

A Hybrid Approach Combining Control Theory and AI for Engineering Self-Adaptive Systems

Ricardo Diniz Caldas
Chalmers | University of Gothenburg
Gothenburg, Sweden
University of Brasília, Brazil
ricardo.caldas@chalmers.se

Arthur Rodrigues
University of Brasília
Brasília, DF, Brazil
arthur.farias@aluno.unb.br

Eric Bernd Gil
University of Brasília
Brasília, DF, Brazil
ericbgil@gmail.com

Genáina Nunes Rodrigues
University of Brasília
Brasília, DF, Brazil
genaina@unb.br

Thomas Vogel
Humboldt-Universität zu Berlin
Berlin, Germany
thomas.vogel@cs.hu-berlin.de

Patrizio Pelliccione
Chalmers | University of Gothenburg
Sweden
University of L'Aquila, Italy
patrizio@chalmers.se

ABSTRACT

Control theoretical techniques have been successfully adopted as methods for self-adaptive systems design to provide formal guarantees about the effectiveness and robustness of adaptation mechanisms. However, the computational effort to obtain guarantees poses severe constraints when it comes to dynamic adaptation. In order to solve these limitations, in this paper, we propose a hybrid approach combining software engineering, control theory, and AI to design for software self-adaptation. Our solution proposes a hierarchical and dynamic system manager with performance tuning. Due to the gap between high-level requirements specification and the internal knob behavior of the managed system, a hierarchically composed components architecture seek the separation of concerns towards a dynamic solution. Therefore, a two-layered adaptive manager was designed to satisfy the software requirements with parameters optimization through regression analysis and evolutionary meta-heuristic. The optimization relies on the collection and processing of performance, effectiveness, and robustness metrics w.r.t control theoretical metrics at the offline and online stages. We evaluate our work with a prototype of the Body Sensor Network (BSN) in the healthcare domain, which is largely used as a demonstrator by the community. The BSN was implemented under the Robot Operating System (ROS) architecture, and concerns about the system dependability are taken as adaptation goals. Our results reinforce the necessity of performing well on such a safety-critical domain and contribute with substantial evidence on how hybrid approaches that combine control and AI-based techniques for engineering self-adaptive systems can provide effective adaptation.

KEYWORDS

hybrid, self-adaptive software, control theory, optimization

ACM Reference Format:

Ricardo Diniz Caldas, Arthur Rodrigues, Eric Bernd Gil, Genáina Nunes Rodrigues, Thomas Vogel, and Patrizio Pelliccione. 2020. A Hybrid Approach Combining Control Theory and AI for Engineering Self-Adaptive Systems. In *IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '20)*, October 7–8, 2020, Seoul, Republic of Korea. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3387939.3391595>

1 INTRODUCTION

There is an increasing trend in the use of control theory to guide the software engineering design process of self-adaptive systems (SAS) in order to provide dynamic adaptation with guarantees of trustworthiness and robustness [16, 34, 41]. Control theory has established solid techniques for designing controllers that enforce controlled (or managed) systems to behave as expected. “These controllers can provide formal guarantees about their effectiveness, under precise assumptions on the operating conditions” [16].

At runtime, control systems are subject to inputs not known at design-time. Based on that, they are designed to respond within acceptable boundaries to dynamic environmental interactions. The analysis of whether the controlled system is operating accordingly demands metrics for comparing the performance of various control systems [30]. A self-adaptive system designed through control theory must be able to quantitatively guarantee its convergence in time to reach the adaptation goal (i.e., setpoint), and on its robustness (i.e., convergence to the setpoint) in the face of errors and noise [16, 34]. Additionally, due to high level of uncertainty in self-adaptive systems, the control model following a control-theoretical approach requires constant evolution.

Therefore, in the control-theoretical design of a self-adaptive software system, it is paramount to support decision-making procedures that can identify appropriate adaptation strategies by efficiently exploring the solution space. To avoid the problem of exhaustive analysis of large adaptation spaces when planning an adaptation, artificial intelligence (AI) techniques have been used, among others: to determine at runtime subsets of adaptation options from a

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SEAMS '20, October 7–8, 2020, Seoul, Republic of Korea
© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-7962-5/20/05...\$15.00
<https://doi.org/10.1145/3387939.3391595>

large adaptation space through classification and regression [31]; to assess and reason about black-box adaptation decisions through on-line learning [9, 10]; or to find Pareto-optimal configurations of the system with respect to the mission goals of mobile robots through machine learning [21]. Moreover, control-based approaches for self-adaptive systems often employ adaptive controllers to address inaccuracies of the control model or radical changes of the environment/controlled system by updating the model using filters (e.g., Kalman or Recursive Least Square) [15, 28, 29] and by tuning the controller parameters using feedback or machine learning [12, 27].

Although such significant contributions in the literature have promoted solid mathematical foundation of control theory for building SAS, there is still the need of understanding (i) how AI could aid the parameter optimization of controllers and (ii) how to apply AI-based techniques in the search of adaptation solutions in control-theoretical SAS strategies. Moreover, “there are various hurdles that need to be tackled to turn control theory into the foundation of self-adaptation of software systems” [40]. Software engineers are not usually knowledgeable of the mathematical principles on control theory. Therefore, there is a need to turn control theory (and its guarantees) into a scientific foundation for engineering self-adaptation [40].

In this paper, we aim at filling those gaps by intertwining control theory and AI in a two-phase optimization approach to support the engineering of control-based SAS. Our controller, hereafter named system manager, contains the mechanisms for deciding how the self-adaptive system must behave in order to achieve the desired goals for the managed system. Due to the gap between high-level requirements specification and the internal knob behavior of the managed system, a hierarchically composed components architecture seek the separation of concerns towards a dynamic solution [3]. Therefore, our system manager is further decomposed into two other components: a strategy manager and a strategy enactor. Each of these components are equipped with AI-based optimization techniques, namely regression model combined with the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [6], for respectively: (i) synthesizing adaptation strategies through high-level reasoning upon the model that represents the managed system, and (ii) enforcing actions through control-theoretical principles to ensure the adaptation strategies are applied and properties of interest guaranteed. By these means, we contribute with a hierarchical and adaptive system manager that relies on optimization at all levels of the decision-making process towards a more efficient adaptation mechanism, while also improving the self-adaptation loop in terms of control theoretic properties: stability, overshooting and steady-state error.

For evaluation purposes, we experimented with a healthcare domain prototype constructed for operating volatile environments. The prototype played the role of the target system. The Body Sensor Network (BSN) prototype was implemented following the Robot Operating System (ROS) framework and consists of a set of distributed sensors that collects the patients vital signs and forwards them to a centralized unit for further processing and reporting. We evaluated the ability of our approach to optimize the adaptation space and to improve the self-adaptation loop w.r.t. the control theoretic properties. Results show that we have found nearly optimal solutions for space and time through our hybrid approach. Not to

mention, the overshoot and steady-state error were below 3% with 100% of stability.

The remainder of this paper is organized as follows: Section 2 presents a glimpse over control theoretical principles and over evolutionary algorithms used in this work. Section 3 presents the core contribution of our work. Section 4 presents the evaluation of our proposal on the BSN prototype. Section 5 presents the most related literature work to ours. Finally, Section 6 presents the conclusion and the directions we intend to pursue as our future work.

2 BACKGROUND

In this section, we discuss the background of our work, feedback control loop and evolutionary algorithms.

