

FEMOSAA: Feature Guided and Knee Driven Multi-Objective Optimization for Self-Adaptive Software

TAO CHEN, Department of Computing and Technology, Nottingham Trent University, UK, and CERCIA, School of Computer Science, University of Birmingham, UK

KE LI, School of Computer Science and Engineering, University of Electronic Science and Technology of China, China, and Department of Computer Science, University of Exeter, UK

RAMI BAHSOON, CERCIA, School of Computer Science, University of Birmingham, UK

XIN YAO, Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China, and CERCIA, School of Computer Science, University of Birmingham, UK.

Self-adaptive software (SAS) can reconfigure itself to adapt to the changing environment at runtime, aiming for continually optimizing conflicted non-functional objectives, e.g., response time, energy consumption, throughput and cost *etc.* In this paper, we present Feature guided and knEe driven Multi-Objective optimization for Self-Adaptive softwAre (FEMOSAA), a novel framework that automatically synergizes the feature model and Multi-Objective Evolutionary Algorithm (MOEA), to optimize SAS at runtime. FEMOSAA operates in two phases: at design time, FEMOSAA automatically transposes the engineers' design of SAS, expressed as a feature model, to fit the MOEA, creating new chromosome representation and reproduction operators. At runtime, FEMOSAA utilizes the feature model as domain knowledge to guide the search and further extend the MOEA, providing a larger chance for finding better solutions. In addition, we have designed a new method to search for the knee solutions, which can achieve a balanced trade-off. We comprehensively evaluated FEMOSAA on two running SAS: one is a highly complex SAS with various adaptable real-world software under the realistic workload trace; another is a service-oriented SAS that can be dynamically composed from services. In particular, we compared the effectiveness and overhead of FEMOSAA against four of its variants and three other search-based frameworks for SAS under various scenarios, including three commonly applied MOEAs, two workload patterns and diverse conflicting quality objectives. The results reveal the effectiveness of FEMOSAA and its superiority over the others with high statistical significance and non-trivial effect sizes.

CCS Concepts: •**Software and its engineering** →**Software performance; Search-based software engineering;**

Additional Key Words and Phrases: feature model, search-based software engineering, multi-objective evolutionary algorithm, multi-objective optimization, self-adaptive system, performance engineering

ACM Reference format:

Tao Chen, Ke Li, Rami Bahsoon, and Xin Yao. 2018. FEMOSAA: Feature Guided and Knee Driven Multi-Objective Optimization for Self-Adaptive Software. *ACM Trans. Softw. Eng. Methodol.* 1, 1, Article 1 (January 2018), 48 pages.

DOI: 10.1145/3204459

This work is supported by the Ministry of Science and Technology of China (Grant No. 2017YFC0804003), Science and Technology Innovation Committee Foundation of Shenzhen (Grant No. ZDSYS201703031748284), and EPSRC (Grant Nos. EP/J017515/01 and EP/K001523). The co-corresponding author: Tao Chen (t.chen@cs.bham.ac.uk), Ke Li (keli.genius@gmail.com) and Xin Yao (xiny@sustc.edu.cn).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s). 1049-331X/2018/1-ART1 \$

DOI: 10.1145/3204459

1 INTRODUCTION

Self-Adaptive Software (SAS) is a special type of software that is capable of adapting and reconfiguring itself at runtime, through a set of known features (e.g., CPU cap, thread pool size and cache size, *etc*), according to the changing environment [17]. One major goal of SAS is to continually optimize multiple and often conflicting non-functional objectives, e.g., response time versus energy consumption, throughput versus cost, *etc*. However, given the dynamic and uncertain nature of running software, it is difficult to fully specify all possible conditions and their adaptation solutions at design time. Thus, designing an efficient and effective runtime optimization approach is necessary, yet challenging. Depending on the complexity of SAS, software engineers have exploited various search algorithms, e.g., exact or stochastic search, for continually finding the optimal (or near-optimal) adaptation solution for SAS at runtime [23][51][12][42][41][15][13].

To optimize SAS at runtime using the search algorithms, there are two crucial challenges: (i) firstly, it is difficult to effectively and systematically convert the SAS design to the context of search algorithm while considering the right encoding of features in the representation of optimization, e.g., using only the features that contribute to different aspects of the variability of SAS. Here, the features might be categorical or numeric, where the former refers to those with distinct characteristics, e.g., the *Cache* feature is 'on' or 'off'; the latter denotes those that can be quantified, measured and sorted, e.g., the size of *maxThreads*. Furthermore, it is difficult to effectively and systematically handle the features' dependencies, e.g., one can change *Cache Mode* only if the *Cache* feature is 'turned on'. Dependency can become even more complex in the presence of numeric features, e.g., in Tomcat [2], the size of *maxThreads* should not be less than the size of *minSpareThreads*. Those conversion tasks are non-trivial as the design of SAS can be complex and most search algorithms cannot handle dependency constraints in nature. (ii) Secondly, optimizing multiple conflicting objectives and managing their trade-offs are complex and challenging, especially for SAS runtime. This is attributed to the huge number of alternative adaptation solutions and the required efficiency for the found solution to be effective. Moreover, the dynamic and uncertain nature of SAS further complicates the conflicting relations between objectives, rendering the trade-off surface difficult to be explored. Those challenges, when not appropriately addressed, can result in compromised quality, unacceptable running overhead and imbalanced trade-off in SAS runtime optimization.

Most existing work fails to handle the first challenge as they have relied on a manual and/or incomplete conversion of the SAS design into the search algorithm's context [42][1][22][51], which renders the process expensive, non-systematic and error-prone. Moreover, the feature dependencies are often ignored, wasting the valuable function evaluations on invalid solutions at SAS runtime while providing no guarantee on finding the valid ones. Inspired by the applications of search algorithms to Software Product Line problems [45], researchers [23][41] have combined the feature model [33] with search algorithms to optimize SAS at runtime, considering categorical dependencies. However, numeric features are ignored and a solution often encodes all the features using a simple binary representation. This might lead to the *curse of dimensionality*, and thereby entailing unnecessary complexity at SAS runtime. Further, existing approaches cannot prevent wasteful exploration of invalid solutions and difficult to handle the dependencies related to numeric features.

For the second challenge above, exact search [23] [9], with the helps of objective aggregation (e.g., a weighted sum), has been exploited for SAS runtime optimization. However, modern SAS often exhibits high variability, leading to an explosion of the search space of all possible solutions and rendering the problem intractable. Henceforth, exact search may fail to scale at runtime. In contrast, stochastic search, particularly Evolutionary Algorithms that are widely applied in Search-Based Software Engineering (SBSE), tends to be naturally robust in solving problems with extremely

high number of alternatives and thus appealing for SAS optimization [29]. Those algorithms, when properly tailored, can lead to approximate and near-optimal solutions for complex software engineering problems with reasonable running time as of minutes, if not seconds [31]. Furthermore, stochastic search has proven to be effective for many real-time systems [22][27][51][12]. Often, existing approaches rely on single-objective evolutionary algorithm to optimize SAS by simply transforming a multi-objective problem into an aggregated single-objective one [42][27]. While objective aggregation might be preferable for some contexts, it has been shown that there are cases where assigning weights to different objectives is a non-trivial task for software engineers and the aggregation can hardly maintain a good diversity of the solutions [29]. To alleviate this issue, studies [1][22] [51] have used NSGA-II [20], a popular Multi-Objective Evolutionary Algorithm (MOEA), to optimize SAS without using the weighted aggregation; they have shown that MOEA can find more convergent and diverse solutions in the trade-off surface than optimizing via objective aggregation. However, NSGA-II has a coarse diversity preservation mechanism that is unable to provide well distributed solutions in certain cases [52]. Therefore, it is desirable to have a general framework that can easily work with different MOEAs for optimizing SAS without suffering the limitation from one specific algorithm. In addition, given the fact that MOEAs produce a set of non-dominated solutions, there is no established method for the SAS to choose an appropriate one for adaptation at runtime, entailing the risk of imbalanced trade-offs.

To address the aforementioned challenges and limitations, this paper presents Feature guided and knEe driven Multi-Objective optimization for Self-Adaptive softwAre (FEMOSAA), a novel framework that automatically synergizes the feature model and a given MOEA, to optimize SAS at runtime. Specifically, our contributions include:

- We rely on the feature model to represent the design of a given SAS with explicit considerations of numeric features and their dependencies. In FEMOSAA, we provide an automatic and systematic approach to transpose a given design of SAS, expressed as a feature model, into the MOEA's context at design time. Further, such transposition extends the internal structure of MOEAs in order to improve their ability to search for better adaptation solutions at SAS runtime. Notably, we contribute to the following in the transposition approach:
 - (1) To tailor the problem to be more suitable for SAS runtime, we discard the lengthy binary encoding. Instead, our approach identifies the *elitist features* from the feature model to encode an elegant and polyadic chromosome representation in the MOEA. By elitist features, we refer to those that cannot be removed in the optimization without damaging the original variability of SAS while minimizing the length of encoding. The benefit of such encoding is that (i) it is intuitive, simpler and enable direct dependency extraction and (ii) reducing the number of genes helps to greatly shrink the search space and simplify the dependency constraints, which also improves the quality of the solutions found while shortening the running time of MOEA.
 - (2) To better guide the search and avoid exploring invalid solutions, our approach extracts the feature dependencies with respect to the elitist features. Then, these dependencies are injected into the basic mutation and crossover operators of the MOEA to create new dependency aware operators. These operators can systematically steer the MOEA to focus on exploring the valid solutions of SAS, creating a larger chance to find better ones.
- Without loss of generality, we design FEMOSAA in such a way that it can be seamlessly integrated with different MOEAs¹ to optimize SAS at runtime. The elitist features and extracted dependencies, as processed by the transposition approach at design time, are used to guide

¹In addition to MOEAs, FEMOSAA also works with single-objective evolutionary algorithms in which case the knee selection method would be deactivated.

the running behaviors of a given MOEA for SAS runtime optimization. In this work, we run FEMOSAA with three fundamentally distinct yet widely-used MOEAs in the literature, i.e., MOEA based Decomposition with STable-Matching model (MOEA/D-STM) [36], Non-dominated Sort Genetic Algorithm-II (NSGA-II) [20] and Indicator Based Evolutionary Algorithm (IBEA) [53].

- To achieve a balanced trade-off in SAS optimization, FEMOSAA identifies knee solutions automatically from the final non-dominated set. The knee solutions often imply well balanced trade-offs, such that any improvement on one objective of a knee will cause relatively severe degradations on others.
- We conduct comprehensive experiments on two running SAS: one is a highly complex SAS that consists of the *eBay*-like RUBiS benchmark [43] and a set of real-world adaptable software (i.e., Apache Tomcat [2], MySQL [40], Ehcache [3] and Xen [48]) under the realistic FIFA98 workload trace [5]; another is a service-oriented SAS that can be dynamically composed by various services. We compare FEMOSAA with four of its variants (e.g., without dependency aware operators) and three other state-of-the-art frameworks (i.e., DUSE [1], PLATO [42] and FUSION [23]) under various scenarios, including three commonly applied MOEAs (i.e., MOEA/D-STM, NSGA-II and IBEA) and two different workload patterns² (i.e., read-write and read-only) along with diverse conflicting quality objectives. The experiments reveal the effectiveness of FEMOSAA and its superiority over the others when optimizing conflicting objectives for SAS, with statistically significant results and non-trivial effect sizes.

The contributions have clear impact on the synergy between software engineering for SAS and evolutionary computation as FEMOSAA combines the strengths from both fields. Unlike many SBSE work that simply formulates the software engineering problem as a classic optimization problem for some MOEAs, our deeper synergy takes one step further by automatically and dynamically extracting the domain information of SAS to extend the internal structure of MOEA, improving its search ability. As a result, to control and exploit the power of MOEAs, the software engineers of SAS only need to provide the feature model when using FEMOSAA, without being an expert on MOEA. In addition, FEMOSAA improves MOEA and provides insights for MOEA researchers to design better algorithms for SAS, since the identified elitist features and their dependencies serve as the engineers' systematic domain knowledge by which we can reduce the search space and better guide the search, providing a larger chance for finding better solutions.

The remainder of this paper is organized as follows: Section 2 illustrates a detailed motivating example of SAS. Section 3 presents the background and the extended notions of numeric features in the feature model. Section 4 gives an overview of FEMOSAA. Section 5 illustrates our approach that transposes a feature model to MOEA. Section 6 presents how the internal structure of existing MOEAs can be extended to combine with our dependency aware operators and knee selection. Experimental results, verifiability and threats to validity are discussed in Section 7. Finally, Sections 8 and 9 present related work and conclusion respectively.

2 A DETAILED MOTIVATING SCENARIO OF SELF-ADAPTIVE SOFTWARE

While our work can be applied to different contexts that demand runtime adaptation, we draw on a representative and realistic SAS to motivate and illustrate the need. As shown in Figure 1, like many SAS, the SAS example consists of two parts: an adaptable software that is being managed at runtime, and an engine that controls the adaptation. Additionally, the SAS contains a complex software stack consisting of RUBiS³ [43], Apache Tomcat [2], Ehcache [3] and MySQL [40], running

²Different workload patterns will create diverse behaviors of the SAS.

³An eBay like software application with 26 services.

