A Framework for Evaluating Quality-Driven Self-Adaptive Software Systems

Norha Villegas, Hausi Müller, Gabriel Tamura, Laurence Duchien, Rubby Casallas

► To cite this version:

Norha Villegas, Hausi Müller, Gabriel Tamura, Laurence Duchien, Rubby Casallas. A Framework for Evaluating Quality-Driven Self-Adaptive Software Systems. SEAMS 2011, ACM SIGSOFT / IEEE TCSE, May 2011, Honolulu, Hawaii, United States. pp.80-89, 10.1145/1988008.1988020 . inria-00578337

HAL Id: inria-00578337

<https://hal.inria.fr/inria-00578337>

Submitted on 21 Mar 2011

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

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

A Framework for Evaluating Quality-Driven Self-Adaptive Software Systems

Norha M. Villegas
Hausi A. Müller
Dept. of Computer Science
University of Victoria
Victoria, Canada
{nvillega,hausi}@cs.uvic.ca

Gabriel Tamura
Laurence Duchien
INRIA Lille-Nord Europe
University of Lille 1
Lille, France
First.Lastname@inria.fr

Rubby Casallas
Dept. of Computer Science
University of Los Andes
Bogotá, Colombia
rcasalla@uniandes.edu.co

ABSTRACT

Over the past decade the dynamic capabilities of self-adaptive software-intensive systems have proliferated and improved significantly. To advance the field of self-adaptive and self-managing systems further and to leverage the benefits of self-adaptation, we need to develop methods and tools to assess and possibly certify adaptation properties of self-adaptive systems, not only at design time but also, and especially, at run-time. In this paper we propose a framework for evaluating quality-driven self-adaptive software systems. Our framework is based on a survey of self-adaptive system papers and a set of adaptation properties derived from control theory properties. We also establish a mapping between these properties and software quality attributes. Thus, corresponding software quality metrics can then be used to assess adaptation properties.

Categories and Subject Descriptors

D.2.8 [Software Engineering]: Metrics—*complexity measures, performance measures*

Keywords

Software adaptation properties, software adaptation metrics, assessment and evaluation of self-adaptive systems, software quality attributes, application of control theory, engineering of self-adaptive systems, run-time validation and verification

1. INTRODUCTION

Over the past decade, self-adaptation has increasingly become a fundamental concern in the engineering of software systems to reduce the high costs of software maintenance and evolution and to regulate the satisfaction of functional and extra-functional requirements under changing conditions. Even though adaptation mechanisms have been widely investigated in the engineering of dynamic software systems,

their application to real problems is still limited due to a lack of methods for validation and verification of complex, adaptive, nonlinear applications [25].

After an exhaustive analysis of self-adaptive approaches, we concluded that adaptation properties and the corresponding metrics are rarely identified or explicitly addressed in papers dealing with the engineering of dynamic software systems. Consequently, without explicit adaptation properties it is impossible to assess and certify adaptive system behavior. In light of this, evaluation techniques, such as run-time validation and verification, are needed to advance the field.

To leverage the capabilities of self-adaptive systems, it is necessary to validate adaptation mechanisms to ensure that self-adaptive software systems function properly and users can trust them. To address this problem we propose a framework for evaluating self-adaptive systems, where adaptation properties are specified explicitly and driven by quality attributes, such as those defined in [2]. Our framework provides (i) a set of dimensions useful to classify self-adaptive systems; (ii) a compendium of adaptation properties for control loops (i.e., in terms of the controller and the managed system); (iii) a mapping of adaptation properties to quality attributes to evaluate adaptation properties; and (iv) a set of quality metrics to evaluate adaptation properties and quality attributes.

To define adaptation properties for our framework, we analyzed existing self-adaptive approaches and investigated properties used in control theory. We then established a mapping between adaptation properties and software quality attributes. Then, we identified a set of metrics used to evaluate software quality attributes. The mapping of adaptation properties to quality attributes and their corresponding quality metrics constitute our framework. The actual evaluation of a self-adaptive system involves both the managed system (i.e., process) and the managing system (i.e., controller).

Borrowing properties and metrics from control theory and re-interpreting them for software-intensive self-adaptive systems is not a trivial task because the semantics of the concepts involved in adaptation for control theory are different than those for self-adaptive software. Moreover, existing self-adaptive software approaches do not generally address adaptation properties explicitly. Yet another important challenge is that, in general, self-adaptive software systems are nonlinear systems [11].

Metrics to evaluate feedback control systems depend on the properties that result from the relationship between the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SEAMS '11 Honolulu, Hawaii, USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

control objectives, the target system’s measured outputs, the disturbances affecting the system and how the target system is considered in the adaptation strategy [10]. Our re-interpretation results from the analysis of several representative approaches and strategies that have been proposed to achieve behavior modification in a managed system. From the identified relationship and the analysis of these strategies we identified two main dimensions to classify and evaluate self-adaptive software. These dimensions arise from the way the strategy addresses (i) the managed system (i.e., the system to be controlled), and (ii) the controller itself. In the managed system dimension, we identified two groups. In the first one, the *control paradigm*, the managed system’s behavior is modeled, evaluated and influenced without affecting its internal structure; in the second group, the *software engineering or planning paradigm*, it is the managed system’s structure that is modeled and modified to influence the system’s behavior. In the controller dimension we identified four types of control actions: (i) continuous signals that affect behavioral properties of the managed system; (ii) discrete operations that affect the computing infrastructure of the managed system; (iii) discrete operations that affect the processes of the managed system; and (iv) discrete operations that affect the managed system’s software architecture. Hence, the nature of the adaptation strategy—*structural or behavioral*—and the way how the managed system is affected by the controller—*control actions*—define the classification of adaptive systems in our analysis.

Some previous papers have addressed the evaluation of self-adaptive software. In [16], Meng proposed a mapping of fundamental concepts from control theory to self-adaptive software systems. In his vision, the fundamental properties to evaluate self-adaptation are stability and robustness. These two properties are analyzed and characterized in terms of what they imply for the programming paradigms, architectural styles, modeling paradigms, and software engineering principles. However, his evaluation model is descriptive and not being applied to any approach. In the taxonomy proposed by Salehie and Tahvildari, several representative projects addressing the adaptation of software systems were surveyed in terms of a set of adaptation concerns: *how, what, when and where* [21]. They also proposed a hierarchical view of self-* properties and discussed their relationship with quality factors of software systems [20]. However, the scope of their work did not include the identification of metrics for evaluating self-adaptive software systems in light of the identified adaptation properties and quality attributes.

Our contribution in this paper differs from the aforementioned in the following aspects. First, we provide a reconciled definition of a more comprehensive list of properties found in control theory and contrast them in the analyzed approaches of self-adaptive software systems. Second, none of the studied contributions presents a comprehensive and unified list of adaptation properties applicable to software systems. Third, we provide a valuable foundation to evaluate adaptation properties in terms of quality attributes and metrics as widely practiced in the engineering of software systems.

