Providentia: Using Search-Based Heuristics to Optimize Satisficement and Competing Concerns between Functional and Non-Functional Objectives in Self-Adaptive Systems

Kate M. Bowers^a, Erik M. Fredericks^a, Reihaneh H. Hariri^a, Betty H. C. Cheng^b

^aOakland University, Rochester, MI USA ^bMichigan State University, East Lansing, MI USA

Abstract

In general, a system may be subject to a combination of functional requirements (FRs) that dictate behavior and non-functional requirements (NFRs) that characterize how FRs are to be satisfied. NFRs also introduce crosscutting concerns that may be difficult to predict, where the degree of satisfaction (i.e., satisficement) of one NFR may be impacted by the satisficement of one or more FRs/NFRs. In particular, self-adaptive systems (SASs) can modify system configurations or behaviors at run time to continuously satisfy FRs and NFRs. This paper presents **Providentia**, a search-based technique to optimize the satisficement of NFRs in an SAS experiencing various sources of uncertainty. Providentia explores different combinations of weighted FRs to maximize NFR/FR satisficement. Experimental results suggest that **Providentia**-optimized goal models significantly improve the satisficement of an SAS when compared with manually- and randomly-generated weights and subgoals. Additionally, we apply a hyper-heuristic (Providentia-SAW) to balance the contribution of NFRs, FRs, and the number of adaptations and further improve the Providentia technique. We apply Providentia and Providentia-SAW to two case studies in different application domains involving a remote data mirroring network and a robotic vacuum controller, respectively.

Keywords: search-based software engineering, non-functional requirements, self-adaptive systems, evolutionary computation, optimization

Preprint submitted to Journal of Systems and Software

December 13, 2019

1. Introduction

A self-adaptive system (SAS) provides adaptation strategies for performing reconfigurations at run time to address unexpected issues that arise as a result of uncertainty (e.g., adverse environmental conditions or unanticipated issues in the system itself) [1, 2]. For example, a smart vacuum can be modeled as an SAS, where a reconfiguration is modeled as updating the cleaning path or navigation strategy as the vacuum encounters an obstacle, such as a chair. Each reconfiguration performed by the system can incur a cost (e.g., computation time, memory resources, etc.) associated with initializing and performing the adaptation. The SAS will use these adaptation strategies to select an optimal configuration that enables requirements to be continuously satisficed (i.e., degree of satisfaction) [3]. Generally, an SAS is governed by functional requirements (FRs) that focus on a specific function or feature of the system and can be mathematically quantified to monitor satisficement [4]. FRs in a smart vacuum may include maintaining battery power above 5% or avoiding obstacles detected by a sensor. Introducing non-functional requirements (NFRs) makes the adaptation selection process more difficult as NFRs specify properties and/or characteristics about system operations, tend to be qualitative, and may not be easily mathematically quantifiable (e.g., specifying resiliency and efficiency) [4, 5]. An example NFR for a smart vacuum SAS may be to optimize performance by cleaning as much dirt as quickly as possible, in contrast to an FR that mandates the vacuum to clean at least 50% of the room. Quantifying NFRs often relies on domain knowledge and may not be optimal given the changing environmental conditions that an SAS must address [5]. Therefore, this paper describes **Providentia** and **Providentia-SAW**, search-based techniques performed at design time that automatically determine an optimal set of FRs to support each NFR in an $SAS.^{1}$

Current techniques to satisfy NFRs in SASs do not offer concrete numerical values to be evaluated at run time. In the KAOS goal modeling framework, NFRs are incorporated as behavioral or soft goals [4]. A KAOS soft goal describes preferences of system behaviors that tend to be qualitative in nature, thereby making the determination of an optimal reconfiguration strategy more challenging [4]. In contrast, KAOS FRs can be quantified via utility functions and provide a concrete numerical basis for comparisons be-

¹A preliminary version of Providentia was presented at SSBSE 2018 [6].

tween reconfiguration strategies [3]. NFRs in the iStar framework are also modeled as soft goals and use the ++/+ or --/- operators to respectively indicate that an NFR makes/helps or breaks/hurts an FR [7]. The iStar operators are also qualitative and can be challenging to use in an SAS, as qualitative descriptions are not necessarily as easy to use when making an adaptation compared to quantitative descriptions. Similar to Providentia, the Analytic Hierarchy Process (AHP) decomposes NFRs into one or more weighted FRs using a prioritization schema [8]. However, prioritizations in an SAS may change drastically or even be inapplicable at a given instance in time (e.g., most or all requirements have equal priority) as the system experiences various forms of uncertainty due to changing environmental and system conditions. In a smart vacuum SAS, environmental uncertainty can be described as the positioning of obstacles in a room, stairs or other changes in elevation, the amount of dirt to be cleaned, and the number of water puddles to avoid. Uncertainty with regard to system conditions relates to sensor failures, damage that occurs in an obstacle collision, and motor degradation in the wheels. Each of these conditions are subject to change and the vacuum must be able to satisfy its requirements for each combination of uncertainties.

This paper describes Providentia [6] and Providentia-SAW (an extension), two techniques that address the challenges of quantifying and analyzing NFRs at run time in SASs. We introduce Providentia within the context of SASs to minimize the number of reconfigurations performed and achieve optimal weighted combinations that maximize requirement satisficement. Non-SASs can also apply Providentia to obtain optimal weighted combinations for each NFR given that the utility functions are adjusted accordingly. However, the effectiveness of Providentia is limited in non-SASs as the system configurations are static and cannot be reconfigured at run time. Therefore, we limit our scope to SASs. Providentia is a design-time technique that takes into account uncertainty from the environment and the system itself and optimizes FR/NFR relationships, where each relationship contributes to quantifying NFR objectives at run time. Each FR is associated with a utility function that specifies a mathematical expression of requirement satisficement [3]. Each NFR comprises a combination of one or more FRs using a linear-weighted sum to indicate the relative impact that an FR has in contributing to the satisficement of the NFR's objectives [8]. For example, the NFR to maximize performance uses four FRs (i.e., Achieve 50% clean, Achieve cleaning efficiency, Achieve cleaning effectiveness, and Maintain safety) to represent performance objectives that might otherwise

be more difficult to quantify. **Providentia** explores different combinations of weights at design time to find an optimal linear-weighted expression that makes the system more robust to various forms of uncertainty at run time. **Providentia-SAW**, in contrast, is a hyper-heuristic [9] approach to adjust the weights of the linear-weighted sum to respond to changing environmental conditions.

Providentia is a search-based evolutionary technique that assesses the system's run-time behavior via an executable system specification that is subjected to randomly-generated sources of uncertainty. The search process identifies optimal goal model configurations, namely the set of FRs and their corresponding weights for a given NFR, to maximize FR/NFR satisficement. **Providentia** uses a genetic algorithm [10] as a search heuristic, where the search space is the weight of each FR set for a given NFR, and the output is a set of optimal weight assignments that results in the highest satisficement of the NFR when faced with uncertainty. The optimal weight assignments determined by **Providentia** are then applied to the SAS at run time. By evaluating traditionally soft goals with FR metrics during execution, the SAS is able to perform online reconfigurations in response to both NFR and FR objectives, where traditionally only FR objectives are mainly considered. Furthermore, the SAS can perform better at run time by optimizing the weighted contributions of FRs to each NFR, as a requirements engineer may not be able to foresee the impact of random sources of uncertainty when determining the weight assignments at design time.

Since it may be difficult to achieve an optimal weighting scheme between NFRs, FRs, and the number of adaptations, this paper extends Providentia [6] by introducing a stepwise adaptation of weights (SAW) hyper-heuristic [11, 12, 13] to optimize the overall fitness value of the SAS for a given set of environmental conditions. Rather than manually selecting weights for the overall fitness function of the system, either according to preference or empirical evidence, Providentia-SAW more accurately determines an optimal set of weights that better guide Providentia's search procedure. As Providentia determines the optimal set of FRs and the weighting scheme for each NFR, Providentia-SAW runs in tandem to determine the weights assigned to balance NFR fitness values (i.e., satisficement), FR fitness value. Note, the term *weights* for Providentia apply to the NFR utility functions, whereas Providentia-SAW *weights* apply to the fitness sub-functions that comprise the Providentia algorithm.

We illustrate the effectiveness and domain independence of Providentia and Providentia-SAW with two case studies: a remote data mirroring (RDM) network and an intelligent robotic vacuum. The RDM is an industry-provided application that replicates and disseminates messages to each RDM within the network [14, 15]. The RDM performs dynamic reconfigurations in response to uncertainty due to dropped or delayed messages, sensor noise, and unexpected server and network link failures. Results from our preliminary work [6] have shown that **Providentia**-optimized goal models result in significantly higher fitness values compared to goal models with manually- and randomly-assigned FR weights. Furthermore, results also indicated that the number of FR violations was significantly reduced when Providentia was used. The second case study of a smart vacuum system (SVS) demonstrates the application of Providentia in a different domain. The SVS is an autonomous robotic vacuum modeled as an SAS and tasked with cleaning a given environment, where adaptations are performed at run time to switch between different configuration modes [16, 17]. Experimental results suggest that the SVS goal model optimized with **Providentia** performs better than SVS goal models with manually- and randomly-selected FR weights. Finally, results from both the RDM and SVS case studies indicate that the overall fitness can be further improved with Providentia-SAW.

Extensions. This paper extends an earlier description of the Providentia technique along several dimensions. First, we introduce Providentia-SAW that makes use of a second level of abstraction to explore the trade-offs between non-functional, functional, and adaptive requirements. Then we apply Providentia to a second case study, an autonomous robotic vacuum (SVS), to further demonstrate the technique's effectiveness in a different application domain. Providentia-SAW is also applied to this second case study. Finally, this paper includes additional details about the overall techniques and expands the related work discussion.

The remainder of this paper is organized as follows. Section 2 provides relevant background information on SASs, the RDM and SVS, goal modeling, NFRs, utility functions, and genetic algorithms. Section 3 presents the **Providentia** approach and introduces the integration with **Providentia**-SAW. Section 4 provides the experimental results of applying **Providentia** and **Providentia-SAW** to the RDM and SVS applications, respectively. Following, Section 5 presents a discussion of the results and threats to validity. Section 6 overviews the work related to **Providentia** and **Providentia-SAW**, and finally, Section 7 summarizes the findings and overviews future work.