Feedback Control Loop. Self-adaptation is typically realized by a feedback loop [4], which originates from control engineering [2]. As shown in Fig. 1, the *plant* is the subsystem to be controlled by a control signal from the *controller*. A reference (setpoint) is defined for a property of interest observed on the plant. The goal is that the observed property is sufficiently close to this reference despite the disturbance affecting the plant. For this purpose, the difference between the observation and reference is fed back to the controller that uses this error to decide about the value of the control signal.




Figure 1: Block diagram of a closed feedback control loop.

A self-adaptive system that uses a feedback control loop to adjust its behavior is a closed-loop system [20, 33]. Such a system can have different responses to a perturbation, that is, how the property of interest observed on the plant develops over time after the perturbation. Figure 2 illustrates a typical response to a sudden variation in the observed property. Various control metrics exist to verify how well the system is responding and if it satisfies requirements previously defined with expert knowledge (cf. [16, 20]). In this work we focus on extracting four metrics to analyze the response of the system: *stability*, *overshoot*, *settling time*, and *steady-state error* [16, 20].

Stability refers to the system reaching a steady-state equilibrium, that is, the observed property converges to a specific value, ideally the setpoint, and stays inside a previously defined stability margin around this convergence point. An overshoot denotes that the property exceeds the setpoint in the transient state, which should be typically avoided. It is measured as maximum peak value of the property relative to the setpoint. The settling time is the time required for the system to reach the steady-state equilibrium. Finally, the steady-state error measures how far from the setpoint the system converges, that is, the steady-state difference between the setpoint and observed property.

In a feedback loop, one of the most popular controllers being used is the proportional-integral-derivative (PID) controller while often

the derivative part is omitted due to difficulties in its tuning, leading to the use of PI controllers (cf. [1]). Moreover, such controllers are also widely used for adapting software systems [34]. Based on this observation, we also use a PI controller throughout this work. The continuous equation of a PI controller is shown in Equation 1, where $u(t)$ is its control action and $e(t)$ the error, both at time t . In this equation we have three constants, which are the proportional gain K_p , the integral gain K_i , and the integral window IW . K_p multiplies the current measured error directly, whereas K_i multiplies the integral of the error over the time window IW . These constants can be defined empirically or using a tuning method of preference.

$$u(t) = K_p e(t) + K_i \int_{j=t-IW+1}^t e(j) \quad (1)$$

A discrete approximation of Equation 1 is shown in Equation 2, which is adapted from the work of Stankovic et al. [39].

$$u(t) = K_p e(t) + K_i \sum_{j=t-IW+1}^t e(j) \quad (2)$$

In this equation, the integral term is transformed into a sum of errors over the time window IW . Since a PI controller is also called periodically, the time window IW is a tunable parameter, which we consider in this work as the number of error measurements to be considered in the sum. This means that in the time instant t the errors from $e(t - IW + 1)$ to $e(t)$ will be used in the integral term.

Evolutionary Algorithms. Evolutionary algorithms originate from computational intelligence, a subfield of artificial intelligence, and refer to meta-heuristic optimization techniques that imitate biological evolution [26]. A widely-used technique is the *Non-dominated Sorting Genetic Algorithm II (NSGA-II)* [6] that has been applied to various software engineering problems [19], and self-adaptive systems to optimize the configuration of the managed system [17, 37]. NSGA-II is a multi-objective algorithm producing a set of Pareto-optimal solutions to a given optimization problem. Imitating biological evolution, it evolves a population of possible solutions encoded as chromosomes by crossover, mutation, and

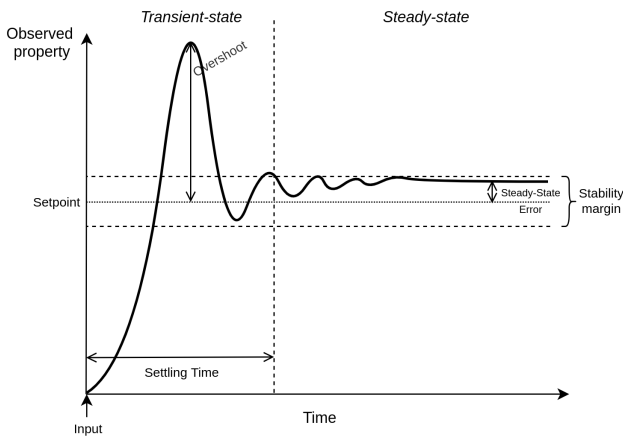


Figure 2: Typical response of a system to a sudden variation on the observed property.

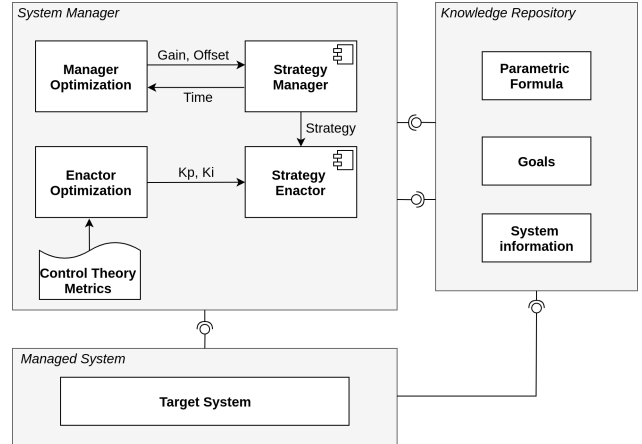


Figure 3: An architectural overview of the hybrid approach.

selection. The evolution is guided by a fitness function that encodes the objectives of the optimization and that evaluates how well a solution satisfies these objectives. Solutions with a higher fitness are preferably selected for further evolution steps in the next generation of the population. Thus, the evolution converges to solutions with higher fitness. Moreover, NSGA-II promotes diversity to obtain Pareto-optimal solutions with different trade-offs between the objectives.

3 A HYBRID APPROACH COMBINING CONTROL THEORY AND AI

In this paper, we propose an AI-based optimization approach to support the engineering of control-based SAS. Following principles of control theory, our system manager decides how the system must behave in order to achieve the desired goals for the managed system. Moreover, our system manager is enhanced with integrated AI-based optimization techniques, namely regression at all levels of the decision-making process towards a more efficient adaptation mechanism for (i) high-level reasoning upon the model that represents the managed system to synthesize adaptation strategies (cf. Section 3.1) and (ii) enforcing actions through control-theoretical principles ensuring the adaptation strategies respect the properties of interest: stability, overshooting and steady-state error (cf. Section 3.2). In our approach, an adaptation strategy is a rule that guides the system behavior. It consists of a goal to be reached, a condition that defines whether the goal is reached, and one or more actions to change the system behavior.

To provide a solution that architecturally promotes the separation of concerns in a hierarchical fashion [3], the proposed architecture relies on two horizontal layers being the *System Manager* and the *Managed System*, and a third and vertical layer being the *Knowledge Repository*. The corresponding architectural overview of our approach is shown in Fig. 3.

The *Managed System* layer contains the *Target System* being the software system to be controlled. The *Knowledge Repository* is orthogonal to the other layers and responsible for maintaining the persistent information, that is, a *parametric formula*, the *goals*, and

system information. A parametric formula evaluates a quality of service (QoS) property of the target system at runtime. For details of how such a formula is created, we refer to our previous work [38]. The goals are the objectives to be achieved by the system with respect to functional or non-functional requirements. The system information is a set of events and data about the status of the target system collected from the running system.

The *System Manager* is responsible for controlling the *Target System* and realizes the combination of control-based self-adaptation and artificial intelligence (AI), which makes our approach hybrid and improves the self-adaptation process.

For this purpose, the system manager has four components, two realizing self-adaptation in a hierarchical fashion and two realizing AI-based optimization to reconfigure the self-adaptation process. This results in a hierarchical and adaptive system manager.

