

What Can Control Theory Teach Us About Assurances in Self-Adaptive Software Systems?

Marin Litoiu¹, Mary Shaw², Gabriel Tamura³,
Norha M. Villegas³, Hausi A. Müller⁴, Holger Giese⁵,
Romain Rouvoy⁶, and Eric Rutten⁷

¹ York University, Canada

`mlitoiu@yorku.ca`

² Carnegie Mellon University, USA

`mary.shaw@cs.cmu.edu`

³ Universidad Icesi, Colombia

`gtamura@icesi.edu.co`, `nvillega@icesi.edu.co`

⁴ University of Victoria, Canada

`hausi@cs.uvic.ca`

⁵ Hasso Plattner Institute for Software Systems Engineering, Germany

`holger.giese@hpi.uni-potsdam.de`

⁶ University of Lille 1, INRIA Nord Europe, France

`romain.rouvoy@univ-lille1.fr`

⁷ INRIA Grenoble Rhône-Alpes, France

`eric.rutten@inria.fr`

Abstract. Self-adaptive software (SAS) systems monitor their own behavior and autonomously make dynamic adjustments to maintain desired properties in response to changes in the systems' operational contexts. Control theory provides verifiable feedback models to realize this kind of autonomous control for a broad class of systems for which precise quantitative models can be defined. Recent MAPE-K models, and variants such as the hierarchical ACRA and similar others, address a broader range of tasks, but they do not provide the inherent assurances that control theory does, as they do not explicitly identify the properties that reliable controllers should have. These properties, in general, result not from the abstract models, but from the specifics of control strategies, which precisely these models fail to analyze. We show that, even for systems too complex for direct application of classical control theory, the abstractions of control theory provide design guidance that identifies important control characteristics and raises critical design issues about the details of the strategy that determine the controllability of the resulting systems. This in turn enables careful reasoning about whether the control characteristics are in fact achieved. In this chapter we examine the control theory approach, explain several control strategies illustrated with examples from both domains, classical control theory and SAS, and show how the issues addressed by these strategies can and should be seriously considered in the design of self-adaptive software systems. From this examination we distill challenges for developing principles that may serve as the basis of a control theory for self-adaptive software systems.

1 Introduction

This chapter explores the contributions that classical feedback loops, defined by control theory, can make to self-adaptive software (SAS) systems, particularly to their assurances. To do this, we focus on the abstract character of classical feedback loops—how they are formulated, the assurances they provide, and the analysis required to obtain those assurances. Feedback loops, in one form or another, have been adopted as cornerstones of software-intensive self-adaptive systems, as documented in [25] and later on in [8,30]. Building on this, we study the relationships between feedback loops and the types of assurance they can provide to SAS systems. We focus particularly on the conceptual rather than implementation level of the feedback model. That is, we concentrate on the relationships among desired properties granted by the feedback loop (e.g., stability, accuracy, settling time, efficient resource use), the ways computational processes can be adjusted, the choice of control strategy, and the quality attributes (e.g., responsiveness, latency, reliability) of the resulting system.

More concretely, this exploration includes, on the one hand, how feedback loops contribute to providing assurances about the behavior of the controlled system and, on the other hand, how assurances improve the realization of feedback loops in SAS. To set the stage for the discussion of concrete challenges identified in this exploration, we first review the major concepts of traditional control theory and engineering, then study the parallels between control theory [4,34] and the more recent research on feedback control of software systems (i.e., MAPE-K loops) [25,21] in the realm of SAS. We establish a common vocabulary of key concepts, such as the control objective or setpoint, the disturbances that may affect the system, control responsiveness, and the control actions used by the controller to adapt the system. By discussing several types of control strategies, backed up by concrete examples, we illustrate how the control theory approach can direct attention to decisions that are often neglected in the engineering of SAS [8]. From the analysis on these concrete examples, we posit key challenges for assurance research related to the application of feedback loops in self-adaptive software.

The contribution of this chapter is twofold. First, it provides a comprehensive, and easy to understand, explanation of the application of control theory concepts to the engineering of SAS systems. This is valuable for both control engineering researchers interested in applying control engineering in new domains, and SAS researchers and engineers interested in taking advantage of control theory concepts and techniques for the engineering of SAS systems. Second, the chapter discusses assurance challenges faced by SAS designers and extrapolates control theory questions, concepts, and desirable properties into the SAS systems realm.

The remaining sections of this chapter are organized as follows. Section 2 revisits the MAPE-K loop. Section 3 explains control theory concepts, including their analogy with the corresponding elements in SAS systems. Section 4 extends these models to adaptive and hierarchical control. Section 5 presents an overview of the application of control theory to the self-adaptation of software

systems, focusing on the model and controller definitions. Section 6 discusses classic control strategies and the assurances they provide about the quality of the control, based on control theory properties; it also analyzes the design of controllers with desirable control theory properties. Section 7 discusses relevant assurance challenges in the design of SAS systems, extrapolated from those of control theory. Finally, section 8 summarizes and concludes the chapter.

2 Self-Adaptive Software (SAS) Systems

The necessity of a change of perspective in software engineering has been discussed during the last decade by several researchers and practitioners in different software application domains [35,11,10]. In particular, Truex *et al.* posited that software engineering has been based on an the incorrect assumption that software systems should support rigid and immutable business structures and requirements, have low maintenance, and fully satisfy their requirements from the initial system delivery [45,32]. In contrast to this “stable” vision, the engineering of self-adaptive software (SAS) originates from a different paradigm based on continuous analysis, dynamic requirements negotiation and incomplete requirements specification.

SAS systems adjust their own structure or behavior at runtime to regulate the satisfaction of functional, behavioral and non-functional requirements that vary over time, for instance when affected by changes in business goals or the system’s execution environment [28,29,10,38,43]. When system requirements evolve in this way, the system must support both short-term adaptation mechanisms to ensure satisfaction of current requirements and also long-term evolution of the requirements [36,32]. The former requires continuous adjustments of system parameters to satisfy current requirements, and the latter requires runtime system analysis to direct long-term evolution [8,10]. Furthermore, assurance mechanisms, both at design time and runtime, must be implemented to guarantee required properties [44].

2.1 The MAPE-K Loop

Inspired by the human autonomous nervous system, IBM researchers proposed the *autonomic element* as a building block for developing self-managing and autonomic software systems. They synthesized this adaptation in the form of the so-called Monitoring-Analysis-Planning-Execution and shared Knowledge (MAPE-K) loop model, as depicted in Fig. 1. The purpose of this model is to develop autonomous control mechanisms to regulate the satisfaction of dynamic requirements, specifically in software systems [25,21,22]. An autonomic manager not only implements the adaptation phases, but also provides the interfaces (i.e., manageability endpoints) required to sense the environment, and to effect changes on the target system (i.e., the system to be adapted, also referred as the managed application) or to manage other autonomic managers.

The responsibilities of each phase of the MAPE-K loop depicted in Fig. 1 are defined as follows:

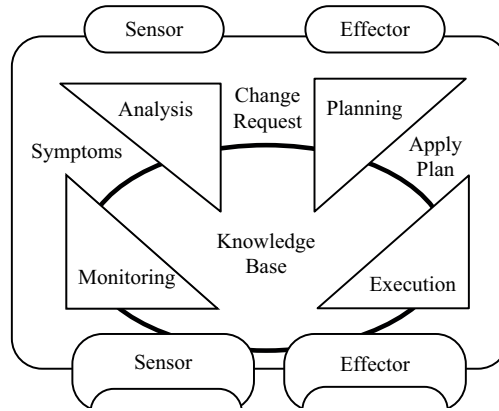


Fig. 1. The MAPE-K loop [25]

- *Monitoring.* In this phase sensors keep track of changes in the managed application’s internal variables corresponding to desired properties (e.g., measured QoS data), the properties to be assured on the adaptation mechanism (e.g., stability), and the external context (i.e., measured from outside the managed application). Based on these changes, monitors must notify relevant context events to the analyzer. Relevant context events are those derivable from system requirements (e.g., from QoS contracts).
- *Analysis.* Based on the high-level requirements to fulfill, and the context events notified by monitors from the monitoring phase, in this phase the analyzer determines whether a system adaptation must be triggered and notifies the planner accordingly. This would occur, for instance, when the notified events signal changes that (may) violate the reference control inputs.
- *Planning.* In this phase, once notified with a reconfiguration event from the context analyzer, the planner selects a strategy to fulfill the new requirements, using the accumulated knowledge in the *knowledge base*. By applying the selected strategy, the planner computes the necessary control actions to be instrumented in the managed software system as discrete operations that are then interpreted by the executor in the next phase.
- *Execution.* In this phase, after receiving a reconfiguration plan, the executor interprets each of the discrete operations specified in the plan and effects them on the managed software system. This implies translating or tailoring the reconfiguration actions to the ones supported by the platform that executes the managed software system.

3 Feedback Control

Control theory depends on reference control points of system behavior and corresponding explicit mathematical model specifications. Some researchers advocate the idea that control theory is applicable as is to SAS systems and propose a process for realizing this idea [17]. Basically, the idea is to find the right mathematical model for the software system behavior and apply the corresponding treatment depending on whether the model is linear or non-linear. This direct application of the methods of classical process control might be tractable for simple cases, but most practical adaptive systems are too complex to apply these methods directly. The approach of [17] only considers the case where process control applies directly. Here, we take a more nuanced position—we posit that understanding the basics of classical control theory provides a foundation for a more abstract view, based upon which important design guidelines, conceived especially for controlling software systems, can be generated and that might otherwise be overlooked. That is, we argue that control theory provides a point of view that enables a designer to design more complex self adaptive systems more effectively.

This section follows [23] and introduces the concepts and vocabulary of control theory. In particular, it discusses factors that affect the quality of control that can be achieved. At each stage, it discusses how control algorithms affect the behavior of systems and it identifies more abstract design tasks that are extensions of the basic theory. We focus here on discrete control, that is, on models with discrete time steps.

In a simple feedback system, a process P has a tuning parameter u (e.g., a knob) that can be manipulated by a controller C to change the behavior of the process, and a tracked metric y that can be sensed in some way. The system is expected to maintain the tracked metric y at a reference level y_r , as illustrated in Fig. 2. At each time increment of the system, C compares the value of y (subject to possible signal translation) to the desired value y_r . The difference is the tracking or control error, and if they are sufficiently different, the controller changes parameter u to drive the process in such a way as to reduce the tracking error.¹ Somewhat confusingly, the tuning parameter u is often called the “input” and the tracked metric y is often called the “output.” It is important to remember that these are the “control” input and output, not the input and output of whatever computation P is performing.