2. Background

This section provides relevant background information on SASs, the case studies, NFRs, utility functions, and genetic algorithms.

2.1. Self-Adaptive Systems

The explosion of the number of possible combinations of system and environmental parameters often inhibits an engineer's ability to fully enumerate each combination [18, 19]. System requirements and objectives may also change following deployment, potentially requiring numerous software updates or patches. An SAS provides an approach for enabling continuous requirements satisfaction by dynamically adapting the system's configuration and/or behavior at run time [2, 20, 21]. As such, the RDM and SVS applications has been modeled as SASs to address uncertainty in the environment and the system itself [21].

Uncertainty. Given the exponential number of system and environmental combinations [18, 19], coupled with the possibility that software requirements and/or models will change following deployment (potentially requiring new software or bug fixes), it is difficult to accurately predict or model all situations an SAS may face throughout its lifetime. An SAS provides an approach for continuous requirements satisfaction by enabling selfreconfiguration at run time to mitigate such issues [2, 20]. SASs are generally guided by a run-time feedback loop such as MAPE-K, comprising *monitoring*, *analyzing*, *planning*, and *executing* components, linked together by common *knowledge* of the system and its elements [1]. This feedback loop enables an SAS to change its configuration and/or its behavior, at run time, to better mitigate current operating conditions or manifested uncertainties.

While many forms of uncertainty exist [22], we focus on *known unknowns* and *emergent behaviors/feature interactions*. Known unknowns tend to deal with knowledge of the system's operating conditions, where data may change unexpectedly, be inaccurate, or be in an unanticipated state [23, 24]. Emergent behaviors/feature interactions occur when multiple subsystems interact, introducing unexpected or possibly dangerous new behaviors that were not explicitly considered at design time [25].

Requirements Monitoring. While requirements monitoring is not specific to the SAS domain, monitoring does feature prominently in the MAPE-K loop for providing feedback to the SAS adaptation engine [1]. Specifically, an SAS that monitors requirements can self-reconfigure in the event that a

requirement is violated or unsatisfied. To quantify requirements at run time, utility functions are often used to provide a mathematical quantification [26] that indicates the degree to which requirements are satisfied. The Rainbow framework uses utility functions at an architectural level [27] and was our initial inspiration for leveraging utility functions, however in this work, we apply them at the requirements level [26].

2.2. Case Studies

This section provides an overview of the two case studies used to demonstrate the **Providentia** technique: the RDM and the SVS.

2.2.1. Remote Data Mirroring

RDM is a technique to protect data by minimizing data loss and maximizing the availability of data [14, 15]. The RDM technique disseminates data replicates to other servers (i.e., data mirrors) in physically remote locations. Each network link between data mirrors is associated with an operational cost. Furthermore, each link has a throughput, latency, and loss rate to collectively measure the performance and reliability of the RDM as a whole. The RDM must optimize the number of links to send messages between data mirrors efficiently without exceeding the budget.

Requirements may become unsatisfied in the face of various forms of uncertainty, such as unexpectedly dropped or delayed messages, random network link or data mirror failures, and noise in the network links or data mirror sensors. An RDM can be modeled as an SAS [28], where reconfigurations can change the network topology (e.g., a minimum spanning tree or a redundant topology) as well as the manner of data propagation among nodes to ensure that requirements are continuously satisfied. The reconfiguration strategies involve modifying the status of data mirrors impacted by uncertainty. The status of a data mirror can be active (i.e., can send and receive messages), passive (i.e., cannot send messages but can receive messages), or quiescent (i.e., cannot send or receive messages).

2.2.2. Smart Vacuum System

Smart vacuums such as iRobot's Roomba² are available in the consumer market to clean dirt by navigating across a room and around obstacles autonomously without guidance from the user. The SVS is an open-source

²See http://www.irobot.com

simulation of a Roomba. The SVS operates by using input from sensors (e.g., bumper sensors and motor sensors) to plan a path in a given area and follow the path to clean its local environment. For the purposes of our simulation, the SVS contains bumper sensors to detect when the robot collides with an item (e.g., a wall or leg of a chair), cliff sensors to prevent the robot from damaging itself by falling down a flight of stairs, and motor sensors to provide feedback on the movement of the robot (e.g., velocities of the wheels and power modes for the suction). A controller uses the sensor data to plan an optimal cleaning path and minimize battery consumption.

The SVS can be modeled as an SAS to reconfigure multiple modes during run time when faced with various forms of uncertainty (e.g., noisy sensor data, amount and location of dirt within a room, obstacle encounters, etc.) [16, 17, 21]. For example, the SVS can change the pathfinding algorithm, power modes with regard to movement and suction, and obstacle avoidance measures at run time to maintain requirements satisfaction.

2.3. Goal-Oriented Requirements Modeling

Goal-oriented requirements modeling (GORE) is an approach that uses goals to model the behaviors of the system [4]. A goal is a desired system behavior achieved through interactions with agents, where an agent is a system component that performs specific actions according to goals. A requirement is a goal that interfaces with a single agent. An expectation is a requirement where the agent is in the environment, compared to agents within the system itself. Requirements and goals can be further classified as functional and non-functional. FRs specify *what* services are to be provided while NFRs specify *how* the FRs are to be satisfied.

GORE uses a directed acyclic graph, where each node represents a goal or requirement and each edge represents a goal/requirement refinement [4]. KAOS and iStar extend GORE by adding additional goal refinements [4, 7, 29]. Figure 1 presents a KAOS goal model of the RDM application, and Figure 2 presents a KAOS goal model of the SVS application.³ KAOS uses ANDand OR-refinements, where an AND-refined goal is satisfied only when all of its subgoals are satisfied (e.g., Goal (A) in Figure 1 is satisfied only if both Goals (B) and (C) are also satisfied) and an OR-refined goal is satisfied when

³This work does not use the KAOS formal refinement infrastructure.



Figure 1: RDM goal model.



Figure 2: SVS goal model.

at least one of its subgoals is satisfied (e.g., Goal (G) in Figure 1 is satisfied if either Requirements (Q) or (R) are satisfied). Furthermore, KAOS FRs can be classified as invariant or non-invariant. An invariant goal, denoted by the keywords "Maintain" or "Avoid," must always be satisfied. If any invariant goal is unsatisfied then the system fails. For example, the safety-related goals in Figure 2 are denoted as invariant goals to emphasize the safety of the SVS. A non-invariant goal, denoted by the keyword "Achieve," may be temporarily unsatisfied due to uncertainty. For example, the path planning goals and requirements of the SVS in Figure 2 (e.g., Goals (C), (H), (I), (O), (P), and (Q)) are non-invariant goals to demonstrate that the SVS may temporarily have a less-than-optimal path to clean a room.

2.4. Non-Functional Requirements

NFRs specify quality constraints on a system, such as performance or reliability [30], and are often difficult to quantify given their subjective nature. Furthermore, NFRs may introduce cross-cutting concerns given the broad impact on the overall system [30]. For example, the RDM may have one NFR to maximize performance that keeps as many data mirrors in an active state as long as possible. However, there may be a second NFR to minimize power consumption that puts data mirrors in a quiescent state, directly contradicting the performance NFR. Such cross-cutting concerns introduce further complexity in measuring the satisfaction of NFRs. Although other approaches to quantify requirement satisficement have been introduced [26, 27], such models often require detailed knowledge of both the system and its environment that may not always be possible with the wide impact of NFRs.

Therefore, Providentia uses FRs already defined as part of the system to quantify NFR objectives. Figure 3 shows a sample NFR for the RDM application to Minimize [Power], where many factors may impact power consumption (e.g., Goals (A), (E), (I), (V), and (W) from Figure 1). For the purposes of this paper, Figure 3 is shown separately from Figure 1 but is intended to be an extension of the FR goal model rather than a separate NFR goal model. We use NFRs to represent all non-functional goals and requirements and FRs to represent all functional goals and requirements. The cloud node in Figure 3 represents a single NFR. Each parallelogram node represents a goal or requirement/expectation that a requirements engineer designates to have an impact in the satisficement of NFR7. Each edge depicts the relative contribution, or weighted sum value, of each FR to NFR7. Note that the sum of the weights is 1.0. Similarly, Figure 4 shows a sample NFR for the SVS application to Minimize [Cost], where cost refers to avoiding damage to obstacles (e.g., damaging the legs of a dining room chair) and avoiding damage to the smart vacuum itself (e.g., hitting an obstacle that damages a sensor or falling off a ledge).



Figure 3: RDM NFR7: Minimize [Power].



Figure 4: SVS NFR1: Minimize [Cost].

2.5. Utility Functions

A utility function can be used to calculate the satisficement of FRs in SASs [26, 31, 32]. The utility function for a specific, single requirement is evaluated and returns a utility value. A utility value of 1.0 indicates the highest degree of satisfaction and a utility value of 0.0 indicates the lowest degree of satisfaction. Utility values between 0.0 and 1.0 indicate the degree of satisficement for a given requirement [3]. Equation (1) shows a utility

function associated with Goal (V) from the RDM application in Figure 1, where n indicates the number of passive data mirrors.

$$util(goal_V) = \begin{cases} 1.0 & \text{if } n == 0\\ f(x) & \text{if } 0 < n < 20\% \text{ of total nodes} \\ 0.0 & \text{if } n \ge 20\% \text{ of total nodes} \end{cases}$$
(1)

Goal (V) is considered to be completely satisfied if no data mirrors are in a passive state (i.e., can receive but not send messages) and evaluates to a utility value of 1.0. Goal (V) completely fails if more than 20% of all data mirrors are in a passive state. Otherwise, if the number of passive data mirrors is greater than 0% but less than 20%, then the f(x) value is linearly determined. For example, if 10% of the data mirrors are in a passive state, then Goal (V) will have a utility value of 0.5.

2.6. Genetic Algorithms

A genetic algorithm is a heuristic used to search a space of solutions to find an optimal result [10]. The evolutionary process of a genetic algorithm generates a population of candidate solutions, performs crossover and mutation operations, and evaluates the fitness of each solution within the population. The evolutionary process repeats until the specific number of generations, or iterations, is reached. We next describe each of these activities.

