

Software Architecture for the Cloud – a Roadmap towards Control-Theoretic, Model-Based Cloud Architecture

Claus Pahl¹ and Pooyan Jamshidi²

¹ IC4 & Lero, School of Computing, Dublin City University

² Department of Computing, Imperial College London

Abstract. The cloud is a distributed architecture providing resources as tiered services. Through the principles of service-orientation and generally provided using virtualisation, the deployment and provisioning of applications can be managed dynamically, resulting in cloud platforms and applications as interdependent adaptive systems. Dynamically adaptive systems require a representation of requirements as dynamically manageable models, enacted through a controller implementing a feedback look based on a control-theoretic framework. We argue that a control-theoretic, model-based architectural framework for the cloud is needed. While some critical aspects such as uncertainty have already been taken into account, what has not been accounted for are challenges resulting from the cloud architecture as a multi-tiered, distributed environment. We identify challenges and define a framework that aims at a better understanding and a roadmap towards control-theoretic, model-based cloud architecture – driven by software architecture concerns.

Keywords: Cloud Computing, Control Theory, Adaptive System, Software Architecture, Microservice, Model-Based Controller, Uncertainty.

1 Introduction

Adapting systems to changing requirements is often a necessity to guarantee on-going correct and satisfying performance. Self-adaptive systems are systems that are able to adjust their behaviour in response to their perception of the environment and the system itself [4]. The software engineering community has approached this from the requirements engineering perspective [10], but has recognised the need for software architecture to play a major role in a solution.

Requirements need to have a representation at runtime to allow self-adaptive systems to interact with the environment, i.e., reflect this through models that also link in the decision-making process necessary to change the underlying system itself [2, 6, 3]. Dynamically adaptive systems require a representation of requirements as dynamically manageable models, enacted through a controller implementing a feedback look based on a control-theoretic framework [1].

The cloud is moving towards a distributed, often federated architecture of many individual cloud services [9], providing resources as services in a tiered

fashion. The configuration, deployment and provisioning of application architectures can be managed dynamically as a response to changes in requirements and changes in the execution platform environment, resulting in cloud platforms and the applications in them as interdependent adaptive systems. Microservices are emerging as a new architectural style, aiming at realising software systems as a package of small services, each deployable on a different platform. These run in their own process while communicating through lightweight mechanisms without any centralized control³. We argue that a cloud-specific control-theoretic, model-based architectural framework is needed. While critical aspects such as uncertainty have been investigated [5, 7, 8] for the cloud, what has not been accounted for are the challenges resulting from the cloud architecture as a multi-tiered, distributed environment for increasingly fragmented application architectures.

We identify the challenges and define a conceptual framework. The target is a roadmap towards control-theoretic, model-based cloud architecture in which software architecture concerns play the central role.

2 Cloud Architecture – Definition and Scenario

Our view on cloud systems from an architectural perspective addresses the key shortcomings of the current discussion of control-theoretic approaches to adaptive systems, and cloud in particular. We will also argue for a model-based approach to controller definition later on as well. The cloud allows the distributed, tiered deployment of software. The underlying architecture links infrastructure and platform providers with the software applications running in them. Software is usually logically architected in a layered format, but in the cloud mapped onto (virtualised) physical tiers.

- Logical layers organise code. Typical layers include presentation, business logic and data management and storage. However, this does not imply that the layers run on different computers or in different processes.
- Physical tiers are about the location of the application execution. Tiers are places where layers are deployed and where layers run.

The cloud services provided as infrastructure-as-a-services (IaaS), platform-as-a-service (PaaS) or software-as-a-service (SaaS) realise these tiers, albeit in a virtualised form accessed through services.

A further complication arises through clouds as distributed, often federated systems, even if providing the same or similar services, will operate differently. Interaction between the layers, but also horizontally is possible and necessary, which we capture in the following architectural scenario in Figure 1.