The *Strategy Manager* synthesizes adaptation strategies by reasoning on the target system using the parametric formula, the goals defined by the stakeholders and the current status of the observed property. This synthesis process and especially the time needed for the synthesis of adaptation strategies is monitored by the *Manager Optimization* component. Based on the measured time, this component optimizes the configuration of the strategy manager by providing optimal values for two parameters (*Granularity* and *Offset*) that influence the synthesis process. The goal of the optimization is to reduce the required time and thus, the efficiency of the synthesis process. The strategy manager provides a synthesized strategy to the *Strategy Enactor*.

The strategy enactor enforces an adaptation strategy in a closed-loop fashion by implementing a feedback control loop on top of the target system. Accordingly, the strategy enactor implements a PI controller as an actuation mechanism for enforcing behavior at the target system, which is guided by conditions and goals. This self-adaptation process is monitored by the *Enactor Optimization* component, that computes various *Control Theory Metrics* based on the *System Information* collected from the running target system. Moreover, the component uses the metrics to tune the strategy enactor adaptation parameters. Upon the metrics and system information, the optimizer tunes the PI controller parameters (K_p and K_i). The goal of this optimization is to improve the adaptation quality with respect to the control theory metrics.

As a result of combining each, the strategy manager and the strategy enactor with an optimization, the goals of the system are guaranteed to be constantly achieved with optimal adaptation search time and adaptation quality, in a limited search space. Both optimizations are realized by AI-based techniques of model learning (regression / curve fitting) and meta-heuristic search (evolutionary algorithms), while the strategy enactor is based on control (PI controller). By this combination, the overall controller (i.e., the system manager) implements a hybrid approach of control and AI.

We evaluate our approach with the Body Sensor Network (BSN)¹, being an exemplar for a target system that monitors and analyzes the health of patients. The BSN is composed of a set of distributed sensors to monitor vital signs of patients and to forward these signs to a centralized processor for analysis. To avoid overload of the centralized processor, the sampling rate of individual sensors can

be adapted, which effects the amount of data sent to and analyzed by the processor.

In the following we detail the strategy manager and strategy enactor of our hybrid approach emphasizing their optimization.

3.1 Strategy Manager and its Optimization

Responsible for high-level reasoning, the strategy manager must be capable of performing expensive decision-making computation. Identifying feasible and effective adaptation strategies is particularly difficult as the size of the solution space grows exponentially with the number of individual adaptation options. For instance, for architectural adaptation the reasoning has to take the variability of each component (e.g., alternative components or parameters of the component) and the composition of components into account to evaluate whether the overall architectural configuration satisfies the goals of the system.

In previous work [38] we have developed a transformation framework for specifying requirements hierarchically in a goal model, translating the goal model to parametric formulae with symbolic model checking, and evaluating these formulae at runtime to express probabilities over the fulfillment of the goals by the running system. Thus, these formulae provide guarantees for satisfying the goals while their evaluation is time efficient (constant computational time) since a costly model exploration as in typical model checking is avoided. The goal modeling and the translation of the goal model to parametric formulae are implemented in the GODA-MDP tool² [38]. A parametric formula is essentially an algebraic equation that relates the probability of successfully achieving a system's overall goal (i.e., the root goal of the goal model) to the combined probabilities of lower-level tasks that contribute towards realizing the goal. Therefore, a formula can be used in the analysis and planning stages of a feedback loop to solve a satisfiability problem, where an appropriate combination of successful local tasks would lead to satisfying a required global property. Searching for such a combination corresponds to the synthesis of adaptation strategies since the combination corresponds to a configuration that can be enacted on the target system. Although using the parametric formula reduces the complexity and time for reasoning, such formulae are still not free from problem of combinatory state-space explosion.

The pseudocode of Algorithm 1 details the search process to solve a parametric formula captured in a *model*. The search process has two input parameters, the *granularity* and *offset* that are used for reducing or widening the solution/search space. Lines 2-4 capture the monitoring of a set of QoS properties P of the target system, the calculation of the current QoS property p_{curr} , and the calculation of the *error*, that is, how far from the setpoint p_{ref} the current property value p_{curr} is. Then, lines 5-24 represent the search for a combination of independent terms (p_x), for each p_i in the set of all properties P , that would lead to the goal (p_{ref}). To do so, lines 6-7 reset all p_x into the value defined by the *offset* and the set of new values (P_s) is plugged into the model. If the error turns to be positive, then the value of each of the properties in P_s should be incremented with the error multiplied by a factor dictated by the granularity *gran* until they, one by one, reach values that approximate the

¹Access <http://bodysensornetwork.herokuapp.com> for an executable artifact.

²<https://pistar-goda.herokuapp.com/>

Algorithm 1 - Strategy Manager Procedure

Input: $gran, offset$
Output: $strategy$

```

1: procedure SEARCH FOR STRATEGY
2:    $P \leftarrow MonitorSystem$ 
3:    $p_{curr} \leftarrow Apply(P, model)$ 
4:    $error \leftarrow p_{ref} - p_{curr}$ 
5:   for each  $p_i$  in  $P$  do
6:      $P_s \leftarrow ResetToStartingPoint(P, offset)$ 
7:      $p_{new} \leftarrow Apply(P_s, model)$ 
8:     if  $error > 0$  then
9:       do  $p_i \leftarrow p_i + gran \cdot error$  until  $p_{new} > p_{ref}$ 
10:      where  $p_{new} \leftarrow Apply(P_s, model)$ 
11:      for each  $P_j$  in  $P_s - p_i$  do
12:        do  $p_j \leftarrow p_j + gran \cdot error$  until  $p_{new} > p_{ref}$ 
13:        where  $p_{new} \leftarrow Apply(P_s, model)$ 
14:      end for
15:    else
16:       $p_i \leftarrow p_i - gran \cdot error$  until  $p_{new} < p_{ref}$ 
17:      where  $p_{new} \leftarrow Apply(P_s, model)$ 
18:      for each  $p_j$  in  $P_s - p_i$  do
19:        do  $p_j \leftarrow p_j - gran \cdot error$  until  $p_{new} < p_{ref}$ 
20:        where  $p_{new} \leftarrow Apply(P_s, model)$ 
21:      end for
22:    end if
23:     $Strategies \leftarrow Strategies + P_s$ 
24:  end for
25:  return  $strategy$  from  $Strategies$ 
26: end procedure

```

dependent value to the goal (see lines 8-15). On the other hand, the properties in P_s are decremented by the same factor (lines 15-22). At line 23, the modified set of P_s is concatenated into the valid set of $Strategies$. In the end, the best or a sufficient strategy is chosen among the valid ones (line 25).

In this search process, the procedure *Apply* matches the independent variables of the model/formula to each respective value, and performs the calculation that results in a value for the dependent variable, that is, the procedure evaluates the formula. Thus, it is used for quantitative reasoning on the system state. As a result, the algorithm relies on performing a search within the combination of all independent variables that would lead the system into reaching the goal. The search itself increases/decreases the value of each independent variable at a time, by gradually changing its value with a step dictated by a factor of *granularity* · *error*.

The granularity and the offset parameters are fundamental to determine whether the strategy manager will be able to find a combination of values (i.e., to converge to a solution) that satisfies the goal of the system, and to find such a combination in time. The smaller the granularity is, the algorithm will more likely find a solution, even though it would take more time to do so. In contrast, the offset broaden the search space boundaries, so the bigger it is the more values will the search go through until it reaches the end line. Consequently, we need a method to enhance the choice of such

parameters, and this is where we employ AI-based optimization.