For instance, assume the process P is a household thermostat to keep a room warm at a reference level of $y_r = 23^\circ C$, and the controller C is a rudimentary *on/off* decision mechanism. Correspondingly, the tuning parameter u in P is a simple switch *on/off*. The tracked metric y is the temperature measured from a strategically located sensor in the room. Regularly, C compares the current value of y to the desired value y_r . If y is less than y_r , C makes u to be in the

¹ If appropriate, the process P can be described by a model and the controller C might use that model to determine changes of u .




Fig. 2. The feedback loop control model.

on position. Of course, in a next cycle, whenever y is greater or equal than y_r , C makes u to turn *off*, and so on.

Control systems are especially important when processes are subject to unpredictable disturbances, illustrated in Fig. 2 as “External Disturbances.” External disturbances are inputs in the Process P that affect the controlled output and they cannot be controlled. For our household thermostat example, external disturbances are the external variations in temperature, together with the room temperature changes created by the sudden opening of the room door and windows. In computing environments, external disturbances examples include, among others, variable loads such as number of users or request arrival rates, and variable hit rates on caches. Control systems can also be important when the computation in the process itself is unpredictable and its accuracy or performance can be tuned by adjusting parameters of the computation itself.

3.1 Correspondences between Feedback Control and MAPE-K

An important difference between the feedback loop from control theory and the MAPE-K loop is the complexity of each constituent component and of the overall loop. For example, in the former, the control actions are atomic operations for physical actuators (e.g., resistors and motors) with clear causality between inputs and outputs, whereas in the latter, they are sequences of discrete plans (thus called reconfiguration plans) with long lags and uncertain consequences.

Another difference is that a control feedback loop is at a more abstract level and focuses on inputs, outputs, states, and external disturbances and the quantitative relationships among those variables.

However, the basic principles are the same and, in the SAS domain, it is useful to identify the analogues of the control elements, ensuring that they can be used effectively for control. For example:

- What corresponds to the tracked metric (or output)? That is, what characteristics of the process are we trying to control? Can we measure them directly? If not, how are we going to infer their values. For example, is there a proxy value readily available? How accurately does the proxy value reflect the true value?

- What corresponds to the reference level? That is, how do we characterize the desired behavior of the system? Is it a specific target? Is it a limit value, so that y should always be less than the limit (e.g., the temperature should approach the boiling point in a cooking pot but never reach it), or are values on either side of the target acceptable (e.g., the cooking pot temperature should be 75°C)? Is it a function of something else? If it is a complex function of other system parameters, how do we analyze interactions between the control discipline and the inputs to the function?
- What corresponds to the tuning parameters of the process and the way of modifying them? That is, in what ways can the controller modify the behavior of the process? Are these ways sufficient to maintain control of the process? (old-style automotive cruise control systems failed this test, because only controlled the accelerator, so they could only speed up the car, not slow it down).
- What effect will have on the process a given change in the control input? In control theory, this is called the system identification problem. The more precisely we know at least this effect, the better our ability to achieve the response we want. For complex SAS systems a full specification is often impractical. If so, it is more important to know whether increasing u will increase or decrease y , than it is to know exactly what the size of the effect will be.
- What are the sources of the external disturbances? What is their nature and how do we characterize them? What is their frequency and amplitude? In control theory, feedback loops are designed also to compensate for the effect of disturbances, within some bounds in amplitude and frequency. Can we characterize those bounds for SAS?
- How does the controller decide how much to change the control input variables? Control theory offers a rich space of control disciplines with well-understood effects on the controlled system, explained later in Sect. 6.1. What are the analogous theories for SAS systems?

Figure 3 depicts structural correspondences between the feedback loop and the MAPE-K loop (Autonomic Manager). The *reference level* (y_r) of desired properties is fed by the system administrator to the analysis phase of the MAPE-K loop. In this phase the analyzer evaluates the difference between the *tracked metric* y , which is captured by the sensors of the Autonomic Manager. The analysis phase communicates the adaptation decision to the planning phase by sending a change request that corresponds to the *control error*. The planning and execution phases correspond to the *controller*, which generates an adaptation plan and executes the adaptation of the managed system, *the process*, through the effector using the *control input* u (i.e., the adaptation commands). Finally, sensors may play the role of the *transducer*.

The application of control theory to the engineering of self-adaptive applications has been actively studied by researchers in the SAS realm [8,10,30]. In particular, recent works have studied the usage of control theory to guide the design

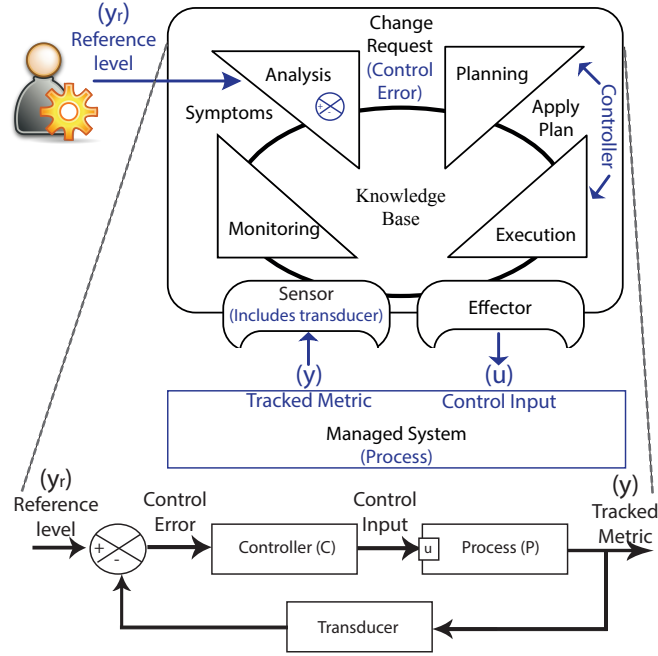


Fig. 3. Correspondences between the feedback loop and the MAPE-K Loop

of MAPE-type systems, in different aspects of self-adaptation from the design of the adaptation mechanism to the assurance of its properties [46,48,44,49,50] To continue advancing in this direction, an important step forward is to support reasoning about the systems by finding the answers to the questions listed above about the correspondences between the feedback loop model from control theory and the MAPE-K loop model from SAS. These questions apply to a wide range of SAS systems, not simply to low-level control systems.

4 Adaptive and Hierarchical Control

For systems that vary in time or face a wide range of external disturbances, it is not possible to design one controller that faces all those changes. For these cases, we need to design a set of controllers or a parameterized controller. When the current controller is not efficient anymore, we change it or retune it. When the controller is updated while the system runs, we call it adaptive control. Adaptive control requires additional logic that monitors the efficiency of the controller and, when some conditions are met, retunes the controller to adapt it to the new situation. The control community has dealt with adaptive control for decades according to Åström [4] since 1961. Åström also provides a good working definition for adaptive control: “An adaptive controller is a controller with adjustable parameters and a mechanism for adjusting the parameters.” This

definition implies two (or more) layered control loops. Home setback thermostats are simple instances of adaptive control. A setback thermostat, useful to reduce overall household energy consumption, gives the user the option of changing the temperature reference value y_r automatically as a function of time and day of the week. In this way, control strategies adapt while the system is running according to changes in the reference temperature. In the SAS area, the use of adaptive feedback loops was first suggested in [8]. One possible realization of an adaptive controller is depicted in Fig. 4. There are two feedback loops: the reactive feedback loop, at the bottom, that takes the output y and brings it to the comparator and further to the controller; and the adaptation loop at the top that estimates a dynamic set of models of the software (Model Estimators) passes the models (Models) to a Controller Tuning/Synthesizing component. This component uses a current Model to either recalculate the parameters of the controller or switch to a different controller.




Fig. 4. Model Identification Adaptive Control

Adaptive control is hierarchical, with at least two layers of feedback loops. Hierarchical control, in addition to adaptive control, can be used for complex systems and complex controllers. The robotics community in particular dealt with hierarchical control since the early eighties—they called it hierarchical intelligent control system (HICS). In the software engineering community, most notably Kephart and Chess [25,22] introduced autonomic systems and autonomic computing reference architecture (ACRA).

The autonomic computing reference architecture (ACRA), depicted in Fig. 5, provides a reference architecture as a guide to organize and orchestrate SAS (i.e., autonomic) systems using autonomic managers (MAPE-K loops). SAS systems based on ACRA are defined as a set of hierarchically structured controllers. Using ACRA, software policies and assurances can be implemented into layers where system administrator (manual manager) policies control lower level policies. System operators have access to all ACRA levels.

Hierarchical control for SAS was further specialized by Kramer et al in a three layer architecture for SAS [26]. In this architecture, each layer has a different time scope and authority.




Fig. 5. The Autonomic Computing Reference Architecture (ACRA) [22]

5 Control Theory Applied to Self Adaptive Software - An Overview

Over the last decade, there have been many efforts to use control theory to design feedback loops for self-adaptive software systems. Hellerstein *et al.*, in their 2004 book [21], introduced several examples where control theory has been applied for controlling the threading level, memory allocation or buffer pool sharing in commercial products such as IBM Lotus Notes and IBM DB2. Other researchers and practitioners have published results on control theory for computer power control [27], thread and web cluster management [1], admission control, video compression [16], performance and denial of service attack mitigation, to name just a few examples.

In the examples we have seen so far, the authors use feedback loops to control quantitative metrics from the categories of performance, cost, energy, reliability. In these control examples, the methodology and the mathematical apparatus follow the control theory methodology and revolves initially around two basic concepts: the model of the controlled subsystem (i.e., the open loop model from the point of view of the controlled metric), and the controller to close the loop, as follows.

5.1 The Open Loop Model

The first major task in applying control theory to SAS is to create a model for the quality attributes to be controlled. Since this model focuses only on the controlled or managed system, we call it open loop model. The model is

quantitative and captures, in a very specific format, the relationships among software components, and between software and environment. The model can be constructed analytically by writing the equations of the modelled phenomena or by following an experimental methodology called system identification[40,31]. In general, defining the model starts by identifying the outputs that need to be controlled, y ; the disturbances, p , that affect the outputs; the control variables (or commands), u , that can control the outputs of the system and compensate for the effects of external disturbances; the internal state variables, x (x can be seen as an intermediate link between u and y); The external disturbances, p , drive the system towards undesired states. For example, an external load increase (increase in number of users) might increase the utilization(state) of a server and therefore increase the response time(output). Note that the perturbations p do not affect the response time directly, but through the intermediate state that we call utilization. To compensate the effects of the external load, a feedback loop designer should engineer a control input that changes the utilization of server in the opposite direction (like transferring some of the load on another server). Note that u, x, p, y are vectors, not scalars, that means that a system has many input, output, state or external disturbance variables. In control terminology, this is defined as MIMO(multiple input, multiple output) system.