Let us illustrate a common problem. An infrastructure server might have the capacity to deal with 100 user applications at the same time, but the workload might temporarily reduce significantly. Load balancing would allow the system architecture to be adapted and applications relocated to one server, thus scaling

³ <http://martinfowler.com/articles/microservices.html>

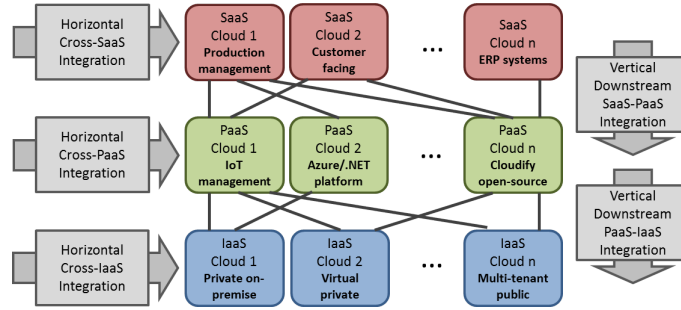


Fig. 1. Tiered and Distributed Cloud Architecture.

down the deployment of servers. Here, the system reacts to external factors – the reduced load – and adapts the configuration to reduce the costs (a non-functional requirement) while still maintaining adequate performance (also a non-functional requirement). Two observations emerge.

- Workload and QoS dictate the adaptation. Cost and quality as drivers for decisions – i.e., decisions are made based on non-functional requirements.
- In the cloud as a tiered architecture, where user applications might run on third-party provided infrastructure servers. Factors that influence here down-scaling as the adaptation include (i) application performance at the user tier/layer and (ii) system workload at the infrastructure tier/layer.

Other scenarios here could involve changing non-functional requirements rather than changing environment factors. The performance requirement might need to be tightened, resulting in an up-scaling of the infrastructure.

Recently, microservice architectures have been discussed, which aim to break up application architectures into independently deployable services that can be rapidly deployed to any infrastructure resource as required. Microservices are independently deployable, usually supported by a fully automated deployment and orchestration framework. They require the ability to deploy often and independently at arbitrary schedules, instead of requiring synchronized deployments at fixed times. The microservice deployment and orchestration across the vertical and horizontal dimensions of the cloud are central architecture concerns. Clouds provide a management tool for their flexible deployment schedules and provisioning orchestration needs, particularly, if these are to be PaaS-provisioned.

3 Dynamic Requirements and Models

As the example above has indicated, both requirements (the user-facing tiers) and the platform (the infrastructure-facing tiers) can change dynamically. What is needed first is a review of modelling concerns for this context. Drivers of change are often requirements to maintain quality-of-service at the user end to maintain within the limits (non-functional requirements) possible stated in a

service-level agreement. In [4], a number of model dimensions are identified that help to frame the adaptivity problem:

- Goals as system objectives: evolution, flexibility, multiplicity, dependency
- Change captures causes of adaptation: source, type, frequency, anticipation.
- Mechanism implements adaptation: type, autonomy, organisation, scope, duration, timeliness, triggering
- Effects define adaptation impact: criticality, predictability, overhead, resilience

A challenge here is the mapping of requirements to the underlying architecture. The solution is a control loop, based on control-theoretic foundations [10], but importantly, the layering of the application architecture onto the tiered cloud. The run-time representation of requirements in the form of application requirements and cloud infrastructure models needs to provide model manipulation and access features to allow introspection and reasoning about these models [2, 6].

A specific challenge is the uncertainty that arises in the interaction between models and the system architecture – the latter possibly at different tiers/layer, all interacting with one another along their interfaces, cf. Figure 1. Respective models that capture uncertainty and can map this as actions within the control loop are needed. The models themselves need to reflect the adaptation approach, requiring to capture the non-functional properties, but more significantly allow prediction and reasoning to take place in an environment prone to uncertainty.