Manager Optimization. In the following, we discuss the optimization of the search process conducted by the strategy manager to synthesize adaptation strategies. The optimization aims for identifying suitable values for the *granularity* and *offset* parameters (cf. Algorithm 1). The optimization of the strategy manager is a single-objective parameter optimization problem since the goal is to minimize the time required by the strategy manager to synthesize an adaptation strategy. This problem is solved by the optimization pipeline depicted in Figure 4. This pipeline consists of the following four steps:




Figure 4: Optimization pipeline for strategy manager.

- 1. Collect Data:** We collect data relating the time to find a solution (t_s) with values for the granularity ($gran$) and *offset* parameters. The variable t_s represents the time taken by the strategy manager to converge to the setpoint. The data is obtained from executing our approach without the optimization in place. Thus, we execute our approach in different scenarios, that is, with different static choices of parameter values for granularity and *offset*.
- 2. Learn Model:** From the collected data $\langle time, gran, offset \rangle$, we learn a model using the curve fitting method. The learned model describes the behavior of the strategy manager as a function $t_s(gran, offset)$. Thus, we can use the model to estimate the time to synthesize an adaptation strategy given concrete values for the granularity and *offset* parameters. An example for such a model is $time = 677 * e^{gran * (-3000)} + e^{offset * (-2874)}$.
- 3. Find Optimal Configuration:** We use the learned model as an input to an evolutionary algorithm, particularly NSGA-II, to find optimal values for the granularity and *offset* parameters in order to minimize the time to synthesize an adaptation strategy. For the evolutionary algorithm, a candidate solution is encoded as a chromosome or a vector two components, one component for each parameter. The fitness function is the minimization of the time. The fitness of a specific candidate solution, that is, a concrete pair of two real values for the two parameters, corresponds to computing the learned function with this pair of values. Moreover, the optimization is constrained by the problem domain, as well as the range of the granularity and *offset* parameters. The optimization performed by evolutionary algorithm results in an optimal or near-optimal configuration (parameter setting) for the strategy manager. To illustrate, let's bring back our previous example model, $time = 677 * e^{gran * (-3000)} + e^{offset * (-2874)}$. Initially, a population of one hundred individuals is generated randomly, having the target values and/or input parameters limited by a given range. An individual is composed of a combination of genes and a calculated target value. In NSGA-II, the genes are represented by the input parameters, in this case *gran* and *offset*. The target value in this step is the *time* variable. Better ranked solutions, i.e., solutions with the best fitness values, are observed in individuals with a lower *time* to find a solution. For

the optimization of both, strategy manager and strategy enactor stages, we have set the algorithm to process a hundred generations. Moreover, we have applied some standard operators like Probability Matching (PM), Simulated Binary Crossover (SBX), and Tournament Selector based on dominance comparison (Pareto).

4. Apply optimal configuration: Finally, we apply the optimal or near-optimal configuration, that is, a pair of values (*gran*, *offset*), to the strategy manager affecting the performance of the synthesis of adaptation strategies.

Delving into the optimization pipeline, we start describing the first part of the *learning*³ process: the *manager optimization*. As mentioned before, the manager behavior is ruled by two features: (i) *gran*, that is the granularity with which the strategy manager will search for new solutions, and (ii) *offset*, representing the starting point of the search within the solution space. Both attributes are crucial to determine a third variable of interest, the (iii) *time to solution*, i.e., how long the manager will take to find a suitable adaptation strategy. From a naive perspective, these values are often defined in an uninformed fashion, before the system operation, and remain unchanged despite the characteristics of the system, or the particularities of the domain. Our proposal tweaks this logic by tracing the strategy manager behavior from its data collected offline. Knowing the fact that the target variable depends on the *gran* and *offset* values, we conduct a curve fitting process to come up with models that explain the *time to solution* variable in terms of the aforementioned parameters. Such a process returns a mathematical function that has the best fit to the data points collected from previous executions of the strategy manager. If the learned function being the output of the curve fitting process is simple, it is possible to solve the optimization problem mathematically. Otherwise, we solve this problem using an evolutionary algorithm, like the NSGA-II.

We should note that our approach has been conceived to provide informed choices about granularity, offset, K_p and K_i at design time. Therefore, the amount of data used for training is not a major concern. A better description of the system behavior allows the engineers to make better choices of control theoretical metrics. It may be the case that the system behavior changes in face of uncertainty. Clearly, learning online from data and continuously feeding the system with new control theoretical metrics may leverage the strength of our approach. However, adopting an online learning approach brings additional complexity to the pipeline, which is out of scope in this work.

3.2 Strategy Enactor and its Optimization

The strategy enactor is responsible for enforcing an adaptation strategy on the target system. The strategies are synthesized by the strategy manager and must be read, interpreted, and enforced by the strategy enactor. Under the guidance of the active strategy, the enactor continually evaluates whether the target system behaves as it is intended to. Otherwise, it adapts the system towards achieving its goals. Whenever the enactor cannot enforce the active strategy,

³Strictly speaking, in this stage the optimization modules do rely on previous adaptation data to learn the behavior and optimize the strategy manager parameters. However, we reckon that the use of term *learning* to describe curve fitting methods and evolutionary algorithms is still disputable.

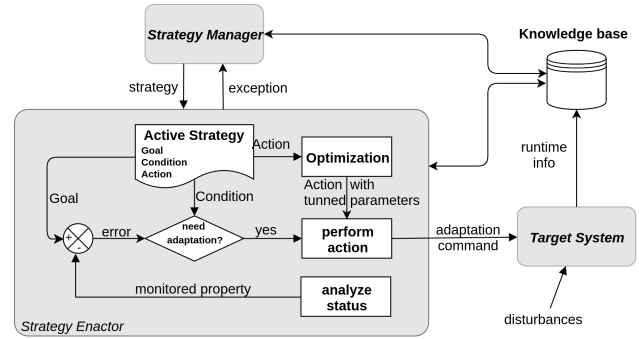


Figure 5: Overview of the Strategy Enactor.

an exception is propagated to the strategy manager in an attempt to receive a new and more adequate adaptation strategy from the strategy manager.

Moreover, the strategy enactor realizes a negative feedback loop. In this sense, the enactor continuously monitors the system QoS properties, analyzes them, and checks whether its status is compliant to the goal demanded by the active strategy. As previously discussed, a strategy is defined by a **goal**, a **condition**, and a **set of actions** to achieve a desired configuration of the target system.

Figure 5 depicts the configurable feedback loop for enforcing strategies on the target system. The strategy enactor collects information from the target system and analyzes the QoS property of interest. If the monitored property status does not correspond to the desired goal under the defined condition then the enactor enforces one or more actions. Whereas the actions are optimized accordingly. In this context, we rely on control theory to implement the function that determines the control signal by using a PI controller. In this work, we use the BSN exemplar as a target system, whose knobs are the sampling rates of the sensors, which are adapted based on the reliability of the BSN’s central processing component affected by the load on this component.

The synthesized adaptation strategy guides the parameters of the closed loop and how the adaptation is executed. To this extent, the PI controller is an instance of the element that enforces the control signal on the target system, characterized as the “perform action” in Fig. 5. Thus, it receives as input the difference between the current state of the observed system property and the goal of the adaptation strategy (error), and multiplies the error by factors of the gains K_p and K_i coming from the optimized action, according to Eq. 2. In the case of our running example, the PI controller would trigger an adaptation command with a description of the calculated new sampling rate for a sensor (e.g., the thermometer of the patient). Depending on the values of K_p and K_i , the increments or decrements on the actuation variable could not guide the property into converging to the setpoint, they could request too much effort from the system, or they could take too long to reach the setpoint. Foremost, the lack of information on a model that relates K_p , K_i , stability, overshoot and steady-state-error demands an optimization of K_p and K_i towards optimal performance.