2.6.1. Population Generation

A genetic algorithm may start with a randomly-generated set, or population, of candidate solutions in the first generation. The population of individuals represents the number of candidate solutions to be evaluated through the evolutionary process. For each generation, evolutionary operators such as crossover, mutation, and selection are applied to each individual to generate new members of the population and evaluate fitness. Ideally, the most fit individuals are preserved during the evolutionary process, with the best performing individual being considered an optimal result.

2.6.2. Crossover and Mutation

Crossover and mutation are evolutionary operators that generate new individuals (i.e., children) during the evolutionary process. While there are many different types of crossover and mutation, we use two-point crossover and single-point mutation. Two-point crossover selects two individuals, either randomly or as a result of tournament selection [10], and then selects two random indices to denote cut points. Two new children are then generated by combining the genetic material from both parents, where genomes are swapped between the two cut points. Single-point mutation selects a single index from a candidate parent and then randomly mutates the gene at that location, creating a new child. Ideally, crossover preserves the best characteristics of its parents, and mutation introduces diversity to ensure that the search procedure does not become "stuck" at a local optimum.

2.6.3. Fitness Evaluation

Populations are evaluated according to fitness functions, where each individual within a population is assigned a fitness value. Fitness values determine those individuals that performed better than others. Often, genetic algorithms use a form of tournament selection where a specific number of the highest performing individuals are used in the next generation. In this way, the genetic algorithm is guided toward an optimal value by carrying over the best solutions through each generation but explores different combinations via crossover and mutation to obtain a global optimal solution.

2.6.4. Hyper-heuristic Algorithms

Hyper-heuristic algorithms are used to search a space of meta-heuristics [9]. Rather than exploring the problem space for a solution (i.e., a metaheuristic), hyper-heuristics search among all the solutions to the problem space. Hyper-heuristics operate at an abstraction level above meta-heuristics [33]. For example, the SVS faces a variety of problems (e.g., cliff detection, object collision, water spots to avoid, etc.). **Providentia** is a meta-heuristic that searches for the most effective solution (i.e., combination of FR weights for each NFR) to satisfy the system requirements to the highest degree. However, satisfying the system requirements to the highest degree is a balance between satisfying FRs, NFRs, and minimizing the number of adaptations performed. Therefore, we apply a hyper-heuristic to search for a solution *among the solutions* that maximizes overall system satisficement. For the purposes of this paper, we classify **Providentia-SAW** as a hyper-heuristic that further improves the **Providentia** meta-heuristic genetic algorithm.

3. Approach

This section introduces Providentia, our technique that analyzes the weighted contributions of FRs to each NFR in the SAS requirements specification and/or goal model. Providentia determines an optimal combination

of weights that yields the highest overall satisfaction of the entire goal model (i.e., including NFR and FR satisficement). The **Providentia** technique is intended to be an add-on feature to make a system more robust to uncertainty. **Providentia** assumes that the system is fully operational and contains models of FRs, NFRs, and mechanisms to measure how well those requirements are being satisfied.

The following sections outline the base technique using the RDM case study [6] as a motivating example and extends previous work to include the SVS case study. First, we discuss **Providentia** with respect to expected inputs and outputs. Next, we describe the process of automatically optimizing the input goal model using a genetic algorithm. Finally, we introduce the **Providentia-SAW** technique to balance the satisficement of FRs, NFRs, and the number of adaptations.

3.1. Providentia: Assumptions, Inputs, and Outputs

Providentia requires four inputs: (1) a goal model of the SAS, consisting primarily of FRs but may include NFRs defined by a requirements engineer, $^{4}(2)$ a list of NFRs with a set of FR suggestions for each, (3) a set of utility functions to measure the satisficement of FRs, and (4) an executable specification with defined sources of uncertainty (for the RDM and SVS case studies, we use environmental and system uncertainty). Each of these inputs are expected to be previously defined by one or more engineers. The sources of uncertainty are parameters (e.g., sensor values, number of dropped messages in a network, etc.) whose values change during runtime that may cause the input FRs and NFRs to become unsatisfied, and therefore trigger a reconfiguration. The **Providentia** technique explores interactions between requirements as well as system behavior while subjected to uncertainty that an engineer may not be able to foresee when designing the system. The output of Providentia is a goal model with optimized FR/NFR relationships. Note that the **Providentia** technique is only as good as the accuracy of the input data (i.e., we assume the set of FRs for a given NFR is accurate and that the executable specification is not missing any sources of uncertainty).

Goal model. A KAOS goal model provides Providentia with a specification of the FRs and NFRs in the SAS. For example, Figures 1 and 2 demonstrate the RDM and SVS goal models, respectively.

⁴For the purposes of this paper, the term "requirements engineer" refers to the authors who performed the manual choices for the NFRs.

Utility functions. Each FR in the input KAOS goal model shall have a corresponding utility function to evaluate SAS requirements at run time [31, 32]. A utility function maps the FR to a utility value within [0.0, 1.0] to represent the degree to which the FR is satisfied. A requirements engineer is expected to provide a set of utility functions that correspond to the input goal model. A sample utility function is demonstrated in Equation (1).

Applicable set of FRs mapped to NFRs. A requirements engineer must provide, in addition to the goal model, a set of FRs that may have an impact in the satisficement of an NFR, with the set of NFRs comprising the non-functional properties intended for the system. For example, Figure 3 for the RDM case study shows an example of an applicable set of FRs, namely Goals (A), (E), (I), (V), and (W), that seem most relevant in minimizing power consumption. However, the requirements engineer may consider expanding the set of five FRs to include Goals (B), (K), (M), (O), (P), and (U) from Figure 1.

Executable specification. An executable specification or simulation is required in order for **Providentia** to evaluate the overall fitness value of the SAS and determine an optimal combination of FR/NFR weights. We use a simulation of both the RDM and SVS case studies. The RDM is an industry-provided simulation shared with the authors and the SVS is an ongoing part of the authors' research lab. The specification provides the utility values to measure how well the requirements were met at a given instance in time during the simulation so that **Providentia** can determine where to further improve requirement satisficement. The executable specification should also include known sources of uncertainty anticipated by the requirements engineer, both in the environment and the system itself (e.g., liquid spills, sensor noise, etc.). This information better guides the search process at design time and makes the system more robust to known and unknown sources of uncertainty at run time. Moreover, the executable specification must include adaptation mechanisms for the NFRs, should they be violated. The requirements engineer is expected to identify and implement the adaptation mechanisms. Providentia does not change any reconfigurations but rather tries to minimize the number that occur, as each system reconfiguration incurs a cost (e.g., computation time, memory used, etc.). For example, if an NFR for maximizing performance is violated, then a reconfiguration strategy (e.g., reconfiguring the RDM network overlay) must be defined in addition to those adaptations already defined for FR violations.

The executable specification acts as a simulation environment where the system is subjected to various randomized forms of uncertainty (e.g., where the RDM experiences a random number of dropped messages, or random connections are lost). The executable specification is not limited to any particular tool, programming language, or environment. Providentia observes how the system behaves and the degree to which FRs and NFRs are satisfied while experiencing adverse conditions. Providentia takes note of the weighted combinations that fulfill both FRs and NFRs to the highest degree and returns the weight assignments. The executable specification can be used as input to **Providentia** if the specification effectively simulates system behaviors with internal and environmental uncertainty.

Output. The output of Providentia is (1) an NFR goal model integrated with the input FR goal model, (2) for each NFR, a set of FRs that collectively contribute to the satisficement of the NFR, and (3) an optimal weight value assigned to each FR. Note that a weight value of 0.0 for an FR indicates that the FR did not at all contribute to the satisficement of the NFR. For example, the initial set of applicable FRs given for NFR7 from the RDM application in Figure 1 were Goals (A), (B), (E), (I), (K), (M), (O), (P), (U), (V), and (W). Providentia determined the most optimal weights for each goal respectively to be as follows: B: 0.237144, E: 0.241000, K: 0.007185, M: 0.373794, O: 0.049442, U: 0.067218, V: 0.024216. The weight assignments to Goals (B), (E), and (M) indicate that maintaining costs below the budget, keeping the minimum number of links active, and achieving an accurate measure of the workload most significantly impact the overall power consumption of the RDM. Goals (K), (O), (U), and (V) minimally impacted the satisfaction of power requirements. Goals (A), (E), (I), (P), and (W) have weights 0.0 as they did not contribute to minimizing power consumption. The utility function for NFR7 is calculated by summing the products of the weight and FR utility value for each nonzero FR.

3.2. Providentia Technique

This section describes the details of **Providentia**, comprising a genetic algorithm [10] to search for optimal FR weights that contribute to each NFR in an SAS. Figure 5 illustrates a data flow diagram of the **Providentia** technique with each step further described in detail.

(1) Define Solution Structure. Each of Providentia's candidate solutions (i.e., the weights of the FRs associated with an NFR) is encoded in a genome shown in Figure 6. The genome as a whole consists of all NFRs



Figure 5: Data flow diagram of Providentia technique [6].



Figure 6: Providentia sample genome.

in the SAS. For example, the RDM case study has 7 NFRs. A single NFR is referred to as a sub-genome, denoted by the bold border in Figure 6. A sub-genome is broken down into genes (i.e., corresponding weights for one NFR), where a single gene is a weight corresponding to a single FR. The sum of weights within one sub-genome (i.e., within one NFR) must be equal to 1.0.

(2) Configure Search Process. The search process of a genetic algorithm is controlled by the population size, number of generations, crossover rate, mutation rate, and selection rate. Based on empirical evidence on convergence rates [6], we specify the following configuration parameters: population size 20, 50 generations, 25% crossover rate, 50% mutation rate. With regard to the selection rate, we use the tournament selection approach [10] where the three individuals with the highest utility values become parents of the next generation. Larger values for population and generation sizes were considered (e.g., 25-50 population individuals and 50-100 generations) but optimal convergence was discovered on average at the specified values.