4 Measurement, Prediction and Uncertainty

The state of a system is characterised by a range of non-functional properties that need to be aligned with non-functional requirements. Due to the layering, mapping and managing these across layers, but also within one layer, is challenging. In general, we need to measure at different layers and map between the different tiers in the cloud.

- The upper level represents the application service-level qualities.
- The lower level are the loads of infrastructure resources that run the service.

Furthermore, there is a mapping of the infrastructure loads into a cost model – which can of course be a major driver of adaptation decisions.

Measurement and Uncertainty. Ideally, system state attributes can be reliably measured. However, the cloud adds a high degree of uncertainty here [10]:

- Uncertainty Level 1: general confidence about the shape of the future, but some key variables do not have precise values.
- Uncertainty Level 2: there are a variety of possible future scenarios, that can be listed and are mutually exclusive and exhaustive.
- Uncertainty Level 3: it is feasible to construct future scenarios, but these are mere possibilities and are unlikely to be exhaustive.
- Uncertainty Level 4: it is not even possible to frame possible future scenarios.

Uncertainty emerges from various sources in cloud systems – as uncertainty from different interpretations and decisions in the adaptation definition process or as uncertainty arising from possible different, distributed monitoring systems resulting in partially unreliable and incomplete data [8]:

- **Uncertainty in Adaptation Definition.** Adaptation policies need a careful determination of thresholds. This relies on a users knowledge of system behaviour and how resources are managed. Therefore, the accuracy of policies remains subjective, making the effect of adaptations prone to uncertainty. Unpredictable changes in environment or application demand may require adaptation models to be continuously re-evaluated and revised.
- **Uncertainty in Dynamic Resource Provisioning.** Acquiring and releasing virtual resources in the cloud is not instantaneous. A cloud controller uses the platform services to initiate the acquisition process and has to wait until resources are available. During this time, which may take minutes for VMs, the cloud application is vulnerable to workload increases, causing uncertainty.
- **Uncertainty in Monitoring Data.** The cloud controller needs to continuously monitor the state of the application as well as of the resources in which the application is deployed in order to timely react to load variations. Monitoring involves a distribution of data collected by measurement-specific probes or sensors, which are not immune to measurement deviations (so-called sensory noise). This sensory noise is another source of uncertainty, as it results in oscillations that may affect how the controller allocates resources.

Formal Models for Uncertainty. Models captures the state, its behaviour and the adaptation rules. Models also reflect how we deal with uncertainty in dynamic systems. The dynamics of a system are often based on state models, describing sequences of possible actions as a protocol. In [2], a Markovian model is used (DTMC – Discrete Time Markov Chains; alternatives could include continuous time models), formalising specific properties in logics such as a probabilistic logics [3] to reason in uncertain spaces – in an uncertain space, the probability of the next state is included in the model.

Others propose fuzzy logic [8], where fuzziness is expressed as a varying, non-binary truth value. This allows the uncertainty of a system situation to be expressed through a membership functions on fuzzy sets. For instance, a fuzzification of adaptation rules [8] can be done. As an example, qualitative values for infrastructure workload and service performance (such as 'very low' or 'very high' for workload) are presented as membership functions in a fuzzy set model, resulting in smoother controller responses.

Analysis and Prediction – Cross-Tier Mapping and Uncertainty. Unreliable or incomplete data causes uncertainty, which can be alleviated to some extent by prediction. Furthermore, the delay in providing resources, as discussed above, also makes prediction a suitable approach. Two aspects emerge:

- Analysing measure system data allows us to predict behaviour, reducing uncertainty and increasing the robustness of the adaptation.

- Prediction also helps to link the layers and tiers in the architecture, as for instance infrastructure tier metrics can be used to predict service-level quality. Prediction captures dependencies and becomes a link between the tiers.

Through predication and analysis of monitored data, we can e.g. identify stable quality utilisation patterns. We can map infrastructure workload patterns for CPU, storage and network utilisation at the infrastructure tier to service-level performance patterns, thus linking models (here pattern-based) across tiers [12].