Enactor Optimization. The designed PI controller is not sufficient to guarantee optimal performance. Therefore, we employ an AI-based optimization technique to tune the controller’s parameters with respect to the desired control-theoretical properties of *settling time*, *overshoot*, and *steady-state error*. In this work, we focus on the *overshoot* and *steady-state error*. These parameters are the proportional gain K_p and the integral gain K_i of the PI controller. Thus, the optimization problem is to find optimal values for K_p and K_i that improve the self-adaptation in terms of *overshoot* and *steady-state error*. These two control properties should be minimized resulting in a multi-objective optimization problem.




Figure 6: Optimization pipeline for strategy enactor.

The optimization pipeline for the strategy enactor is shown in Fig. 6. Except for slight differences, this pipeline and the employed techniques are the same as for the strategy manager (cf. Fig. 4). Considering Fig. 6, the control-theoretical metrics are collected from the behavior of the system during execution. From this data we get the data points relating K_p and K_i to settling time, overshoot, and steady-state error. Using this data, we adopt a curve fitting strategy to come up with three mathematical functions that describe the relationship between the parameters K_p and K_i and each control-theoretical metric: overshoot (OS) and steady-state error (SSE) (see second step in Fig. 6). The type of the curve (e.g., linear, polynomial, exponential, etc.) depends on the relation revealed by the data points. This is illustrated in Fig. 7 for the BSN case study.



Figure 7: Model learning by curve fitting the behavior of the strategy enactor.

Once we have these functions, it is possible to feed the NSGA-II with them in order to find values of K_p and K_i that minimize the overshoot and steady-state error. Since we have now a multi-objective optimization problem, the meta-heuristic will not return

just a single pair of optimal values for K_p and K_i , but a set of Pareto-optimal solutions, i.e., several pairs of (K_p, K_i) . These results are optimal or near-optimal solutions trading off the two objectives, overshoot and steady-state error as depicted in Fig. 8.

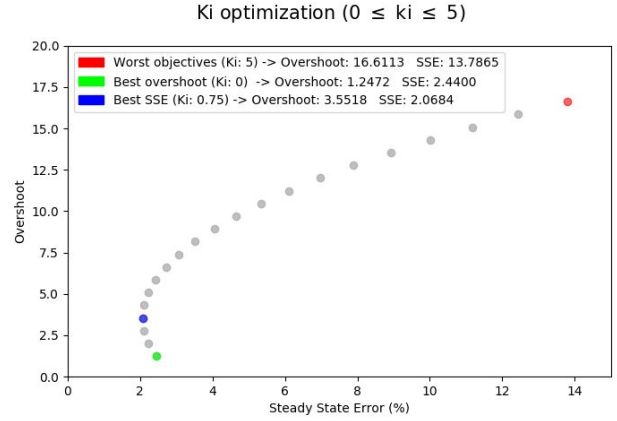


Figure 8: Multi-objective optimization example

Finally, we update the parameters of the PI controller employed in the strategy enactor with K_p and K_i values that optimize the gains that lead to lower and more balanced overshoot and steady-state error.

4 EXPERIMENTAL RESULTS

Our proposal is evaluated through the execution⁴ of an actual prototype from the healthcare domain (BSN) as the target system alongside an implementation⁵ of the system manager components upon ROS. The evaluation itself follows the Goal-Question-Metric methodology [5]. Our focus is on answering the guiding questions, from a broader perspective (i) how could AI aid the parameter optimization controllers, and (ii) how to apply AI-based techniques in the search for adaptation solutions in control-theoretical SAS strategies. Thus, we narrow down the questions towards an operational scenario without loss of generality, see Table 1.

G1: Improve time to find an eligible adaptation	
Questions	Metrics
Q1: How much time does the optimization pipeline take to find a solution?	Duration [s]
G2: Improve quality of the adaptation	
Question	Metric
Q2: How efficient is the proposed learning and optimization?	CT metrics

Table 1: Goal-Question-Metric

The first goal (G1) stems from the fact that the proposed AI-based optimization pipeline aims at providing support to the search for

⁴Configurations for experiments: CPU 2x Intel(R) Core(TM) i7-8550U CPU @ 1.80GHz, 8192MB RAM, Ubuntu 18.04.3 LTS, GNU C Compiler version 7.4.0

⁵<https://www.github.com/lesunb/bsn>

an optimal adaptation task by reducing time to solution at runtime. Therefore, to support **G1**, we collect the duration of the optimization pipeline for both components, Strategy Manager and Strategy Enactor. The optimization pipeline relies on the identification of parameters through the collection of control theoretical metrics from QoS attributes of the system at runtime. Thus, with **G2**, we want to guarantee that our approximation does lead to efficient solutions. By this means, we measure the efficiency of our method by the reached control theoretical metrics with optimized parameters configuration and the best solutions given with the exhaustive search algorithm.

Under the light of the GQM, the experiment is performed in complementary phases (i) time and space evaluation of the search for the optimal adaption, and (ii) evaluation of the adaptation quality with distinct configurations. Thus, following sections further describe the experimental setup, the operational scenario under which the experiments were conducted, and finally the results.

4.1 Time and space to find an eligible adaptation

In this phase, we have first instrumented the pipeline with a chronometer to measure the duration of the learning model step and the finding an optimal configuration step, from both strategy manager instance and strategy enactor instances.

For the strategy manager, the time to perform the curve fitting (time to solution, gran and offset) was in average 0.3099 seconds per execution, whereas the curve type was classified as exponential. The NSGA-II took in average 0.0023 seconds per execution. Totaling in 0.3123 seconds per execution for the strategy manager pipeline.

For the strategy enactor, four curve fittings had to be realized (i) Kp-SSE, (ii) Kp-Overshoot, (iii) Ki-SSE and (iv) Ki-Overshoot. For (i) the algorithm took 0.01295 seconds to fit the model into a quadratic function, (ii) it took 0.0498 seconds to fit into an exponential, (iii) it took 0.0127 seconds to fit into a quadratic and (iv) it took 0.0075 seconds to fit into a linear model. Furthermore, the NSGA-II was executed twice (v) Kp-SSE Kp-Overshoot and (vi) Ki-SSE & Ki-Overshoot. For (v) the algorithm took 5.3947 seconds and (vi) it took 4.896 seconds. Totaling in 10.374099032 seconds per execution for the strategy manager pipeline.

Our execution of NSGA-II makes ten thousand function evaluations. However, this amount can be reduced to one thousand without great quality loss if the execution time of the pipeline needs to be improved.

4.2 Quality of the adaptation

We have deployed and executed the prototype under a scenario where at least one adaptation was triggered for each five minutes execution. Then, 50 randomly selected combinations of K_p and K_i , namely configurations, were executed. The 50 configurations are a product of the combination of two sets $K_p = [60,150]$ with steps of 10, $K_i = [0.2,1]$ with steps of 0.2 and $IW = 5$. Once we select an optimal configuration with the AI technique, another scenario is performed, totaling 51 configuration scenarios. The configuration selected by the optimization is $K_p = 114.0$, $K_i = 0.75$ and $IW = 5$.

The scenario was derived from the running example, where the stakeholders wish to maintain the system's reliability level at 95%. The maintenance process is held by the system manager, as follows.