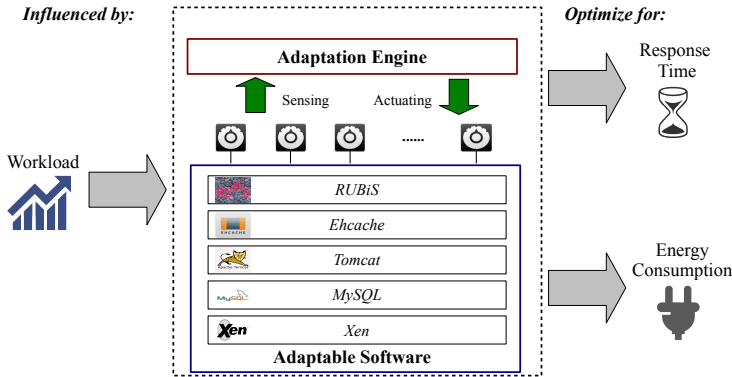


Fig. 1. An example of SAS.

on the virtualization hypervisor Xen [48]. The RUBiS benchmark serves as a representative of many real-world software applications that offers diverse functionalities and services to many end-users concurrently. We can see from Figure 1, as the case of most practical software applications, the SAS’s software stack contains many off-the-shelf real-world software as described above. Each of the software supports various control features, which, together with those from the other software in the stack, can be changed dynamically on-the-fly to influence the runtime behaviors of the software system. Example of the control features includes the number of threads, the memory allocation and enabling/disabling cache mechanism, *etc.* By design, all the possible configurations of control features form the search space, or variability of the SAS. As the workload changes, the SAS is capable of adapting the features at runtime to optimize for various non-functional quality attributes, e.g., response time. To achieve such goal, thanks to the rapid development of search algorithms, through which the SAS is often designed as to continually search for the combination of feature configurations that lead to the optimal (or near optimal) quality at runtime. However, to effectively and efficiently engineer the SAS in this way is challenging for the following reasons:

Encoding the Features from the SAS Design. Consider a complex SAS which contains many features and configurations, systematically and generically choosing the right features and encode them into the representation of search algorithm is difficult. To optimize the SAS at runtime, such representation defines the fundamental search space of the problem to be explored, therefore the encoding of features could have positive or negative impact on the search ability of a potential search algorithm. Given that some features in the SAS design do not contribute to the SAS’s variability or they can represent the same aspect of variability [6], existing work [23] [41] that simply encodes all features in a binary format is unnecessary. Suppose a feature model with 100 features, binary representation can easily create a search space of 2^{100} and this, as we will show in Section 7.4.1, can negatively affect the adaptation quality and overhead.

Handling the Dependencies in the SAS Design. Many widely-used exact and stochastic search algorithms, e.g., MOEA, are not designed to handle dependencies constraints. This makes the treatment of dependencies difficult especially when the dependencies in SAS come in a mixture of categorical dependencies, e.g., *Cache Mode require Cache*, and numeric ones, e.g., $maxThreads \geq minSpareThreads$. As we will show in Section 7.4.2, those dependencies, when ignored [42] [1] [27] or incorrectly handled [23] [41] (as in existing work), can degrade the adaptation quality.

Explosion of the Search Space. Modern SAS often has high variability leading to an explosion of the search space. For example, the original design of the SAS shown in Figure 1 has a search space of more than a billion, which we will elaborate in details at Section 3.3.

Trade-off on the Conflicting Objectives. SAS often exhibits multiple conflicting quality objectives, e.g., response time versus energy consumption and throughput versus cost, which need to be optimized simultaneously and trade-off needs to be made to comprise some of them. In general, many existing approaches [42] have assumed that the relative importance of objectives can be correctly quantified as numeric weights, which has been found to be difficult in some cases [29]. Those weights, when inappropriately specified and expressed, would inevitably create negative impact on the search process and result in unwanted bad adaptation quality. It is even more difficult to achieve balanced trade-off.

These difficulties motivate our work, which automatically synergizes the feature model of SAS and a given MOEA, creating feature guided MOEA with knee selection, to optimize SAS at runtime.

Algorithm 1 General algorithmic process of MOEA

Input: given mutation rate r_m , crossover rate r_c and the maximum number of evaluation $eval_{max}$, which is often equivalent to the size of population \times the maximum number of generations

Output: a set of optimized non-dominated solutions

```

1: start evolution
2:  $P = \emptyset$ 
3:  $eval = 0$ 
4: for  $i = 1$  to  $P_{size}$  do
5:    $S = \text{GETRANDOMSOLUTION}()$ 
6:    $\text{EVALUATEFITNESS}(S)$ 
7:    $eval = eval + 1$ 
8:    $P = P + S$ 
9: end for
10: while  $eval < eval_{max}$  do
11:    $P_0 := \emptyset$ 
12:   while  $|P_0| \leq P_{size}$  do
13:      $parents := \text{DOMATINGSELECTION}(P)$ 
14:      $offspring := \text{DOCROSSOVER}(parents, r_c)$ 
15:     for each solution  $S$  in  $offspring$  do
16:        $\text{DOMUTATION}(S, r_m)$ 
17:     end for
18:      $\text{EVALUATEFITNESS}(offspring)$ 
19:      $eval := eval + |offspring|$ 
20:      $P_0 := P_0 \cup offspring$ 
21:   end while
22:    $P := P \cup P_0$ 
23:    $\text{DOSURVIVALSELECTION}(P, P_{size})$ 
24: end while
25: return  $\text{GETNONDOMINATEDSOLUTIONS}(P)$ 
26: end evolution

```

3 BACKGROUND AND PRELIMINARIES

3.1 Multi-Objective Evolutionary Algorithm (MOEA)

Evolutionary algorithm, a stochastic search-based meta-heuristic, has been widely accepted as a major approach for solving multi-objective optimization problems [19], in which case it is also known as MOEA. In MOEA, the population contains a set of solutions (individuals), each of which is represented by a fixed-length thread-like chromosome carrying different values at each gene. As shown in Figure 2 and Algorithm 1, the evolutionary search of MOEA starts after the initialization of the population (line 2 to 9). During the search process, the elite information can propagate from the parents to the offspring via some random and probabilistic reproduction operations (i.e., crossover and mutation) upon the mating parents chosen from the mating selection procedure. Inspired by the *survival of the fittest* rule from the evolutionism, the survival selection preserves the high quality individuals, having superior fitness values, to the next iteration (generation), as shown from line 10 to 24. The evolution process repeats until a stopping criteria, e.g., a predefined function evaluation threshold, is satisfied. The major difference between MOEA and the classic




Fig. 2. The general workflow of MOEA.




Fig. 3. Pareto optimal and knee solutions.

single-objective evolutionary algorithm lies in the mating and survival selection mechanisms. In particular, instead of finding a single optimal (or near optimal) solution as in the single-objective evolutionary algorithm, MOEA aims to find a set of non-dominated solutions⁴ that approximates the *Pareto Front* with a good convergence and uniform distribution (line 25). Notably, for every solution in the non-dominated set, any improvement of an objective will result in a degradation for at least one other objective.

Generally, the existing MOEAs can be divided into the following three categories according to the survival selection mechanisms:

- *Decomposition-based method*: The MOEA decomposes the original multi-objective optimization problem into several single-objective optimization subproblems by linear or non-linear aggregation methods [38]. Then, it uses a population-based technique to solve these subproblems in a collaborative manner. MOEA/D [52], MOEA/D-STM [36] and NSGA-III [18] are the representative algorithms of this sort.
- *Pareto-based method*: The MOEA uses Pareto dominance relation as the primary selection criterion to push the solutions towards the Pareto front as close as possible. In the meanwhile, it employs some density estimation techniques, e.g., the crowding distance [20] and the clustering analysis [54], to maintain the population diversity. The representative algorithms are NSGA-II [20], SPEA2 [54] and PAES [34] etc.
- *Indicator-based method*: Here, sophisticated performance indicators are designed to measure the overall quality of a solution set. The representative algorithm is IBEA [53], which transfers the multi-objective optimization problem into a new single-objective one that aims to find the optimal set of solutions with respect to a given indicator.

3.2 Knee Solutions

The MOEA generates a set of non-dominated solutions that approximate the Pareto front. However, not every non-dominated solutions can lead to balanced trade-off for SAS runtime optimization. Indeed, the most common purpose of MOEA is to search and visualize a set of non-dominated solutions that are as close to the true Pareto front as possible. Then, a human decision maker can pick whichever solution that s/he prefers. However, there is no such a human available in the SAS optimization problem. Therefore, a method is required to pick a sole solution from the resulted set of non-dominated solutions to execute adaptation.

A simple Pareto optimal front is shown in Figure 3 where the two objectives should be minimized. Clearly, solutions near the edges strongly favor one objective over the other but there is a visible bulge around the middle, which is the knee region. Those solutions in the knee region (or simply knee solutions) are characterized by the fact that a small improvement in either objective will cause

⁴A solution dominates another if it has at least one objective better than another while all other objectives are not worse than another. Non-dominated solutions denote those solutions that are not dominated by any other solutions in the set.

a large deterioration in the other. In case where the human intervention is limited while the two objectives are equally important; or it is difficult to correctly weight them (which is common for SAS), the knee solutions are more balanced than the others and they are almost the most preferable ones. This is because the knee solutions achieve a good sense of compromise, while moving the solution in any direction from the knee region would create a bias towards an objective, leading to imbalanced adaptation results. Finding the knee solutions is challenging because real-world runtime SAS problems may not pose a perfect convex objective surface as shown in Figure 3.

3.3 Feature Model with Numeric Features

The feature model [16], expressed as the tree structure, is a widely used notation for software engineers to represent the functional variability of a software [6]. In the feature oriented domain analysis, the feature model is particularly important for expressing the possible variations under which a software system can operate in order to improve functional and non-functionary quality [33]. In this perspective, features define the prominent or distinctive aspects between different variations of a software system [33], which range from high-level architectural elements (an entire component) to low-level configurations (a specific parameter).

In the context of SAS, the inherited concept of a feature model allows it to define the extent to which the SAS is able to adapt at runtime, i.e., a range of variations that the SAS can achieve. Given such nature, there has been some successful attempts that apply the feature model to design SAS [23][41]. Therefore, to correctly exploit the feature model for SAS, the software engineer must identify (i) the variations of different features that are supported by the SAS; and (ii) the dependency constraints that determine the validity of a given variation (adaptation solution). However, while the feature model is useful to express the variability of SAS, i.e., the search space of the adaptation decision making problem, it does not correlate the effects of those variations to the concerned quality attributes. Therefore, in this work, we exploit additional system model to evaluate how a variation can affect the quality of SAS, as we will discuss in Section 7.2.

Figure 4 shows an example of a feature model for one of the SASs we study in this paper⁵. As we can see there are four types of in-branch relation between a feature and its parent:

- *Optional* refers to the feature might be deselected, e.g., *Cache*.
- *Mandatory* denotes core features, which cannot be deselected, e.g., *Thread Pool*.
- *XOR* represents the feature in a group such that exactly one group member can be selected, e.g., *Cache Mode*.
- *OR* means a group that at least one group member needs to be selected, e.g., *Cache Size*.

When a feature is selected, it means that such a feature is ‘turned on’; similarly, deselection of a feature refers to it is ‘turned off’. Selecting a feature implies that its parent should be selected too. In this work, we call a feature **deselectable** if it has *Optional*, *OR* or *XOR* relation to its parent; or **conditionally deselectable** if it has *Mandatory* relation to its parent but there exist deselectable ancestors. On the other hand, common cross-branch relations include:

- F_i *require* F_j means the former can only be selected if the latter is selected.
- F_i *exclude* F_j denotes two features are symmetrically mutually exclusive.
- F_i *at-least-one-exist* F_j is an implied relation between the members of an *OR* group. It represents the same notion as that of *OR*.

⁵In this paper, we use graphical figure of the feature model for more intuitive presentation. In practice, the feature model might be expressed in XML or conjunctive normal form, which can be parsed and analyzed directly by FEMOSAA.

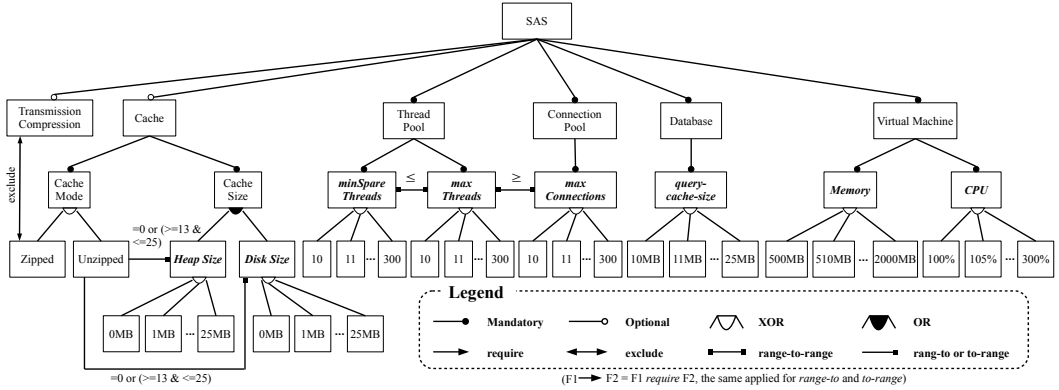


Fig. 4. The feature model for the example SAS shown in Figure 1. (The numeric features are shown in bold and italic letters. A zipped cache mode means the cached data is compressed, thereby costing smaller amount of memory; otherwise, it is unzipped. CPU % denotes to what extent the SAS can consume CPU capacity where each 100% means a CPU core, e.g., 200% means two cores; 150% means one full core and 50% of another)

- F_i *at-least-one-require* F_j is an implied relation between a member of an *OR* group and another external feature, which has *require* to the root of the said *OR* group. It means F_i can only be selected if at least one member of the *OR* group, which F_j belongs to, is selected.