The remainder of this paper is organized as follows. Section 2 presents our proposed model to characterize and classify self-adaptive software according to the dimensions mentioned above. Section 3 presents the analysis of selected adaptive systems based on the characterization model pre-

sented in Sect. 2. Section 4 presents a compendium of the metrics and properties found in the analyzed approaches with their corresponding definitions. This section also presents our proposed mapping between adaptive properties and quality attributes and how this mapping can be used to evaluate adaptive systems. Section 5 discusses different aspects of our analysis and some challenges to be addressed for the evaluation of self-adaptation. Finally, Section 6 concludes the paper.

2. A MODEL TO CHARACTERIZE SELF-ADAPTIVE SOFTWARE

In this section we propose a model consisting of eight analysis dimensions to characterize self-adaptive software. This model constitutes a foundation for evaluating self-adaptive systems. For each of the analysis dimensions, the model considers a set of standardized classification options. These options resulted from combining classification attributes from recognized authoritative sources (e.g., Software Engineering Institute (SEI)) with those found in the set of papers that we analyzed. For instance, the set of options for the analyzed quality attributes as observable adaptation properties was identified mainly using the taxonomy proposed by an SEI study [2]. This taxonomy provides a comprehensive characterization of software quality attributes, their concerns, factors that affect them and methods for their evaluation.

For each analysis dimension, we include the relevant options available from control theory, as follows.

- **Adaptation goal.** This is the main reason or justification for the system or approach to be self-adaptive. Adaptation goals are usually defined through one or more of the self-* properties, the preservation of specific quality of service (QoS) properties, or the regulation of non-functional requirements in general.
- **Reference inputs.** The concrete and specific set of value(s) and corresponding types that are used to specify the state to be achieved and maintained in the managed system by the adaptation mechanism, under changing conditions of system execution. Reference inputs are specified as (a) single reference values (e.g., a physically or logically-measurable property); (b) some form of contract (e.g., quality of service (QoS), service level agreements (SLA), or service level objectives (SLO)); (c) goal-policy-actions; (d) constraints defining computational states (according to the particular proposed definition of state); or even (e) functional requirements (e.g., logical expressions as invariants or assertions, regular expressions).
- **Measured outputs.** The set of value(s) (and corresponding types) that are measured in the managed system. Naturally, as these measurements must be compared to the reference inputs to evaluate whether the desired state has been achieved, it should be possible to find relationships between these inputs and outputs. Furthermore, we consider two aspects on the measured outputs: how they are specified and how monitored. For the first aspect, the specification, the identified options were (a) continuous domains for single variables or signals; (b) logical expressions or conditions for contract states; and (c) conditions expressing states of system malfunction. On the other side,

the options for monitoring are (a) measurements on physical properties from physical devices (e.g., CPU temperature); (b) measurements on logical properties of computational elements (in software e.g., request processing time; in hardware e.g., CPU load); and (c) measurements on external context conditions (e.g., user's profiles, weather conditions).

- **Computed control actions.** These are characterized, in the monitor, analyze, plan, execute, and knowledge (MAPE-K) loop context, and in particular by the nature of the output of the adaptation *planner* or controller [12]. These outputs affect the managed system to have the desired effect. The computed control actions can be (a) continuous signals that affect behavioral properties of the managed system; (b) discrete operations affecting the computing infrastructure executing the managed system (e.g., host system's buffer allocation and resizing operations; modification of process scheduling in the CPU); (c) discrete operations that affect the processes of the managed system directly (e.g., processes-level service invocation, process execution operations—halt/resume, sleep/respawn/priority modification of processes); and (d) discrete operations affecting the managed system's software architecture (e.g., managed system's architecture reconfiguration operations). The nature of these outputs is related to the extent of the intrusiveness of the adaptation mechanism with respect to the managed system and defines the extent of the adaptation mechanism with respect to exploiting the knowledge of either the structure or the behavior of the managed system in the adaptation process.
- **System structure.** Self-adaptive systems have two well-defined subsystems (although possibly un-distinguishable): (i) the adaptation controller and (ii) the managed system. One reason for analyzing controller and managed system structures is to identify whether a given approach implements the adaptation controller embedded with the managed system. Another reason is to identify the effect that the separation of concerns in these two subsystems has in the achievement of the adaptation goal. The analyzed approaches can be grouped into two sets: (i) those modeling the structure of the managed system to influence its behavior by modifying its structure; and (ii) those modeling the managed system's behavior to influence it directly. We consider the behavior model and the structure model as part of the system's structure. The identified options for the controller structure are variations of the MAPE-K loop with either behavioral or structural models of the managed system: (a) feedback control, that is, a MAPE-K structure with a fixed adaptation controller (e.g., a fixed set of transfer functions as a behavior model of the managed system); (b) adaptive control: a MAPE structure extended with managed system's reference and identification models of behavior (e.g., tunable parameters of controller for adaptive controllers—model reference adaptive control (MRAC) or model identification adaptive control (MIAC)); (c) reconfigurable control: MAPE-K structure with modifiable controller algorithm (e.g., rule-based software architecture reconfiguration controller). For the target system structure, the identified options are: (a) non-modifiable structure (e.g., monolithic system); and (b) modifiable structure with/without reflection capabilities (e.g., reconfigurable software components architecture). It is worth noting that not all options for system structure can be combined with any options for computed control actions. For instance, discrete operations affecting the computing infrastructure executing the managed system could be used to improve the performance of a monolithic system, whereas discrete operations affecting the managed system's software architecture would not make sense.
- **Observable adaptation properties.** By adaptation property we mean a quality (or characteristic) that is particular to a specific adaptation approach or mechanism. A quality can be a specific attribute value in a given state or a characteristic response to a known stimulus in a given context. Thus, observable adaptation properties are properties that can be identified and measured in the adaptation process. Given that we distinguish between the controller and the managed system in any self-adaptive system, we analyze observable adaptation properties also in both, the controller and the managed system. The identified observable properties in the controller are (a) stability; (b) accuracy; (c) settling-time; (d) small-overshoot; (e) robustness; (f) termination; (g) consistency (in the overall system structure and behavior); (h) scalability; and (i) security. For the managed system, the identified observable properties result from the adaptation process: (a) behavioral/functional invariants; (b) quality of service conditions, such as performance (latency, throughput, capacity); dependability (availability, reliability, maintainability, safety, confidentiality, integrity); security (confidentiality, integrity, availability); and safety (interaction complexity and coupling strength). Our proposed definitions for these properties are given in Sect. 4.2.
- **Proposed evaluation.** For the analyzed approaches, we used this element to identify the strategies proposed to evaluate themselves. Amongst the most used evaluation mechanisms are the execution of tests in real or simulated execution platforms, and the illustration with example scenarios.
- **Identified metrics and key-performance indicators (KPIs).** This element was identified from the analyzed approaches, the definition of metrics and KPIs that were used to measure the adaptation's variables of interest.