The system manager must continuously monitor the target system's reliability status during runtime. Moreover, the prototype was developed to satisfy complementary tasks with single responsibility components. For that reason, the reliability is locally calculated for each component and composed through a model that relates their probability of success, i.e. the reliability formula from previous works [38]. To keep track of the system goals, the system manager systematically analyzes whether the monitored reliability status fulfills them. Where in case of violations to the adaptation strategy conditions the system managers adapts the target system to maintain the reliability level accordingly to the stipulated goal.

In the experimental sessions, we simulated a scenario where multiple sensors flood the central processor with a higher rate than what it can handle. On top of that, noise is generated by an external agent that triggers random failures on the sensors preventing them from forwarding data to the central processor. The simulated failures reduce the number of messages received by the central processor at a random factor, placing another layer of uncertainty on the system execution. To cope with the disturbances into the reliability status, the system manager delivers reconfiguration signals containing increments/decrements to the sensors' sampling rate that could give more or less time for the central processor to maintain a always flowing processing queue.

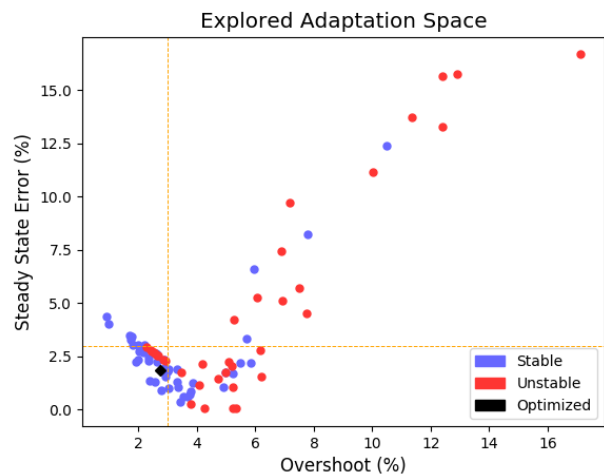
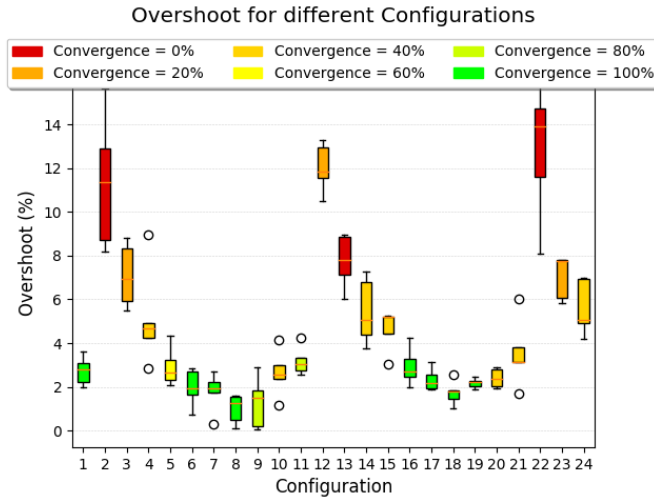


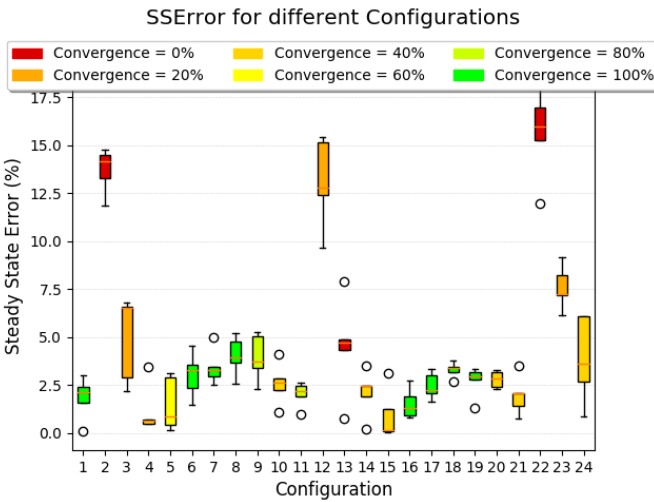
Figure 9: Explored adaptation space.

The SAS was instrumented to log information regarding the quality of the adaptations w.r.t the control theory metrics. The outcoming data is processed accordingly to each evaluation and is presented in Figures 9 and 10.

Figure 9 shows the results of the adaptation space explored by the naïve (in red and blue) and our approach (in black). Our approach was able to meet the nearly optimal adaptation space for the threshold of steady state error and overshoot of 3%. Figures 10a and 10b provides a bigger picture of the experiments over the



(a) Overshoot for distinct configurations.



(b) Steady-state-error for distinct configurations.

Figure 10: Boxplots for the sample of 24 configurations

50 configurations and the optimized one. For the sake of space, we refer only to the first 24 configurations, where configuration 1 refers to our approach. Results for the overshoot and steady-state error show that our approach managed to perform among the best results, reaching out convergence point of 100%. In other words, our approach not only managed to optimize steady-state and overshoot but also the stability of our BSN. By these means, the self-adaptation loop of the BSN from the control theoretic principles can be considered quite robust for the scenarios evaluated.

4.3 Threats to Validity

Internal validity. The main threat to internal validity is the data that we collected from the scenarios of executing our approach and that we used for the model learning and the optimization. Different scenarios may lead to different data and thus, to different

results of our approach. Moreover, we rely on specific techniques for the model learning (curve fitting) and optimization (NSGA-II). Other techniques could be used, which may lead to different results. In this context, we have not tuned the techniques, especially the hyper-parameters of NSGA-II, to avoid any bias concerning BSN. **External validity.** Although we present a generic self-adaptation approach that combines AI and control-theoretical principles, our evaluation focuses on the BSN case study with a single non-functional requirement (reliability) as the adaptation concern. Thus, we cannot generalize our results to other concerns (e.g., performance, security, or costs) and target systems, which calls for further experiments. **Construct validity.** The main threat is the correctness of the implementation of our approach and the BSN used for the evaluation. Concerning our approach, at least two authors of this paper have reviewed the implementation and checked the plausibility of the evaluation results based on the experience they have with the BSN from earlier work [32, 38]. With this experience we are confident about the validity of the BSN that is publicly available.

5 RELATED WORK

As stated by Weyns [40], the use of control theory in the design of self-adaptive software can bring several benefits since it allows providing analytical guarantees for several system properties such as stability and robustness. Consequently, the idea of using control-based approaches to achieve self-adaptation for software systems has been widely studied [34], and the use of control theory as a well-suited solution for self-adaptation to systematically meet the adaptation goals despite a certain degree of uncertainty of the system and environment has been recognized [8, 11, 15, 16].

Control-based self-adaptation approaches typically propose controller synthesis to automatically construct a controller for managing the software system’s adaptation needs [13, 14, 35]. The resulting controllers are related to the Strategy Enactor in our work. These approaches to synthesize these controllers could be applied for the Strategy Enactor whose focus is on using a control-theoretical controller. However, our work presents a further module called Strategy Manager that synthesizes an adaptation strategy to which the controller subsequently adheres. As previously stated, the combination of these two modules is the essence of our work, where self-adaptations is realized at two different levels: we first use a global goal, set by users, to come up with an adaptation strategy, and then we adapt the target system’s parameters based on goals, used as setpoints, conditions and actions defined in the strategy.

Moreover, we use AI-based learning and optimization techniques in our work to reduce the adaptation space at the strategy manager level and to optimize the controller’s parameters at the strategy enactor level. For the latter level, the focus is on meeting acceptable system behavior with respect to control properties such as overshoot and steady-state error. In this context, other approaches use learning, online or offline, and other statistical methods to reduce the adaptation space, often called configuration space, in order to find a configuration for the system in available time [9, 10, 18, 21–23, 31]. The main difference to our work is that improve efficiency of synthesizing adaptation strategies by optimizing the search process (in terms of granularity and offset) exploring the adaptation

space. Thus, we optimize the exploration of the search space rather than explicitly reducing the space.