In the general case, with f and g non-linear discrete *vector* functions, the open loop model can be described as:²

$$x(k+1) = f(x(k), u(k), p(k)) \quad (1)$$

$$y(k) = g(x(k), u(k)) \quad (2)$$

where k is the current discrete time and $k+1$ means the “next time increment.” How big is the difference between k and $k+1$ depends on what exactly the model tries to capture. Equation 1 projects the state forward, as a function of the current state x , some control commands u and disturbances/perturbations, while Equation 2 expresses the output as a function of the same parameters.




Fig. 6. A discrete non-linear model

The above model formulation, as depicted in Fig. 6, is general enough to describe any quantitative aspect of a system that might need adaptation, including

² In many cases the models include explicitly the modelling and measurement noise; to keep the presentation simple, in this chapter we do not represent noise explicitly.

aspects of SAS. When constructing the above model for a software system, a special attention should be paid to identifying (or engineering) \mathbf{u} , the control inputs or commands. Those are variables in the system that control engineers pay special attention to and they are a *sine qua non* condition for designing and engineering a feedback loop. The commands \mathbf{u} are control knobs, that is software parameters, scripts, programs, interfaces through which a system administrator or an automation program can change a software configuration. They are always designed with some knowledge about the possible external disturbances. Load inputs to the systems, such as workloads, are frequently modeled as external disturbances or perturbations, \mathbf{p} , because they are out of the control of the SAS designer. Their effect is compensated by designing control inputs \mathbf{u} , engineered into the system: increase the capacity of the server, scale out by replicating servers and distributing the load, changing the data flows inside the system, changing the threading levels, etc.

In many cases, the non-linear model can be simplified or approximated with a linear one, such as the one below. In this model we assumed an additive external perturbation:

$$x(k+1) = A * x(k) + B * u(k) + F * p(k) \quad (3)$$

$$y(k) = C * x(k) + D * u(k) \quad (4)$$

where A, B, C, D, F are constant matrices that can be determined experimentally for a specific deployments and under particular perturbations³.

Example 1: software queues modeling (performance view). Many software entities can be modelled as queues (semaphores, critical sections, thread and connection pool containers, web services, servers, etc.). Although the performance metrics of those entities can be captured with non-linear models, many authors prefer to linearize the models around an equilibrium point and express them with the canonical model (3)-(4). Below, there is an example of such model for a web server (taken from [21] and [12]). In equations (5)-(6), x_1 and x_2 represent the server utilization and the memory usage, respectively; the commands, u_1 and u_2 represent the number of HTTP connections and the keep-alive intervals; and y represents the response time of the server. A_{ij} , B_{ij} , where $i, j = 1, 2$ and C_1 are constants that have to be determined experimentally for a specific server deployment. Intuitively, the model says that the future memory usage and the server utilizations are a function of the current usage, the number of connections, and the length of interval they are kept alive.

$$\begin{bmatrix} x_1(k+1) \\ x_2(k+1) \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix} + \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \begin{bmatrix} u_1(k) \\ u_2(k) \end{bmatrix} \quad (5)$$

$$y(k) = [C_1 \ 0] \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix} \quad (6)$$

³ Here we assume a time invariant case, but those values can depend on time as well.

Similar models have been presented for various entities that contain queues [41,42] [20,16,2]. They were derived from the dynamics of the queues and from experiments. For example, Solomon *et al.* [42] use a control theory specific model identification methodology to experimentally identify A, B, C , and D for a threading pool.

5.2 The Controller

Equations (3)-(4) suggest that if we want to keep y at a predefined value y_r (the goal), we accomplish that by computing a command u_r (feed-forward control). In practice, that is impossible for two main reasons: (a) the model (2) is an inaccurate representation of the real system; (b) there are external disturbances, p , that affect the system. A feedback loop implies adding a new software component, a controller (or Adaptive Manager) as depicted in Fig. 7, that is fed back with information from the controlled software system.




Fig. 7. Controller and Feedback Loop.

The controller and the feedback will compensate for modelling errors and for external disturbances. Furthermore, they can also stabilize unstable systems and achieve additional control criteria. The feedback can be taken from the state variables (state feedback) or from the output (output feedback). The state feedback is the most effective, since it assumes an understanding of inner workings and dependencies of the software system. Output feedback is more practical since output variables are easier to measure. For observable systems, a state feedback can be constructed from output feedback, using estimators as for example Kalman filters [24]. When using estimators, we measure y ; estimate x based on the model, the observability matrix and the measurements; and then compute a state based controller. The controller can vary from a simple constant matrix to a complex set of differential (or difference) equations and can be synthesized to achieve some design goal, such as: stability of the entire system, perturbations rejection across wide ranges (robust control), or optimization. While the controller design goals can be diverse, there are standard design techniques that have well defined procedures and solutions.

The interesting aspect of the feedback controller is that the overall equations of the closed loop system (or closed loop model) can be inferred from the open loop model and the equations of the controller. For example, consider the controller a matrix K that needs to be determined. The equations of the closed

loop system, having y_r as input and y as output, can be described as follows (for simplicity, $F=0$):

$$x(k+1) = A_c * x(k) + B_c * y_r(k) \quad (7)$$

$$y(k) = C_c * x(k) + D_c * y_r(k) \quad (8)$$

where $A_c = (A - BK)$; $B_c = B$; $C_c = C$; $D_c = D$.

Having the equations of the closed loop system allows the designer to shape the behavior of the states and therefore of the outputs by choosing the right K . At the same time, a model will enable the analysis and verification that the system has indeed the desired behavior.

5.3 Feedback Control Behavior

Independent of having a model for the open or closed loop, a controller for a SAS achieves its goals by successive approximations. It is hard to know exactly how a given system will respond to changes in the control inputs (i.e., the system dynamics). Control algorithms operate by making successive corrections that are designed to reduce the tracking error. Several factors lead to this strategy, all of them derived from the uncertainties about the exact system dynamics: delays in system response (e.g., startup time to bring a new server online), lags in system response (e.g., the time it takes to clear a backlog), environmental changes and events (e.g., the execution of the garbage collector), and variability in the interaction of the controlled process with its environment (e.g., randomness in arrival rates to a queue, or processing times for the jobs in a queue). Feedback loops provide some robustness in the face of these uncertainties. However, the consequence of their use is that the process approaches the reference level in a series of steps. For example, in the case of the home thermostat, the controller turns the furnace *on* and evaluates how the temperature goes with respect to the reference value. Depending on the results, the thermostat controller turns the furnace *off* and evaluates the results again. This process is continuously repeated to reach the desired temperature. The design of the control algorithm determines whether the path to achieving the reference value is a series of overshoots above and below the desired value or a gradual approach to the reference value. The former strategy may get close to the reference value faster, but at the cost of oscillating behavior (imagine riding in a car that entered a new speed zone and made drastic changes up and down in speed, first overshooting and then undershooting to get to the new speed limit). The latter strategy may take longer to reach the reference value, but with much smoother behavior and less resource wasting. These are the aspects of concern when analyzing control properties and respective assurances.

In the next section, based on the model, the controller and the control behavior definitions, we analyze different control strategies and the assurances they

provide with respect to control properties, the analysis of the open model properties, the design (or synthesis) of the controller, and then the analysis of the resulting closed loop model.

6 Assurances

Most physical processes, and many computational processes (especially complex ones), do not respond immediately to changes in their control inputs. Further, most systems do not respond so precisely that an exact correction can be made in a single step. Therefore, the design of a feedback system, whether it be a simple feedback loop or a complex SAS, must take into account the way the controlled process responds to changes in the control inputs. The characteristic responses constitute the basis on which system properties are built, and the assurances correspond to what (and in what extent) can be guaranteed about these properties.

This section focuses on *assurances* provided by feedback loops, and their correspondences in SAS domains; in other words, we focus on what we can be certain of when using feedback loops, and explore the correspondences of these assurances in SAS domains. Of course, ideally, it would be desirable that just by using feedback loops we would have granted assurances on desirable properties and optimum behavior in the controlled software applications. For instance, for regulating the throughput of the web service for computing definite integrals it would be desirable to have a controller that always (i) corrects errors fastly; (ii) reaches the reference level with precision and minimum error; (iii) spends the resources at the minimum required; and (iv) reaches stability after achieving the reference level. However, obtaining such a controller for a SAS system in general is not easy and constitutes precisely the main challenge.

In the following subsections we discuss different control strategies and the assurances they provide with respect to desirable behaviors. Then, we analyze the open model properties, the conditions that can help to design controllers that provide assurances on desirable properties, and then the analysis of the resulting closed loop model properties. Even though we do not provide a complete answer to the main challenge, this analysis helps to understand better the implied problems and sets the basis for defining the principles on which a software controlling theory should emerge.

6.1 Classic Control Strategies

From the discussions of previous sections and the questions posed in Sect. 3.1 we can infer that decisions about the control strategy are absolutely critical to the design of SAS. Classical control theory operates in settings that are simple enough that mathematical analysis is usually possible. SAS systems, on the other hand, are usually too complex to characterize completely. This makes the

design and validation of the control strategy even more important for this kind of systems.⁴

This section reviews common control strategies and discusses their application to SAS. For this application, we use an extremely simplified version of the WolframAlpha⁵ web services as the software to be controlled. Our web service only computes the integral of elementary functions (i.e., functions with exponentials, logarithms, radicals, trigonometric functions, and the four basic arithmetic operations). The service is implemented using an exact method for computing definite integrals of a given elementary function, that is, by computing the symbolic integration of the function and evaluating the result on the integration limits. This method is slow (i.e., worst case is exponential, even though for many cases is polynomial [7]), but exact.

6.1.1 On/Off Control. The simplest type of control is for a system where the control inputs of the managed system are either turned *on* or *off*. The most common example is the old-style household thermostat, where the thermostat turns the furnace *on* or *off*, as described previously (cf. Sect. 3). This simple strategy has significant drawbacks: it oscillates constantly (and hence is often called “bang-bang” control). This problem can be moderated by introducing hysteresis, that is, by delaying the control response until a modest overshoot has occurred.

