessed Jan 2019
38. Puppet Labs: Marionette collective. <http://docs.puppetlabs.com/mcollective/>. Accessed Jan 2019
39. Woods, D.: On infrastructure at scale: a cascading failure of distributed systems. <https://medium.com/@daniel.p.woods/on-infrastructure-at-scale-a-cascading-failure-of-distributed-systems-7cff2a3cd2df>. Accessed Jan 2019

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

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