3. ANALYSIS OF SELF-ADAPTIVE APPROACHES

To validate our model for evaluating self-adaptive systems, we analyzed over 20 published approaches dealing with such systems. Of course, developing the model and analyzing the subject systems was an iterative process. The results of this analysis process are summarized in Tables 1 and 2 below. Note that two of the dimensions (i.e., columns *Adaptation Properties* and *Metrics* in Table 1) are detailed in Sect. 4. Table 1 presents the characterization of selected adaptive approaches based on the evaluation model proposed in Sect. 2.

Table 1: Applying the characterization model to selected adaptive approaches

Approach	Adaptation Goal	Reference Inputs	Measured Outputs	Control Actions	System Structure	Adaptation Properties	Evaluation	Metrics
Appley et al. Océano [1]	Self-management Self-optimization	Contracts: SLAs	SLOs/Logical properties of computational elements	Discr. operations affecting computational infrastructure	Adaptive control/Modifiable structure (reflection)	Stability, settling time, small overshoot and QoS:Dependability (scalability, maintainab.), performance (scalability)	Based on settling time to test performance of the adaptation process	Active connections per server, response time, output bandwidth, throttle rate, admission rate, active servers
Baresi and Guinea [3]	Self-recovery	Contracts: SLAs, funct. req. (logical and regular expressions)	SLOs/Logical properties of computational elements	Discr. operations affecting the process of the managed system	Adaptive control/Non-modifiable structure	None/Behavioral, QoS:Dependability (safety, integrity, availab., reliab.)	Functional and reliability-based tests	Reliability $\geq$ 0.95 (last two hours)
Candea et al. [4]	Self-recovery	Contracts: SLOs-QoS	Malfunction conditions/Logical properties of computational elements	Discr. operations affecting the process of the managed system	Adapt. control/Modifiable structure (reflection)	Small overshoot, settling time/QoS:Dependability (availab.)	Recovery-based tests	Availability, downtime (cf. Table 5)
Cardellini et al. MOSES [5]	QoS preservation	Contracts: QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the process of the managed system	Reconfigurable control/Modifiable structure (reflection)	Accuracy/QoS:Performance (latency), Dependability (reliab., cost).	Based on the accuracy of the adapt. strategy	resp. time, exec. cost, reliability
Dowling and Cahill KComponents [6]	Self-management QoS	Contracts: SLOs-QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Reconfigurable control/Modifiable structure (reflection)	Robustness, scalability/QoS:Performance (throughput, capacity)	None	Load cost of components
Ehrig et al. [7]	Self-healing	Goal actions	Malfunction conditions/Logic. properties of comput. elem.	Discr. operations affecting the process of the managed system	Adaptive control/Non-modifiable structure	Termination/QoS: Dependability (reliab.)	Formal verification of properties. Running example	None
Floch et al. MADAM [8]	QoS preservation Self-configuration	Contracts: SLAs-SLOs-QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Adaptive control/Modifiable structure (reflection)	Scalability/QoS: Performance, Dependability (reliab., maintainab.)	Simulated environment to evaluate scalability	None
Garlan et al. Rainbow [9]	Self-repairing	Contracts: SLAs-SLOs-QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Adaptive control/Modifiable structure (reflection)	None/QoS: Performance (latency)	Running examples to evaluate effectiveness	None
Kumar et al. MWare [13]	Self-management Self-configuration Self-optimization	Contracts: QoS; policy actions	SLOs/Logical properties of computational elements, external context	Discr. operations affecting the managed system's soft. arch.	Adaptive control/Modifiable structure (reflection)	Setting time, small overshoot/QoS:Performance (throughput, capacity)	Execution tests on real scenarios	Business utility function (cf. Table 5)
Léger et al. [14]	QoS preservation Self-configuration	Constraints defining computational states	Malfunction conditions/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Reconfigurable control/Modifiable structure (reflection)	Consistency (atom., isol., durab.)/QoS: Dependability (availab., reliab.)	Running examples to evaluate performance	None
Mukhiya and Glinz CASA [17]	QoS preservation Self-configuration	Contracts: QoS	SLOs/Logical properties of computational elements, external context	Discr. operations affecting the managed system's soft. arch.	Reconfigurable control/Modifiable structure (reflection)	Consistency/QoS: Performance	Running examples to evaluate performance	None
Parekh et al. [18]	QoS preservation	Single reference value	SLOs/Logical properties of computational elements	Continuous signals affecting behavioral properties	Feedback control/Non-modifiable structure (mathematical model)	Stability, small overshoot/QoS: Performance (throughput, capacity)	Running examples on Lotus Notes	Offered load
Sicard et al. [22]	Self-management self-healing	Constraints defining computational states	Not explicit monitoring phase	Discr. operations affecting the managed system's soft. arch.	Feedback control/Modifiable structure (reflection)	None/QoS: Dependability (reliab., availab.)	Simulated experiments to evaluate performance	Availability (cf. Table 5)
Solomon et al. [23]	Self-optimization	Contracts: QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Adaptive control/Modifiable structure (reflection)	Accuracy/QoS: Performance	Running example to test accuracy	None
Tamura et al. SCeSAMe [24]	QoS preservation	Contracts: SLOs-QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the managed system's soft. arch.	Reconfigurable control/Modifiable structure (reflection)	Termination, consistency/QoS properties (contract)	Formal properties probed with theorems. Running example	None
White et al. Autonomic JBeans [26]	Self-management Self-healing Self-optimization Self-protection	Contracts: SLOs-QoS	SLOs/Logical properties of computational elements	Discr. operations affecting the process of the managed system	Adaptive control/Non-modifiable structure	Setting time/QoS: Performance (throughput), Dependability (availab.)	Empirical experiments to probe development effort savings	Average response time

Table 2 summarizes the characterization and classification of the studied self-adaptive systems.

Self-adaptive approaches range from pure control theory approaches to pure software engineering-based approaches with many hybrid approaches in-between. In control theory-based approaches, control actions are continuous signals that affect behavioral parameters of the managed system. The structure of the managed system in these approaches is generally non-modifiable while its behavior is modeled mathematically [18]. In contrast, software engineering-based approaches are characterized by implementing discrete control actions that affect the managed system’s software architecture (i.e., the system structure). In these approaches the adaptation is supported by a model of the managed system’s structure and reflection capabilities that allow the modification of the structure [1, 6, 8, 9, 13, 14, 17, 22, 24]. In hybrid adaptive systems, control actions are generally discrete operations that affect either the computing infrastructure executing the managed system or the set of processes comprising the managed system. Usually, the structure of the managed system is non-modifiable [3, 4, 5, 7, 26]. It might be useful to classify hybrid approaches further. For instance, we classified the approach proposed by Solomon et al. [23] between hybrid and software engineering-based approaches, given that their control actions affect the architecture of the managed system but the analyzer is based on a behavioral model of the managed system to decide when to adapt. These predictive mechanisms use control engineering techniques (i.e., Kalman Filters) to estimate adaptation parameters that require outputs that are not measurable on the actual managed system.