To illustrate the application of the *on/off* control to our web service, assume we deploy the control method implementation both on a main server and on a spare server. Assume also that the required reference level is to maintain a given throughput of definite integral computation requests per time unit. Whenever the main server can service the requests maintaining the required throughput, the controller makes u to turn *off* the spare server. Otherwise, the controller makes u to turn it *on* and redirects requests to it. However, independent of whether the number of requests grows in greater magnitudes, the action of this type of controller is clearly limited to activating the spare server.

6.1.2 Proportional Control. This type of control acts by making the size of the adjustment on the control input proportional to the size of the tracking error. In the simple case, the constant of proportionality is called the gain. The proportional gain determines the ratio of output response to the tracking error. For instance, if the error term has a magnitude of 10, a proportional gain of 2 would produce a proportional response of 20. In general, increasing the proportional gain will increase the speed of the control system response. Even though this strategy means a fast response to eliminate the tracking error, a shortcoming of proportional control is that it generally is unable to eliminate the tracking

⁴ We often describe “self-adaptation” in terms of strategic changes to maintain the system behavior as close as possible to the reference value. Note that a system can also self-adapt by changing the control discipline. This is another reason to make explicit design decisions about the control strategy.

⁵ <http://www.wolframalpha.com>

error entirely, a phenomenon called proportional droop. Thus, one critical design task is gain estimation. Gains too high result in excessive oscillation, possibly never reaching convergence toward the reference level, while gains too low result in sluggish performance.

Revisiting the thermostat example, modern high-efficiency furnaces run at several power levels⁶: the level 1 is energy-efficient but produces lower temperatures (i.e., has lower gain), and the others gradually produce more heat at the expense of efficiency (i.e., have higher gains). For an automated furnace with eleven internal power levels (0 being the Off position, 10 the maximum power), covering a range between $18^{\circ}C$ and $27^{\circ}C$ the thermostat's ordinary operation uses the energy-efficient level in most cases, but at the expense of slower adjustment of temperatures in the living area. Assuming the thermostat has a proportional gain of 3, and the temperature error is of $3^{\circ}C$, the controller would produce a command control for selecting the ninth power level position in order to correct the temperature rapidly; notice that being this almost the maximum power level, maintaining it for too long would lead to significant overshoots, and therefore, correspondingly large oscillations.

For the application of this type of control to our web service, assume we deploy the implementing method on the servers of a cluster computing infrastructure, and that the required reference level is the same throughput as the defined for the *On/Off* control example. In this case, we estimate the gain as $2/100$ times the tracking error. That is, in short, and roughly speaking, the resulting command control on this strategy is $u(k+1) = \frac{2}{100} * e(k)$, being $e(k) = y(k) - y_r(k)$. Thus, if the error reaches a difference of 200 in the throughput, the controller response is to activate four spare servers and distribute the service requests among them with the purpose of maintaining the required throughput. Nonetheless, despite the fast corrective action, small tracking errors (i.e., below 50 in this case) produce no correction (i.e., evidencing proportional droop).

6.1.3 Integral Control. As analyzed in the proportional control, independent small tracking errors can produce null corrective actions. However, the accumulation of these errors sooner or later should produce a corrective action, possibly small, but significant enough. The accumulated errors is essentially the integral of the error (or sum for discrete systems), so this is called integral control. However, given that this strategy applied alone naturally results in very slow responses, a common control strategy is to combine an integral control term with a proportional one. This combination provides a fast response action, with elimination of remnant small tracking errors.

For the thermostat example, assume the reference level is set to $25^{\circ}C$, the controller having an integral coefficient of $3/4$, and the current tracked metric being $24.5^{\circ}C$. Having an integral term alone would yield a command control of $0.5 * 3/4 = 0.375$, that is, selecting the Off position for the furnace. After 3 more cycles with the same error (i.e., with null corrective actions) a first corrective action is produced by setting the furnace on the first position ($0.375 * 3 = 1.125$).

⁶ <https://www.ecobee.com>

However, as this power level setting is too low to achieve and maintain $25^{\circ}C$, the controller will have to keep accumulating errors until reaching the eighth power level, in which, after some more cycles, it will stabilize (recall from the initial specifications that the tenth power level achieves $27^{\circ}C$). Once achieved the reference level, the temperature error will be diminished to near zero, contributing in this amount to the integral term. Indeed, the main contribution of the integral term, besides eliminating remnant tracking errors, is to maintain achieved levels of control (i.e., the eighth power level) at a given point, in which other control terms (e.g., the proportional one) would be null. Thus, this term models inertial load effects inherent of physical systems.

For the application of the integral control to our web service, assume the same deployment and required reference level as in the corresponding proportional control example. Also, let's define the integral control term as $I(k+1) = I(k) + \frac{1}{5} * e(k)$, that is, $u(k+1) = I(k+1) = \sum_{j=1}^k \frac{1}{5} * e(j)$. Thus, an independent small tracking error $e(k)$ of 2 units produce no immediate corrective actions. Nonetheless, if this error is maintained, after three control cycles the accumulated error is of 6 units, and the integral control term reaches $6/5$. Then, the controller response is to activate one ($6/5 = 1.2$) spare server and distribute the service requests.

6.1.4 Derivative Control. It is important in many cases to predict or anticipate the software behavior trends. If we would know that in near future the tracking error is going to increase, then we would take a proactive action now. In many cases, performing the command u , which in many cases is a workflow, might take several minutes. This is called an execution lag. In these cases, the command will affect the output y not instantaneously but with a lag of several minutes. It is better in these situation to anticipate the *trend*. The simplest way to accomplish it is to determine the sign and magnitude of the error derivative (or difference for discrete systems, $e(k) - e(k-1)$), and calculate u as a function of this difference.

Applied to the thermostat example, assume the same reference level of $25^{\circ}C$, the controller having a derivative coefficient of $3/4$, and the current tracked metric being $24.5^{\circ}C$. After some minutes, the error trend will become worst, and thus, a derivative term alone would yield a command control of $(e(k) - e(k-1)) * 3/4$. When the error difference reaches $4/3$, the controller produces the first corrective action by setting the furnace on the first position ($(4/3) * (3/4) = 1$), which, depending on the control cycle duration, will correct the error trend the more or the less. In absence of other control terms, this power level is too low to achieve and maintain $25^{\circ}C$. Moreover, as this controller reacts to error differences, trying to nullify the error trend, it will never reach the reference level set.

For the application of this type of control to our web service, assume the same deployment and required reference level as in the corresponding proportional control example. Also, let's define the derivative control command as $u(k+1) = \frac{1}{100} * (e(k) - e(k-1))$. Notice that, given that this strategy is based on the error

trend, it should be used to complement other strategy: even if the error is large, but the trend is negative (e.g., $e(k) - e(k-1) = -100$) indicating that the error is diminishing, the controller response is to deactivate one spare server. However, if we combine it with the proportional strategy in our example, the command control would be $u(k+1) = \frac{2}{100} * e(k) + \frac{1}{100} * (e(k) - e(k-1))$. In this case, if the error presents a difference of 200 in the throughput, but the error difference trend is -100, the corrective action is to activate 3 spare servers (not 4, as the proportional control alone would yield in this case).

Example 2: A PID Controller for Cluster Utilization. Figure 8 shows a PID controller presented by Gergin *et al.* [18] for the control of a software cluster. The controlled variable is the cluster utilization, y , and the command, u , is the number of software replicas within the cluster. A PID controller computes the number of replicas u as a function of error $e = y - y_r$, where y_r is the setpoint utilization for the cluster. The PID controller amplifies the error (K_p term), integrates all past errors (K_i term) and anticipate the direction of the error (K_d term). The coefficients K_p, K_i, K_d are determined experimentally or based on the system model in such a way that some quality control metrics (overshoot, rising and settling time, steady errors) are achieved.




Fig. 8. A feedback loop with a PID controller for distributing the web service load in a cluster.

6.2 Control Theory to the Rescue of the MAPE-K Loop

From the above discussion, it is apparent that each control strategy poses different advantages and disadvantages (i.e., implying different properties). Moreover, some of the strategies can be combined, as in Example 2, being it usual to search for combinations that exploit the advantages of ones to compensate the disadvantages of others, given a particular control problem.

In addition to the ones discussed, there are many other control disciplines and strategies. However, the presented examples should be sufficient to show how critical it is to design the control strategy of an SAS system very carefully. It is important to notice that no other software engineering discipline forces attention to these matters. For instance, the well-known MAPE-K model, proposed for implementing autonomic computing and self-management computing

systems, does not forcefully raise these design issues. Of course, the design of the planning element of the model should address these matters, but its specification is purely functional, being concerned fundamentally with *what* is required in order to modify the system behavior. Our discussion argues for balancing the concerns, considering even more importantly the longer term effects of the proposed behavior modifications, that is, the *properties* implied by the chosen control strategy, over the specific technique or method for producing the modification.

Any discussion about the properties of the closed loop system should start with a discussion about the properties of the open loop system. Only certain properties of the open loop system will allow us to design and analyze the closed loop properties.

In this setting, it is important to notice that a properly achieved property effectively becomes an assurance for the desired system behavior. Even though the MAPE-K loop—and its further refinements—can be considered as an important contribution for the SAS engineering, we strongly believe that the lessons learned by control theory in the assurance of desired properties is the direction we need to follow in order to consolidate it as a milestone. The remainder of this section makes the case for this claim by showing how control considerations provide a basis for reasoning about the control properties of systems.

6.3 Properties of the Open Loop Model

One of the most important reasons for having a model is to study the properties of the system it models. In the open loop model, the most important properties are stability, observability and controllability.

6.3.1 Stability. In simple terms, stability means that for bound inputs (commands or perturbations), the system will produce bound state and output values. In the case of Example 1 (cf. Sect. 5.1), a stable system might mean that for keep-alive connection intervals of 10 minutes (u_2) and for a 100 connections (u_1), the response time of the server is guaranteed to be between 1 and 2 seconds. Formally, the stability is proven by necessary and sufficient conditions and, given a model, one can use Bode and Nyquist plots [3] to study the stability. Examples of stability studies in control of software and computing systems have been presented in [21,2]. Unfortunately, perturbations are the enemy of stability in open loop SAS, because the system is not set up to recognize how much the perturbations affect the system. If the open SAS is not stable, it can be stabilized through the design of controller. However, by doing a stability analysis of the open system, we understand the source of instability and design the controller appropriately. For example, it is known that a source of instability is a saturated queue. It can lead to performance instability, faults and crashes. In a complex SAS, we have to know which queue has the potential to be saturated by external disturbances so we design the controller to generate u that will avoid saturations.