(3) Evaluate NFR Models. Once the search process has been appropriately configured, each individual candidate solution is input to the executable specification and evaluated to compare against other individuals within the population. The overall fitness function for the SAS is shown in Equation (2) and evaluates the satisfaction of FRs, NFRs, and the number of adaptations (NA) into a single fitness value.

$$FF = \begin{cases} \alpha_{NFR} * FF_{NFR} + \alpha_{FR} * FF_{FR} + \alpha_{NA} * FF_{NA} & \text{iff invariants true} \\ 0.0 & \text{otherwise} \end{cases}$$

(2)

Equation (2) is a linear-weighted sum, where α_{NFR} , α_{FR} , and α_{NA} are further optimized by the **Providentia-SAW** technique (**Providentia-SAW** is described in Section 3.3). The weights must cumulatively sum to a value of 1.0. The terms FF_{NFR} and FF_{FR} represent fitness sub-functions. Note that if any invariant goals or requirements are unsatisfied, then the system is considered to be failed and evaluates its fitness value to be 0.0. Although many other approaches exist to combine fitness sub-functions, we find that a linear-weighted sum balances competing concerns between requirement satisficement and adaptations adequately for this experiment.

The fitness sub-function to evaluate the satisficement of NFRs is shown in Equation (3).

$$FF_{NFR} = \frac{\sum_{i=1}^{|NFRs|} utility_value_{NFR_i}}{|NFRs| * timesteps}$$
(3)

The numerator in Equation (3) represents the sum of utility values for all NFRs in the SAS. For example, the RDM application contains seven NFRs, resulting in a numerator evaluation of: $utility_value_{NFR_1} + utility_value_{NFR_2} + ... + utility_value_{NFR_7}$. The denominator of Equation (3) is the number of NFRs in the SAS (e.g., seven for the RDM) multiplied by the number of timesteps run in the executable specification.

The utility value calculation of each NFR from the numerator in Equation (3) is shown in Equation (4), where n refers to the number of the NFR (e.g., NFR_1 , NFR_2 , etc.).

$$utility_value_{NFR_n} = \sum_{i=1}^{|FR_{NFR_n}|} utility_value_{FR_i} * weight_{FR_i}$$
(4)

Equation (4) calculates a utility value for one NFR by evaluating the utility function of each supporting FR (e.g., Equation (1) supplied by the input utility functions), multiplying the utility value with its corresponding weight, and then summing the results for all supporting FRs for a given NFR. The term $weight_{FR_i}$ represents the weights present in **Providentia**'s genes and is the focus of the technique's optimization.

The fitness sub-function for the FF_{FR} term from Equation (2) is shown in Equation (5).

$$FF_{FR} = \frac{\sum_{i=1}^{|FRs|} utility_value_{FR_i}}{|FRs| * timesteps}$$
(5)

Equation (5) sums all of the utility values for the functional requirements and goals of the SAS in the numerator and divides by the product of the number of FRs and the number of timesteps in the executable specification.

Finally, the equation to minimize the number of adaptations performed by the SAS is shown in Equation (6), where |adaptations| refers to the total number of reconfigurations performed by the SAS and |faults| reports the total number of adverse conditions introduced within the executable specification.

$$FF_{NA} = 1.0 - \frac{|adaptations|}{|faults|} \tag{6}$$

Equation (6) defines the term FF_{NA} from Equation (2). A minimized number of adaptations performed by an SAS reduces the overall disruption of the system.

(4) Select NFR Models. Providentia uses tournament selection to select the genomes with the highest fitness value among k genomes to be used as a parent for the next generation. The remaining individuals not chosen in the population are removed from consideration.

(5) Generate NFR Models. When the most fit individuals from the population are selected, Providentia then performs crossover and mutation to obtain new individuals for the following generation. Providentia uses two-point crossover that selects two points on a genome within an NFR (i.e., crossover does not occur between different NFRs) and swaps the genes with a second genome, creating two new candidate solutions. Furthermore, Providentia performs single-point mutation that selects a single gene (i.e., FR weight) and randomly generates a value within $\pm 20\%$ of its original value. Note that although values are modified via crossover and mutation, the weights must still be normalized within a particular NFR/sub-genome.

Steps (3) - (5) are applied iteratively until the number of generations is reached. When the genetic loop concludes, **Providentia** returns a set of optimized weights of FRs for each NFR.

3.3. SAW Integration

This paper extends the original Providentia technique [6] with SAW optimization [13, 34] to explore a search space containing an optimal weighting scheme that balances the competing concerns between FR satisficement, NFR satisficement, and the number of SAS adaptations. Our preliminary work used an empirically chosen weighting scheme where the terms from Equation (2) were set as follows: $\alpha_{NFR} = 0.375$, $\alpha_{FR} = 0.375$, and $\alpha_{NA} = 0.25$. However, the chosen weights may not have considered all possible values in the search space, given the uncertainty surrounding an SAS in terms of its environment and configured parameters. Therefore, we apply the hyperheuristic SAW to gradually update the weighting scheme in tandem with Providentia, denoted as Providentia-SAW, to explore how and whether overall SAS fitness may be improved over the course of system execution.

Providentia uses the *online* SAW method of optimization described in the original SAW paper [34]. Step (3) in Figure 5 is augmented with the SAW technique. Rather than randomly generating the weights, a requirements engineer can provide initial values (i.e., seed) for Providentia-SAW to

use in the first generation. The fitness sub-functions are evaluated every *fifth* generation and the sub-function with the lowest fitness value (i.e., in comparison to the other fitness sub-functions in the equation) has its corresponding weight increased to guide the search process for an optimal combination of weights. By increasing the weight of the least fit sub-function, the search process ensures that the weighting scheme balances all concerns appropriately to achieve the highest overall fitness value. The generation number to evaluate the fitness sub-functions (i.e., five) was determined empirically. Other numbers, both higher and lower, were simulated, however optimal results were obtained when evaluating every fifth generation. This generation value balances the competing concerns of adapting too frequently, where the search space does not have enough time to consider optimal solutions, and adapting too infrequently, where more optimal solutions are never explored.

4. Experimental Results

This section describes the experimental setup and results for both the RDM and the SVS case studies. The RDM is an industry-provided case example and the SVS was created by the authors' research lab. First, we introduce the experiment setup parameters that we used and the types of adaptations performed for each application. Next, we present our results for each application to compare the performance of randomly-generated, manually-selected, and Providentia-derived FR/NFR weights. We then describe how Providentia-SAW enabled further optimization towards the satisficement of SAS requirements.

4.1. RDM Study

This section describes the experimental setup and results of the RDM application.

4.1.1. RDM Experimental Setup

The RDM application is modeled as a completely-connected graph. Each node of the graph represents an RDM. Each edge of the graph represents a network link. For each trial, system and environmental parameters were randomized based on a model previously presented by Keeton *et al.* [14, 15]. For example, the randomized system parameters include a random number of RDMs (i.e., within [15,30]) and a random number of valid messages (i.e., [100,200]) inserted into RDMs at random timesteps. Each message is required to be replicated to all other RDMs. The RDM simulation was performed over 300 timesteps. In addition to the 23 FRs presented in Figure 1, we also examine seven NFRs specific to the RDM that are next presented in Table 1.

NFR	Set of FRs			
NFR1: Maximize	(A), (B), (C), (E), (F), (G), (H), (I), (L),			
[Reliability]	0.4, 0.2, 0.2, 0.0, 0.0, 0.1, 0.0, 0.1, 0.0,			
	(M), (N), (R), (T), (U), (V), (W)			
	0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0			
NFR2: Maximize	(A), (C), (D), (E), (F), (G), (H), (L), (M),			
[Throughput]	0.0, 0.6, 0.0, 0.0, 0.0, 0.1, 0.1, 0.0, 0.0,			
	(N), (O), (P), (Q), (R), (S), (T), (U), (V),			
	0.0, 0.0, 0.0, 0.0, 0.1, 0.0, 0.1, 0.0, 0.0			
	(W)			
	0.0			
NFR3: Maximize	(C), (D), (H), (I), (K), (M), (Q), (S), (U),			
[Speed]	0.2, 0.1, 0.3, 0.2, 0.0, 0.0, 0.0, 0.2, 0.0,			
	(V), (W)			
	0.0, 0.0			
NFR4: Maximize	(A), (B), (D), (G), (H), (L)			
[System Security]	0.4, 0.2, 0.2, 0.1, 0.1, 0.0			
NFR5: Maximize	(C), (G), (H), (Q), (R), (S), (T), (W)			
[Secure Communication]	0.5, 0.3, 0.1, 0.0, 0.0, 0.0, 0.0, 0.1			
NFR6: Maximize	(C), (D), (H), (R), (T)			
[Message Security]	0.2, 0.2, 0.0, 0.3, 0.3			
NFR7: Minimize	(A), (B), (E), (I), (K), (M), (O), (P), (U),			
[Power]	0.3, 0.0, 0.1, 0.1, 0.0, 0.0, 0.0, 0.0, 0.0			
	(V), (W)			
	0.3, 0.2			

Table 1: RDM NFRs with sets of FRs and manually-derived weights.

We compared and evaluated different combinations of FRs and their supporting weights for every NFR. The seven NFRs were derived using three different techniques: (1) FR weights generated by random search [35], (2) manually-selected weights assigned by a requirements engineer, and (3) **Providentia**-optimized weights. Note that although a requirements engineer initially selects a subset of FRs for each NFR, all three techniques may disable (but not add) one or more FRs by setting the corresponding weight to 0.0, effectively allowing limited flexibility in the selection of FRs as well as the weight. Table 1 shows the initial sets of FRs for each NFR chosen by a requirements engineer with the manually-derived weights listed below. For instance, NFR6 uses the utility functions from Goals (C), (D), (H), (R), and (T) to calculate its own utility function in aggregate. Specifically, Equation (7) demonstrates the utility function for NFR6:

$$util(NFR_6) = \alpha_C * util(goal_C) + \alpha_D * util(goal_D) + \alpha_H * util(goal_H) + \alpha_R * util(goal_R) + \alpha_T * util(goal_T)$$
(7)

Each specified NFR has a similar utility function to determine its utility value at run time. Note that **Providentia** optimizes the α values on a per-NFR basis (i.e., each set of FR weights is optimized by **Providentia** and normalized to ensure they sum to 1.0 to ensure that the associated utility function is also normalized). To support NFR feedback within the SAS decision loop [1], additional reconfiguration strategies were implemented for the RDM, where the new reconfiguration strategies were defined based on the requirements engineer's knowledge of the system. For instance, reconfiguration as a result of an NFR3 violation resulted in an internal search for a new network overlay. Finally, we configured the genetic algorithm as shown in Table 2.