All those relations constitute to the dependency chain(s) in the model. As in Figure 4, the number of features in the above example is 1151 with a search space of more than a billion.

To better incorporate the feature model with SAS and simplify the design, we distinguish **categorical features** and **numeric features**. We define numeric features as: *a feature is numeric if it has more than one child in its XOR group and all its children can be quantified by real numbers*. For example, in Figure 4, the *Memory* is clearly a numeric feature. Otherwise, the feature is categorical, e.g., *Cache Mode*. Similarly, a dependency is numeric as long as it is linked to numeric features and it involves quantitative comparisons. As in Figure 4, we propose the following cross-branch numeric dependencies for the engineers to specify in their design:

- *range-to-range*. This is associated with two numeric features and it can be expressed as, e.g., F_i *range-to-range* F_j ($F_i < F_j$), meaning that F_i 's selected child in its *XOR* needs to be smaller than that of F_j . It can be easily translated into categorical dependency: $F_i < F_j$ simply means F_i 's *XOR* child C would have *exclude* dependency on each F_j 's *XOR* child, which is larger than or equal to the value of C . Other quantitative comparisons (e.g., $>$) can be also applied.
- *to-range*. This constrains a categorical feature F_i (dependent) with respect to a numeric feature F_j (main), e.g., F_i *to-range* F_j ($F_j < 10$), meaning that F_i can only be selected if F_j 's selected child in its *XOR* falls in the given range, as expressed by the mathematical formula. This can be translated to categorical dependency such that F_i would have *exclude* dependency on each of F_j 's *XOR* children that are not in the range.
- *range-to*. This is the inverse of *to-range* dependency where a numeric feature (dependent) is constrained by a categorical feature (main).

Clearly, numeric dependencies can only be cross-branched while categorical ones exist on both in-branch and cross-branch. When a dependency is associated with one categorical feature and one numeric feature (i.e., *to-range* and *range-to*), we call it hybrid dependency which is a special case of numeric dependency. Note that numeric features might have all types of dependencies but categorical features cannot be linked to the *range-to-range* numeric dependency.

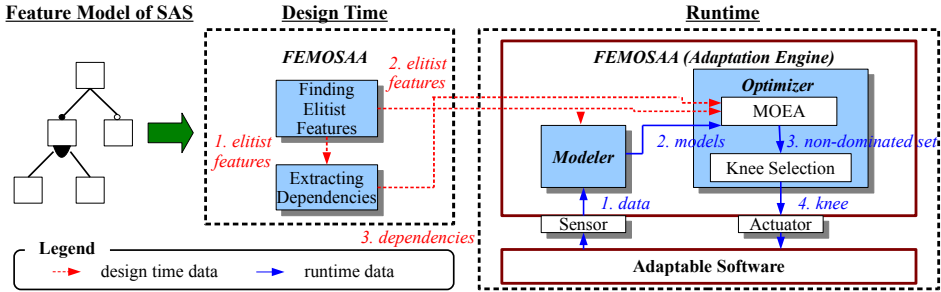


Fig. 5. The architecture of FEMOSAA.

3.3.1 The benefits of explicitly considering numeric features. As mentioned, given that the feature model is discrete and statically defined at design time, it is possible to convert those numeric features and their dependencies into categorical ones without affecting the original variability of SAS. However, explicitly considering numeric features in the feature model will introduce the following benefits in terms of both design time analysis and runtime optimization in FEMOSAA:

- Explicitly considering the numeric features provides simpler and more intuitive design of the feature model as numeric features can be interpreted directly by the software engineers.
- Converting the numeric features into categorical ones will unnecessarily complicate the feature model, which can implicitly induce the software engineers to design the feature model in a way that the children of numeric features would need to be encoded as genes. As mentioned, this will largely increase the number of solution variables in the optimization, leading to the curse of dimensionality. Therefore, explicitly considering numeric features can provide us with the foundation to design novel and simpler encoding of chromosome representation in MOEA, as we will show in Section 5.1.
- Explicitly considering the numeric features results in less number of dependencies in contrast to the case where the numeric features are converted into categorical ones. As we will show in Section 5.2, this simplifies our dependency extraction process for injecting the dependencies into mutation and crossover operators of MOEA. In addition, less number of dependencies implies simpler dependency structure, i.e., a dependent feature has less number of main features, which in turn reduces the running overhead of our dependency aware operators at runtime.

4 FEMOSAA OVERVIEW

As shown in Figure 5, a SAS generally consists of two parts: an adaptable software that is managed at runtime, and an engine that controls the adaptation. The adaptable software could be a software stack that contains different inter-connected software or middleware.

Our FEMOSAA framework is deployed as the adaptation engine and it operates on both design time and runtime of the SAS. At design time, FEMOSAA analyzes and transposes the feature model of SAS, which is provided by the software engineers, to the context of MOEA. The transposition at first identifies the elitist features (see Section 5.1), which are passed to the process for extracting the dependency to accommodate with the selected features (step 1), as we will explain in Section 5.2. With the help of FEMOSAA, those elitist features and dependencies are stored and will be used directly by the MOEA at runtime (step 2 and 3). Given that only the elitist features would be encoded into the chromosome representation of MOEA, the identified elitist features can be used as the objective functions' inputs, and can serve as the indication of which sensors/actuators to use or to implement (step 2).

FEMOSAA has two main components at runtime: a *Modeler* which contains the objectives (fitness) functions that build the correlation between features and quality attributes. Those objectives functions can be created using analytical models [44], simulation [26] or machine learning [10] [11] [14] in which they might be updated on-the-fly using the data from sensors; and an *Optimizer* that realizes the MOEA (extended by our knee selection), which is guided by the transposed information from the feature model, to find a single optimized solution for adaptation via actuators (see Section 6).

Given the uncertain and dynamic environment, these two components constitute the feedback loop that continually adapts the SAS towards better quality, e.g., improve response time. The adaptation cycle starts from monitoring status of the SAS and the environment (step 1), which is then used to update the objective functions and model (step 2). Next, the feature guided MOEA optimizes and searches for a set of non-dominated solutions based on the updated objective functions (step 3), after which the knee selection selects the most balanced one for adaptation (step 4). The optimization can be triggered either by the violations on the quality requirements or, as what we did in this work, by a fixed frequency, e.g., at each point in time. Note that we consider the execution order of a solution as a separate issue from the optimization. Thus, given a valid and optimized solution, we assume that the valid order of execution, with respect to the dependency, is enforced in the actuators through analyzing the dependencies in the feature model.

5 TRANSPOSING FEATURE MODEL OF SELF-ADAPTIVE SOFTWARE TO MOEA AT DESIGN TIME

In this section, we present an automatic and systematic approach as part of FEMOSAA that transposes a feature model into MOEA's context. At design time, the approach finds the elitist features from the model, by which we refer to those that cannot be removed in the optimization without damaging the original variability of SAS while minimizing the length of encoding, to form chromosome representation; and then extracts the feature dependency with respect to the elitist features. Such information will be used at runtime to guide the evolutionary optimization.

To guarantee correctness of the transposition, it is imperative to ensure that the feature model has been fully tested and verified by existing tools [6]. Henceforth, this ensures that faults, e.g., dead features, false options and contradictory relations, have been already dealt with before the transposition. The verification of a feature model is beyond the scope of this work, however. Unlike our work, the dependency related to numeric features is not treated explicitly in existing testing tools. However, as discussed in Section 3.3, the numeric (and hybrid) dependency can be easily transferred into the categorical dependency, which can be then tested directly. We also assume that all possible children (including 0) of numeric features are discretized and predefined. It is worth noting that discretizing the numeric features is the firstly step to remove the unnecessary complexity of our SAS optimization problem, this is because many real-world features are often discrete and/or can be customized based on software engineers' knowledge, e.g., it can be known that changing memory allocation by less than 1MB does not affect the behaviors and quality of SAS, henceforth, instead of considering the memory feature as a continuous feature, the possible child features of the memory feature can be discretized at every 1MB.

While FEMOSAA is generic and can be applied on any cases as long as the feature model and MOEAs are involved, in the following, we specify the transposition approach in FEMOSAA for the general cases but refer to a concrete example for more intuitive illustration where appropriate. Specifically, in Section 5.1, we introduce the approach to identify the elitist chromosome representation of a SAS's feature model. Subsequently, in Section 5.2, we illustrate how the related




Fig. 6. The growing process in SAS's feature model.

dependency chains and the value trees can be extracted (Section 5.2.1) and merged (Section 5.2.2), according to the genes identified in Section 5.1.

5.1 Finding Elitist Features for Chromosome Representation

5.1.1 Growing the Feature Model Tree. Deselectable features in a feature model often do not explicitly indicate the 'on' and 'off' features as children, but they are important information for us to parse and understand the full variability of the model. Hence, to correctly transpose the feature model, we firstly grow the feature model tree for disclosing the hidden information inferred from the deselectable features. As illustrated in Figure 6, this is achieved by adding children representing *On* and/or *Off* to any given feature *F* in the feature model using the following steps in order:

- **G-1.** If *F* is a leaf feature that has *OR* relation to its parent, we then add two children representing *On* and *Off* in a *XOR* group to *F*. This explicitly states that in such case, the leaf *F* can have two mutually exclusive options, which is important to our encoding.
- **G-2.** If *F* is a leaf feature that has neither *OR* nor *XOR* relation to its parent, we then add one child representing *On* in a *XOR* group to *F*. This ensures that every feature has the option of 'on' (and translate them into branches to be parsed by G-3), except those with *OR* nor *XOR* relation to its parent, as the former has been considered in G-1 while the latter's 'on' option can be expressed by the parent.
- **G-3.** If *F* is a branch feature that has *Optional*, *OR* or *XOR* relation to its parent, we then add one child representing *Off* in a *XOR* group to *F* and to the descendants of *F* that are branch features (if they do not currently have child representing *Off*). This ensures that both the deselectable and conditionally deselectable features expose the option of 'off'.

After growing the tree, the added features and the steps that create them are shown in Figure 7.

5.1.2 Identifying Genes from the Feature Model Tree. We have now obtained a model with no hidden information, the next phase is to find the elitist features for genotype encoding in MOEA, creating an **elitist chromosome representation**. Intuitively, following the grown tree, our approach *encodes a feature *F* as gene in the chromosome, if and only if, it is the parent of a XOR group, which contains more than one group member*. Hence, *F*'s children within the *XOR* group constitute its set of alternative optional values to be chosen in MOEA, subject to the constraints in dependencies. Drawing on this, the representation can be simplified in three aspects without affecting the original variability:

- (1) Eliminating features whose variability can be expressed by their parent, i.e., those with *XOR* relations to the parent, e.g., the variability of *CPU*'s children can be represented by itself.

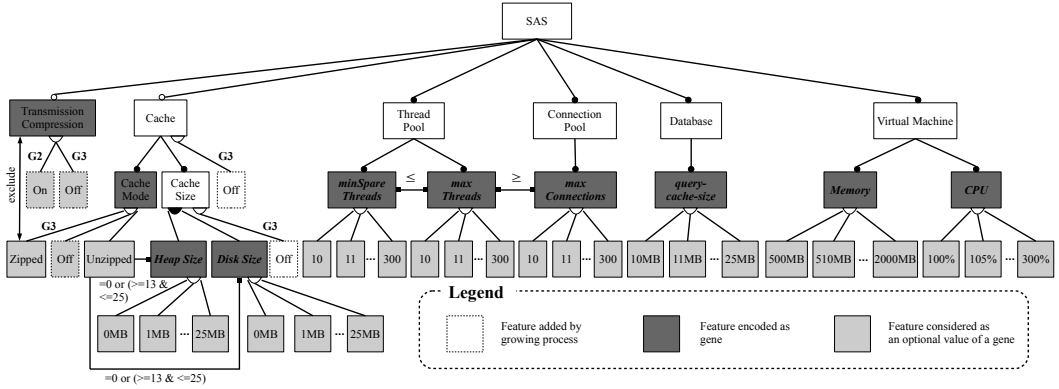


Fig. 7. The example SAS’s feature model after the process of finding elitist features.




Fig. 8. The resulted chromosome representation of the SAS studied in form of genes (G_1, G_2, \dots, G_{10}).

- (2) Eliminating features whose variability can be expressed by their descendants, e.g., the variability of *Cache* feature can be represented by the combination of *Cache Mode* and *Cache Size* features; *Cache Mode* can be represented by *Heap Size* and *Disk Size*.
- (3) Eliminating the features that have no implication on the variability, e.g., the *Thread Pool* is always mandatory. This, however, does not mean that we simply remove all mandatory features (as in [31]); instead, our approach retains those mandatory features that with a *XOR* group of children as they would often help us to considerably reduce the number of genes even more, as explained in (1) above.

From now on, we call those features, which are chosen to be encoded in the chromosome, as **genes**. It is easy to see that numeric features are always chosen as genes. As shown in Figures 7 and 8, there are 10 features being considered as genes in the example feature model of SAS⁶. To make the model informative, we prune the features that are the only member in their corresponding *XOR* group. For all genes, if they select *Off* or 0 as their value then it means they are deselected; any other values mean that they are selected. Note that when a gene selects 0, it implies that the numeric value is 0 and that the feature is ‘turned off’, which will have no further effects on the SAS.

In this way, the elitist chromosome representation is polyadic, elegant and free of unnecessary information (e.g., some unneeded relations to parent of a feature), which is otherwise unavoidable in the classic binary representation. This, as we will show in Section 7.4.1, can bring non-trivial benefits for the optimization quality and runtime overhead.

⁶The other features, which are not genes, can be fixed to *On* in the SAS




Fig. 9. The vertical analysis for extracting in-branch dependencies in SAS's feature model with respect to the elitist genes.