6.3.2 Observability. Observability is the property of the model that allows to find, or at least estimate, the internal state variables of a system from the output variables. This property is important from a pragmatic point of view. In a real system, it is impossible, hard or impractical to measure all state variables. On the other hand, the commands and outputs are easier to measure. By knowing the model of the system in the form of equations (3)-(4), if the matrices A and C have are well behaved, one can compute x as a function of y . Formally, a model is observable if the observability matrix

$$O = [C \ C A \ C A^2, \dots, C A^{n-1}] \quad (9)$$

has the rank n , where n is the number of the state variables. Examples of web service observability from a performance point of view can be found in [9]. Examples of how to estimate software performance parameters for applications deployed across multi-tiers and using Kalman filters are presented in [52].

6.3.3 Controllability (or Reachability). The concept of controllability (or state controllability) describes the possibility of driving the open system to a desired state, that is, to bring its internal state variables to certain values [9]. Of course, the system can be driven to a certain state by changing the command variables u . In the case of Example 1 (cf. Sect. 5.1), we should be able to find the values of u_1 and u_2 that will bring the server utilization (x_1) and memory usage (x_2) to 50%, for example. Like observability, the formal proof is given by the structure of the matrices that make the equation (3)-(4): the system is controllable if the controllability matrix,

$$S = [B \ AB \ A^2B, \dots, A^{n-1}B] \quad (10)$$

has the rank n , where n is the number of state variables.

If observability is not a necessary condition for designing a controller for SAS, the Controllability property is. Even we do not have a explicit open loop model, a qualitative analysis should be done. If the the external disturbances, their frequency and amplitudes are known, do we have enough control inputs to compensate those perturbations?

6.4 Complex Open SAS and Model Composition

It is always the case that software systems are made of many components, interconnected together. The components might have different life cycles and be under different administrative domains. It is therefore imperative to answer the question: What are the properties of the overall open system (controllability, observability, stability) if we know the properties of the individual components?. If the components are described by models, models like the ones presented in equations (3)-(4), then we have a method to answer that question. Models like (3)-(4) can be composed, that is, if they are connected in series, parallel or through a feedback loop, a resultant composed model can be derived from the individual

models. The composed model will have the same structure as the one presented in equations (3)-(4), but the matrices A, B, C, D will be different but determined analytically from the individual models. Two services connected in series means that the output of the first one is the input of the second one (inputs and outputs have the control theoretic meaning). Parallel composition means that the same input reaches two different models and the output of the models will be summed up and use (eventually) as an input in other model. Feedback composition refers to a composition in which the output of a model is brought back and subtracted from the input of another model, as illustrated in equations (7)-(8).




Fig. 9. Model composition for web services.

Example 3: Web service composition. In [41], the authors showed, based on web service performance case studies, how web services models can be composed in series and in parallel. Figure 9 illustrates the composition process: a set of web services are interconnected in an workflow; if we consider their performance, then the inputs (commands), u , are arrival rates and the outputs, y , are throughputs. For services in series, the throughput of the service on the left becomes the arrival rate for the service on the right. In Figure 9a, WS_2 and WS_3 are connected in parallel, which means $u_2 = u_3 = y_1/2$ and $u_4 = y_2 + y_3$. By composing the models using their property of linearity, we can get a model of the entire system that is similar to the one presented in equations (3) and (4). It is interesting to note that by composing web service models, the properties of the individual models (stability, controllability, observability) are not necessarily transferred to the composed model [9]. Model composition through feedback loops was briefly described by equations (7)-(8).

6.5 Properties of the Closed Loop Model

Once an open SAS has been analyzed and the designer has a good understanding of its stability, observability and controllability, a controller can be designed. When an explicit model of the open loop is available, the closed loop model can be synthesized mathematically to achieve the properties the designer wants. As equations (7) and (8), suggest, A_c can be shaped and therefore the properties of closed loop model can be shaped. In general, the controller is designed to achieve some goals. The most frequent ones are described below.

6.5.1 Stability refers to whether control corrections move the open system state, over time, toward the reference stepping. A system is unstable if the control causes overcorrections that never decrease, or that increase without limit. Instability can be introduced by making corrections that are too large in an attempt to achieve the reference level quickly. This leads to oscillating behaviors in which the system overshoots the reference value alternately to the high side and the low side. The opposite problem from stability is sluggish performance: if the control corrections are too small, it may take a very long time for the system to reach the reference value. In designing a control algorithm, the goal is to make the largest control correction that does not make the system unstable. Of course, the more we know about the way the system responds to changes in the tuning parameters (i.e., the dynamics of the system), the easier it is to accomplish this goal.

In a model of the open loop exists, the controller is designed to correct any instability of the controlled software system, caused by perturbations or intrinsic properties of the software. The stability property then is defined for the entire composed system, (considering y_r or p as input and y as output). Commonly, the stability is designed in the frequency domain (Z or S transforms) and implies solving an algebraic equation that will give the controller equations. Recent work on controller design for stability is presented by Cortellesa *et al.* [2] for a performance case study.

6.5.2 Robustness or Robust Stability. A special type of stability in control theory is the *robuststability*. Robust stability means that the closed loop system is stable in the presence of external disturbances, model parameters and model structure uncertainties. If we refer to models (3)-(4), external disturbances uncertainties refer to wide variations of p in amplitude and frequency. Consider an application in which the disturbance is the external load, that is the number of users and the frequency of user requests. The number of users can vary from 100 to 1 million and their requests frequency from 0.1 to 40 requests per second. Can we design one single controller that maintain the performance of the system (response time) below a certain threshold in the presence of those large variations? Model parameter uncertainties in equations (3)-(4) refer to errors in identifying A,B,C,D,F. How can we design a controller that can achieve stability when there are large errors in those parameters? Model structure uncertainties refer to working with models (3)-(4) that miss dynamics of the system (like ignoring important state variables). Can the controller provide stability in this situation? In general, can we design one static controller that assure stability within quantifiable uncertainties bounds?

In control theory, there are several design techniques that help a designer achieve robust stability (e.g. H_∞ , loop shaping, quantitative feedback theory or QFT, etc.). When SAS can be described by a system of equations (1)-(2) or (3)-(4), then we should consider those control techniques. For more complex SAS, it is probably feasible to consider Adaptive and/or Hierarchical Control

where multiple controllers are synthesized or tuned dynamically to face wide disturbances.

6.5.3 Performance In addition to stability, a controller has to achieve certain performance metrics: rise time, overshoot, settling time, steady error or accuracy [3,6], as illustrated in Example 2. For a system subject to a PID controller, such as the one treated in that example, all of these metrics depend on the selected values for coefficients K_p , K_i , K_d . Figure 10 indicates these metrics, and illustrates the cluster utilization of a web service when perturbed by a workload change.




Fig. 10. Quality metrics for a controlled system [6]: rise time (A); overshoot (B); settling time (C); steady error or accuracy (D). Fine-tuned values for a PID controller, K_p , K_i , K_d , result in a smooth behavior reaching the reference level with accuracy and acceptable settling time (curve labeled (E)). Large values for K_p (and in lesser extents K_i and K_d), produce corresponding large oscillations around the reference level (curve labeled (F)). Small values usually result in a sluggish behavior (curve labeled (G)).

Assume the reference level is the horizontal continuous line in the figure. The perturbation, that affects the system at the origin of the time axis, drops the utilization to close to zero. A good controller (PID in this case) will compute a command u (the number of replicas in the cluster) that will bring the web service behavior back to normal in a short time (rise time) by varying the number of replicas in the cluster. Very likely, the system will be under- or over-provisioned, but the over- and undershooting must be kept small since it can affect other attributes, like response time and stability. Naturally, a few oscillations around the reference level are acceptable, but after a time (settling time), preferably

short, the cluster utilization must reach a stable behavior (cf. curve labeled (E) in the figure). The figure also presents two examples of unstable and unacceptable control behavior: one which keeps oscillating around the reference level, and one which never reaches it (cf. curves labeled (F) and (G), respectively).

6.5.4 Optimality or Linear Quadratic Regulator (LQR). In many cases a controller has many goals, especially when there are multiple reference points, multiple input variables and multiple outputs. For example, a cruise controller should track the set speed but also minimize the consume of fuel. Usually the goals are conflicting. In control, the common way to treat this is to define an objective function across different goals. The goals can have different weights. Most commonly, the goals are defined across the state, output and control variables in the form of a quadratic function. The role of the controller is to find the commands u that minimize this quadratic function.

This type of controllers with multiple goals are especially important for SAS systems, especially when the goal of the controller is a combination of conflicting goals: maintaining performance within certain bonds, optimization of the use of resources for cost and energy, maximizing user experience, etc. For example, Ghanbari *et al.* [20], for a SAS described by equations similar to (3)-(4) defined a quadratic function over the commands u (number of service instances) and the deviations from the setpoint y_r (or SLA violations) for a cost and performance optimization in cloud, with Q and R matrices that penalize some elements of y and u :

$$J = \sum_{t=0}^k (y - y_r)^T Q (y - y_r) + u^T R u \quad (11)$$

They showed that a controller can be designed to optimize the sum of two conflicting goals (performance and cost).

In control theory, there are algorithms that assure the optimality (or sub-optimality) of the solution and those can applied to SAS that can be described by equations (3)-(4). However, the definition of the objective function, practical ways to establish the weights of the goals, are subject of further studies.

6.5.5 Look-Ahead Optimality or Model Predictive Optimization (MPC). In control theory, MPC optimizes a set of goals over a future time horizon H . This makes sense for economic goals, where the revenue, profit, etc.. are calculated over a time horizon. The idea behind these controllers is to make strategic decisions, that is to make a current decision by taking into account the long term goals. In other words, a decision may be suboptimal now but may prove optimal in long term. The goal of the controller in this case is to find $u(k), u(k+1), \dots, u(k+H)$, with k the current time that optimizes a function similar to LQR over the future horizon H .