According to our proposed spectrum, most approaches were identified as software engineering-based and hybrid adaptive systems. With respect to the adaptation goal, we did not identify a relationship with our proposed spectrum. Thus, it is possible to address any of the adaptation goals along the entire spectrum. Concerning reference inputs, most approaches use contracts as the way to specify reference values for the adaptation goal and the corresponding measured outputs. All approaches that explicitly addressed monitoring monitor logical properties of computational elements (internal context), while two of them take into account the external context [13, 17]. Regarding the controller structure, all approaches, except [18] and [22] that implement a simple feedback loop, implement either adaptive or reconfigurable control. Finally, the most common evaluation mechanism is the implementation of running examples based on simulated environments.

4. MEASURING ADAPTATION PROPERTIES

Our evaluation of a self-adaptive system has two aspects. The first one concerns the evaluation of desired properties for the managed system. In our analysis we focused only on desired properties that correspond to quality attributes of software systems. The second one relates to desired properties of the controller of the adaptation process.

In this section we present a set of properties and metrics useful to evaluate adaptation. For the identification of desired properties on the managed system, we based our analysis on the taxonomy of quality attributes for software systems proposed by SEI researchers [2]. For properties re-

Table 2: Characterization Summary

Characteristic	Count [List of Approaches]
Spectrum Classification	
Control Engineering	1 [18]
Hybrid	5 [3, 4, 5, 7, 26]
Hybrid-Software	1 [23]
Software Engineering	9 [1, 6, 8, 9, 13, 14, 17, 22, 24]
Adaptation Goal Specification	
Contract-based	14
Monitoring Mechanisms	
Monitor internal context	15
Monitor external context	2 [13, 17]
Non specified	1 [22]
Controller’s Structure	
Feedback control	2 [18, 22]
Adaptive control	9 [1, 3, 4, 7, 8, 9, 13, 23, 26]
Reconfigurable Control	4 [5, 6, 14, 17, 24]
Managed System’s Structure	
Non-modifiable	4 [3, 7, 18, 26]
Modifiable with reflection	12 [1, 4, 5, 6, 8, 9, 13, 14, 17, 22, 23, 24]
Adaptation Properties	
Settling time	4 [1, 4, 13, 26]
Small overshoot	4 [1, 4, 13, 18]
Scalability	3 [1, 6, 8]
Stability	2 [1, 18]
Accuracy	2 [5, 23]
Termination	2 [7, 24]
Consistency	3 [14, 17, 24]
Robustness	1 [6]
Security	0
Quality Attributes	
Performance	10 [1, 5, 6, 8, 9, 13, 17, 18, 23, 26]
Dependability	7 [1, 3, 4, 5, 7, 8, 26]

lated to the controller, we based our analysis on the *SASO properties* identified by Hellerstein et al. in the application of control theory to computing systems [10], and other properties identified in self-adaptive software systems surveys [7, 14, 16, 17].

Furthermore, we classified the identified adaptation properties according to *how* and *where* they are observed. Concerning how they are observed, some properties can be evaluated using static verification techniques while others require dynamic verification and run-time monitoring. We use the term *observed* as some properties are difficult to measure, despite the fact that controllers are designed to preserve them [24]. With respect to where, properties can be evaluated on the managed system or the controller. On the one hand, some properties to evaluate the controller are observable on the controller itself or on both, the controller and the managed system; however, most properties can only be observed on the managed system. On the other hand, properties to evaluate the managed system are observable only on the managed system. In both cases, the environment that can affect the behavior of the controller or the managed system is also a factor worth considering.

Based on Sect. 2 and the analysis presented in Sect. 3,

this section presents the foundations of a framework for the evaluation of self-adaptation, where properties observable on the managed system, either to evaluate the controller or the managed system, can be evaluated in terms of quality attributes. After analyzing several approaches on self-adaptive software systems and identifying adaptation properties, we propose a process for evaluating self-adaptation with which software engineers should be able to (i) identify required adaptation goals (i.e., the quality attributes that drive the adaptation of the managed system); (ii) identify adaptation properties to evaluate the controller, this must include the identification of properties that are observable on the controller, the managed system, or both; (iii) map quality attributes used to evaluate the managed system to properties that evaluate the controller but are observable on the managed system; and (iv) define metrics to evaluate properties observable on the managed system and the controller.

4.1 Quality Attributes as Adaptation Goals

If we intend to evaluate an adaptive software system we need to identify the motivation to build it—the *adaptation goal*. In general, adaptation can be motivated by the need of continued satisfaction of functional and regulation of non-functional requirements under changing context conditions. Nevertheless, as most analyzed contributions focus on non-functional factors, we based our analysis on software systems whose adaptation goals are motivated by quality concerns. Moreover, characteristics of self-adaptive systems, such as self-configuring or self-optimizing, can be mapped to quality attributes. Following this idea, Salehie and Tahvildari discussed the relationships between autonomic characteristics and quality factors such as the relationship between self-healing and reliability [20].

Our main contribution is the application of quality attributes to evaluate self-adaptive software systems, as quality attributes are commonly used to evaluate desirable properties on the managed system. More importantly, we propose a mapping between quality factors and adaptation properties. This in fact introduces a level of indirection for evaluating adaptation properties that are not directly observable on the controller. The quality attributes that we analyzed in the approaches are the ones introduced in Sect. 2. In this subsection we present the definitions of the selected quality attributes, as well as the citations of the analyzed contributions whose adaptation goals are related to these quality attributes.

- **Performance.** Characterizes the timeliness of services delivered by the system. It refers to responsiveness, that means the time required for the system to respond to events or the event processing rate in an interval of time. Identified factors that affect performance are latency, the time the system takes to respond to a specific event [5, 9]; throughput, the number of events that can be completed in a given time interval—beyond processing rate as the desired throughput must also be observed in time sub-intervals [6, 13, 18, 26]; and capacity, a measure of the amount of work the system can perform [6, 13, 18].
- **Dependability.** Defines the level of reliance that can justifiably be placed on the services the software system delivers. Adaptation goals related to dependability are availability, readiness for usage [1, 3, 4, 14, 22,

26]; reliability, continuity of service [3, 7, 8, 13, 14, 22]; maintainability, capacity to self-repair and evolve [1, 4, 8, 22]; safety (from a dependability point of view), non-occurrence of catastrophic consequences from an external perspective (on the environment) [3]; confidentiality, freedom from unauthorized disclosure of information; integrity, non-improper alterations of the system structure, data and behavior [3].