5.2 Extracting Feature Dependency for Guiding Evolution

5.2.1 Analyzing and Refactoring the Dependency. While the identified elitist chromosome representation can naturally prevent violation of the XOR relation, it does not contain any information about the other dependency constraints (*require*, OR , etc). This issue is non-trivial as leaving it without treatment could result in a high possibility of exploring invalid solutions, which negatively affect the quality of adaptation. To this end, our next step in the transposition is to extract and analyze the dependency chain(s) to accommodate with the genes, so that they can be injected into the mutation and crossover operators of MOEA to prevent the search from exploring the invalid solution. Here, a single dependency between two genes represents the constraint on the dependent gene with respect to the conditions of main gene. The extracted dependencies and their imposed constraints are shown in the Table 1, which will be discussed in Section 5.2.2. Specifically, we distinguish two categories of dependency: *in-branch* and *cross-branch*.

Extracting in-branch dependency chain(s) aims to handle the constraints introduced by *Optional*, *Mandatory*, OR and XOR relations with respect to the genes. To achieve this, the features' in-branch dependencies are extracted in both vertical and horizontal directions while considering all the four relations.

Vertical analysis for extracting in-branch dependency helps to ensure that the in-branch relation between feature and parent are captured. As shown in Figure 9, for any feature F in the original feature model, we conduct the following vertical analysis:

- **VA-1.** If F is a gene and it is deselectable (*Optional*, OR or XOR to its parent) or conditionally deselectable (*Mandatory* to its parent but has deselectable ancestors), then for each path from F , the closest descendant gene D_g of F would have *require* dependency on F , as D_g cannot be selected without the presence of F . Additionally, if D_g has *Mandatory* relation to its parent and the path between F and D_g does not contain deselectable features, then F would also have *require* on D_g , as both features need to be selected at the same time.
- **VA-2.** In addition to VA-1, if F has XOR relation to its parent and it is a gene, then F would have a *require* on its parent, denoted as $P_g = \alpha$ (F 's parent P_g would always be a gene as ensured by our gene identification process), where α is the reference of F in P_g ; similarly, $P_g = \alpha$ would also have *require* on F , as both features need to be selected at the same time. On the other hand,




Fig. 10. The horizontal refactoring for extracting in-branch dependencies in SAS's feature model with respect to the elitist genes.

if F has χOR relation to its parent but it is not a gene, then for each path from F , the closest descendant gene D_g of F would have its own *require* on $P_g = \alpha$. Under the same case, if D_g has *Mandatory* relation to its parent and the path between F and D_g does not contain deselectable features, then $P_g = \alpha$ would also need to have *require* on D_g .

- **VA-3.** In addition to **VA-1**, if F has OR relation to its parent, we find the closest deselectable ancestor of F , denoted as A , if such a A does exist. Now, if at least one ancestor of F is gene; or F 's parent is neither deselectable nor conditionally deselectable; or there exist at least one closest descendant gene, D_g , of a path from A , such that D_g has *Mandatory* relation to its parent and there is no deselectable features in the path between D_g and A , then this means that the OR group, for which F 's parent is the root, needs to select at least one group member. Thus, unless there already exist an *at-least-one-exist* dependency, F (if F is a gene) or its closest descendant genes, each of which follows different paths (if F is not a gene), would have *at-least-one-exist* on (i) the other closest descendant genes of F if it is not gene; (ii) those sibling genes of F in the same OR group; and (iii) the closest descendant genes, each of which follows different paths from the F 's siblings that are not genes but in the same OR group as F .

The horizontal refactoring, on the other hand, is to ensure the elimination of some features does not mislead the dependencies implied by the original variability. Suppose that F is a feature in the original feature model and that A is the closest deselectable ancestor of F , if such a A does exist. Now, assuming that A is not a gene, and that there is no gene on the path from A to F , we then conduct the following horizontal refactoring, illustrations can be found in Figure 10:

- **HR-1.** If F has *Optional* relation to its parent then we do nothing, even if it is a gene. This is because the selection of F does not affect A 's closest descendant genes, each of which follows the other paths from A .
- **HR-2.** If F is a gene that has *Mandatory* relation to its parent and it is conditionally deselectable, then for each path from A , the closest descendant gene D_g of A (excluding F itself) would have *require* on F . This can ensure when D_g is selected, F would be also selected.
- **HR-3.** If F has OR relation to its parent and it does not have *at-least-one-exist* dependency for the group, then for each path from A (except the paths that pass through F 's OR group), the closest descendant gene D_g of A would have *at-least-one-require* on F , if F is a gene; or on those closest descendant genes of F , each of which follows different paths from F , if F is not a gene. Hence, when D_g is selected, at least one member of the OR group of F (or their closest descendant genes) would be also selected.
- **HR-4.** If F has χOR relation to its parent then we do nothing, even if it is a gene. This is because our gen identification process ensures that the parent of F would always be a gene, which also express the selection of F .




Fig. 11. The refactoring for extracting cross-branch dependencies in SAS's feature model with respect to the elitist genes.

After considering the in-branch dependency, we now focus on refactoring the cross-branch dependency. If both sides of a cross-branch dependency are genes, then it can be extracted directly. However, if either side (or both) of the feature is not a gene, then a treatment is needed. Suppose that a feature F is associated with one or more cross-branch dependency and that F is not a gene, we then do the following refactoring as shown in Figure 11:

- **CR-1.** If F is a branch, then its cross-branch dependencies are migrated to those closest descendant genes of F , each of which follows different paths from F . Further, if F is the root of a OR group and it is the main feature in any *require* dependency, then those *require* would be changed to *at-least-one-require*, which are migrated to the member genes of F 's OR group, and to the closest descendant genes, each of which follows different paths from those members that are not genes.
- **CR-2.** If F is a leaf, then its cross-branch dependencies are migrated to the parent of F , denoted as P_g , where the dependency would remain the same but the main gene becomes $P_g = \alpha$, where α is the reference of F in P_g . Here, F would always have XOR relation to P_g , because if F was to have *Mandatory* relation to P_g , then there would be a contradiction as the main feature of a cross-branch dependency is mandatory. In addition, when F is a leaf, our growing process has ensured that F has neither OR nor *Optional* relation to P_g which would always be a gene.

Finally, putting everything together, the extraction which occurs on the model and the extracted dependency chain(s), with respect to the elitist chromosome, are shown in Figure 12 and 13 respectively. The constraint of a dependent gene imposed by a dependency, according to Table 1, can be expressed using a *value tree*, where each leaf is a set of optional values constrained by the corresponding condition in a branch, i.e., the selected values of the main gene. For example, in Figure 14, the value tree for the dependency between *Cache Mode* (G_2) gene and *Transmission Compression* (G_1) gene constrains that the former can only be *Off* or *Unzipped* if the latter is *On*; or any optional values, otherwise. Note that if a gene is not a dependent of any dependency, then it would have a value tree without any branches.

5.2.2 Merging the Dependency. After the extraction, we can see that a dependent gene might have multiple dependencies on the same or different main genes. To construct a combined value tree for a dependent gene, the dependencies, by which it is constrained, need to be merged one by one using set operators (union or intersection) to combine the leaves from their value trees. Table 1 shows what set operators are needed for each dependency type when merging with the others, which is derived from the conjuncture normal form of the related genes.

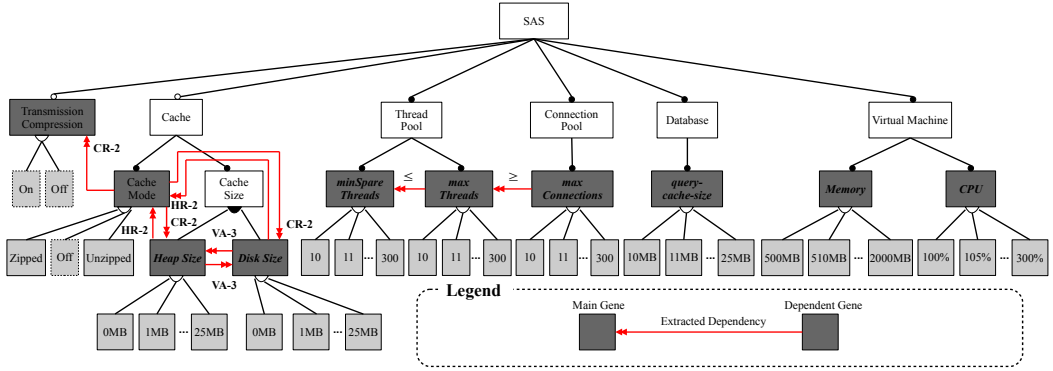


Fig. 12. The example SAS's feature model and the extracted dependency after the process of VA, HR and CR.

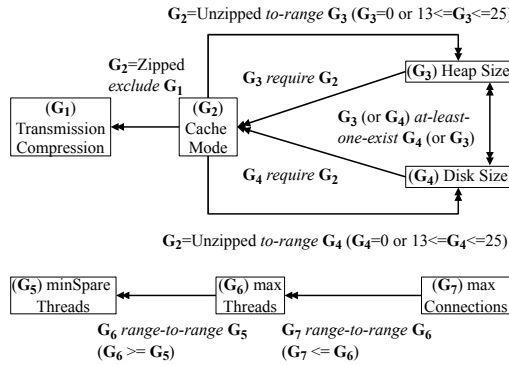


Fig. 13. The extracted dependency chains of genes for the example feature model.

Specifically, for every dependent gene, the merging process has the following steps:

- **Step 1.** If it has two or more dependencies with identical main genes, then the leaves, which are constrained by the same condition in the branches, would be combined directly using the set operations shown in Table 1.
- **Step 2.** If it has dependencies on different main genes, all branch nodes of one single or already combined value tree are replicated and grafted (as a whole) to each right-most branch node of another single or already combined value tree, forming new levels for the newly combined value tree, representing the combinatorial conditions. Then, for the two value trees that were grafted, their leaves, whose original ancestors are now on the same path from root to a right-most node in the newly grafted tree, are combined using the set operations shown in Table 1 to create the new leaf-set.

The process stops when all related dependencies are merged and their value trees are combined, resulting a finally combined value tree. As an example, Figure 14 illustrates the merging process and the finally combined value tree for *Cache Mode* (G_2) gene, which originally contains dependencies on three different main genes. The same merging procedure is repeated for every dependent gene.

To ensure the correctness of merging, in both steps, the dependencies that require union are always merged ahead of the others, leading to a set of partially merged dependencies. Since each

Table 1. The Extracted Dependency Constraints Between two Genes and The Related Set Operations for Merging Dependency.

Dependency (denoted as D)	Constraints on G_i	Merge with Other Dependency D'
G_i require G_j	if $G_j = \text{Off}$ or 0, then $A_s = \{\text{Off}\}$ or $\{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha$ require G_j	if $G_j = \text{Off}$ or 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i require $G_j = \alpha$	if $G_j \neq \alpha$, then $A_s = \{\text{Off}\}$ or $\{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ require $G_j = \alpha_2$	if $G_j \neq \alpha_2$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i exclude G_j	if $G_j \neq \text{Off}$ or 0, then $A_s = \{\text{Off}\}$ or $\{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha$ exclude G_j	if $G_j \neq \text{Off}$ or 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ exclude $G_j = \alpha_2$	if $G_j = \alpha_2$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i at-least-one-require G_j	if $G_j = \text{Off}$ or 0, then $A_s = \{\text{Off}\}$ or $\{0\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) at-least-one-require that related to the same root of OR group as D 's, then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
$G_i = \alpha$ at-least-one-require G_j	if $G_j = \text{Off}$ or 0, then $A_s = A - \{\alpha\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) at-least-one-require that related to the same root of OR group as that of D , then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
G_i at-least-one-exist G_j	if $G_j = \text{Off}$ or 0, then $A_s = A - \{\text{Off}\}$ or $A - \{0\}$. Otherwise, $A_s = A$.	if D' is (or merged from only) at-least-one-exist that related to the same root of OR group as that of D , then $A_s \cup A'_s$. Otherwise, $A_s \cap A'_s$.
G_i range-to-range G_j (e.g., $G_i < G_j$)	if $G_j = \alpha$, then $A_s = \{\alpha_1, \dots, \alpha_n\}$ where $\forall \alpha_n \in A$, and $\forall \alpha_n$ meets the given condition with respect to α , e.g., $\forall \alpha_n < \alpha$, etc.	$A_s \cap A'_s$.
$G_i(R)$ range-to G_j	if $G_j = \text{Off}$ or 0, then $A_s = A - R$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i(R)$ range-to $G_j = \alpha$	if $G_j \neq \alpha$, then $A_s = A - R$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
G_i to-range $G_j(R)$	if $G_j = \alpha$; $\alpha \notin R$, then $A_s = \{\text{Off}\}$ or $\{0\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.
$G_i = \alpha_1$ to-range $G_j(R)$	if $G_j = \alpha_2$; $\alpha_2 \notin R$, then $A_s = A - \{\alpha_1\}$. Otherwise, $A_s = A$.	$A_s \cap A'_s$.

Additionally, when the set operation leads to an empty set, we fix $A_s = \{\text{Off}\}$ or $\{0\}$.

G_i and G_j are dependent and main gene, respectively; A is the entire set of optional values for G_i ; A_s denotes the set of values for G_i , given a selected value of G_j ; α , α_1 , α_2 and α_n denote some selected values for G_i or G_j ; R is a given constrained set of range, e.g., $G_j < 10$, etc; D' is another single or merged dependency for which G_i is the dependent gene; A'_s denotes the set of values for G_i , given the selected value(s) of the main gene(s) in D' .