This type of optimization is specially important to complex SAS that have multiple design goals and have to prove efficient over long periods of time. Examples of such optimizations, mostly in server provisioning for performance,

cost and power can be found in [19,51,27]. The optimization problem for MPC problem looks like:

$$J = \sum_{t=k}^{k+H} (y - y_r)^T Q (y - y_r) + u^T R u \quad (12)$$

6.6 Open Questions

Control theory models provide fundamental assurances that can be leveraged in SAS adaptation mechanisms. The open and closed loop properties, even if not captured in formalized models for SAS, must be considered by SAS designers along the system's engineering life cycle. Villegas *et al.* comprehensively define these properties in the context of SAS systems [46]. Implied questions are, for example:

- Above all, how to determine whether a given SAS will be stable?
- How quickly will the system respond to a change in the reference value? Is this fast enough for the application? Control of a supersonic aircraft requires much faster response times than control of a slow-moving ground vehicle. Are there lags or delays that will affect the response time? If so, are they intrinsic to the system or can they be optimized?
- Are there some bounds of the external disturbances we know and what to design the SAS for?
- Can we design one robust controller to achieve robust stability or use Adaptive or Hierarchical Control?
- How accurately shall the system track the reference value? Is this good enough for the application?
- How much resources will the system spend in tracking the reference value? Can this amount be optimized? What is more important, minimizing the cost of resources or tracking the reference values? Can this tradeoff be quantified?
- It is very likely that multiple control inputs are needed to achieve robust stability. How will the control algorithm select the corrections to the control inputs?

7 Assurance Challenges for Self-Adaptive Software

From the examples presented in the previous sections, it should be clear that the concepts and principles of control theory and the assurances they provide, at least in abstract, can be applied to a large class of problems in SAS. In this application, we must consider at least two scenarios.

The first, in which it is possible to apply control theory *directly* to SAS, that is, by building a mathematical model (e.g., a set of equations (1)-(2)) for the software system behavior, and applying control theory techniques to obtain

properties such as stability, performance, and robustness (or robust stability). These properties automatically grant corresponding assurances about the controlled system behavior. This scenario requires two main necessary conditions: the problem must refer to quantitative metrics, and must be simple enough to be modeled as a linear set of equations relating control inputs, outputs and state variables in the open and controlled system. Nonetheless, there are still no clear guidelines about the limitations of control theory as directly applied to SAS systems in the general case.

The second scenario arises when it is not feasible to build a reasonably precise mathematical model, but instead, it is possible to create an approximated operational or even qualitative model of the SAS behavior. In this case, the formal definitions and techniques of control theory may not apply directly, but their comprehension can guide the sorts of questions the designer should answer and take care of. In other words, even if a mathematical model is not available, control theory provides validation obligations and elements that can be exploited to model the phenomena to be controlled.

In this section we highlight the main challenges a designer faces when dealing with the problem of designing SAS systems, and extrapolate the questions and concepts from the control theory approach into the general case of SAS.

7.1 Modeling Challenges

In both of the aforementioned scenarios, and based on the specification of the system properties to achieve, the main modeling challenges the SAS system designer should solve are:

- How to implement and instrument the control commands u in the software system to be controlled? For instance, Hellerstein *et al.* control the number of active threads in the Apache thread pool by varying the respective parameter [21]. Even though this can appear as natural, it is not clear whether the Apache server was intentionally designed with this parameter to be modified automatically, but rather to be configured manually at installation time. For any other given software system to be controlled, what parameter should be defined, and affecting what components and functionalities? Is the provision of one (or several) parameters the best alternative for implementing the control commands? Other options, very different to varying numerical parameters, are modifying the software structure in critical parts (i.e., using domain-specific design patterns), as in [43]; or using machine-learning strategies to modify the system behavior, as in [13].
- How to obtain an operational understanding of the effects that control commands will have on the controlled system? That is, a variation of one unit on the control parameter affects in what magnitude the number of related components and the behavior of their functionalities?
- Given a set of control commands and the operational understanding of their effects on the system behavior, how to refine these commands to refine also

the granularity of their effects, in order to improve the accuracy controllability?

- How many cycles of control would have to be applied to reach the reference level fastly? how to determine the size of the change to be applied in each control cycle without producing excessive overshooting?

7.1.1 Model Identification. The design of a feedback system, whether it be a simple feedback loop or a complex SAS, must consider the behavioral characteristics of the target system, and the characteristic response of the controlled process to unitary disturbances.

The Core Phenomena to Control. In control theory, differential (or difference) equations are the models on which principles and properties of the phenomena to control are described and analyzed. Depending on the behavioral characteristics of the target system, a controller can be defined to make the system behave as desired. The desired behavior is usually specified by either providing a reference input the system should follow (e.g., PID controllers), or as an optimization problem (e.g., Model Predictive Control) [4,34]. The characteristics of a controller are usually defined by adjusting its parameters, which have special significance to generate the signals that will adapt the system depending on how far measured outputs are from the corresponding reference inputs.

Indeed, controllers are designed as parametric functions that generate the control signals (“knobs”) that drive the target system towards accomplishing its goals. In software systems the identification of the core phenomena to control is typically a complex task. In contrast to physical systems, software systems still lack general methods to model the multi-dimensional and non-linear relationships between system goals and adaptation mechanisms [46,37]; moreover, the problem and solution spaces can differ drastically (e.g., in the web service example for computing integrals of elementary functions, calculating $\int \frac{x}{\sqrt{x^4+10x^2-96x-71}} dx$ takes polynomial time, but only changing the constant 71 by 72, that is, $\int \frac{x}{\sqrt{x^4+10x^2-96x-72}} dx$, takes exponential time, as analyzed in [7]). This simple example illustrates how complex can be the problem space, on which, for SAS systems, tracked metrics depend upon, directly and explicitly.

Furthermore, an adaptation mechanism can reconfigure the software structure by applying domain-specific architectural patterns with the goal of improving the system performance, as realized in [43]. However, it is still an open challenge to model the exact effect of the structural and behavioral elements introduced by a pattern in the system performance.

In general, the analysis of the system model should determine whether the “knobs” have enough power (command authority) to actually drive the system in the required direction. Many other research questions remain open in the identification and modeling of the core phenomena to control in software systems. For example, how to model explicitly and accurately the relationship among system goals, adaptation mechanisms, and the effects produced by controlled variables when the control variables are as complex as a pattern and workflow?

Can we design software systems having an explicit specification of what we want to assure with control-based approaches? Can we do it by focusing only on some aspects for which feedback control is more effective? Can we improve the use of control, or achieve control-based design, by connecting as directly as possible some real physics inside the software systems? How far can we go by modelling SAS systems mathematically? What are the limitations?

Sampling Period. One important decision when building a model is the sampling rate, that is, the frequency with which the states, outputs and commands are monitored and processed. In terms of our models (1)-(2), what is the difference between time k and $k + 1$? The meaning of “1” above was “the next sample” but after how many microseconds or minutes is the next sample? For models such as (1)-(2) and for quantitative software qualities that can be approximated with continuous signals (performance, cost, energy), Nyquist-Shannon sampling theorem [39] can provide some practical guidelines. Simplifying, the theorem says that if the highest signal harmonic one wants to reconstruct from the sampling data has h Hertz, one should sample with a minimum frequency of $2h$ Hertz.

In Zheng *et al.* [53], the authors showed that for the specific performance model they were building, a sampling rate of 5 minutes was appropriate. Note that in [53] the authors were interested in mean response time and server provisioning decisions, which are rare control inputs. For higher frequency models and decisions (like changing the number or threads), that sampling rate is not enough. Further, Solomon *et al.* [42] showed how to find the sampling rate from the cutoff frequency. The cutoff frequency is the frequency of the command u that has no effect on the system (one can imagine that, for the example described by equations (5)-(6), by adding and removing an http connection at very high frequency, in average, has limited effect on the response time because the connection cannot be used). Solomon *et al.* use Bode plots to analyze models like (3)-(4) in the frequency domain and then determine the cutoff frequency. The sampling rate is then twice the cutoff frequency. It is obvious that the sampling rate depends of what we monitor and control, but it is not clear how to determine it in many cases and how to guarantee its correctness. Moreover, in SAS, most events are discrete and makes more sense to use the number of events (like number of failures, number of requests, etc..) instead of a sampling rate.

7.2 Composition and Incrementality: V&V Tasks

Considering the two scenarios we introduced for applying control theory to the engineering of SAS, validation and verification (V&V) acquires significant relevance, especially for the scenarios in which obtaining mathematical models is not feasible. If it is not possible to guarantee desirable properties based on the control theory principles and techniques, at least we should consider introducing validation and verification (V&V) tasks on critical properties, an aspect that is commonly omitted in self-adaptive software proposals [46]. Nonetheless, performing V&V tasks (e.g., model checking) over the entire system—at runtime, to guarantee desired properties and goals, is often infeasible due to prohibitive

computational costs. Therefore, one fundamental requirement for the assurance of SAS systems is for these V&V tasks to be composable and able to be applied incrementally along the adaptation loop, among other conditions, as analyzed in [44].

In these regards, relevant research questions include: which V&V tasks can guarantee which control properties, if any, and in what extent? are stability, accuracy, overshoot, settling-time and other properties composable (e.g., when combining control strategies which independently guarantee them)? what are suitable techniques to realize the composition of V&V tasks? what approaches can we borrow from testing? how can we reuse or adjust them for the assurance of SAS systems? Regarding incrementality: in which cases is it useful? How can incrementality be realized? How to characterize increments, and their relationship to system changes?

7.3 Timing Issues and Lags

As discussed previously, a major challenge in designing a control algorithm is to determine what is the largest control correction that does not make the system unstable. In addition to stability, SAS is affected by specific characteristics of the system dynamics, such as lags (when the system responds only slowly to a change of tuning parameter), delays (when it takes time before a change takes any effect, such as startup time when adding a server), and differences between transient response (when response to a change in environment decays over time) and steady-state response (when the change is enduring). If the open and closed SAS are described by equations like (3), the time lag should be included in the equations. The main problem in SAS is that lags are highly variable and it is hard, if not impossible, to bound.

Another difference is that, in contrast to physical plants, behavior load in software systems sometimes presents no inertial effect. That is, in the room temperature control case, the trend of the difference between the reference level y and the measured output y_r usually augments or diminishes quite slowly. However, in software systems, all of the service requests can disappear instantly, for instance if all the users cancel their requests, or the internet connection of the web server is interrupted. Besides, while most physical processes do not respond immediately to changes in the control parameters, most software systems do not respond so precisely to control changes.

Even in the absence of a mathematical model, it is clear that the designer needs a good understanding of at least the direction of change that will result from a given control input, and preferably a sense of its magnitude. Since SAS systems often have more than one control input, the absence of a mathematical model makes understanding the relations among the inputs, disturbances, and lags more difficult, but no less important.