Parameter	Value
Population size:	20
Number of generations:	50
Crossover rate:	25%
Crossover type:	Two-point
Mutation rate:	50%
Mutation type:	Single-point
Selection:	Tournament selection, $k = 3$
Stepwise adaptation of weights:	Online, every 5^{th} generation

Table 2: Configuration of Providentia genetic algorithm.

Each of the parameters in the genetic algorithm were determined empirically. For example, population sizes of 10, 25, 50, etc. were run in the simulation as well as other values for the remaining parameters. Table 2 shows the values where the genetic algorithm converged. Note that the population size (i.e., the number of individuals) combined with the number of generations results in the number of evaluations *per experimental replicate*. Using the values presented in Table 2, each experimental replicate evaluates 1000 individuals. To ensure statistical significance of our results, we performed 50

experimental replicates, each of which was seeded differently. The p-value was calculated using the Wilcoxon-Mann-Whitney u-test by comparing the results of (1) manually-selected weights vs. randomly-generated weights, (2) Providentia-optimized weights vs. randomly-generated weights, and (3) Providentia-optimized weights vs. manually-selected weights. Manually selected weights indicate a single set of weights chosen by engineers to best represent the contribution of each FR to each NFR. The configurations for the genetic algorithm shown in Table 2 produced the experimental results used in evaluating Providentia and Providentia-SAW.

4.1.2. RDM Experimental Results

This section presents our results from investigating how NFRs impact an SAS. Specifically, we examine how a set of FRs can contribute to the satisfaction of NFRs that are then in turn incorporated into the SAS decision loop [1] to support run-time reconfigurations. We compare and evaluate the impact of applying automatically-selected **Providentia** FR weights with manual and random selection, respectively. For this experiment, all FR weights associated with each NFR are normalized to 1.0. For manual weight selection, we apply the weights as shown in Table 1, where weights of 0.0 indicate that the related FR's utility function is not applied to the calculation of its associated NFR.

With respect to the fitness function in Equation (2), we set $\alpha_{NFR} = 0.375$, $\alpha_{FR} = 0.375$, and $\alpha_{NA} = 0.25$, where these values were selected based on empirical evidence.

For this experiment, we define two null hypotheses. First, $H1_0$ states that "there is no difference in fitnesses achieved by a **Providentia**-optimized goal model and those that are unoptimized." Second, $H2_0$ states that "there is no difference in fitnesses achieved by a **Providentia**-optimized goal model and those manually optimized by a requirements engineer."

Figure 7 shows three boxplots for each NFR in the RDM with average fitness values calculated from randomly-selected FR weights, manuallyselected FR weights, and **Providentia**-optimized FR weight selection. Figure 7 demonstrates that **Providentia** can significantly improve overall NFR fitness than those manually selected by a requirements engineer or selected at random (p < 0.05, Wilcoxon-Mann-Whitney u-test). Table 3 presents the average utility values (μ) and standard deviation (σ) for each NFR, with the optimal value highlighted in gray. These results suggest that **Providentia** can improve overall NFR fitness when an SAS is subject to uncertainty.



Figure 7: NFR fitness experimental results [6].

Providentia also enabled a significant *decrease* in the number of encountered FR violations in comparison to randomly- and manually-defined FR weights (p < 0.05, Wilcoxon-Mann-Whitney u-test), as shown in Figure 8. As requirements violations tend to signify a significant problem with a system, a reduction in run-time violations is an ideal result for an optimization procedure.

Given the results presented in Figure 7, Table 3, and Figure 8, we can reject both $H1_0$ and $H2_0$ and accept our alternate hypotheses that **Providentia** provides a significant improvement over manual and random search.

We next describe the integration of SAW within Providentia to provide further points of optimization.

4.1.3. RDM Results with SAW Integration

We now examine how the fitness function weights (α_{NFR} , α_{FR} , and α_{NA} ; c.f., Equation (2)) impact overall fitness, as the weights were initially selected

NFR	Random	Manual	Providentia
NFR1: Maximize	μ : 0.654	μ : 0.615	μ : 0.905
[Reliability]	σ : 0.325	σ : 0.191	σ : 0.149
NFR2: Maximize	μ : 0.655	μ : 0.666	μ : 0.882
[Throughput]	σ : 0.325	σ : 0.262	σ : 0.153
NFR3: Maximize	μ : 0.875	μ : 0.743	μ : 0.975
[Speed]	σ : 0.207	σ : 0.148	σ : 0.085
NFR4: Maximize	μ : 0.802	μ : 0.736	μ : 0.979
[System Security]	σ : 0.273	$\sigma: 0.177$	σ : 0.085
NFR5: Maximize	μ : 0.621	μ : 0.742	μ : 0.925
[Secure Communication]	σ : 0.273	σ : 0.191	σ : 0.146
NFR6: Maximize	μ : 0.921	μ : 0.919	μ : 0.980
[Message Security]	σ : 0.181	$\sigma: 0.072$	σ : 0.069
NFR7: Minimize	μ : 0.821	μ : 0.758	μ : 0.926
[Power]	σ : 0.270	σ : 0.172	σ : 0.188

Table 3: NFR average utility values and standard deviations [6].



Figure 8: FR violation experimental results [6].

based on empirical evidence and domain knowledge [6]. We apply SAW to dynamically adjust the weights during execution (c.f., Section 3.3). For this experiment, we executed **Providentia** with and without SAW applied. We reuse the configuration of the genetic algorithm as presented in Table 2. Moreover, SAW dynamically updates the weights of the fitness function every fifth generation, where the poorest-performing fitness subfunction's weight is increased and the remaining subfunction weights are normalized to sum to 1.0.

We define an additional null hypotheses for this experiment. $H3_0$ states that "there is no difference between a **Providentia**-optimized goal model with static fitness subfunction weights and a goal model with dynamicallyoptimized fitness subfunction weights." For this experiment, we reuse the static weights defined in the previous section (i.e., $\alpha_{NFR} = 0.375$, $\alpha_{FR} =$ 0.375, and $\alpha_{NA} = 0.25$) as both our static weights and seed weights for **Providentia-SAW**.

Figure 9 presents two boxplots that show the fitness values obtained from goal models optimized with Providentia and Providentia-SAW, respectively. As this figure demonstrates, applying SAW to Providentia results in higher fitness values (p < 0.05, Wilcoxon-Mann-Whitney u-test), suggesting that weights that are dynamically adjusted better reflect the environment and/or configuration of the system. Moreover, these results enable us to reject $H3_0$ and conclude that the fitness function weighting scheme directly impacts overall fitness resulting from monitoring the system.

Next, Figure 10 presents a set of grouped boxplots that demonstrate the average utility value (calculated by the aggregate utility values of each associated FR) for each NFR. As can be seen by this figure, there is no statistical difference that exists between the average utility values of each NFR that were optimized by **Providentia** and **Providentia-SAW**, respectively (p > 0.05, Wilcoxon-Mann-Whitney u-test). This result is interesting in that overall FR fitness was significantly improved (c.f., Figure 9), however there is no improvement for the NFR values. This result suggests that the NFR utility values may be independent of the overall RDM fitness calculation, however reconfigurations that are performed at a coarser-grain (NFRs) can significantly impact and improve the performance of FR utility functions.

Figure 11 shows the average number of *non-invariant* and *invariant* requirement violations that occurred during execution, respectively. Again, there is no significant difference (p > 0.05, Wilcoxon-Mann-Whitney u-test) in the number of violations between Providentia and Providentia-SAW,



Figure 9: Comparison of fitness values between Providentia and Providentia-SAW experiments for the RDM application.

suggesting that adjusting the fitness function weights does not significantly impact the SAS reconfiguration engine for the RDM application.

Lastly, Figure 12 provides a comparison of the starting and ending fitness sub-function weights attained with **Providentia-SAW**. As can be seen from this figure, α_{FR} is maximized and α_{NFR} and α_{NA} tend to be minimized, thereby suggesting that the satisficement of FRs is considerably more important to the performance of the RDM application. This result correlates with those found in Figures 10 and 11 in that **SAW** seems to have a minimal impact on NFR satisficement and violation reduction. However, an improvement in overall performance of the RDM application still presents a significant finding.

We next repeat our experiments on the SVS application to demonstrate the domain independence of Providentia and Providentia-SAW.

4.1.4. SVS Experimental Setup

The SVS (c.f., Section 2.2.2) comprises an autonomous vacuum system tasked with cleaning a desired space *safely*. As such, the SVS must balance competing concerns to satisfy its requirements, including maximizing clean-



Figure 10: Comparison of NFR fitness values between Providentia and Providentia-SAW experiments for the RDM application.

ing efficiency, minimizing needless power consumption, and ensuring the safe operation of the vacuum.

For this experiment, the SVS was simulated in the Open Dynamics Engine,⁵ where its physical appearance (c.f., Figure 13) and behaviors were modeled on observations from the iRobot Roomba vacuum system. The SVS comprises two wheels, a circular body, seven touch sensors for collision detection, two wheel velocity sensors to monitor the status of each individual wheel speed, and a vacuum sensor that monitors the suction capabilities of the robot.

To motivate the need for run-time adaptation, uncertainty was configured in terms of system and environment-based uncertainty. System uncertainty comprised random sensor noise, sensor failures induced at random during execution, and fluctuations in the main controller timing logic (e.g., variations in the amount of time to spend in a particular path plan) to represent concerns that are found in real-time operating systems. Environment uncer-

 $^{^{5}}See http://www.ode.org/.$



Figure 11: Comparison of requirement violations between **Providentia** and **Providentia-SAW** experiments for the RDM application.



Figure 12: Comparison of starting/ending fitness subfunction weights for the RDM application with and without SAW optimization.



Figure 13: Screenshot of SVS simulation environment in the Open Dynamics Engine.

tainty includes the amount and distribution of dirt spread throughout the room, a downward step to avoid, and randomly-placed objects that may either hurt the SVS (e.g., a pole or liquid spill) or the object itself (e.g., a pet or child) that introduce safety concerns.