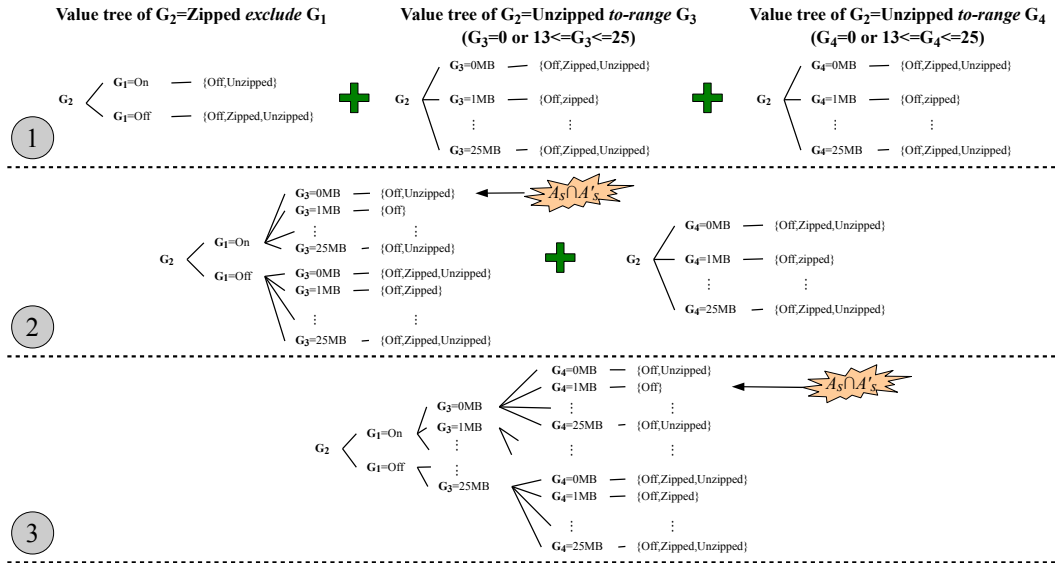


Fig. 14. The example of dependency merging and the combined value tree for the gene G_2 (Cache Mode), which is the dependent gene of three different main genes, i.e., G_1 (Transmission Compression), G_3 (Heap Size) and G_4 (Disk Size). A_s and A'_s denote two sets of values (leaf-set) to be combined, as specified in Table 1.

of them are merged from the identical dependency type, as shown in Table 1, they can then be merged with each others (and the remaining single dependencies) using the same set operators as if they are single dependencies.

It is worth noting that, despite the given feature model is contradictions free, it might still be possible for the combined leaves to have an empty set under some conditions. For example, recall *Disk Size* gene from Figure 13, when its *require* dependency on *Cache Mode* merges with its *at-least-one-exist* on *Heap Size*, then combing the leaves for the condition where *Cache Mode* selects *Off* and *Heap Size* selects *0 MB* would lead to an empty set. This is due to the fact that some dependencies have different priorities: in the above example, *at-least-one-exist* would constrain *Disk Size* only if it is allowed to be selected as indicated by the *require* on *Cache Mode*. In those cases, we fix $\{Off\}$ (or $\{0\}$) as the new leaf-set for the dependent gene's combined value tree, representing the fact that it needs to be 'turned off' and cannot affect the SAS under the corresponding conditions.

In Section 6, we will describe how the elitist chromosome, the extracted dependency chain(s) and the (combined) value trees of genes are seamlessly injected into the MOEA for optimizing SAS at runtime, regardless to the order of changes on genes.

6 FEATURE GUIDED AND KNEE DRIVEN MOEA AT RUNTIME

We have now completed all the design time transposition and analysis in FEMOSAA. Next, to optimize the SAS at runtime using FEMOSAA, we extend MOEA with the domain information gathered from the feature model, creating a feature guided MOEA with knee selection, as shown in Figure 15 and Algorithm 2. In particular, the elitist features from Section 5 form the basic chromosome representation of the solution in MOEA. To avoid exploring invalid solutions, we explicitly inject the extracted dependencies (e.g., Figure 13) and the (combined) value trees of genes (e.g., Figure 14) into the basic mutation and crossover phases of MOEA to create dependency aware

Algorithm 2 The Extended Feature Guided MOEA with Knee Selection

Input: the extracted dependency chain(s) C and the (combined) value trees of the genes VT , in addition to the inputs in Algorithm 1

Output: a single knee solution

```

1: start evolution
2:  $P = \emptyset$ 
3:  $eval = 0$ 
4: for  $i = 1$  to  $P_{size}$  do
5:    $S = \text{GETRANDOMSOLUTION}()$ 
6:    $\text{DODEPENDENCYAWAREMUTATION}(S, 1, C, VT)$   $\triangleright$ mutation rate of 1 means mutating every gene.
7:    $\text{EVALUATEFITNESS}(S)$ 
8:    $eval = eval + 1$ 
9:    $P = P + S$ 
10: end for
11: while  $eval < eval_{max}$  do
12:    $P_0 = \emptyset$ 
13:   while  $|P_0| \leq P_{size}$  do
14:      $parents = \text{DOMATINGSELECTION}(P)$ 
15:      $offspring = \text{DODEPENDENCYAWARECROSSOVER}(parents, r_c, C, VT)$ 
16:     for each solution  $S$  in  $offspring$  do
17:        $\text{DODEPENDENCYAWAREMUTATION}(S, r_m, C, VT)$ 
18:     end for
19:      $\text{EVALUATEFITNESS}(offspring)$ 
20:      $eval := eval + |offspring|$ 
21:      $P_0 = P_0 \cup offspring$ 
22:   end while
23:    $P = P \cup P_0$ 
24:    $\text{DOSURVIVALSELECTION}(P, P_{size})$ 
25: end while
26:  $P = \text{GETNONDOMINATEDSOLUTIONS}(P)$ 
27: return  $\text{GETKNEESOLUTION}(P)$ 
28: end evolution

```

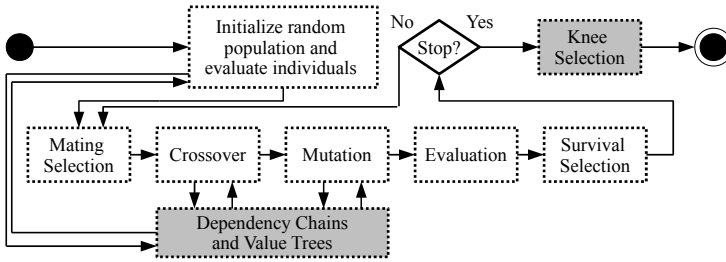


Fig. 15. The workflow of the extended feature guided MOEA with knee selection.

operators (line 6, 15 and 17). Finally, we apply the knee selection proposed in Section 6.4 to identify a single adaptation solution from the set of non-dominated solutions returned by MOEA (line 27).

6.1 Objective Functions

The elitist chromosome representation from Section 5 also helps to define the inputs of the objective (fitness) functions used in the optimization. It is worth noting that FEMOSAA works with a range of quantifiable quality objectives and is agnostic to the actual objective functions, as the framework itself does not rely on any assumptions about the internal structure of those objectives. FEMOSAA is also possible to consider more than two objectives but we use two in this paper to provide more intuitive illustration, since the fundamental principle of multi-objective optimization is the same regardless to the objective number.

The actual objective functions exploited by FEMOSAA can be built using various modeling approaches from the literature, e.g., machine learning based [10][14][11], analytical [44] and simulation based [26], as long as they are compatible with the genes identified by FEMOSAA. In Section 7.2, we will elaborate the concrete objective functions, which are built by using different approaches, for each subject SAS.

6.2 Dependency Aware Mutation Operator

At runtime, to prevent generating invalid offspring when mutating the solutions, the extracted dependency chain(s) and the (combined) value trees of the genes are seamlessly injected into the mutation operation. In this work, we use the boundary mutation operator as the basis, in which each gene might be mutated subject to a mutation rate. Upon mutation of a gene, one of its optional values is randomly assigned. The reason why the boundary mutation operator was chosen is because: (i) it is one of the most commonly used mutation operator; and (ii) it works under discrete optimization problems while allowing to randomly select a value from a predefined value tree, which particularly fits with our SAS optimization problem. However, since the violation of dependency is prevented whenever a gene is changed, it is easy to modify FEMOSAA to work with any other operators that mutate the genes in a similar way as the boundary mutation operator.

As illustrated in Algorithm 3 and Figure 16, our extended dependency aware mutation operator has the following recursive steps:

- **Step 1.** When a gene G (e.g., *Cache Mode*) needs to be mutated as identified by the basic mutation operator or due to its violation, we randomly select a value from its (combined) value tree with respect to the selected values of G 's main genes (e.g., *Transmission Compression*, *Heap Size* and *Disk Size*).
- **Step 2.** Then, we propagate, according to the dependency chain, to G 's dependent genes (e.g., *Heap Size* and *Disk Size*) and we then validate if those genes in the solution violate




Fig. 16. Workflow of dependency aware mutation operator. (The example genes from left to right represent *Transmission Compression*, *Cache Mode*, *Heap Size* and *Disk Size*)

Algorithm 3 Dependency aware mutation operator

Input: given a valid solution S , the extracted dependency chain(s) C and the (combined) value trees of the genes VT

Output: a mutated valid solution (individual) S

```

1: start mutation
2:  $G_{set}$  = the genes that requires mutation as identified by the basic crossover operator based on mutation rate
3: for each  $G$  in  $G_{set}$  do
4:   | MUTATEWITHDEPENDENCY( $G,S$ )
5: end for
6: return  $S$ 
7: end mutation
    
```

Function: MUTATEWITHDEPENDENCY(G,S)

```

1:  $VT_g$  = the (combined) value tree of  $G$  from  $VT$ 
2: randomly choose a new value for  $G$  of  $S$  from  $VT_g$ 
3:  $C_g$  = the chain that contains  $G$  from  $C$ 
4: for each dependent gene of  $G$  ( $G_d$ ) in  $C_g$  do
5:   |  $VT_{g_d}$  = the (combined) value tree of  $G_d$  from  $VT$ 
6:   | validate  $G_d$  based on current conditions in  $VT_{g_d}$ 
7:   | if  $G_d$ 's value in  $S$  is invalid then
8:     | MUTATEWITHDEPENDENCY( $G_d,S$ )
9:   | end if
10: end for
    
```

any dependency using their (combined) value trees. If violation found, we would mutate the corresponding gene and start from **Step 1** for it (e.g., *Heap Size*=5MB in Figure 16), as shown from line 4 to 10 of the MUTATEWITHDEPENDENCY function.

- **Step 3.** The process stops when all the genes, which are identified by the basic mutation operator, have been mutated and there is no further violation found.

Using this operator, we guarantee that the mutation process in MOEA would be better guided, such that only valid adaptation solutions can be explored regardless to the order of mutation. Since the given feature model has been fully validated and exhibits no design errors, there will be at least one state such that all genes can satisfy all dependencies, which in turn, preventing infinite loop in the presence of circular dependencies. As we will show in Section 7.4.2, the dependency aware operators can lead to better quality of adaptations.

6.3 Dependency Aware Crossover Operator

Similar to the mutation process, it is necessary to eliminate invalid offspring when swapping elements of the solutions. To this end, the extracted dependency chain(s) and the (combined) value trees can be injected into the given crossover operator⁷ in the MOEA. In this work, we rely on

⁷In this work, we have used the most common type of crossover operators that takes two parents and produces two offsprings. However, our dependency can be injected with the other types of crossover operators (which are rare) in a similar way.




Fig. 17. Workflow of dependency aware crossover operator. (The example genes from left to right represent *Transmission Compression, Cache Mode, Heap Size and Disk Size*)

Algorithm 4 Dependency aware crossover operator

Input: given two valid parent solutions S_1 and S_2 , the extracted dependency chain(s) C and the (combined) value trees of the genes VT

Output: two new valid solutions (individuals) S_{n1} and S_{n2}

```

1: start crossover
2:  $S_{n1} = S_1$ 
3:  $S_{n2} = S_2$ 
4:  $G_{set} =$  the genes that requires crossover as identified by the basic crossover operator based on crossover rate
5: for each  $G$  in  $G_{set}$  do
6:   | CROSSOVERWITHDEPENDENCY( $G, S_1, S_2, S_{n1}, S_{n2}$ )
7: end for
8: return  $S_{n1}$  and  $S_{n2}$ 
9: end crossover

```

Function: CROSSOVERWITHDEPENDENCY($G, S_1, S_2, S_{n1}, S_{n2}$)

```

1: if  $G$  in  $S_{n1}$  and  $S_{n2}$  has already been swapped then
2:   | return
3: end if
4: swap the values of  $G$  in  $S_{n1}$  and  $S_{n2}$ 
5:  $C_g =$  the chain that contains  $G$  from  $C$ 
6: for each solution  $S$  in  $\{S_{n1}, S_{n2}\}$  do
7:   | for each dependent gene of  $G$  ( $G_d$ ) in  $C_g$  do
8:     | |  $VT_g =$  the (combined) value tree of  $G_d$  from  $VT$ 
9:     | | validate  $G_d$  based on current conditions in  $VT_g$ 
10:    | | if  $G_d$ 's value in  $S$  is invalid then
11:    | |   | CROSSOVERWITHDEPENDENCY( $G_d, S_1, S_2, S_{n1}, S_{n2}$ )
12:    | |   end if
13:    | | end for
14:    | |  $VT_g =$  the (combined) value tree of  $G$  from  $VT$ 
15:    | | validate  $G$  based on current conditions in  $VT_g$ 
16:    | | if  $G$ 's value in  $S$  is invalid then
17:    | |   | for each main gene of  $G$  ( $G_m$ ) in  $C_g$  do
18:    | |   | | CROSSOVERWITHDEPENDENCY( $G_m, S_1, S_2, S_{n1}, S_{n2}$ )
19:    | |   | end for
20:    | |   end if
21:    | | end for

```

the widely used uniform crossover where two genes, each of which from a different parent and both are at the same position in the chromosome, might be swapped subject to a crossover rate. Such a uniform crossover operator was chosen because it mitigates the problem of genes locus, i.e., the ability to explore the search space is less sensitive to the closeness of highly dependent genes (features) in the encoding, which helps to relax extra design requirements of the SAS. However, since the violation of dependency is prevented whenever a pair of genes is swapped, it is easy to modify FEMOSAA to work with any other operators in which each pair of the swapped genes would be always at the same position in the encoding.