7.4 Challenges in Control Strategies Design

It is clear from the above discussion that decisions about the control strategy are absolutely critical to the design of SAS. Classical control theory operates in settings that are simple enough that mathematical analysis is often possible. SAS, on the other hand, are usually too complex to characterize completely, are highly non-linear and time variant. This makes the design of the control strategy even more critical.

We often describe control in terms of strategic changes to the reference value or the level of external disturbances. It is almost impossible in SAS to design a controller that can work well with all possible values of references or disturbances. This is another reason to make explicit design decisions about the control strategy. When the system is time variant (that is, the matrices A, B, C, D, F in equations (3)-(4) are time dependent) or the reference and disturbances change over large ranges, the design goals are achieved by engineering an adaptive controller that is tuned on-line (adaptive control), based on the operational point of the system. Adaptive control can be achieved using various strategies such as Model Identification Adaptive Control (MIAC) or Model Reference Adaptive Control (MRAC) [30]. Those strategies are paramount to SAS.

7.4.1 Modeling External Disturbances and Uncertainties. In control, the feedback loop is designed with a knowledge of the nature of the external disturbances, their characteristics and bounds. Furthermore, the design of the feedback system relies on sensed values about internal state or system behavior. This design must consider explicitly how accurately those sensed values actually represent the true state of the system. In a cruise control example, the control engineer knows that the speed of the car (the controlled output) is altered by friction, road conditions, hills and valleys, and all those can be captured in a model (like the term $F * p(k)$ in equations (3)-(4)). Once a set point is set (e.g, 100km/hour) the controller has to respond adequately (i.e., with certain accuracy and overshoot) to changes in perturbations. These changes refer to both amplitude (the slope of the hill) and frequency (alternations of hills and valleys). Similar considerations are taken into account for thermostat control, autopilots, etc. In SAS, external disturbances are harder to identify and model. In addition, the most likely to change in SAS are the external disturbances, not the reference values, so we have to design feedback loops to respond well to external disturbances. Consider an application deployed in a public cloud or an application using services provided by third parties. Assuming a feedback loop that maintains the response time at a setpoint, what are the external disturbances that affect the response time? We can make some assumptions about application load (number of users, their distributions and requests during a day) but what about the external disturbances affecting the cloud or third party services? Indirectly those disturbances are going to affect our application response time. How do we identify all those disturbances, their amplitude and frequency, how do we model or take their influence into account when we assure our feedback loop?

7.4.2 Complex Reference Values. It is very tempting to set a reference level as a combination of goals. For example, to keep wait time below some threshold and energy consumption below some other threshold. Selecting a composite reference point exposes the risk of creating a situation in which one term of the reference value calls for increasing the tuning parameter and another term calls for decreasing the reference parameter. For example, if both wait time and energy consumption are limited, the former might lead to activating a server while the latter demands deactivating a server. In the case of the home thermostat, maintaining a comfortable temperature and saving energy consumption are clearly two goals that may lead to a conflictive situation. We must at minimum be aware of these conflicts; even better, we should refrain from creating them. In any event, thinking carefully about such conflicts can help avoid or manage the resulting complexity. One way to solve this is to use just one variable as a set point and use an optimization function like LQR to mediate among the other variables.

7.4.3 Synergy Between Control and Software Engineering. Despite the efforts to apply control theory to the engineering of SASs, evidenced for instance in [21,33,8,15,14,16], the assurance of SASs still demands effective ways of integrating control theory and software engineering. Research questions to advance towards this direction include: what are the best practices from both sides that can be exploited in the assurance of SAS systems? How to construct reliable controllers leveraging formal model techniques used in software engineering? Can we define and/or use anti-patterns in control-based assurance of SAS systems? How to apply off-line V&V mechanisms at runtime? How to close the semantic gap between control in physical systems and control in software systems? Can we define a model in the middle of these two worlds? Would these models be domain or problem specific? Can we characterize (at least some of) these problems and corresponding suitable models? How can we educate software engineers in the application of control to the design of software systems?

7.4.4 Management of Viability Zones. A viability zone of a SAS system can be defined as the set of possible states in which the system operation is not compromised with respect to a desired property [5]. Therefore, a SAS system may have more than one viability zone: at least one for the managed system and other for the controller or adaptation mechanism. The level of assurance of a SAS system may then be judged by taking into account the state of the system with respect to its viability zone(s). A viability zone can be characterized in terms of the properties to be satisfied, the quality attributes that describe them and corresponding desired values [47]. Moreover, viability zones can change according to variations in relevant context. Feedback control may also help in the specification of viability zones for SAS systems. Properties to be satisfied define the dimensions of a viability zone. These properties are usually characterized in terms of quality attributes that correspond to reference inputs (y_r) defining the boundaries of viability zones. The goal of adaptation mechanisms (commands u)

is to maintain the managed system within its viability zone(s). Viability zones may be also dynamic, not only because relevant context and desired properties may change over time due to uncertainty, but also when the adaptation goal is optimization.

The definition of viability zones is not a trivial task, particularly because desired properties are usually characterized in terms of several variables (i.e., quality attributes). Relevant research questions in this regard include: how many viability zones are required for the assurance of a particular SAS system? Does each desired property require an independent viability zone? How to manage trade-offs and possible conflicts among several viability zones?

The dynamic nature of viability zones is also a relevant research challenge in the assurance of SAS systems. Particularly, because it implies adjusting the domain coverage not only of design and adaptation realization but also of V&V tasks. Open research questions in this aspect include: what are runtime models that can be used for the incremental and dynamic derivation of software artifacts for implementing V&V tasks? How to maintain the causal connection between viability zones, adapted system, and its corresponding V&V software artifacts? How to adapt these artifacts at runtime?

8 Conclusions

In this chapter, we explored the control theory applications to self-adaptive software (SAS) design. We started with an overview of the SAS and control theory concepts, summarized the main research results of control theory applied to software and computing in general and then we focused on control assurances. We described the properties of open loop models, the control strategies and goals, and then we elaborated on assurance challenges.

There are several conclusions we can draw from this exploration. Control theory can be applied as is to a subset of SAS and for a subset of quality attributes, including performance, cost, reliability, energy consumption. Research results suggest that it is feasible to design a controller for relatively simple use cases, especially for single input/single output systems and for time invariant cases. Even for these metrics the research is still open, and more complex case studies are needed to draw general conclusions. To apply control theory, a specific type of model is needed, linear or non-linear, that captures the outputs, control inputs, states, and eventually external perturbations of the open loop SAS. With this model, a controller can be synthesized.

Nonetheless, many complex SAS are time variant and multidimensional, with multiple input and output control variables. For these cases, more advanced control theory concepts such as Model Identification Adaptive Control (MIAC) are needed, where the model is identified at runtime and a controller is synthesized periodically. In even more cases, an explicit model for SAS cannot be constructed. Nevertheless, the control theory general principles are still a good guideline to follow. The designer needs to understand the properties of open systems: inputs, outputs, perturbations, states, and how these influence each other. Properties

such as stability and controllability should be investigated and understood before starting to design a controller (or adaptive/autonomic manager). The controller should be designed with some goals in mind: optimization, stability, performance and accuracy, and then tested and verified. A final conclusion is that the field of control theory applied to SAS and the related assurance challenges are still open, and constitute important research areas in the engineering of this kind of software systems.

References

1. Abdelzaher, T., Stankovic, J., Lu, C., Zhang, R., Lu, Y.: Feedback performance control in software services. *Control Systems, IEEE* 23(3), 74–90 (June 2003)
2. Arcelli, D., Cortellessa, V., Filieri, A., Leva, A.: Control theory for model-based performance-driven software adaptation. In: *Proceedings of the 11th International ACM SIGSOFT Conference on Quality of Software Architectures*. pp. 11–20. ACM (2015)
3. Åström, K.J., Murray, R.M.: *Feedback systems. An introduction for scientists and engineers* (2008)
4. Åström, K., Wittenmark, B.: *Adaptive Control*. Addison-Wesley series in electrical engineering : control engineering, Addison-Wesley (1995)
5. Aubin, J.P., Bayen, A.M., Saint-Pierre, P.: *Viability theory: new directions*. Springer Science & Business Media (2011)
6. Balzer, B., Litoiu, M., Müller, H., Smith, D., Storey, M.A., Tilley, S., Wong, K.: 4th international workshop on adoption-centric software engineering. In: *Proceedings of the 26th International Conference on Software Engineering*. pp. 748–749. ICSE 2004, IEEE Computer Society, Washington, DC, USA (2004)
7. Bronstein, M.: Integration of Elementary Functions. *Journal of Symbolic Computation* 9(2), 117 – 173 (1990)
8. Brun, Y., Serugendo, G.D.M., Gacek, C., Giese, H., Kienle, H., Litoiu, M., Müller, H., Pezzè, M., Shaw, M.: Engineering self-adaptive systems through feedback loops. In: *Software engineering for self-adaptive systems*, pp. 48–70. Springer (2009)
9. Checiu, L., Solomon, B., Ionescu, D., Litoiu, M., Iszlai, G.: Observability and controllability of autonomic computing systems for composed web services. In: *2011 6th IEEE International Symposium on Applied Computational Intelligence and Informatics (SACI)*. pp. 269–274. IEEE (2011)
10. Cheng, B.H., Lemos, R., Giese, H., Inverardi, P., Magee, J., Andersson, J., Becker, B., Bencomo, N., Brun, Y., Cukic, B., Marzo Serugendo, G., Dustdar, S., Finkelstein, A., Gacek, C., Geihs, K., Grassi, V., Karsai, G., Kienle, H.M., Kramer, J., Litoiu, M., Malek, S., Mirandola, R., Müller, H.A., Park, S., Shaw, M., Tichy, M., Tivoli, M., Weyns, D., Whittle, J.: *Software Engineering for Self-Adaptive Systems: A Research Roadmap*. In: *Software Engineering for Self-Adaptive Systems*. pp. 1–26. LNCS, Springer-Verlag (2009)
11. Dahm, W.: *Technology Horizons a Vision for Air Force Science & Technology During 2010-2030*. Tech. rep., U.S. Air Force (2010), <http://www.af.mil/information/technologyhorizons.asp>
12. Diao, Y., Gandhi, N., Hellerstein, J.L., Parekh, S., Tilbury, D.M.: Using mimo feedback control to enforce policies for interrelated metrics with application to the apache web server. In: *Network Operations and Management Symposium, 2002. NOMS 2002. 2002 IEEE/IFIP*. pp. 219–234. IEEE (2002)