- **Security.** The selected concerns of the security attribute are confidentiality, protection from disclosure; protection from unauthorized modification, integrity; and availability, protection from destruction [3].
- **Safety.** The level of reliance that can justifiably be placed on the software system as not generator of accidents. Safety is concerned with the occurrence of accidents, defined in terms of external consequences. The taxonomy presented in [2] includes two properties of critical systems that can be used as indicators of system safety: interaction complexity and coupling strength. In particular, interaction complexity is the extent to which the behavior of one component can affect the behavior of other components. SEI's taxonomy presents detailed definitions and indicators for these two properties.

4.2 Adaptation Properties

An important part of our contribution is the identification of adaptation properties that have been used for the analyzed spectrum of adaptive systems, from control theory to software engineering, to evaluate the adaptation process. The identified adaptation properties are stated as follows. The first four properties, called *SASO properties*, correspond to desired properties of controllers from a control theory perspective [10]; note that the stability property has been widely applied in adaptation control from a software engineering perspective. The remaining properties in the list were identified from hybrid approaches. Citations attached to each property refer either to papers where the property was defined or to examples of adaptive systems where the property is observed in the adaptation process.

- **Stability.** The degree in that the adaptation process will converge toward the control objective. An unstable adaptation will indefinitely repeat the process with the risk of not improving or even degrade the managed system to unacceptable or dangerous levels. In a stable system responses to a bounded input are bounded to a desirable range [1, 8, 15, 16, 18].
- **Accuracy.** This property is essential to ensure that adaptation goals are met, within given tolerances. Accuracy must be measured in terms of how close the managed system approximates to the desired state (e.g., reference input values for quality attributes) [5, 23].
- **Short settling time.** The time required for the adaptive system to achieve the desired state. The settling time represents how fast the system adapts or reaches the desired state. Long settling times can bring the system to unstable states. This property is commonly referred to as recovery time, reaction time, or healing time [4, 10, 13, 15, 16].

- **Small overshoot.** The utilization of computational resources during the adaptation process. Managing resource overshoot is important to avoid the system instability. This property provides information about how well the adaptation performs under given conditions—the amount of excess resources required to perform the adaptation [1, 4, 13, 15, 18].
- **Robustness.** The managed system must remain stable and guarantee accuracy, short settling time, and small overshoot even if the managed system state differs from the expected state in some measured way. The adaptation process is robust if the controller is able to operate within desired limits even under unforeseen conditions [6, 16].
- **Termination (of the adaptation process).** In software engineering approaches, the planner in the MAPE-K loop produces, for instance, discrete controlling actions to adapt the managed system (cf. Sect. 2), such as a list of component-based architecture operations. The termination property guarantees that this list is finite and its execution will finish, even if the system does not reach the desired state. Termination is related to deadlock freeness, meaning that, for instance, a reconfigurable adaptation process must avoid adaptation rules with deadlocks among them [7, 24].
- **Consistency.** This property aims at ensuring the structural and behavioral integrity of the managed system after performing an adaptation process. For instance, when a controller bases the adaptation plan on dynamic reconfiguration of software architectures, consistency concerns are to guarantee sound interface bindings between components (e.g., component-based structural compliance) and to ensure that when a component is replaced dynamically by another one, the execution must continue without affecting the function of the removed component. These concerns help protect the application from reaching inconsistent states as a result of dynamic recomposition [17]. Léger et al. define this property to complete the atomicity, consistency, isolation and durability (ACID) properties found in transactional systems that guarantee transactions are processed reliably [14]:
 - **Atomicity.** Either the system is adapted and the adaptation process finishes successfully or it is not finished and the adaptation process aborts. If an adaptation process fails, the system is returned to a previous consistent state.
 - **Isolation.** Adaptation processes are executed as if they were independent. Results of unfinished adaptation processes are not visible to others until the process finishes. Results of aborted or failed adaptation processes are discarded.
 - **Durability.** The results of a finished adaptation process are permanent: once an adaptation process finishes successfully, the new system state is made persistent. In case of major failures (e.g., hardware failures), the system state can be recovered.
- **Scalability.** The capability of a controller to support increasing demands of work with sustained performance

using additional computing resources. For instance, scalability is an important property for the controller when it must evaluate an increased number of conditions in the analysis of context. As computational efficiency is relevant for guaranteeing performance properties in the controller, scalable controllers are required to avoid the degradation of any of the operations of the adaptive process in any situation [1, 6, 8].

- **Security.** In a secure adaptation process, not only the target system but also the data and components shared with the controller are required to be protected from disclosure (confidentiality), modification (integrity), and destruction (availability) [2].

Table 3: Classification of adaptation properties according to how and where they are observed.

Adaptation properties	Property verification mechanism	Where the property is observed
Stability	Dynamic	Managed system
Accuracy	Dynamic	Managed system
Settling Time	Dynamic	Managed system
Small Overshoot	Dynamic	Managed system
Robustness	Dynamic	Controller
Termination	Both	Controller
Consistency	Static	Managed system
Scalability	Dynamic	Controller
Security	Dynamic	Controller

4.3 Mapping Adaptation Properties and Quality Attributes

Once the adaptation goal and adaptation properties have been identified, the following step maps the properties of the controller, which are observable at the managed system, to the quality attributes of the managed system. Table 4 presents a general mapping between adaptation properties and quality attributes. These quality attributes refer to attributes of both the controller and the managed system depending on where the corresponding adaptation properties are observed.

According to Tables 3 and 4, SASO properties, including stability, can be verified at run-time by observing performance, dependability and security factors in the managed system. Stability is one of the adaptation properties addressed in Océano, a dynamic resource allocation system that enables flexible SLAs in environments where peak loads are an order of magnitude greater than in the normal steady state. Quality attributes addressed in Océano concern dependability (i.e., availability and maintainability), and performance (i.e., throughput and capacity—scalability) [1]. In the same way, the controller proposed by Parekh et al. to guarantee desirable performance levels (i.e., throughput and capacity) is also concerned with stability as adaptation property [18]. They apply an integral control technique to construct a transfer function that describes the system and the way the behavior of the managed system is affected by the controller. Baresi and Guinea propose a self-recovery system where service oriented architecture (SOA) business processes recover from disruptions of functional and non-functional

Table 4: Mapping adaptation properties to quality attributes

Adaptation Properties	Quality Attributes	
Stability	Performance	Latency Throughput Capacity
	Dependability	Safety Integrity
	Security	Integrity
Accuracy	Performance	Latency Throughput Capacity
Settling Time	Performance	Latency Throughput
Small Overshoot	Performance	Latency Throughput Capacity
	Dependability	Availability Reliability
Robustness	Safety	Interact. Complex. Coupling Strength
	Dependability	Reliability Integrity
Termination	Dependability	Maintainability Integrity
Consistency	Performance	Latency Throughput Capacity
Scalability	Security	Confidentiality Integrity Availability

requirements, to avoid catastrophic events (safety) and improper system state alterations (integrity), and to guarantee readiness for service (availability) and correctness of service (reliability) [3].