The SVS will reconfigure to minimize the impacts of uncertainty, thereby maximizing overall requirements satisficement. As with the RDM, the SVS performs run-time requirements monitoring via utility functions to quantify the performance of each separate requirement, where violations and/or unsatisfactory performance result in a reconfiguration. Possible reconfigurations include changing the current path plan (e.g., from a spiral to random search), updating the power moding strategy (e.g., from full power to reduced power), and instantiating emergency avoidance procedures to bypass a critical object.

For the SVS, we specify three NFRs to guide the system. Table 4 specifies the defined NFRs for the SVS, along with the sets of FRs and respective weights that comprise its aggregate utility function. Note that unlike the RDM application, **Providentia** is not provided FRs with weights of 0.0 to demonstrate its effectiveness in a smaller search space (i.e., fewer FRs) for each NFR.

NFR			Set of FRs
NFR1:	Minimize	[Cost]	(D), (F), (J), (R), (S)
			0.15, 0.3, 0.15, 0.2, 0.2
NFR2:	Minimize	[Time]	(A), (B), (C), (E), (G)
			0.3, 0.2, 0.2, 0.15, 0.15
NFR3:	Maximize	[Performance]	(A), (B), (C), (D)
			0.3, 0.25, 0.25, 0.2

Table 4: SVS NFRs with sets of FRs and manually-derived weights.

4.1.5. SVS Experimental Results

For this experiment, we reuse the experimental setup for the RDM (c.f., Section 4.1.1). Specifically, we examine how **Providentia**-optimized FR weighting schemes compare with weighting schemes that were manually- and randomly-specified (note that, for presentation purposes, we condense our results section to include **SAW** as well as the replication of the **Providentia** experiment in the SVS application domain). As the fitness function introduced in Equation (2) was specific to the RDM application, we now extend the fitness function to be specific to the SVS. The fitness function for the SVS is shown in Equation (8):

$$FF = \begin{cases} \alpha_{NFR} * FF_{NFR} + \alpha_{FR} * FF_{FR} + \\ \alpha_{adaptations} * FF_{adaptations} & \text{iff invariants true} \\ 0.0 & \text{otherwise} \end{cases}$$
(8)

First, we combine these fitness subfunctions with those fitness subfunctions specific to **Providentia**, as previously defined in Equations (3) and (5). As with the RDM, we set $\alpha_{NFR} = 0.375$, $\alpha_{FR} = 0.375$, and $\alpha_{adaptations} = 0.25$. We substitute Equation (6) that is specific to the RDM with Equation (9), specific to the SVS, to minimize the number of adaptations that the SVS experiences at run time:

$$FF_{adaptations} = \frac{1.0}{|adaptations|} \tag{9}$$

Figure 14 presents boxplots that show the fitness values between goal models that were optimized with Providentia-SAW and Providentia, those whose FR weights were manually-selected (Manual), and those whose weights were randomly selected (Random). As can be seen from the plots, optimizing

with Providentia significantly improves overall fitness of the SVS (p < 0.05, Wilcoxon-Mann-Whitney u-test). Moreover, introducing SAW to optimize the fitness sub-functions that guide Providentia (c.f., Equation (8), α_{NFR} , α_{FR} , $\alpha_{adaptations}$) further significantly improves fitness (p < 0.05, Wilcoxon-Mann-Whitney u-test).



Figure 14: Comparison of fitness values between all run types for the SVS application.

Similar to the results presented for the RDM application, Figure 15 demonstrates how the SVS attains significant NFR fitness improvements as Providentia and then Providentia-SAW are applied (p < 0.05, Wilcoxon-Mann-Whitney u-test).

Table 5 and Figure 16 demonstrate that although Providentia-SAW provides higher fitness than Providentia, the number of invariant and non-invariant violations are not significantly increased, therefore suggesting a positive impact to the Providentia technique with minimal negative repercussions.



Figure 15: Comparison of NFR fitness values between Providentia and Providentia-SAW experiments for the SVS application.

	Providentia	SAW
Invariant	μ : 3.592	μ : 3.418
violations	$\sigma: 25.042$	σ : 25.292
Noninvariant	μ : 732.765	μ : 748.941
violations	$\sigma: 36.174$	$\sigma: 45.178$

Table 5: Invariant and noninvariant average violations and standard deviations for SVS application.

Table 6 presents the average utility values (μ) and standard deviation (σ) for each of the SVS NFRs. As with the RDM, we see a significant improvement in average utility for most NFRs (p < 0.05, Wilcoxon-Mann-Whitney u-test), further suggesting that Providentia and Providentia-SAW significantly improve system performance while experiencing uncertainty.



Figure 16: Comparison of violations between Providentia and Providentia-SAW experiments for the SVS application.

NFR	Random	Manual	Providentia	SAW
NFR1: Minimize	μ : 0.972	μ : 0.948	μ : 0.957	μ : 0.951
[Cost]	σ : 0.114	σ : 0.099	σ : 0.200	σ : 0.210
NFR2: Minimize	μ : 0.564	μ : 0.510	μ : 0.833	μ : 0.773
[Time]	σ : 0.363	σ : 0.014	σ : 0.302	σ : 0.342
NFR3: Maximize	μ : 0.489	μ : 0.426	μ : 0.875	μ : 0.886
[Performance]	σ : 0.399	σ : 0.072	σ : 0.281	σ : 0.279

Table 6: NFR average utility values and standard deviations for SVS application.

5. Discussion

This paper has described search-based techniques and their corresponding empirical studies that use quantitative NFR performance information in the SAS feedback loop to support online decision making. To enable this calculation, a set of FRs (either new or existing) are identified from a requirements specification to support quantification of an NFR, where each selected FR is assigned a weight to indicate its relative importance in the satisfaction of the NFR's objectives. **Providentia** was introduced to perform automated optimization of the FR selection and weight definition process, given the large search space that results from this problem. Furthermore, **SAW** was added to **Providentia** to further optimize the weights of the fitness subfunctions that guided the search procedure.

Providentia and Providentia-SAW determine optimal sets of weights to make the system more robust to uncertainty. For example, the uncertainty introduced in the RDM simulation include the percentage of dropped/delayed messages, the percent chance that a server goes down, of a network link severing, sensor fuzz, and more. Uncertainty introduced in the SVS include a randomized chance of each sensor failing and/or fuzzing, the distribution of dirt particles in the room, a downward step in the room, the random instantiation of hazardous objects to the SVS (e.g., liquid, large objects, etc.), and more. The different types of uncertainty faced by the two case studies show the effectiveness and domain independence of Providentia and Providentia-SAW, as the RDM is from an industrial collaborator and the SVS is based on a real-world system.

Experimental results suggest that the introduction of NFRs into the SAS feedback ecosystem, along with **Providentia**'s automated optimization procedure to determine the configuration of each NFR, can produce an SAS that performs better in terms of overall fitness while significantly reducing the number of requirements violations incurred during execution. These results demonstrate that non-functional objectives, while historically difficult to quantify, can be automatically reconfigured and tuned to enable an SAS to deliver optimal performance in the face of uncertainty.

Figure 14 shows the comparison of fitness values (i.e., the values calculated in Equation (8)) for each of the four runs. With the manually-specified weights for the fitness function, **Providentia** performs the best because the GA is able to explore the weights for NFRs *that are separate* from the weights of the fitness function. For manual, random, and **Providentia**, the alpha values in Equation (8) are all equal. **Providentia-SAW** combines the results from **Providentia** (leading to a higher median fitness), but the search space for an optimal fitness value is explored as the overall fitness function is adjusted every fifth generation. **Providentia-SAW** has led to a greater degree of variance yet further improved the results from standalone **Providentia** (i.e., without SAW integration). Additional results demonstrated the significance of adjusting the weights of the fitness subfunctions that guide **Providentia**'s search procedure. While there was no significant impact to the satisficement of NFR utility values, FR utility values were significantly improved when using **Providentia-SAW**, thereby leading to a significant improvement in overall fitness of both SAS applications that were studied. This result suggests that, as was previously discovered [36], a linear-weighted fitness function is subject to the uncertainties imposed by each individual environment and that the statisticallyspecified weights may not generate optimal solutions for all environments considered. As a result, running a hyper-heuristic optimizer can significantly improve fitness in uncertain environments.

In cyber-physical systems, modeling real-world uncertainty is an ongoing problem as simulations often cannot include unexpected issues that in reality a system may face. Such a system would benefit from Providentiaoptimized NFRs as well as Providentia-SAW to balance NFR/FR/adaptation utility functions in unforeseen circumstances.

5.1. Threats to validity.

This paper has presented an extended proof of concept to demonstrate that quantifying NFRs at run time can support the reconfiguration engine in an SAS to improve overall requirements satisficement and minimize violations. As such, we have identified the following internal and external threats to validity for this research.

Internal. First, the derivation of the requirements for both the RDM and SVS applications was manually performed and may not be wholly inclusive of all possible requirements. Moreover, each application was simulated based on the executable specifications and may not exhaustively capture all of the detailed requirements. The manual selection of each FR set to support an NFR is another threat, as the selection may either be too limited or broad as it relies on domain knowledge, which is a common problem in general when identifying NFRs [7, 37].

External. External threats include impacts to the system as a result of unanticipated environmental conditions, unexpected human interaction with systems under execution, and unplanned for changes to the system requirements, thereby invalidating prior optimizations discovered by **Providentia** and **Providentia-SAW**.

Construct. Construct threats include scalability and generalizability concerns. In terms of scalability, optimization heuristics traditionally suffer from

a larger search space. For this paper, the size of the requirements specifications for both the RDM and SVS case studies may be considered small (i.e., 23 FRs and 7 NFRs for the RDM, and 19 FRs and 3 NFRs for the SVS), and as such, the results may not necessarily generalize to larger search spaces. In terms of generalizability, we have demonstrated **Providentia** and **Providentia-SAW** in two application domains: a networking application and a cyber-physical system, both of which are modeled as SASs. As such, it is possible that our techniques do not generalize to non-SAS domains and can be considered as future work for the authors.

We also only explored the genetic algorithm as a search heuristic. As such, other applicable search techniques, such as multi-objective optimization [38], could be used to discover ideal or more globally-optimal solutions. Therefore, an additional threat to validity lies in the search technique, and we plan to explore other such search heuristics in future work.