Similarly to the strategy enactor in our work, existing control-based approaches for self-adaptation use PI or PID controllers that are adapted at run-time (cf. [34]), to compensate inaccuracies of the model or to handle radical changes. Particularly, values of parameters in the model are estimated and updated at run-time based on measurements often processed by filter algorithms such as Kalman or Recursive Least Square [13, 24, 25, 29, 35, 36], or controller parameters are tuned online by relay feedback [12] or machine learning [27], or offline by experiments [7]. In contrast, in our approach the controller parameters are tuned by an evolutionary algorithms identifying (near-)optimal values for them. Evolutionary algorithms have been applied previously to self-adaptive systems in order to find optimal configurations of the target system [17, 37] while we apply them to improve the self-adaptation performed by the strategy manager and strategy enactor.

To the best of our knowledge, our work is a novel approach that combines the use of algebraic models (i.e., the parametric formulae) of a system, which are used to synthesize adaptation strategies based on a non-functional requirement, and control theory, which is used to adapt the target system's parameters systematically based on the previously defined strategy. Also, we combine AI-based optimization with these two approaches to adapt the system in an efficient way while also improving the self-adaptation in terms of the control-theoretical properties: overshoot and steady-state error.

6 CONCLUSION AND FUTURE WORK

In this work, we have proposed a hybrid approach that combines control theory principles with AI techniques to optimize the adaptation process in self-adaptive systems. Using curve fitting methods aligned with a meta-heuristic technique, namely NSGA-II, we were able to i) efficiently synthesize adaptation strategies through high-level reasoning upon the model that represents the managed system, and ii) enforce actions through control-theoretical principles to ensure the adaptation strategies are applied and properties of interest guaranteed. By these means, we contribute with a hierarchical and dynamic system manager that relies on optimization at all levels of the decision-making process towards a more efficient and robust adaptation mechanism. We have evaluated our approach on the BSN prototype implemented in the ROS framework. The evaluation results have shown that our hybrid approach is able to find optimal solutions for the adaptation space while also improving the self-adaptation loop in terms of control theoretic properties: stability, overshooting and steady-state error.

We envision that our future work will be devoted into expanding our technique to accommodate PID controllers as well as exploiting further evolutionary algorithms, while ensuring our approach is also able to optimize the adaptation space at runtime. During the experiments, we have noticed that defining a suitable data window in which the method should train the model is a hard task. A second issue is knowing when the learned model should be applied. Some kind of quality monitor is likely to be incorporated to the architecture in order to orchestrate the learning process. From a performance perspective, this brings us to another challenge, that is building the pipeline in a way that the continuous learning and

adaptation of metrics do not affect negatively the operation of the managed system. We plan to further investigate these issues in a future work. Additionally, for generalization purposes, we plan to conduct case studies with other exemplars from other domains than healthcare as well as to further compare our solution to state-of-the-art adaptive control algorithms with respect to time to solution and robustness.

ACKNOWLEDGMENT

The authors express their utmost gratitude to Léo Moraes and Gabriel Levi (UnB/Brazil) for implementing an accessible version of the BSN for experimentation on SAS. This study was financed in part by the CAPES-Brasil – Finance Code 001, through CAPES scholarship. This work was also partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, the FLASH project (GR 3634/6-1) funded by the German Science Foundation (DFG), and the EU H2020 Research and Innovation Prog. under GA No. 731869 (Co4Robots). The authors also acknowledge financial support from Centre of EXcellence on Connected, Geo-Localized and Cybersecure Vehicle (EX-Emerge) funded by Italian Government under CIPE resolution n. 70/2017 (Aug. 7, 2017). Finally, we thank CNPq for partial support under grant number 306017/2018-0.

REFERENCES

- [1] Kiam Heong Ang, Gregory Chong, and Yun Li. 2005. PID control system analysis, design, and technology. *IEEE transactions on control systems technology* 13, 4 (2005), 559–576.
- [2] K.J. Åström and R.M. Murray. 2010. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press. http://www.cds.caltech.edu/~murray/amwiki/index.php/Main_Page
- [3] Victor Braberman, Nicolas D'Ippolito, Jeff Kramer, Daniel Sykes, and Sebastian Uchitel. 2015. MORPH: A Reference Architecture for Configuration and Behaviour Self-adaptation (*CTSE 2015*). ACM, New York, NY, USA, 9–16.
- [4] Yuriy Brun, Giovanna Di Marzo Serugendo, Cristina Gacek, Holger Giese, Holger Kienle, Marin Litoiu, Hausi Müller, Mauro Pezzè, and Mary Shaw. 2009. Engineering Self-Adaptive Systems through Feedback Loops. In *Software Engineering for Self-Adaptive Systems*, Betty H. C. Cheng, Rogério de Lemos, Holger Giese, Paola Inverardi, and Jeff Magee (Eds.). Springer, 48–70.
- [5] Victor R Basili, Gianluigi Caldiera, and H Dieter Rombach. 1994. The goal question metric approach. *Encyclopedia of software engineering* (1994), 528–532.
- [6] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
- [7] David Desmeurs, Cristian Klein, Alessandro Vittorio Papadopoulos, and Johan Tordsson. 2015. Event-Driven Application Brownout: Reconciling High Utilization and Low Tail Response Times. In *Proceedings of the 2015 International Conference on Cloud and Autonomic Computing (ICCAC '15)*. IEEE, 1–12.
- [8] Yixin Diao, Joseph L Hellerstein, Sujay Parekh, Rean Griffith, Gail Kaiser, and Dan Phung. 2005. Self-managing systems: A control theory foundation. In *12th IEEE International Conference and Workshops on the Engineering of Computer-Based Systems (ECBS'05)*. IEEE, 441–448.
- [9] Ahmed Elkhodary, Naem Esfahani, and Sam Malek. 2010. FUSION: A Framework for Engineering Self-tuning Self-adaptive Software Systems (*FSE '10*). ACM, New York, NY, USA, 7–16.
- [10] N. Esfahani, A. Elkhodary, and S. Malek. 2013. A Learning-Based Framework for Engineering Feature-Oriented Self-Adaptive Software Systems. *IEEE Transactions on Software Engineering* 39, 11 (Nov 2013), 1467–1493.
- [11] Antonio Filieri, Carlo Ghezzi, Alberto Leva, and Martina Maggio. 2011. 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*. IEEE Computer Society, 283–292.
- [12] Antonio Filieri, Carlo Ghezzi, Alberto Leva, and Martina Maggio. 2012. Reliability-Driven Dynamic Binding via Feedback Control. In *Proceedings of the 7th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '12)*. IEEE, 43–52.
- [13] Antonio Filieri, Henry Hoffmann, and Martina Maggio. 2014. Automated design of self-adaptive software with control-theoretical formal guarantees. In *Proceedings*