After the injection, as shown in Algorithm 4 and Figure 17, our dependency aware crossover operator uses the following recursive steps:

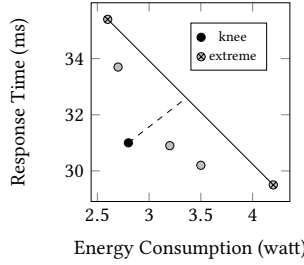


Fig. 18. Finding knee solution(s) in the non-dominated set.

- **Step 1.** When a gene G (e.g., *Disk Size*) needs to be swapped as identified by the basic crossover operator or due to violation, we swap it in the offspring if it has not been swapped already.
- **Step 2.** We propagate, according to the dependency chain, to G 's dependent genes (*Heap Size* and *Cache Mode*). If a dependent gene in an offspring violates the dependency, we then attempt to swap the gene by repeating from **Step 1** for it (e.g., *Heap Size=0MB* in Figure 17), as shown from line 7 to 13 of the CROSSOVERWITHDEPENDENCY function.
- **Step 3.** Next, we check if G in an offspring violates dependencies; if violation exists, we then attempt to swap all the main genes of G (e.g., *Cache Mode* and *Heap Size*, but this is not needed in the example shown) by repeating from **Step 1** for each of them, as shown from line 14 to 20 of the CROSSOVERWITHDEPENDENCY function.
- **Step 4.** The process terminates when there is no dependency violation or all genes in the parents have been swapped.

In this way, we guarantee that only valid adaptation solutions would be produced, given any order of the crossover. As we will show in Section 7.4.2, the dependency aware operators can lead to better quality of adaptations.

6.4 Knee Selection

As mentioned in Section 3.2, knee points are those solutions that achieve well-balanced trade-off on all the objectives. This is particularly appealing as FEMOSAA aims for the cases where the relative importance between conflicting objectives of SAS is unknown and it is too difficult to quantify them. As we can see in Figure 18, which shows the solutions that minimize both response time and energy consumption, those more balanced knee solutions are likely to be around the visible bulge, representing a good sense of compromise. Intuitively, the knee solutions tend to be the furthest away from all the solutions that have the worst result on each single objective, i.e., the extreme solutions. As a result, finding the knee solutions within a non-dominated set is equivalent to search for the solution(s) that has the largest general distance to the extreme solutions in the set.

To achieve this, we developed a knee selection method for selecting single solution from the set returned by MOEA. As shown in Figure 18, given the final non-dominated set obtained by the feature guided MOEA, we at first construct a line ℓ that connects the extreme solutions holding the worst value at each single objective. Then, we calculate the perpendicular distance from each non-dominated solution \mathbf{x} to the ℓ :

$$d(\mathbf{x}, \ell) = \begin{cases} \frac{|\epsilon|}{\sqrt{a^2+b^2}}, & \text{if } \epsilon < 0 \\ -\frac{|\epsilon|}{\sqrt{a^2+b^2}}, & \text{otherwise} \end{cases} \quad (1)$$

where $\epsilon = a \times f(\mathbf{x}) + b \times g(\mathbf{x}) + c$ and the parameters a , b and c can be identified through the extreme solutions. In particular, $\epsilon < 0$ means that \mathbf{x} is on the left side of ℓ ; otherwise, \mathbf{x} is on

the right side of ℓ . Clearly, solutions on the left side of ℓ are preferable than those on the right when considering a minimization problem. The solution(s), which has the largest perpendicular distance to ℓ , is the knee solution(s) that we are seeking. When there are multiple knee solutions, we randomly select one for adaptation. As we will show in Section 7.4.3, the knee selection can lead to more balanced trade-offs on the quality of adaptations. Note that, by replacing ℓ with a surface or a hyperplane, we can easily extend our knee selection for more than two objectives.

6.5 The Concrete MOEAs

Without loss of generality, FEMOSAA can work easily with a wide range of MOEAs. In this work, we run FEMOSAA with three distinct MOEAs (i.e., MOEA/D-STM, NSGA-II and IBEA), each of which is a widely-used representative of its own category as explained in Section 3.1. In the following, we briefly explain their principles and ideas, while the details are beyond the scope of this paper.

- NSGA-II [20]—As one of the most popular MOEAs, NSGA-II at first uses the non-dominated sorting to divide the population into several non-domination levels. Solutions in the first several levels have a higher priority to survive to the next iteration. If the size of the current non-dominated solution set exceeds the pre-defined threshold, NSGA-II uses the crowding distance, a density estimation technique, to trim the population.
- IBEA [53]—The basic idea of IBEA is to firstly define the optimization goal in terms of a binary performance measure/indicator, which is then used to guide the survival selection process. In this way, IBEA transfers a multi-objective optimization problem into a new single-objective optimization problem, with respect to the chosen indicator, to facilitate the fitness assignment procedure.
- MOEA/D-STM [36]—MOEA/D is a MOEA framework that combines the mathematical rigour of the classic multi-objective optimization method and the implicit parallelism of evolutionary algorithms in a single paradigm. Different from the classic multi-objective optimization method [38], which can only obtain a single Pareto-optimal solution at a time by aggregating all objectives into a single-objective aggregation function, MOEA/D decomposes the original multi-objective optimization problem into a population of single-objective optimization subproblems. In particular, each subproblem corresponds to a pre-defined weight vector, generated in a systematic manner [35]. Afterwards, MOEA/D uses a population-based technique to solve these subproblems in a collaborative manner. As a recent variant of MOEA/D, MOEA/D-STM achieves a balance between convergence and diversity by modifying the survival selection mechanism of the original MOEA/D. In a nutshell, MOEA/D-STM treats subproblems and solutions as two sets of agents. Each agent has its preferences over the agents on the other side. In particular, subproblems concern convergence while solutions concern diversity. The survival selection process is modelled as a matching process between subproblems and solutions. The stable matching between them finally turns out to be the selection result. Note that the stable matching achieves an equilibrium between the preferences of subproblems and solutions, thus the selection strikes the balance between convergence and diversity simultaneously.

7 EXPERIMENTS AND EVALUATION

By using the actual running SAS that consists of a stack of real-world software, we have conducted comprehensive experiments to evaluate the effectiveness of FEMOSAA by means of comparing it with its variants and the state-of-the-art frameworks under different metrics. The source code of FEMOSAA, the comparative variants and state-of-the-art frameworks, the subject SAS and all the

experiment data can be publicly accessed via GitHub⁸ ⁹. Specifically, our evaluation aims to answer the following research questions:

- **RQ1.** What are the added values of the elitist chromosome representation in contrast to the conventional binary representation?
- **RQ2.** What benefits do the dependency aware operators provide compared against the classical and widely-used operators
- **RQ3.** What benefits does the knee selection mechanism provide in contrast to selecting arbitrary non-dominated solutions?
- **RQ4.** What is the effectiveness of FEMOSAA in contrast to other state-of-the-art search-based frameworks?
- **RQ5.** What is the overhead of FEMOSAA, in terms of execution time, compared to other state-of-the-art frameworks? Which part(s) cause the most overhead? Is the overhead suitable for SAS runtime?

7.1 Verifiability and Methodology

We have conducted experiments on a dedicated server, which runs Ubuntu Linux 14.04 on an Intel i5 2.8GHz Quad Core processor, 4GB RAM. To separate the adaptation engine and the adaptable software, we used Xen v3.0.3 [48] as the hypervisor to create a virtualized environment. We have implemented FEMOSAA as the adaptation engine using Java, JDK 1.6, and it is deployed on the *Dom0* of Xen. For setting up all the MOEAs, we use a population size of 100 for 10 generations as the termination criteria; the mutation and crossover rate are 0.1 and 0.9 respectively. In particular, for IBEA, we use the ϵ -indicator and an archive size of 500; while for MOEA/D-STM, we apply Tchebycheff aggregation for creating subproblems where the number of evenly distributed weight vectors is 100 (i.e., 100 subproblems) and the size of each subproblem's neighborhood is 20. Those settings are either common values or have been tailored for runtime optimization in our cases with respect to quality and overhead. All MOEAs used in the experiments are extended from the jMetal Framework [21]. To mitigate interference caused by the adaptation engine, we used one vCPU and 800MB RAM on *Dom0*.


7.2 The Subject SAS

We deploy two running SAS for our real-time experiments. The two diverse subject SAS aims to examine the generality and applicability of FEMOSAA under different domains. On the micro level, they help to demonstrate how FEMOSAA can be applied to different feature models, dependency structures, environmental factors, dimensions of quality objectives and degrees of objective conflicts:

- **RUBiS-SAS.** This adaptable software is a software stack that contains RUBiS [43], which is a well-known software benchmark that simulates the *eBay* model, and a set of real-world software including Tomcat v6.0.28 [2], MySQL v3.23.58 [40] and Linux kernel v2.6.16.29 running on a configurable guest virtual machine. Ehcache v2.6 [3] is plugged to RUBiS as the cache management module. All the features that can be adapted at runtime are represented by the feature model in Figure 4 with a total of 1151 features and a search space of around 1.3×10^{16} using the elitist chromosome representation (including both valid and invalid solutions). We have also use two distinct workload patterns, a read-write patten and a read-only one, to create diverse runtime behaviors of the SAS. To simulate a realistic time-varying environmental conditions within the capacity of our testbed, we vary the number of clients according to the compressed FIFA98 workload trace [5], as shown in Figure 19a. This setup can generate up to 600

⁸<https://github.com/taochen/ssase>

⁹<https://github.com/JerryI00/Software-Adaptive-System>



(a) The compressed FIFA98 workload for all patterns in *RUBiS-SAS* (from June 7 to July 1, 1998) (b) The mean relative standard deviation of all concrete services' throughput and cost in *SOA-SAS*

Fig. 19. The changing environment of subject SAS.

parallel requests, which offers sufficient dynamic and uncertainty. The workload is generated by another machine using the client emulator provided in *RUBiS*. In general, a heavier workload implies more pressure on the SAS, which reduces the number of effected solutions and thus make the problem harder. The goal is to continually optimize the following two conflicting quality attributes that need to be minimized:

- (1) **Response Time:** In the *RUBiS-SAS*, response time is measured as the elapsed time between a request's arrival and its response. To express the quality of *RUBiS-SAS* by the end of a time point, we calculated the average response time for all monitored requests served in the past time interval.
- (2) **Energy Consumption:** The energy consumed by software systems, which is accounted for 2% of the global carbon emissions in 2007 [39], is increasingly becoming an important quality concern that has clear conflicts with response time. To measure energy consumption, we leveraged PowerAPI [7], a tool that measures the actual energy (watt) incurred by the software's CPU and memory utilization through probing into the sensors of the hardware infrastructure. For each time interval, we computed the average of 30 measurement results of PowerAPI as the energy consumed by the *RUBiS-SAS* for that interval.

These two quality attributes serve as the case of moderate degree of conflict. In other words, it is easier to find solutions that have better response time and energy consumption when the workload is lower. However, as the workload increases, the degree of conflict between them tends to amplify. Formally, the objective functions for *RUBiS-SAS*'s quality attributes are defined as:

$$RT(t+1) = f(G_1(t+1), G_2(t+1), \dots, G_n(t+1), \delta(t)) \quad (2)$$

$$EC(t+1) = g(G_1(t+1), G_2(t+1), \dots, G_n(t+1), \delta(t)) \quad (3)$$

whereby $G_1(t+1), G_2(t+1), \dots, G_n(t+1)$ denote the genes and their selected values in a possible adaptation solution; $\delta(t)$ denotes a set of most recent environmental factors, e.g., workload in this paper. Given an adaptation solution, $RT(t+1)$ and $EC(t+1)$ are the expected fitness values for response time and energy consumption respectively.

In *RUBiS-SAS*, we adopt the machine learning based model from [10][14][11] to build the objective functions for *RUBiS-SAS*, since it relies on few assumptions of the application domain. By using the actual data of quality performance, software status and environment monitored from the SAS, this model can be continually updated at runtime to provide sufficient accuracy [11]. To stabilize the objective functions in this work, we pre-trained these models by monitoring the data from SAS under random workload for 120 intervals, after which the models are gradually and dynamically updated at runtime. This step is similar to the testing phase before the actual production deployment happened in the industry. As we will show for the next subject SAS,

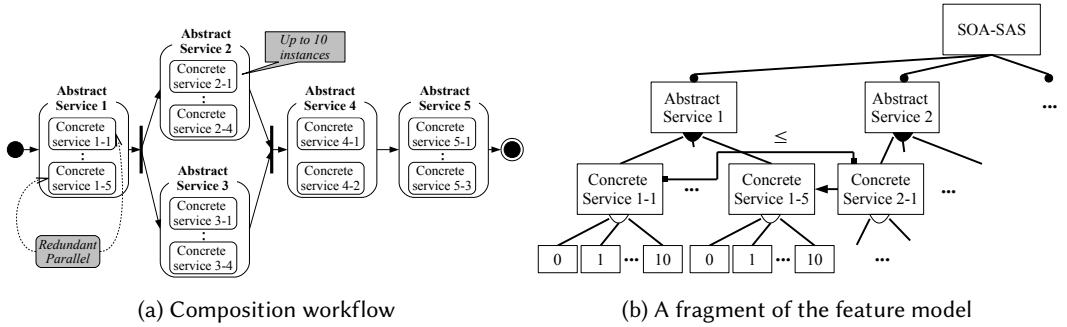


Fig. 20. The architecture of SOA-SAS.

other modeling approaches, e.g., analytical [44] or simulation based [26], can be easily applied as long as they are compatible with the genes identified by FEMOSAA.