We next describe related work that span a number of complementary areas to further highlight our contributions.

6. Related Work

This section presents related work with regard to self-adaptive systems, obstacle mitigation and requirement satisficement, goal modeling, and NFRs.

Self-Adaptive Systems. Self-adaptive capabilities are generally expensive to build, difficult to modify, and are usually specific to a given application [27]. The Rainbow framework generalizes an SAS such that it can be reused in different systems, separating the self-adaptive control infrastructure from the system itself [27]. This separation enables use in legacy systems, localization of problems in separate modules, and software reuse. Rainbow is an architecture-based modeling technique while **Providentia** extends GORE at a higher level. Cheng *et al.* also presented the use of utility functions in SASs with quality NFRs such as performance, cost, and content fidelity [39]. Rather than introducing new metrics to measure quality attributes in a separate architecture, **Providentia** extends the preexisting goal model of FRs and uses metrics already defined by the system to measure NFRs. The use of predefined metrics saves both time and space when performing reconfigurations at run time.

Aceituna *et al.* presented a model to determine if, based on a given requirements model, an SAS can be put into undesired states [40]. Although this model can be useful in evaluating FRs, **Providentia** operates under the assumption that the FRs keep the system in an acceptable state. Moreover, **Providentia** focuses on guiding the SAS behavior based on FRs. Should the NFRs introduced with **Providentia** fail, the system will remain in a functionally-valid, yet less optimal, state.

Bencomo *et al.* use a Bayesian definition of surprise to measure the degrees of uncertainty that cause a self-adaptive system to deviate from expected behavior [41]. Their approach uses dynamic decision networks that use probability to determine the satisficement of an NFR based on a system's decision. Initial probabilities are either estimated or derived based on past statistical performance. **Providentia** uses utility functions to determine requirement/goal satisficement and violations, rather than probability, to measure uncertainty.

A recent direction for SASs involves Complex Event Processing (CEP) systems [42]. A CEP system analyzes event streams and detects specific events or patterns. Weisenburger *et al.* address the difficulties in making CEP systems self-adaptive and the need to identify conditions that trigger adaptations [42]. CEP systems can directly benefit from **Providentia** as CEP systems typically do not allow developers to specify metrics on quality attributes. **Providentia** provides quantifications on ambiguous goals and requirements that can be optimized in SASs and therefore can be used to evaluate CEP systems to determine the specific conditions that require an adaptation.

In addition to the numerous approaches for representing NFRs in goal models and SASs, Aspect-Oriented Requirements Engineering presents an approach to identify and specify cross-cutting concerns in separate modules, or aspects [43]. Yu *et al.* demonstrate that aspects can be identified in goal-oriented requirement analysis using both FRs and NFRs [44]. Similar to Providentia, Gray *et al.* elevate cross-cutting concerns to be represented with FRs in a goal model, in contrast to Providentia that includes NFRs [45]. As shown in the performance NFR of the RDM case study, cross-cutting concerns are prevalent in generic NFRs. DeVries and Cheng use evolutionary computation to automatically detect unwanted feature interactions [46], in contrast to Providentia and Providentia-SAW that introduce additional feature interactions by adding new non-functional requirements to a goal model. Bisbal and Cheng also explore unwanted feature interactions due to non-functional conflicts due to shared resources [47]. Note that Providentia and Providentia-SAW do not add unwanted feature interactions but rather evaluate the feature interactions between FRs with NFRs.

Obstacles and Requirements. Obstacle mitigation is a strategy for identifying and resolving obstacles to goal satisfaction. van Lamsweerde *et al.* [4, 48] have described a set of strategies for obstacle mitigation, however this approach does not specify to what degree of non-satisfaction that NFRs can become without impacting or degrading the overall system. **Providentia** can be used to supplement these strategies by extending the non-functional goal model in KAOS, and moreover, automatically optimize the FR/NFR weighting scheme.

Requirements monitoring is an approach for quantifying requirements at run time for use in detecting and mitigating obstacles as the system executes, including a monitoring framework developed by Feather *et al.* [49]. Sawyer *et al.* [50] have posited that requirements can be promoted to live run-time entities for use in self-adaptation feedback loops, with a notable example being the SAS MAPE-K feedback loop [1]. However, these approaches mainly focus on FRs. **Providentia**-optimized NFRs are also intended to be used in adaptation decisions at run time, supported by a set of well-defined FRs and their accompanying utility functions.

Goal Modeling. Many approaches similar to Providentia use goal modeling to address dependencies between FRs [51] or represent NFRs as soft goals [7, 52]. Other approaches use probabilistic methods to improve NFR/FR satisficement [53, 54] or optimize SAS satisficement [55, 56]. However, Providentia focuses solely on NFR/FR dependencies to optimize FR and NFR satisficement in an SAS without prior knowledge of system performance that most probabilistic methods require. Furthermore, Providentia does not use early-phase requirements engineering or high-level abstraction [57, 37], but rather focuses on a run-time model used by an SAS.

Non-Functional Requirements. Other techniques have been introduced to quantify NFRs, generally representing NFRs as soft goals [5, 58, 59]. Although both the RDM and SVS case studies use the KAOS goal modeling framework, Providentia is independent of any framework (e.g., NFR Framework, iStar, and KAOS). In contrast to modeling NFRs as soft goals, the weighted approach enables greater flexibility for an SAS to use in finding an optimal reconfiguration strategy at run time. Salehie *et al.* use a Goal-Action-Attribute Model (GAAM) and an automated weighting scheme called Analytic Hierarchy Process to prioritize NFRs [8]. However, priorities may shift due to uncertainty and requirement interactions at run time. Therefore, Providentia uses a genetic algorithm to optimize goal and weight selection instead of prioritization to make the goal model more robust to uncertainty at run time. Contributing work has decomposed NFR behaviors into monitored patterns rather than explicit requirements in a goal model [60] and used quantifiable metrics to represent NFRs separately from the FR goal model [61]. **Providentia** monitors requirements at run time and does not separate NFRs from the goal model of FRs, as separating NFRs and FRs may prevent the requirements engineer to identify cross-cutting concerns in NFRs.

7. Conclusion

This paper described **Providentia**, a design-time approach to automatically quantify NFRs at run time. **Providentia** uses a genetic algorithm to determine an optimal weighting scheme of FRs to describe each NFR, where optimal results yield the highest overall fitness. To illustrate the effectiveness of **Providentia**, we used an industry-provided RDM application that distributes messages across a network while experiencing random sources of uncertainty. We extended previous work [6] to use a second case study to evaluate the effectiveness of the technique in a different domain.

To further improve **Providentia**, we incorporated **SAW** to optimize the weighting scheme of the fitness subfunctions for both case studies. Experimental results suggest that the **Providentia**-optimized systems are more robust against system and environmental uncertainty and fulfills its requirements to a higher degree when compared to the systems without **Providentia**.

Future directions for this research include extending the Providentia search procedure to execute at run time, applying both techniques to a realworld cyber-physical system, and evaluating other search techniques that may better handle competing concerns in fitness calculation than a linearweighted sum (e.g., multi-objective optimization). We also intend to work towards open-sourcing the RDM and SVS applications, respectively.

Acknowledgements

This work has been supported in part by grants from the NSF (CNS-1657061, CNS-1305358, and DBI-0939454), the Michigan Space Grant Consortium, the Comcast Innovation Fund, Oakland University, Ford Motor Company, General Motors Research, the Air Force Research Laboratory (AFRL) under agreement numbers FA8750-16-2-0284 and FA-8750-19-2-0002,

and Michigan State University. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and do not necessarily represent the opinions of the sponsors.

References

- J. Kephart, D. Chess, The vision of autonomic computing, Computer 36 (2003) 41 - 50.
- [2] P. McKinley, S. Sadjadi, E. Kasten, B. H. C. Cheng, Composing adaptive software, Computer 37 (2004) 56 – 64.
- [3] L. Chung, B. Nixon, E. Yu, J. Mylopoulos, Non-functional requirements, Software Engineering (2000).
- [4] A. van Lamsweerde, Requirements Engineering: From System Goals to UML Models to Software Specifications, Wiley, 2009.
- [5] A. Yrjönen, J. Merilinna, Extending the nfr framework with measurable nonfunctional requirements, in: NFPinDSML@ MoDELS.
- [6] K. M. Bowers, E. M. Fredericks, B. H. Cheng, Automated optimization of weighted non-functional objectives in self-adaptive systems, in: International Symposium on Search Based Software Engineering (SSBSE 2017), Springer, pp. 182–197.
- [7] E. S. K. Yu, Towards modelling and reasoning support for early-phase requirements engineering, in: Proceedings of the Third IEEE International Symposium on Requirements Engineering, pp. 226–235.
- [8] M. Salehie, L. Tahvildari, Towards a goal-driven approach to action selection in self-adaptive software, Software: Practice and Experience 42 (2012) 211–233.
- [9] E. Burke, G. Kendall, J. Newall, E. Hart, P. Ross, S. Schulenburg, Hyper-heuristics: An emerging direction in modern search technology, in: F. Glover, G. A. Kochenberger (Eds.), Handbook of Metaheuristics, volume 57, Springer US, 2003, pp. 457–474.
- [10] J. H. Holland, Adaptation in Natural and Artificial Systems, MIT Press, Cambridge, MA, USA, 1992.
- [11] B. Craenen, A. E. Eiben, Stepwise adaption of weights with refinement and decay on constraint satisfaction problems, in: Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation, Morgan Kaufmann Publishers Inc., pp. 291–298.