We can implement a prediction technique for the same workload and performance prediction context, based on simple and double exponential smoothing to smoothen outliers and to anticipate trends. Here the aim is the robustness of the prediction and overall adaptation process (by looking ahead in vulnerable moments when the system is about to change).

5 Control Theory and Controller Architecture

Control theory and control engineering can be applied to build self-adaptive systems. Control theory can help to build the models and the reasoning about them to inform the decision making [4]. Decision making is a multi-objective process [10]. Constructing a utility function that involves all stakeholders (such as end-users and the providers of the various tiers of the system in question) is a challenging task [11]. This utility function is implemented by the cloud controller. This construction of utility function (the model) and the controller is a process involving the following steps [1]: identify goals, identify knobs (measure), devise model and design controller, complemented by validation and verification steps.

A key property of this controller is robustness. Robustness tells how resilient the controller is against noise and uncertainty. Prediction, as discussed above, is in addition to a proper calibration of the model a contributor to robustness. Prediction across layers has already addressed the challenges arising from the tiered cloud. architecture. Techniques such as horizontal scaling can deal with the distribution dimension at each tier.

All concerns need to be managed by a control loop. Often, the MAPE-K model is utilised [2], cf. Fig. 2, as the structure of a controller: *Monitor* application and environment (in control-theoretic terms disturbances such as workload). *Analyse* the input data and detect any possible violation. *Plan* corrective actions in terms of adding resources or removing existing unutilized ones. *Execute* the plan according to a specific platform. Utilise a shared *Knowledge* (model).

It is the task of the controller to synchronise models with run-time architecture [10]. The model part of the controller needs to be implemented and integrated with the cloud architecture in order to allow a model-driven cloud control of non-functional aspects [1].

Controller construction still faces a number of problems [5], including uncertainty, synthesize controllers, heterogeneity, unpredictable workloads, resource bottlenecks, multi-tier applications, multi-cloud resources and scalability. We have already discussed uncertainty and unpredictability. The last few points in-

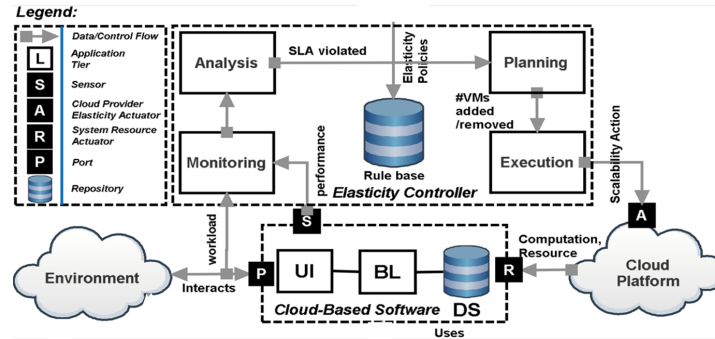


Fig. 2. MAPE-K Control Loop for the Cloud.

dicating the importance of the cloud as a problem from the software architecture perspective, i.e., an architecture onto which the concerns need to be projected:

- Measurement: the controller integrates different models representing the application and infrastructure models at the different tiers – vertical dimension.
- Actuating/executing: typically within a tier, but across services, e.g., based on scalability actions as adaptations – the horizontal dimension.

Uncertainty [1, 7] could also be addressed by reducing the dependency on human stakeholders. Here, machine learning can serve to learn adaptation rules rather than relying on uncertain, possibly erroneous or inconsistent user input. Again, the software architecture perspective can clarify this. A suitable architecture would add a meta-model layer on top of the MAPE-K control loop, representing the learning loop on the models. Models can provide prediction and the feedback loop can correct it, e.g., a queuing model provides how much resources are needed to guarantee an SLA. Since the model is not precise, then it can be augmented with a feedback to correct the error, called feedforwarding.