To emulate the behaviors of running software system, we run the SAS, empowered with FEMOSAA, under the entire FIFA98 workload trace where the sampling interval is 120s for a total of 102 timesteps, leading to around 5 hours per experiment run including the emulated end-users’ thinking time. For all experiments, we trigger an optimization run by the end of each interval.

- **SOA-SAS.** This is a Service Oriented Architecture (SOA) based adaptable software derived from [49]. At the highest level, it is composed by 5 abstract services connected in parallel or in sequence, as shown in Figure 20. Each abstract service can select up to 2 or 5 concrete services, which are redundant parallel and associated with different quality values. Each concrete service could have up to 10 concurrent instances with the same quality of throughput and cost. Those abstract services, concrete services and their replicas are all considered as features of SOA-SAS (e.g., a fragment in Figure 20), creating a total of 221 features in the feature model with a search space of around 5.6×10^{18} using the elitist chromosome representation (including both valid and invalid solutions). We have also placed various categorical and numeric dependencies on the feature model. To simulate dynamic and uncertainty under time-varying environmental conditions, we amended the throughput and cost of certain concrete services at each timestep, i.e., changing their diversity level according to Gaussian distribution, as shown in Figure 19b. Generally, more diverse concrete services implies more sparse adaptation solutions, which reduces the number of effected solutions and thus make the problem harder. Here, we consider the following conflicting objectives to be maximized and minimized for the entire composition:
 - (1) **Throughput:** In the SOA-SAS, throughput of each concrete service is its maximum capacity under normal operation, expressed as number of requests per second. The throughput of the entire composition is calculated via an aggregate analytical model as shown in Eq. 4.
 - (2) **Cost:** Each concrete service, when utilized, comes with a monetary cost. The cost of the entire composition is again calculated via an aggregate analytical model as shown in Eq. 5.

As suggested in [49], the basic throughput and cost for each concrete service in SOA-SAS were specified following two different Gaussian distributions such that a concrete service with higher throughput would have higher cost too. In SOA-SAS, we leverage analytical model to define the objective functions:

$$T(t + 1) = \min\{a \in A : \sum_{i=1}^a G_i(t + 1) \times T_i(t)\} \tag{4}$$

$$C(t+1) = \sum_{i=1}^n G_i(t+1) \times C_i(t) \quad (5)$$

where $G_i(t+1)$ is the i th gene (a concrete service) and its selected value (the actual number of instances for that concrete service) in a possible adaptation solution for the total of n genes. $T_i(t)$ and $C_i(t)$ are the related throughput and cost value, respectively. a refers to the genes that is associated with a given abstract service, denoted by A . Since the analytical model is customizable, we have forced the objective functions to produce a worst possible fitness value for any invalid solutions, creating a stronger pressure for them to be eliminated in the classic, non-dependency aware approaches.

Similar to *RUBiS-SAS*, there is a total of 102 timesteps and we trigger an optimization run by the end of each interval.

As discussed, FEMOSAA finds elitist features and extracts their dependencies at design time, after which the outcomes are passed to the MOEA and knee selection for runtime optimization.

7.3 The Metrics

We consider the following metrics to evaluate various aspects of FEMOSAA:

- **Individual Quality Attribute:** In addition to the detailed plots of the results, we also report on the Geometric Mean (GM) for all the observed objective values over 102 intervals, as GM tends to be more resilient to the outliers than the arithmetic mean while still can be influenced by those outliers. This is important because outliers on the achieved quality are common in real-world scenarios and thus they cannot be eliminated, but we need to prevent them from strongly dominating the overall result.
- **Aggregate Quality of SAS:** To compare the overall quality for both objectives and provide an overall assessment, we apply a modified Hypervolume (HV) [55][37] for two objectives and Euclidean Distance (ED) [50] on GMs to assess the balance in trade-off and the extent to which both objectives are optimized, respectively. The GMs are normalized before computing those metrics so that they range from 0 to 1. Hence, the modified HV is computed as:

$$HV = \prod_{i=1}^n \left(1 - \frac{GM_i - GM_{i,b}}{GM_{i,w} - GM_{i,b}}\right) \quad (6)$$

and ED can be calculated by:

$$ED = \frac{1}{n} \times \sqrt{\sum_{i=1}^n \left(\frac{GM_i - GM_{i,b}}{GM_{i,w} - GM_{i,b}} - 0\right)^2} \quad (7)$$

where n is the number of objectives and GM_i is the GM for the i th objective; $GM_{i,b}$ and $GM_{i,w}$ are the best and worst GMs that we observed for the i th objective, respectively. Notably, we have converted the maximizing objective (e.g., throughput) into a minimizing one by inverting the results.

- **Percentage of Valid Solutions Found:** For all the intervals, we also compare the average percentage of valid adaptation solutions found in the final population.
- **Running Overhead:** We report on the mean running overhead, in terms of the execution time, over all the intervals in the experiment runs.

To confirm statistical significance of the comparisons on the quality attributes, we performed *Wilcoxon Signed-Rank* test (two-tailed) for all comparisons between FEMOSAA (or FEMOSAA-N) and the others, as our data does not follow Gaussian distributions. We use 95% as the confidence interval ($\alpha = 0.05$), which means that, if the test produces a p that is smaller than 0.05, then we can

reject the null hypothesis H_0 , which states that the given two approaches cannot be statistically distinguished when optimizing a quality attribute of the SAS. The effect size for each test is also reported and we follow the categories in [32] to measure the meaningfulness of effect size.

7.4 Effectiveness of FEMOSAA

To ensure generality of the evaluation, we run each of the three MOEAs for optimizing the SAS under each case which results in a total of 9 scenarios. Further, to evaluate the effectiveness of FEMOSAA in each scenario, we compare FEMOSAA with numbers of its variants:

FEMOSAA-K—This is similar to our FEMOSAA except that it relies on the classic, non-dependency aware operators. Therefore, if the final population contains invalid solutions, it automatically filters them and work on the valid ones only. When no valid solutions found in the final non-dominated set, it corrects the invalid solutions via dependency aware mutation operator. FEMOSAA-K aims to examine whether the specifically tailored dependency aware operators can create benefit over the classic operators that are widely-used in SBSE.

FEMOSAA-D—Another variant of FEMOSAA without knee selection. Hence, in the final non-dominated set, one solution is randomly selected for adaptation. FEMOSAA-D aims to examine the importance of considering knee in runtime SAS optimization. Notably, the use of randomized baseline is strongly recommended by existing SBSE community [4].

FEMOSAA-N—This variant neither considers dependency nor knee selection in the evolutionary optimization. Hence, it uses the same ad hoc strategies from FEMOSAA-D and FEMOSAA-K. FEMOSAA-N is designed to evaluate the combinatorial benefit of dependency aware operators and the knee selection.

FEMOSAA-0/1—A baseline variant, built from FEMOSAA-N, that exploits the conventional binary chromosome representation using all the features from the feature model. Thus, it works under a search space of 2^{1151} and 2^{221} for the two subject SAS, as opposed to the search space of 1.3×10^{16} and 5.6×10^{18} when using the elitist chromosome representations. FEMOSAA-0/1 neither considers dependency nor knee in the evolutionary optimization. Note that if more than one members in a XOR group are selected, we then randomly choose one among them to create a computable solution for the fitness functions. FEMOSAA-0/1 aims to evaluate whether the elitist chromosome representation outperforms the commonly used binary one.

We report the results for all the cases from Figures 21 to 29, we also plot the achieved quality for all timesteps with respect to the environmental conditions in Figures 30 and 31, as well as the results of an example optimization run in Figure 32. Note that, in Figures 30 and 31, the area near the bottom-left line of the cube is the ideal area; the closer the points converge to that line, the better the overall result is.

7.4.1 Evaluating the elitist chromosome representation. To evaluate the effectiveness of our elitist chromosome representation in contrast to the conventional binary representation, we firstly compare FEMOSAA with FEMOSAA-0/1 for all cases on *RUBiS-SAS*, as shown in Figures 21 to 26. Clearly, we see that FEMOSAA largely outperform FEMOSAA-0/1 on both response time and energy consumption with $p < 0.05$ and non-trivial effect sizes. Further, FEMOSAA yields better HV and ED values for all cases. In Figure 30, we also note that, in contrast to FEMOSAA-0/1 under all cases, the quality achieved by FEMOSAA tends to be much more convergent to the left-bottom line of the cube even under heavy workload, meaning that it leads to better quality results and more balanced trade-off. For *SOA-SAS*, we can observe similar results on throughput and cost, as shown in Figures 27 to 29 and Figure 31.

However, the above comparison only demonstrates the combinatorial effectiveness of elitist chromosome representation, the dependency aware operators and the knee selection, while it is not

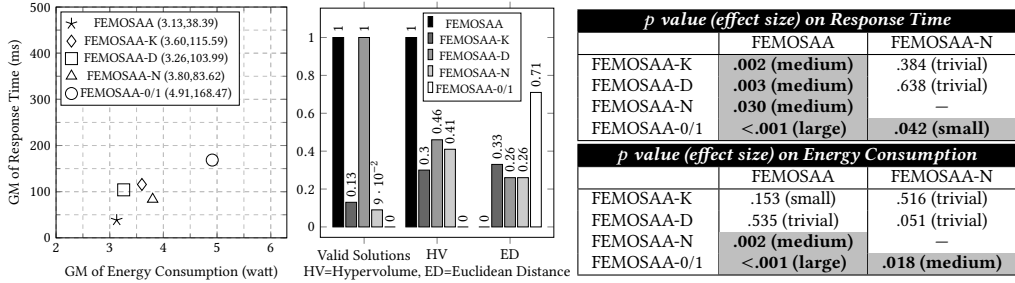


Fig. 21. Results with MOEA/D-STM under read-write pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)




Fig. 22. Results with NSGA-II under read-write pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)




Fig. 23. Results with IBEA under read-write pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)

clear whether the superiority over the conventional binary representation is truly resulted from the elitist chromosome representation. To answer this, we then compare FEMOSAA-N with FEMOSAA-0/1 for all cases on RUBiS-SAS, as shown in Figures 21 to 26. We see that, again, FEMOSAA-N yields much better results on both quality attributes than FEMOSAA-0/1 under all cases while such improvements are statistically significant on at least one attribute with non-trivial effect sizes. It also achieves better HV and ED results. Further, although it is less than 15%, FEMOSAA-N does produce more valid solutions than FEMOSAA-0/1 which cannot identify any valid solution and thus affect the quality of optimization. Next, we take a more detailed comparison between FEMOSAA-N and FEMOSAA-0/1 through Figure 30, which reveals that, for all cases, the results of FEMOSAA-N are more closed to left-bottom line of the cube when the workload is low, e.g., less than around 500 requests/min. For SOA-SAS, we observe even better results: FEMOSAA-N largely outperforms FEMOSAA-0/1 on all the metrics, and the comparisons on quality attributes exhibit $p < 0.05$ and large effect sizes for all cases, as shown in Figures 27 to 29 and Figure 31. Those improvements of




Fig. 24. Results with MOEA/D-STM under read-only pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)




Fig. 25. Results with NSGA-II HV under read-only pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)




Fig. 26. Results with IBEA under read-only pattern on RUBiS-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)

elitist chromosome representation over the conventional binary representation are mainly due to the fact that it fundamentally reduces the search space without affecting the original variability of SAS, so that the search of MOEA has a larger chance to identified the ideal solutions, leading to more valid solutions and better quality. In addition, as we will show in Section 7.6, the elitist chromosome representation can significantly shorten the running time of MOEAs.

Notably, as shown in Figure 30 for RUBiS-SAS, when the workload increases (i.e., more than 500 requests/min), we can see that the achieved quality between FEMOSAA-N and FEMOSAA-0/1 are barely different. This is because the number of good and effected solutions tends to be limited when the workload is heavy, causing the benefits of reducing search space introduced by the elitist chromosome representation less obvious. Those results indicate that, overall, the elitist chromosome representation can better guide the search towards ideal solutions but such improvement tends to blur as the workload increases. Similarly, for SOA-SAS in Figure 31, we see

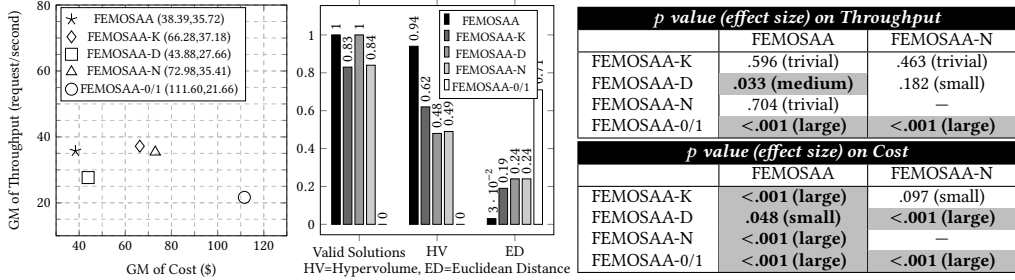


Fig. 27. Results with MOEA/D-STM on SOA-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)




Fig. 28. Results with NSGA-II on SOA-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)





Fig. 29. Results with IBEA on SOA-SAS. (GM denotes Geometric Mean. The significant statistics of comparisons, i.e., $p < 0.05$, are highlighted and shown in bold.)

that although FEMOSAA-N creates better results than FEMOSAA-0/1 across different levels of diversity on the concrete services, the improvement tends to degrade under high diversity.