Accuracy is addressed by Cardellini et al. in the MOSES framework using adaptation policies in the form of directives to select the best implementation of the composite service according to a given scenario [5]. MOSES adapts chains of service compositions based on service selection using a multiple service strategy. It has been tested with some experiments to observe the behavior of the adaptation strategy in terms of its accuracy (i.e., how close the managed system is to the adaptation goal). Multiple adaptation goals were used for these tests. The self-optimizing mechanism for business processes proposed by Solomon et al. also deals with accuracy as an adaptation property [23]. Their approach is based on a simulation model to anticipate performance levels and make decisions about the adaptation process. They developed a tuning algorithm to keep the simulation model accurate. The algorithm compensates for the measurement of actual service time to increase the accuracy of simulations by modeling errors, probabilities and inter-arrival times. Then it obtains the best estimate for these data such that a square root error between the simulated and measured metrics is minimized.

Settling time and small overshoot are addressed as adaptation properties in Océano [1], the framework for developing autonomic Enterprise Java Bean (EJB) applications proposed by White et al. [26], and the self-recovery approach

based on microreboots proposed by Candea et al. [4]. In Océano settling time is measured in terms of the time required for deploying a new processing node including the installation and reconfiguration of all applications and data repositories. In White's framework, settling time is evaluated in terms of the average response time required for autonomic EJBs to adapt. Finally, Candea's approach applies a recursive strategy that reduces mean time to repair (MTTR) by means of recovering minimal subsets of a failed system's components. If localized, minimal recovery is not enough, progressively larger subsets are recovered. In the control-based approach, to ensure SLOs proposed by Parekh et al., small overshoot is addressed by avoiding that control values (e.g., MAXUSERS) are set to values that exceed their legal range. Root-locus analysis is used to predict the valid values of the maximum number of users. They divided the valid range into three regions to decide when control values reach undesirable levels. Based on empirical studies, they analyzed properties of the transfer function to predict the desired range of values [18].

One of the addressed adaptation properties in the self-management approach for balancing system load proposed by Dowling and Cahill is robustness [6]. They aim to realize a robust controller by implementing adaptation an mechanism via decentralized agents that eliminate centralized points of failure.

Termination can be verified using static mechanisms. Ehrig et al. proposed a self-healing mechanism for a traffic light system to guarantee continuity of service (reliability) by self-recovering from predicted failures (integrity) [7]. In their adaptive solution, termination is addressed by ensuring deadlock-freeness in the managed system by statically checking the self-healing rules in such a way that the self-repairing mechanism never inter-blocks traffic lights in the same road intersection.

Scalability is also an adaptive property in K-Components, the agent-based self-managing system proposed by Downling and Cahill [6]. Scalability is addressed by evolving the self-management local rules of the agents. Another approach where scalability is addressed as an adaptation property is Madam, the middleware proposed by Floch et al. for enabling model-based adaptation in mobile applications [8]. Scalability is a concern in Madam for several reasons. First, its reasoning approach might result in a combinatorial explosion if all possible variants are evaluated; second, the performance of the system might be affected when reasoning on a set of a concurrently running applications competing for the same set of resources. They proposed a controller where each component (e.g., the adaptation manager) can be replaced at run-time to experiment with different analysis approaches for managing scalability.

Security was not addressed as an adaptation property by any of the self-adaptive systems analyzed. However, we propose the use of SEI's definition of security as a quality attribute and its corresponding quality factors to evaluate security on the controller [2]. As presented in Table 3, security of the controller should be evaluated independently of the managed system. This means that ensuring security at the managed system does not guarantee security with respect to the adaptation mechanism.

4.4 Adaptation Metrics

Adaptation metrics provide the way of evaluating adap-

tive systems with respect to particular concerns of the adaptation process [16, 19]. Thus, metrics provide a measure to evaluate desirable properties. For instance, metrics to evaluate control systems measure aspects concerning the SASO properties (i.e., stability, accuracy, settling time, and small overshoot).

To characterize the evaluation of adaptive systems, we analyzed the variety of self-adaptive software systems to identify adaptation properties (i.e., at the managed system and the controller) that were evaluated in terms of quality attributes (cf. Sects. 3 and 4.1). Just as the evaluation of most properties is impossible by observing the controller itself, we propose the evaluation of these properties by means of observing quality attributes at the managed system. To identify relevant metrics, we characterized a set of factors that affect the evaluation of quality attributes such as speed, memory usage, response time, processing rate, mean time to failure, and mean time to repair [15, 2]. These factors are an essential part of the metrics used to evaluate properties of both the controller and properties of the managed system [19].

The evaluation of MOSES, the QoS-driven framework proposed by Cardellini et al. to adapt service-oriented business processes, is based on the following metrics to measure performance and reliability. Expected response time (Ru), the average time needed to fulfill a request for a composite service; expected execution cost (Cu), the average price to be paid for a user invocation of the composite service; and expected reliability (Du), the logarithm of the probability that the composite service completes its task for a user request [5].

For Océano, the following metrics were defined to evaluate dependability factors (e.g., availability) and performance factors (e.g., scalability in terms of throughput and capacity) [1]: active connections per server—the average number of active connections per normalized server across a domain; overall response time—the average time it takes for any request to a given domain to be processed; output bandwidth—the average number of outbound bytes per second per normalized server for a given domain; data base response time—average time it takes for any request to a given domain to be processed by the back-end data base; throttle rate (T)—a percentage of connections disallowed to pass through Océano on a customer domain; admission rate—the complement of the domain throttle rate ($1 - T$); and active servers—the number of active normalized-servers which service a given customer domain.

Average response time is a common metric used to evaluate performance in several adaptive approaches, such as the framework to develop autonomic EJB applications proposed by White et al. [26]. In K-Components, the self-adaptive component model that enables the adaptation of software components to optimize system performance, a load balancing function on every adaptation contract uses a cost function to calculate its internal load cost and the ability of its neighbors to handle the load [6]. This cost function is defined as the addition of the advertised load cost and internal cost of the component (i.e., calculated as the estimated cost to handle a particular load type).

In the control-based approach proposed by Parekh et al., to achieve performance service level objectives the length of the queue of the in-progress client requests is the metric defined as control offered load, the load imposed on the server

by client requests [18]. Baresi and Guinea proposed a metric to control reliability on the adaptation of business processes based on BPEL [3]. In their approach, reliability is calculated as the number of times a specific method responds to within two minutes over the total number of invocations. They also defined a KPI based on this metric such that reliability must be greater than 95% over the past two hours of operation.