6 Conclusion

The cloud is a distributed, multi-tiered platform onto which layered, modular software application architectures are mapped. The virtualisation of the cloud resources causes this to be an adaptive system, that is, however, subject to uncertainty and other challenges. Our contribution is the discussion from a software architecture perspective and to propose a roadmap towards a model-based control-theoretic solution that defines some core contributors to future solutions.

- models for uncertainty, allowing prediction and enforcing robustness in a control-theoretic framework
- a model-driven multi-tier cloud controller to manage layered built from easily deployable microservices
- adapting the architectural configuration in the cloud, but also re-architecting the application for the cloud

There is a need for a controller framework that addresses the layered architecture of an application mapped onto tiered cloud resource services through a set of linked models for robust control-theoretic uncertainty management.

Acknowledgments

This work was supported by Science Foundation Ireland grant 13/RC/2094 to Lero (www.lero.ie) and by the Irish Centre for Cloud Computing and Commerce (IC4), a Technology Centre funded by Enterprise Ireland and the IDA.

References

1. K. Angelopoulos N. D'Ippolito I. Gerostathopoulos A. Hempel H. Hoffmann P. Jamshidi E. Kalyvianaki C. Klein F. Krikava S. Misailovic A.V. Papadopoulos S. Ray A.M. Sharifloo S. Shevtsov M. Ujma A. Filieri, M. Maggio and T. Vogel. Software engineering meets control theory. In *10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems SEAMS'2015*, 2015.
2. L. Baresi and C. Ghezzi. A journey through smscom: self-managing situational computing. *Computer Science – Research and Development*, 28(4):267–277, 2013.
3. K. Chan, I. Poernomo, H. Schmidt, and J. Jayaputera. A model-oriented framework for runtime monitoring of nonfunctional properties. In *Quality of Software Architectures and Software Quality*. 2005.
4. R. De Lemos, H. Giese, H.A. Müller, M. Shaw, J. Andersson, M. Litoiu, B. Schmerl, G. Tamura, N.M. Villegas, T. Vogel, et al. Software engineering for self-adaptive systems: A second research roadmap. In *Soft Eng for Self-Adaptive Syst II*. 2013.
5. S. Farokhi, P. Jamshidi, I. Brandic, and E. Elmroth. Self-adaptation challenges for cloud-based applications: A control theoretic perspective. In *10th International Workshop on Feedback Computing 2015*, 2015.
6. C. Ghezzi, L.S. Pinto, P. Spoletini, and G. Tamburrelli. Managing non-functional uncertainty via model-driven adaptivity. In *Intl Conf on Soft Eng*, 2013.
7. M.U. Iftikhar and D. Weyns. Assuring system goals under uncertainty with active formal models of self-adaptation. In *Companion Proceedings of the 36th International Conference on Software Engineering*. ACM, 2014.
8. P. Jamshidi, A. Ahmad, and C. Pahl. Autonomic resource provisioning for cloud-based software. In *Intl Symp on Software Engineering for Adaptive and Self-Managing Systems*, SEAMS'14, 2014.
9. C. Pahl. Containers and clusters for edge cloud architectures - a technology review. In *Intl Conference on Future Internet of Things and Cloud*, FiCloud'15, 2015.
10. P. Sawyer, N. Bencomo, J. Whittle, E. Letier, and A. Finkelstein. Requirements-aware systems: A research agenda for re for self-adaptive systems. In *International Requirements Engineering Conference RE'2010*, pages 95–103, 2010.
11. A. van Hoorn, M. Rohr, A. Gul, and W. Hasselbring. An adaptation framework enabling resource-efficient operation of software systems. In *Proceedings of the Warm Up Workshop for ACM/IEEE ICSE 2010*, WUP '09. ACM, 2009.
12. L. Zhang, Y. Zhang, P. Jamshidi, L. Xu, and C. Pahl. Workload patterns for quality-driven dynamic cloud service configuration and auto-scaling. In *International Conference on Utility and Cloud Computing UCC'2014*, 2014.