- of the 36th International Conference on Software Engineering. ACM, 299–310.
- [14] Antonio Filieri, Henry Hoffmann, and Martina Maggio. 2015. Automated multi-objective control for self-adaptive software design. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*. ACM, 13–24.
- [15] Antonio Filieri, Martina Maggio, Konstantinos Angelopoulos, Nicolás D'ippolito, Ilias Gerostathopoulos, Andreas Berndt Hempel, Henry Hoffmann, Pooyan Jamshidi, Evangelia Kalyvianaki, Cristian Klein, et al. 2015. Software engineering meets control theory. In *Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. IEEE Press, 71–82.
- [16] Antonio Filieri, Martina Maggio, Konstantinos Angelopoulos, Nicolás D'ippolito, Ilias Gerostathopoulos, Andreas Berndt Hempel, Henry Hoffmann, Pooyan Jamshidi, Evangelia Kalyvianaki, Cristian Klein, Filip Krikava, Sasa Misailovic, Alessandro V. Papadopoulos, Suprio Ray, Amir M. Sharifloo, Stepan Shevtsov, Mateusz Ujma, and Thomas Vogel. 2017. Control Strategies for Self-Adaptive Software Systems. *ACM Trans. Auton. Adapt. Syst.* 11, 4, Article 24 (Feb. 2017), 31 pages.
- [17] Erik Fredericks, Ilias Gerostathopoulos, Christian Krupitzer, and Thomas Vogel. 2019. Planning as Optimization: Dynamically Discovering Optimal Configurations for Runtime Situations. In *13th International Conference on Self-Adaptive and Self-Organizing Systems (SASO '19)*. IEEE, 1–10.
- [18] Ilias Gerostathopoulos, Christian Prehofer, and Tomas Bures. 2018. Adapting a system with noisy outputs with statistical guarantees. In *Proceedings of the 13th International Conference on Software Engineering for Adaptive and Self-Managing Systems*. ACM, 58–68.
- [19] Mark Harman, S. Afshin Mansouri, and Yuanyuan Zhang. 2012. Search-Based Software Engineering: Trends, Techniques and Applications. *ACM Comput. Surv.* 45, 1, Article Article 11 (2012), 61 pages.
- [20] J. Hellerstein, Y. Diao, S. Parekh, and D. Tilbury. 2004. *Feedback Control of Computing Systems*. John Wiley Sons, Inc., USA.
- [21] Pooyan Jamshidi, Javier Cámara, Bradley Schmel, Christian Kästner, and David Garlan. 2019. Machine Learning Meets Quantitative Planning: Enabling Self-adaptation in Autonomous Robots (SEAMS '19). IEEE Press, Piscataway, NJ, USA, 39–50.
- [22] Pooyan Jamshidi and Giuliano Casale. 2016. An uncertainty-aware approach to optimal configuration of stream processing systems. In *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE, 39–48.
- [23] Pooyan Jamshidi, Miguel Velez, Christian Kästner, Norbert Siegmund, and Prasad Kawthekar. 2017. Transfer learning for improving model predictions in highly configurable software. In *Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. IEEE, 31–41.
- [24] Cristian Klein, Martina Maggio, Karl-Erik Arzén, and Francisco Hernández-Rodríguez. 2014. Brownout: Building More Robust Cloud Applications. In *Proceedings of the 36th International Conference on Software Engineering (ICSE 2014)*. ACM, 700–711.
- [25] Cristian Klein, Alessandro Vittorio Papadopoulos, Manfred Dellkrantz, Jonas Dürango, Martina Maggio, Karl-Erik Arzén, Francisco Hernández-Rodríguez, and Erik Elmroth. 2014. Improving Cloud Service Resilience Using Brownout-Aware Load-Balancing. In *Proceedings of the 2014 IEEE 33rd International Symposium on Reliable Distributed Systems (SRDS '14)*. IEEE, 31–40.
- [26] Rudolf Kruse, Christian Borgelt, Christian Braune, Sanaz Mostaghim, Matthias Steinbrecher, Frank Klawonn, and Christian Moewes. 2016. *Computational Intelligence: A Methodological Introduction* (2nd ed.). Springer.
- [27] Palden Lama and Xiaobo Zhou. 2013. Autonomic Provisioning with Self-Adaptive Neural Fuzzy Control for Percentile-Based Delay Guarantee. *ACM Trans. Auton. Adapt. Syst.* 8, 2, Article Article 9 (2013), 31 pages.
- [28] Martina Maggio, Henry Hoffmann, Alessandro V. Papadopoulos, Jacopo Panerati, Marco D. Santambrogio, Anant Agarwal, and Alberto Leva. 2012. Comparison of Decision-Making Strategies for Self-Optimization in Autonomic Computing Systems. *ACM Trans. Auton. Adapt. Syst.* 7, 4, Article Article 36 (2012), 32 pages.
- [29] Martina Maggio, Cristian Klein, and Karl-Erik Arzén. 2014. Control strategies for predictable brownouts in cloud computing. *IFAC Proceedings Volumes* 47, 3 (2014), 689–694. 19th IFAC World Congress.
- [30] Katsuhiko Ogata. 2001. *Modern Control Engineering* (4th ed.). Prentice Hall PTR, Upper Saddle River, NJ, USA.
- [31] F. Quin, D. Weyns, T. Bamelis, S. Singh Buttar, and S. Michiels. 2019. Efficient Analysis of Large Adaptation Spaces in Self-Adaptive Systems using Machine Learning. In *2019 IEEE/ACM 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. 1–12.
- [32] Arthur Rodrigues, Ricardo Diniz Caldas, Genaina Nunes Rodrigues, Thomas Vogel, and Patrizio Pelliccione. 2018. A learning approach to enhance assurances for real-time self-adaptive systems. In *2018 IEEE/ACM 13th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. IEEE, 206–216.
- [33] Maziar Salehie and Ladan Tahvildari. 2009. Self-Adaptive Software: Landscape and Research Challenges. *ACM Trans. Auton. Adapt. Syst.* 4, 2, Article Article 14 (2009), 42 pages.
- [34] Stepan Shevtsov, Mihaly Berekmeri, Danny Weyns, and Martina Maggio. 2018. Control-Theoretical Software Adaptation: A Systematic Literature Review. *IEEE Trans. Softw. Eng.* 44, 8 (Aug. 2018), 784–810.
- [35] Stepan Shevtsov and Danny Weyns. 2016. Keep it simplex: Satisfying multiple goals with guarantees in control-based self-adaptive systems. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 229–241.
- [36] Stepan Shevtsov, Danny Weyns, and Martina Maggio. 2017. Handling New and Changing Requirements with Guarantees in Self-Adaptive Systems Using SimCA. In *Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '17)*. IEEE, 12–23.
- [37] Seung Yeob Shin, Shiva Nejati, Mehrdad Sabetzadeh, Lionel C. Briand, Chetan Arora, and Frank Zimmer. 2019. Dynamic Adaptive Network Configuration for IoT Systems: A Search-based Approach. *CoRR abs/1905.12763* (2019). arXiv:1905.12763 <http://arxiv.org/abs/1905.12763>
- [38] Gabriela Félix Solano, Ricardo Diniz Caldas, Genaina Nunes Rodrigues, Thomas Vogel, and Patrizio Pelliccione. 2019. Taming Uncertainty in the Assurance Process of Self-Adaptive Systems: A Goal-Oriented Approach. In *Proceedings of the 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '19)*. IEEE, 89–99.
- [39] Jack A Stankovic, Chenyang Lu, Sang Hyuk Son, and Gang Tao. 1999. The case for feedback control real-time scheduling. In *Proceedings of 11th Euromicro Conference on Real-Time Systems. Euromicro RTS'99*. IEEE, 11–20.
- [40] Danny Weyns. 2017. Software engineering of self-adaptive systems: an organised tour and future challenges. *Chapter in Handbook of Software Engineering* (2017).
- [41] Danny Weyns. 2018. Engineering Self-Adaptive Software Systems - An Organized Tour. In *FAS*W@SASO/ICAC*. IEEE, 1–2.