13. Elkhodary, A., Esfahani, N., Malek, S.: Fusion: a framework for engineering self-tuning self-adaptive software systems. In: *Procs. of 18th ACM Intl. Symp. on Foundations of Software Engineering*. pp. 7–16. FSE '10, ACM (2010)
14. Filieri, A., Ghezzi, C., Leva, A., Maggio, M.: Self-adaptive software meets control theory: A preliminary approach supporting reliability requirements. In: *Proceedings of the 2011 26th IEEE/ACM International Conference on Automated Software Engineering*. pp. 283–292. IEEE Computer Society (2011)
15. Filieri, A., Ghezzi, C., Leva, A., Maggio, M.: Reliability-driven dynamic binding via feedback control. In: *Software Engineering for Adaptive and Self-Managing Systems (SEAMS), 2012 ICSE Workshop on*. pp. 43–52. IEEE (2012)
16. Filieri, A., Hoffmann, H., Maggio, M.: Automated design of self-adaptive software with control-theoretical formal guarantees. In: *Proceedings of the 36th International Conference on Software Engineering*. pp. 299–310. ACM (2014)
17. Filieri, A., Maggio, M., Angelopoulos, K., D'Ippolito, N., Gerostathopoulos, I., Hempel, A.B., Hoffmann, H., Jamshidi, P., Kalyvianaki, E., Klein, C., Krikava, F., Misailovic, S., Papadopoulos, A.V., Ray, S., Sharifloo, A.M., Shevtsov, S., Ujma, M., Vogel, T.: *Software Engineering Meets Control Theory*. In: *Proceedings of 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2015)*. pp. 71–82. IEEE Press (2015)
18. Gergin, I., Simmons, B., Litoiu, M.: A decentralized autonomic architecture for performance control in the cloud. In: *2014 IEEE International Conference on Cloud Engineering (IC2E)*. pp. 574–579. IEEE (2014)
19. Ghanbari, H., Litoiu, M., Pawluk, P., Barna, C.: Replica placement in cloud through simple stochastic model predictive control. In: *2014 IEEE 7th International Conference on Cloud Computing (CLOUD)*. pp. 80–87. IEEE (2014)
20. Ghanbari, H., Simmons, B., Litoiu, M., Iszlai, G.: Feedback-based optimization of a private cloud. *Future Generation Computer Systems* 28(1), 104–111 (2012)
21. Hellerstein, J.L., Diao, Y., Parekh, S., Tilbury, D.M.: *Feedback control of computing systems*. John Wiley & Sons (2004)
22. IBM Corporation: *An Architectural Blueprint for Autonomic Computing*. Tech. rep., IBM Corporation (2006)
23. Janert, P.K.: *Feedback Control for Computer Systems*. O'Reilly Media, Inc. (2013)
24. Kalyvianaki, E., Charalambous, T., Hand, S.: Self-adaptive and self-configured cpu resource provisioning for virtualized servers using kalman filters. In: *Proceedings of the 6th international conference on Autonomic computing*. pp. 117–126. ACM (2009)
25. Kephart, J.O., Chess, D.M.: The Vision of Autonomic Computing. *IEEE Computer* 36(1), 41–50 (2003)
26. Kramer, J., Magee, J.: Self-managed systems: an architectural challenge. In: *FOSE '07: 2007 Future of Software Engineering*. pp. 259–268. IEEE Computer Society, Washington, DC, USA (2007)
27. Kusic, D., Kephart, J.O., Hanson, J.E., Kandasamy, N., Jiang, G.: Power and performance management of virtualized computing environments via lookahead control. *Cluster computing* 12(1), 1–15 (2009)
28. Laddaga, R.: Guest Editor's Introduction: Creating Robust Software Through Self-Adaptation. *IEEE Intelligent Systems* 14(3), 26–29 (May 1999)
29. Laddaga, R.: Active software. In: *Proceedings of the First International Workshop on Self-adaptive Software*. pp. 11–26. IWSAS 2000, Springer-Verlag New York, Inc., Secaucus, NJ, USA (2000)

30. de Lemos, R., Giese, H., Müller, H.A., Shaw, M., Andersson, J., Baresi, L., Becker, B., Bencomo, N., Brun, Y., Cukic, B., Desmarais, R., Dustdar, S., Engels, G., Geihs, K., Goeschka, K.M., Gorla, A., Grassi, V., Inverardi, P., Karsai, G., Kramer, J., Litoiu, M., Lopes, A., Magee, J., Malek, S., Mankovskii, S., Mirandola, R., Mylopoulos, J., Nierstrasz, O., Pezzè, M., Prehofer, C., Schäfer, W., Schlichting, R., Schmerl, B., Smith, D.B., Sousa, J.P., Tamura, G., Tahvildari, L., Villegas, N.M., Vogel, T., Weyns, D., Wong, K., Wuttke, J.: Software Engineering for Self-Adaptive Systems: A Second Research Roadmap. In: de Lemos, R., Giese, H., Müller, H., Shaw, M. (eds.) *Software Engineering for Self-Adaptive Systems 2*. LNCS, vol. 7475, pp. 1–32. Springer (2013)
31. Lennart, L.: Perspectives on system identification. *Annual Reviews in Control* 34(1), 1–12 (2010)
32. Müller, H., Villegas, N.: Runtime evolution of highly dynamic software. In: Mens, T., Serebrenik, A., Cleve, A. (eds.) *Evolving Software Systems*, pp. 229–264. Springer Berlin Heidelberg (2014)
33. Müller, H., Pezzè, M., Shaw, M.: Visibility of control in adaptive systems. In: *Proceedings of the 2nd international workshop on Ultra-large-scale software-intensive systems*. pp. 23–26. ACM (2008)
34. Murray, R.M.: Control in an information rich world: report of the panel on future directions in control, dynamics, and systems. SIAM (2003)
35. Northrop, L., Feiler, P., Gabriel, R., Goodenough, J., Longstaff, T., Kazman, R., Klein, M., Schmidt, D., Sullivan, K., Wallnau, K.: *Ultra-Large-Scale Systems—The Software Challenge of the Future*. Tech. rep., Carnegie Mellon University Software Engineering Institute (2006), <http://www.sei.cmu.edu/uls>
36. Oreizy, P., Medvidovic, N., Taylor, R.N.: Runtime Software Adaptation: Framework, Approaches, and Styles. In: *Proceedings 30th International Conference on Software Engineering (ICSE 2008)*. pp. 899–910 (2008)
37. Patikirikorala, T., Colman, A., Han, J., Wang, L.: A systematic survey on the design of self-adaptive software systems using control engineering approaches. In: *Proceedings of the 7th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2012)*. pp. 33–42. IEEE Press (2012)
38. Salehie, M., Tahvildari, L.: Self-adaptive Software: Landscape and Research Challenges. *ACM Transactions on Autonomous and Adaptive Systems* 4, 14:1–14:42 (May 2009)
39. Shannon, C.: Communication in the presence of noise. *Proceedings of the IEEE* 86(2), 447–457 (Feb 1998)
40. Söderström, T., Stoica, P.: *System identification*. Prentice-Hall, Inc. (1988)
41. Solomon, B., Ionescu, D., Litoiu, M., Iszlai, G.: Autonomic computing control of composed web services. In: *Proceedings of the 2010 ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2010)*. pp. 94–103. ACM (2010)
42. Solomon, B., Ionescu, D., Litoiu, M., Iszlai, G., Prostean, O.: Measurements and identification of autonomic computing processes. In: *Computational Intelligence for Measurement Systems and Applications (CIMSAs)*, 2010 IEEE International Conference on. pp. 72–77. IEEE (2010)
43. Tamura, G., Casallas, R., Cleve, A., Duchien, L.: QoS Contract Preservation through Dynamic Reconfiguration: A Formal Semantics Approach. *Science of Computer Programming (SCP)* 94(3), 307–332 (2014)
44. Tamura, G., Villegas, N.M., Müller, H.A., Sousa, J.P., Becker, B., Pezzè, M., Karsai, G., Mankovskii, S., Schäfer, W., Tahvildari, L., Wong, K.: Towards Practical Runtime Verification and Validation of Self-Adaptive Software Systems. In:

Software Engineering for Self-Adaptive Systems 2. LNCS, vol. 7475, pp. 108–132. Springer (2013)

45. Truex, D.P., Baskerville, R., Klein, H.: Growing Systems in Emergent Organizations. *Communications of the ACM* 42(8), 117–123 (1999)
46. Villegas, N., Müller, H., Tamura, G., Duchien, L., Casallas, R.: A Framework for Evaluating Quality-Driven Self-Adaptive Software Systems. In: *Proceedings of 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2011)*. pp. 80–89. ACM (2011)
47. Villegas, N.M.: Context Management and Self-Adaptivity for Situation-Aware Smart Software Systems. Ph.D. thesis, University of Victoria (2013)
48. Villegas, N.M., Tamura, G., Müller, H.A., Duchien, L., Casallas, R.: DYNAMICO: A Reference Model for Governing Control Objectives and Context Relevance in Self-Adaptive Software Systems. In: *Software Engineering for Self-Adaptive Systems 2. LNCS*, vol. 7475, pp. 265–293. Springer (2013)
49. Vromant, P., Weyns, D., Malek, S., Andersson, J.: On Interacting Control Loops in Self-Adaptive Systems. In: *Proceedings of 6th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2011)*. pp. 202–207. ACM (2011)
50. Weyns, D., Schmerl, B., Grassi, V., Malek, S., Mirandola, R., Prehofer, C., Wuttke, J., Andersson, J., Giese, H., Göschka, K.M.: On patterns for decentralized control in self-adaptive systems. In: *Software Engineering for Self-Adaptive Systems II*, pp. 76–107. Springer (2013)
51. Zhang, Q., Zhu, Q., Zhani, M.F., Boutaba, R., Hellerstein, J.L.: Dynamic service placement in geographically distributed clouds. *IEEE Journal on Selected Areas in Communications* 31(12), 762–772 (2013)
52. Zheng, T., Woodside, M., Litoiu, M.: Performance model estimation and tracking using optimal filters. *IEEE Transactions on Software Engineering* 34(3), 391–406 (2008)
53. Zheng, T., Yang, J., Woodside, M., Litoiu, M., Iszlai, G.: Tracking time-varying parameters in software systems with extended kalman filters. In: *Proceedings of the 2005 Conference of the Centre for Advanced Studies on Collaborative research (CASCON)*. pp. 334–345. IBM Press (2005)