In the adaptive middleware proposed by Kumar et al., a business value KPI is defined in terms of factors, such as the priority of the user accessing the information, the time of day the information is being accessed, and other aspects that determine how critical the information is to the enterprise [13]. For this, they defined a utility function as a combination of some of these factors: $utility(e_{gj-k}) = f(\sum d_{ni}, \min(b_{ni}), b_{gj-k})$, where $i|e_{ni} \in M(e_{gj-k})$. The business utility of each edge (e_{gj-k}), which represents data streams between operators that perform data transformations, is a function of the delay d_{ni} , the available bandwidth b_{ni} of the intervening network edges e_{ni} , and the required bandwidth b_{gj-k} of the edge e_{gj-k} .

In the self-healing approach based on recursive micro-reboots proposed by Candea et al., availability is evaluated in terms of mean time to recover ($MTTR$) [4]. To evaluate the availability of the system they defined two metrics, availability ($A = MTTF/(MTTF + MTTR)$) and downtime of unavailability ($U = MTTR/(MTTF + MTTR)$), where $MTTF$ is the mean time for a system or subsystem to fail (i.e., the reciprocal of reliability), $MTTR$ is the mean time to recover, and A is a number between 0 and 1. U can be approximated to $MTTR/MTTF$ when $MTTF$ is much larger than $MTTR$. Similarly, Sicard et al. define a metric for availability in terms of $MTTR$ [22].

Table 5: Metrics to evaluate quality attributes of a subject self-adaptive system being analyzed

Quality Attributes	Approach	Metrics
Performance	MOSES [5]	Response Time (Ru) Execution Cost (Cu)
	Océano [1]	Response Time Output Bandwidth Throttle Rate Admission Rate
	EJB Framework [26]	Response Time
	K-Comp. [6]	Load Cost
	Parekh et al. [18]	Offered Load
	Kumar et al. [13]	$f(\sum d_{ni}, \min(b_{ni}), b_{gj-k})$
Dependability	MOSES [5]	Expected Reliability (Du)
	Océano [1]	Active Connections per Server Active Servers
	Baresi and Guinea [3]	$ResponseFrequency/TimeUnit$
	Reboots [4]	$A=MTTF/(MTTF+MTTR)$ $U=MTTR/(MTTF+MTTR)$
	Sicard et al. [22]	$A=MTTR$

Table 5 summarizes our identified metrics to assess self-adaptive systems. Although these metrics are directly re-

lated to the measurement of quality factors, we expect that these metrics will be useful for evaluating adaptation properties based on our proposed mapping between quality attributes and adaptation properties (cf. Sect. 4). The approach by Reinecke et al. [19] supports our hypothesis. Their metric measures the ability of a self-adaptive system to adapt. They argue that adaptivity can be evaluated using a meta-metric named *payoff* which is defined in terms of performance metrics to measure the effectiveness of the adaptation process. That is, the optimal adaptive system is characterized by the fact that its adaptation decisions are always optimal (i.e., always yield the optimal payoff). To apply their metric it is necessary to (i) identify the adaptation tasks, (ii) define one or more performance metrics on these tasks (i.e., these metrics should reflect the contribution of these tasks toward the adaptation goal), (iii) define a payoff metric in terms of the performance metrics, and (iv) to apply the metric by observing the performance of the system.

5. DISCUSSION

We started the analysis phase for this work with 34 research papers with different proposals for self-adaptive software systems published over the past decade. From this set, 18 were filtered-out mainly because either they presented very generic proposals (i.e., with non-measurable self-adaptive properties) or they did not include enough information in the paper for characterization purposes.

From the analysis of the 16 remaining papers that we presented in the previous sections (and even considering the 18 papers that we filtered out), it is worth noting the prevalent difficulty—or the lack of awareness of—to identify metrics to evaluate self-adaptive software. Nonetheless, some advances have been made based on concepts from control theory and the recognized importance of quality metrics and corresponding measures as the basis for understanding and improving processes. However, as we can conclude from our analysis, there are plenty of opportunities and challenges to be addressed, even when we consider concepts more abstract than metrics, such as the properties of self-adaptation. In the following we outline some of these opportunities and challenges.

First, most of the proposals focus only on self-adaptation mechanisms, not addressing explicitly the level of achievement of adaptation properties nor the adaptation properties themselves. On the one hand, it is known that even though control theory has defined standardized properties that a controller must realize (i.e., the SASO properties), self-adaptive software systems require additional properties, due to their discrete nature (e.g., termination). On the other hand, it is clear from the discussion in previous sections that quality attributes are a plausible option to evaluate some adaptation properties. However, it will be necessary (i) to evaluate if our proposed set of self-adaptation properties defined in Sect. 4.2 is general enough for self-adaptation mechanisms; (ii) to find standardized metrics to measure self-adaptation properties, which could be based on proposals such as the one by Reinecke et al. [19]; and (iii) to analyze whether the proposal of measuring adaptation properties based on quality attributes is meaningful enough, and if they fulfill the conceptual definitions of corresponding properties of, for instance, control theory.

Second, the lack of awareness of adaptation properties as a goal to be measured results in a lack of evaluation meth-

ods and metrics for these properties and for the adaptation mechanisms themselves. However, this trend could be reversed by designing self-adaptive mechanisms with implicit controllable and measurable properties. For some verifiable properties, this can be obtained, for instance, by developing new or using existing formal models as a basis for the self-adaptation process. For adaptation mechanisms with measurable properties, one main challenge is to develop mathematical behavior models based on the architecture itself of the target computing systems to be controlled.

Third, without declared evaluation methods and metrics it is very difficult to compare and to reason about the engineering of self-adaptation; for instance, from our analysis it resulted not possible to identify any measurable relationship between the adaptation goals and the evaluation of the adaptation strategies as such. From the structure point of view, it is clear that decoupling the controller from the managed system, with respect to evaluation, is a first critical step toward to be able to reason and control the properties of dynamic self-adaptation. However, several questions remain: does the system structure (i.e., controller and managed system) have any relationship with the quality of the adaptation approach? Do non-explicit controllers imply undefined adaptation properties? From the behavior point of view and considering the adaptation mechanism as a black-box, how do we compare managed system behavior in the different phases of the adaptation process? Maybe in terms of stability and other properties such as settling time. Under which circumstances and characteristics of the managed system is an adaptation mechanism better than another? There is no available evaluation framework to compare adaptation mechanisms to help answer these questions.

6. CONCLUSIONS

Self-adaptive software evaluates and modifies its behavior to preserve the satisfaction of its functional requirements and regulation of non-functional requirements, under changing context conditions of execution. Researchers devised and proposed many diverse approaches and strategies to modify the behavior of a managed system. In this paper, as a result of our analysis, we proposed a classification of self-adaptive systems spanning the spectrum from control-based to software engineering-based approaches.