- [12] A. E. Eiben, J. K. van der Hauw, Solving 3-sat by gas adapting constraint weights, in: Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC'97), IEEE, pp. 81–86.
- [13] A. E. Eiben, J. K. van der Hauw, Adaptive penalties for evolutionary graph coloring, in: Artifical Evolution, Springer, pp. 95–106.
- [14] M. Ji, A. Veitch, J. Wilkes, Seneca: Remote mirroring done write, in: USENIX 2003 Annual Technical Conference, USENIX Association, Berkeley, CA, USA, 2003, pp. 253–268.
- [15] K. Keeton, C. Santos, D. Beyer, J. Chase, J. Wilkes, Designing for disasters, in: Proceedings of the 3rd USENIX Conference on File and Storage Technologies, USENIX Association, Berkeley, CA, USA, 2004, pp. 59–62.
- [16] N. Bencomo, J. Whittle, P. Sawyer, A. Finkelstein, E. Letier, Requirements reflection: Requirements as runtime entities, in: Proc. of the 32nd ACM/IEEE International Conference on Software Engineering, ACM, Cape Town, South Africa, 2010, pp. 199–202.
- [17] N. Bencomo, A. Belaggoun, Supporting decision-making for selfadaptive systems: from goal models to dynamic decision networks, in: Requirements Engineering: Foundation for Software Quality, Springer, 2013, pp. 221–236.
- [18] B. H. C. Cheng, P. Sawyer, N. Bencomo, J. Whittle, A goal-based modeling approach to develop requirements of an adaptive system with environmental uncertainty, in: Proc. of the 12th International Conference on Model Driven Engineering Languages and Systems, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 468–483.
- [19] J. Whittle, P. Sawyer, N. Bencomo, B. H. C. Cheng, J. Bruel, Relax: Incorporating uncertainty into the specification of self-adaptive systems, in: 17th IEEE International Requirements Engineering Conference (RE'09), pp. 79–88.
- [20] P. Oreizy, M. Gorlick, R. Taylor, D. Heimhigner, G. Johnson, N. Medvidovic, A. Quilici, D. Rosenblum, A. Wolf, An architecture-based approach to self-adaptive software, Intelligent Systems and their Applications, IEEE 14 (1999) 54 –62.

- [21] S. Neema, T. Bapty, J. Scott, Development environment for dynamically reconfigurable embedded systems, in: Proceedings of the International Conference on Signal Processing Applications and Technology. Orlando, FL.
- [22] Y. Li, J. Chen, L. Feng, Dealing with uncertainty: A survey of theories and practices, IEEE Transactions on Knowledge and Data Engineering 25 (2013) 2463–2482.
- [23] C. Chua Chow, R. K. Sarin, Known, unknown, and unknowable uncertainties, Theory and Decision 52 (2002) 127—138.
- [24] N. Esfahani, E. Kouroshfar, S. Malek, Taming uncertainty in selfadaptive software, in: Proceedings. of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering, ACM, 2011, pp. 234–244.
- [25] D. O. Keck, P. J. Kuehn, The feature and service interaction problem in telecommunications systems: A survey, IEEE Transactions on Software Engineering 24 (1998) 779–796.
- [26] A. Ramirez, B. Cheng, Automatically deriving utility functions for monitoring software requirements, in: Proceedings of the 2011 International Conference on Model Driven Engineering Languages and Systems Conference, Wellington, New Zealand, pp. 501–516.
- [27] D. Garlan, S.-W. Cheng, A.-C. Huang, B. Schmerl, P. Steenkiste, Rainbow: Architecture-based self-adaptation with reusable infrastructure, Computer 37 (2004) 46–54.
- [28] A. J. Ramirez, D. B. Knoester, B. H. C. Cheng, P. K. McKinley, Applying genetic algorithms to decision making in autonomic computing systems, in: Proceedings of the 6th International Conference on Autonomic Computing, pp. 97–106. (Best paper award).
- [29] A. Dardenne, A. Van Lamsweerde, S. Fickas, Goal-directed requirements acquisition, Science of computer programming 20 (1993) 3–50.
- [30] L. Chung, B. A. Nixon, E. Yu, J. Mylopoulos, Non-functional requirements in software engineering, Springer Science & Business Media, 2012.

- [31] P. deGrandis, G. Valetto, Elicitation and utilization of applicationlevel utility functions, in: Proc. of the 6th International Conference on Autonomic Computing, ICAC '09, ACM, 2009, pp. 107–116.
- [32] W. Walsh, G. Tesauro, J. Kephart, R. Das, Utility functions in autonomic systems, in: Proceedings of the First IEEE International Conference on Autonomic Computing, IEEE Computer Society, 2004, pp. 70–77.
- [33] A. C. Kumari, K. Srinivas, Software module clustering using a fast multiobjective hyper-heuristic evolutionary algorithm, International Journal of Applied Information Systems 5 (2013) 12–18.
- [34] van der Hauw K, Evaluating and improving steady state evolutionary algorithms on constraint satisfaction problems, Master's thesis, Leiden University, 1996.
- [35] A. Arcuri, L. Briand, A practical guide for using statistical tests to assess randomized algorithms in software engineering, in: Proceedings of the 33rd International Conference on Software Engineering, ICSE '11, ACM, 2011, pp. 1–10.
- [36] E. M. Fredericks, B. DeVries, B. H. C. Cheng, Autorelax: Automatically relaxing a goal model to address uncertainty, Empirical Software Engineering 19 (2014) 1466–1501.
- [37] J. Mylopoulos, L. Chung, B. Nixon, Representing and using nonfunctional requirements: A process-oriented approach, IEEE Transactions on software engineering 18 (1992) 483–497.
- [38] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: Nsga-ii, Evolutionary Computation, IEEE Transactions on 6 (2002) 182 –197.
- [39] S.-W. Cheng, D. Garlan, B. Schmerl, Evaluating the effectiveness of the rainbow self-adaptive system, in: ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems, 2009. SEAMS '09., pp. 132–141.

- [40] D. Aceituna, H. Do, Exposing the susceptibility of off-nominal behaviors in reactive system requirements, in: 2015 IEEE International Conference on Requirements Engineering (RE), IEEE, pp. 136–145.
- [41] N. Bencomo, A. Belaggoun, A world full of surprises: Bayesian theory of surprise to quantify degrees of uncertainty, in: Companion Proceedings of the 36th International Conference on Software Engineering, ACM, pp. 460–463.
- [42] P. Weisenburger, M. Luthra, B. Koldehofe, G. Salvaneschi, Qualityaware runtime adaptation in complex event processing, in: Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, IEEE Press, pp. 140–151.
- [43] A. Rashid, P. Sawyer, A. Moreira, J. Araújo, Early aspects: A model for aspect-oriented requirements engineering, in: 2002 IEEE Joint International Conference on Requirements Engineering, IEEE, pp. 199–202.
- [44] Y. Yu, J. C. Leite, J. Mylopoulos, From goals to aspects: discovering aspects from requirements goal models, in: 2004 IEEE International Requirements Engineering Conference, IEEE, pp. 38–47.
- [45] J. Gray, T. Bapty, S. Neema, D. Schmidt, A. Gokhale, B. Natarajan, An approach for supporting aspect-oriented domain modeling, in: Generative Programming and Component Engineering, Springer, pp. 151–168.
- [46] B. DeVries, B. H. Cheng, Automatic detection of feature interactions using symbolic analysis and evolutionary computation, in: 2018 IEEE International Conference on Software Quality, Reliability and Security (QRS), IEEE, pp. 257–268.
- [47] J. Bisbal, B. H. Cheng, Resource-based approach to feature interaction in adaptive software, in: Proceedings of the 1st ACM SIGSOFT workshop on Self-managed systems, ACM, pp. 23–27.
- [48] A. van Lamsweerde, E. Letier, Handling obstacles in goal-oriented requirements engineering, IEEE Transactions on Software Engineering 26 (2000) 978–1005.
- [49] M. Feather, S. Fickas, A. Van Lamsweerde, C. Ponsard, Reconciling system requirements and runtime behavior, in: Proceedings of the Ninth

International Workshop on Software Specification and Design, 1998, pp. 50–59.

- [50] P. Sawyer, N. Bencomo, J. Whittle, E. Letier, A. Finkelstein, Requirements-aware systems: A research agenda for RE for self-adaptive systems, in: Requirements Engineering Conference (RE), 2010 18th IEEE International, pp. 95–103.
- [51] B. Nagel, C. Gerth, J. Post, G. Engels, Kaos4SOA-extending kaos models with temporal and logical dependencies., in: CAiSE Forum, pp. 9–16.
- [52] P. Giorgini, J. Mylopoulos, R. Sebastiani, Goal-oriented requirements analysis and reasoning in the tropos methodology, Engineering Applications of Artificial Intelligence 18 (2005) 159–171.
- [53] A. Cailliau, A. van Lamsweerde, Runtime monitoring and resolution of probabilistic obstacles to system goals, in: Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, IEEE Press, pp. 1–11.
- [54] L. H. G. Paucar, N. Bencomo, The reassessment of preferences of non-functional requirements for better informed decision-making in self-adaptation, in: Requirements Engineering Conference Workshops (REW), IEEE International, IEEE, pp. 32–38.
- [55] E. Letier, A. van Lamsweerde, Reasoning about partial goal satisfaction for requirements and design engineering, in: Proceedings of the 12th ACM SIGSOFT Twelfth International Symposium on Foundations of Software Engineering, pp. 53–62.
- [56] Z. Yang, Z. Jin, Z. Li, Achieving adaptation for adaptive systems via runtime verification: A model-driven approach, arXiv preprint arXiv:1704.00869 (2017).
- [57] F. Dalpiaz, A. Borgida, J. Horkoff, J. Mylopoulos, Runtime goal models: Keynote, in: 2013 IEEE Seventh International Conference on Research Challenges in Information Science (RCIS), IEEE, pp. 1–11.

- [58] N. Kobayashi, S. Morisaki, N. Atsumi, S. Yamamoto, Quantitative non functional requirements evaluation using softgoal weight., J. Internet Serv. Inf. Secur. 6 (2016) 37–46.
- [59] S. Yamamoto, An approach for evaluating softgoals using weight, in: Information and Communication Technology, Springer, 2015, pp. 203– 212.
- [60] S. Supakkul, T. Hill, L. Chung, T. T. Tun, J. C. S. do Prado Leite, An nfr pattern approach to dealing with nfrs, in: Requirements Engineering Conference (RE), 2010 18th IEEE International, IEEE, pp. 179–188.
- [61] D. Sykes, W. Heaven, J. Magee, J. Kramer, Exploiting non-functional preferences in architectural adaptation for self-managed systems, in: Proceedings of the 2010 ACM Symposium on Applied Computing, ACM, pp. 431–438.