Many studied approaches did not identify nor address adaptation properties. Thus, the evaluation of adaptive systems is generally not addressed explicitly—neither in the controller nor in the managed system. Consequently, since adaptation properties are not identified in many approaches, metrics are not addressed either. Thus, validation mechanisms discussed for these approaches are usually limited to the evaluation of performance properties observed in the managed system, even when the adaptation goal is not related to performance quality attributes.

Future work will focus on the validation of the adaptation properties and their mapping to quality attributes as proposed in this paper through evaluation of existing adaptive systems.

Acknowledgments

This work was funded in part by the National Sciences and Engineering Research Council (NSERC) of Canada, IBM Corporation, CA Inc., Icesi University (Cali, Colombia), and

7. REFERENCES

- [1] K. Appleby, S. Fakhouri, L. Fong, G. Goldszmidt, M. Kalantar, S. Krishnakumar, D. P. Pazal, J. Pershing, and B. Rochwerger. Océano - SLA based management of a computing utility. In *Proceedings 7th IFIP/IEEE International Symposium on Integrated Network Management*, pages 855–868, 2001.
- [2] M. Barbacci, M. H. Klein, T. A. Longstaff, and C. B. Weinstock. Quality attributes. Technical Report CMU/SEI-95-TR-021, CMU/SEI, 1995.
- [3] L. Baresi and S. Guinea. Self-supervising BPEL processes. *IEEE Transactions on Software Engineering*, 99(Preprint), 2010.
- [4] G. Candea, J. Cutler, and A. Fox. Improving availability with recursive microreboots: a soft-state system case study. *Performance Evaluation*, 56(1-4):213–248, 2004. Dependable Systems and Networks - Performance and Dependability Symposium (DSN-PDS) 2002: Selected Papers.
- [5] V. Cardellini, E. Casalicchio, V. Grassi, F. Lo Presti, and R. Mirandola. QoS-driven runtime adaptation of service oriented architectures. In *Proceedings 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*, (ESEC/FSE '09), pages 131–140, New York, NY, USA, 2009. ACM.
- [6] J. Dowling and V. Cahill. Self-managed decentralised systems using k-components and collaborative reinforcement learning. In *Proceedings 1st ACM SIGSOFT Workshop on Self-Managed Systems*, (WOSS '04), pages 39–43, New York, NY, USA, 2004. ACM.
- [7] H. Ehrig, C. Ermel, O. Runge, A. Bucchiarone, and P. Pelliccione. Formal analysis and verification of self-healing systems. In *Fundamental Approaches to Software Engineering*, (FASE '10), volume 6013 of LNCS, pages 139–153. Springer, 2010.
- [8] J. Floch, S. Hallsteinsen, E. Stav, F. Eliassen, K. Lund, and E. Gjørven. Using architecture models for runtime adaptability. *IEEE Software*, 23:62–70, March 2006.
- [9] D. Garlan, S.-W. Cheng, A.-C. Huang, B. Schmerl, and P. Steenkiste. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *IEEE Computer*, 37:46–54, 2004.
- [10] J. L. Hellerstein, Y. Diao, S. Parekh, and D. M. Tilbury. *Feedback Control of Computing Systems*. John Wiley & Sons, 2004.
- [11] J. L. Hellerstein, S. Singhal, and Q. Wang. Research challenges in control engineering of computing systems. *IEEE Trans. on Network and Service Management*, 6(4):206–211, 2009.
- [12] IBM Corporation. An architectural blueprint for autonomic computing. Technical report, IBM Corporation, 2006.
- [13] V. Kumar, B. F. Cooper, Z. Cai, G. Eisenhauer, and K. Schwan. Middleware for enterprise scale data stream management using utility-driven self-adaptive information flows. *Cluster Computing*, 10:443–455, 2007.
- [14] M. Léger, T. Ledoux, and T. Coupaye. Reliable dynamic reconfigurations in a reflective component model. In *Proceedings 13th International Symposium on Component Based Software Engineering*, (CBSE '10), volume 6092 of LNCS, pages 74–92. Springer, 2010.
- [15] C. Lu, J. A. Stankovic, T. F. Abdelzaher, G. Tao, S. H. Son, and M. Marley. Performance specifications and metrics for adaptive real-time systems. In *Real-Time Systems Symposium*, pages 13–23, 2000.
- [16] A. C. Meng. On evaluating self-adaptive software. In *Proceedings 1st International Workshop on Self-Adaptive Software*, (IWSAS' 00), pages 65–74, Secaucus, NJ, USA, 2000. Springer-Verlag New York.
- [17] A. Mukhija and M. Glinz. Runtime adaptation of applications through dynamic recomposition of components. In *Proceedings 18th International Conference on Architecture of Computing Systems*, pages 124–138, 2005.
- [18] S. Parekh, N. Gandhi, J. Hellerstein, D. Tilbury, T. Jayram, and J. Bigus. Using control theory to achieve service level objectives in performance management. *Real-Time Systems*, 23:127–141, July 2002.
- [19] P. Reinecke, K. Wolter, and A. van Moorsel. Evaluating the adaptivity of computing systems. *Performance Evaluation*, 67(8):676–693, 2010. Special Issue on Software and Performance.
- [20] M. Salehie and L. Tahvildari. Autonomic computing: emerging trends and open problems. *SIGSOFT Software Engineering Notes*, 30:1–7, May 2005.
- [21] M. Salehie and L. Tahvildari. Self-adaptive software: Landscape and research challenges. *ACM Transactions on Autonomous and Adaptive Systems*, 4:14:1–14:42, May 2009.
- [22] S. Sicard, F. Boyer, and N. De Palma. Using components for architecture-based management: the self-repair case. In *Proceedings 30th International Conference on Software Engineering*, ICSE '08, pages 101–110. ACM, 2008.
- [23] A. Solomon, M. Litoiu, J. Benayon, and A. Lau. Business process adaptation on a tracked simulation model. In *Proceedings 2010 Conference of the Center for Advanced Studies on Collaborative Research*, (CASCON '10). ACM, 2010.
- [24] G. Tamura, R. Casallas, A. Cleve, and L. Duchien. Qos contract-aware reconfiguration of component architectures using e-graphs. In *Proceedings 7th International Workshop on Formal Aspects of Component Software*, (FACS '10). LNCS, 2010. To appear.
- [25] United States Air Force Chief Scientist (AF/ST). Technology Horizons a Vision for Air Force Science & Technology During 2010-2030. Technical report, U.S. Air Force, May 2010.
- [26] J. White. Simplifying autonomic enterprise Java bean applications via model-driven development: A case study. *The Journal of Software and System Modeling*, pages 601–615, 2005.