In summary, the results conclude that:

Answering RQ1—In contrast to the binary encoding, the elitist chromosome representation helps to produce better optimized quality for SAS with statistically significant results and non-trivial effect sizes on at least one quality attribute. It can also shorten the running time of MOEAs (in Section 7.6). However, the improvements might be hardly observed when the number of effected solutions is limited, e.g., under heavy workload.


7.4.2 Evaluating the dependency aware operators. To evaluate the benefit of dependency aware operators, we firstly compare FEMOSAA with FEMOSAA-N for RUBiS-SAS as shown in Figures 21 to 26. It is easy to see that for all cases, FEMOSAA yields better results on both quality attributes




(a) MOEA/D-STM and read-write pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]




(b) NSGA-II and read-write pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]




(c) IBEA and read-write pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]



(d) MOEA/D-STM and read-only pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]




(e) NSGA-II and read-only pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]




(f) IBEA and read-only pattern [the axis from left to right are Response Time (ms), Energy Consumption (watt) and Workload (requests/min)]


Fig. 30. Measured quality results with respect to workload for all the timesteps for RUBiS-SAS.



(a) MOEA/D-STM [the axis from left to right are **inverted** Throughput (request/second), Cost (\$) and mean relative standard deviation of the candidate concrete services' quality]



(b) NSGA-II [the axis from left to right are **inverted** Throughput (request/second), Cost (\$) and mean relative standard deviation of the candidate concrete services' quality]



(c) IBEA [the axis from left to right are **inverted** Throughput (request/second), Cost (\$) and mean relative standard deviation of the candidate concrete services' quality]

Fig. 31. Measured quality results with respect to changes for all the timesteps for SOA-SAS.

with $p < 0.05$ and non-trivial effect sizes on at least one attribute. FEMOSAA also outperforms FEMOSAA-N in terms of HV, ED and the number of valid solutions. In Figure 30, we can see that the achieved quality of FEMOSAA better converges to the ideal area of the cube on different degrees of workload. Similar observations can be obtained for SOA-SAS as illustrated in Figures 27 to 29 and Figure 31.

However, the comparison above does not indicate whether the improvement is mainly introduced by the dependency aware operators or the knee selection method, thus, we then further compare FEMOSAA with FEMOSAA-K, which omitted the dependency aware operators. In Figures 21 to 26 for RUBiS-SAS, we see that FEMOSAA still exhibits superior quality results under all cases which are statistically significant on at least one attribute with non-trivial effect sizes; it also comes with better HV and ED values while the number of valid solutions are 100% versus less than 20%. In all cases, the achieved results of FEMOSAA, as shown in Figure 30, are more convergent around the ideal area of the cube, especially when the workload is heavy, i.e., more than around 500 requests/min. Similar results can be seen for SOA-SAS as shown in Figures 27 to 29 and Figure 31. But, FEMOSAA-K is better than FEMOSAA on throughput when using MOEA/D-STM, the comparisons are not statistically significant though. This means that the degree of objective conflicts in SOA-SAS is much stronger than the case of RUBiS-SAS.

To gain a better understand about why our dependency aware operators can improve the quality of optimization, in Figure 32, we plot the objective space of all searched, non-dominated and selected valid solution(s) in the final population of an example run. We can clearly see that, under all studied MOEAs, the dependency aware operators (i.e., FEMOSAA) can help to find solutions with better convergence and diversity, when compared with the case where dependencies are ignored (i.e.,

FEMOSAA-N). Such a benefit eventually leads to a better set of non-dominated solutions within which a final one can be selected for adaptation. Fundamentally, this is because the dependency aware operators can better guide the search to prevent MOEAs from exploring the unneeded invalid solutions, which virtually reduces the search space and provides a larger chance to find better solutions.

Interestingly, when we compare FEMOSAA-D with FEMOSAA-N in Figures 21 to 26 for RUBiS-SAS, we did not observe significant differences between them in terms of all the metrics and their comparison has also failed in all statistical significance tests, which differs from our expectation that FEMOSAA-D should achieve statistically better results. The same can be also registered in Figure 30. We believe the reason for this is due to although FEMOSAA-D can guide the search for better solutions, the fact that a solution from the final non-dominated solutions is randomly selected for adaptation has obscured the benefits of preventing the exploration of invalid solutions, as such a solution might be highly imbalanced for the conflicting objectives. On contrary, for SOA-SAS in Figures 27 to 29 and Figure 31, we see that FEMOSAA-D generally outperforms FEMOSAA-N with better HV and ED values. In particular, under NSGA-II and IBEA, FEMOSAA-D is better on both attributes with $p < 0.05$ and non-trivial effect sizes. This implies that for a feature model with relatively more complex dependencies as in the case of SOA-SAS, the dependency aware operators can guide the search process to evolve many highly effected solutions in the final non-dominated set, thus the randomly selected adaptation solution, although might be imbalanced, could still be much better than those of the case when dependencies are omitted.

In conclusion, the results suggest that:

Answering RQ2—In contrast to the classic and widely-used operators, the benefit of dependency aware operators is that they help to find solution with better convergence and diversity, leading to better optimized quality for SAS with statistically significant results and non-trivial effect sizes on at least one quality attribute. Such improvement can be more obvious when the number of effected solutions is limited, e.g., when the workload is heavy. However, applying dependency aware operators without ensuring balance of the selected adaptation solution might obscure its effectiveness.

7.4.3 Evaluating the knee selection. We have already shown that the combination of dependency aware operators and knee selection can lead to results that outperform the case where both of them are omitted, we now evaluate whether the knee selection method itself can introduce benefit for runtime optimization of SAS. From Figures 21 to 26 for RUBiS-SAS, when comparing FEMOSAA with FEMOSAA-D in which knee selecting has been omitted, we see that FEMOSAA achieves superior HV, ED and better results on both quality attributes. The comparisons are statistically significant with non-trivial effect sizes on at least one attribute for 5 out of 6 cases. In Figure 30, we can see that the results of FEMOSAA are more balanced, i.e., points are more converge to the bottom-left, and such improvement is more obvious when the workload is heavy under all cases. For SOA-SAS, we observed similar outcomes overall as illustrated in Figures 27 to 29 and Figure 31, except that FEMOSAA-D has better throughput than FEMOSAA under NSGA-II due to the strong conflicts of the objectives in SOA-SAS. These results show that our knee selection method in FEMOSAA contributes to the overall quality of adaptation by automatically selecting the solution, which has balanced trade-off on the objectives (i.e., a good sense of compromise), to execute the adaptation. The resulted quality achieved by those knee solutions are naturally the most preferable ones, when the relative preferences between objectives are unknown or it is too difficult to quantify them (which is common for SAS).

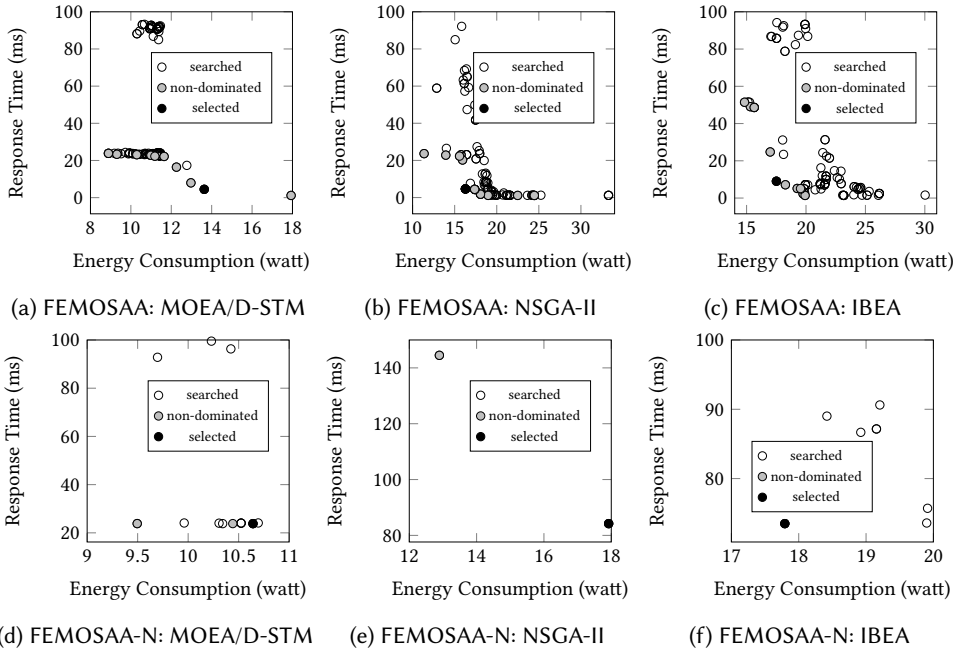


Fig. 32. Objective space of searched, non-dominated and selected valid solution(s) on a timestep of *RUBiS-SAS*.

For all the studied MOEAs in Figure 32 of *RUBiS-SAS*, we can see that, when comparing with randomly selecting an adaptation solution from the non-dominated set (i.e., FEMOSAA-N), incorporating knee selection can indeed select more balanced solution for adaptation (i.e., FEMOSAA) as it is closer to the bulge near the bottom-left area of the objective space.

Nevertheless, unlike our expectation, FEMOSAA-K and FEMOSAA-N do not yield statistically significant differences as shown in Figures 21 to 26 and Figure 30 for *RUBiS-SAS*. This is because the effectiveness of knee selection is fundamentally dependent on the quality of searched valid solution; henceforth, as we have shown in the previous section, when the search process wastes efforts to explore invalid solutions, the quality of the solutions in the final population might be compromised, which would negatively affect the benefit that could have been introduced by our knee selection method. Due to the same reason, for *SOA-SAS* in Figures 27 to 29 and Figure 31, we see that although FEMOSAA-K slightly outperforms FEMOSAA-N on both attributes under all the cases, the majority of the differences (5 out of 6) are not significant statistically.

Overall, the results indicate that:

Answering RQ3—In contrast to the randomly selected non-dominated adaptation solution, the knee selection helps to find a more balanced solution for adaptation, leading to better optimized quality for SAS with statistically significant results and non-trivial effect sizes on at least one attribute. Such improvement can be more obvious when the number of effected solutions is limited, e.g., when the workload is heavy. However, applying knee selection without ensuring the quality of searched valid solution can obscure its effectiveness.

7.5 Comparing FEMOSAA with State-of-the-art Frameworks

To further evaluate the effectiveness of FEMOSAA, we compare it with the following state-of-the-art search-based frameworks from the literature:

DUSE¹⁰ [1]—A representative framework that optimizes SAS using NSGA-II. Since it does not consider dependency and knee in the optimization, we adapt the ad hoc strategies applied by FEMOSAA-D and FEMOSAA-K. As DUSE relies on manual transposition, we have used elitist chromosome representations to ensure fair comparison and to eliminate bias in reproducibility.

PLATO [42]—An approach that applies weighted sum of objectives and Genetic Algorithm, a popular single-objective evolutionary algorithm, in the optimization. It also does not consider dependency in the optimization and thus we use the same ad hoc strategy as FEMOSAA-K. We specify equal weights on the objectives to find the single best solution. As PLATO relies on manual transposition, we have used elitist chromosome representations to ensure fair comparison and to eliminate bias in reproducibility.

FUSION [23]—A well-known feature model based framework that formulates the SAS optimization as an Integer Programming problem, which assumes aggregate objective function and is resolved by an exact algorithm (we use *Branch-and-Bound* in the experiments). FUSION uses binary representation of the solutions and it considers categorical dependency only, thus when no valid solution found, we then fix the violations on numeric dependency. To avoid unacceptable execution time, we forcibly return the best solution so far when it hits a time threshold, i.e., 40s.

Since the source codes of these state-of-the-art frameworks are not openly accessible, we have reproduced the implementation of the optimization algorithms, which are core to these frameworks, following the exact guidance and setups as mentioned in the work. Relevant open-sourced frameworks (e.g., jMetal [21]) are exploited when possible. In particular, we have reproduced the representation of the solution and dependency handling with respect to these algorithms as realized by the state-of-the-art frameworks. The other parts, e.g., the modeling component, are assumed to be identical to FEMOSAA in all experiments. We apply the same set of metrics, statistical test and method for categorizing effect size as discussed earlier.

As shown in Figure 33, for both workload patterns on *RUBiS-SAS*, we can see that FEMOSAA with MOEA/D-STM, NSGA-II and IBEA obtain better quality results than the other state-of-the-art frameworks. They also yield better results in terms of HV and ED. The statistical tests and effect size for the comparisons have been shown in Table 2, in which we can see that the improvements obtained by FEMOSAA with all three MOEAs over the other state-of-the-arts frameworks have been statistically significant in general, resulting $p < 0.05$ with non-trivial effect sizes on at least one quality attribute. The superiority of FEMOSAA when compared with DUSE again conforms that the combination of elitist representation, the dependency aware operators and knee selection can guide the MOEA to achieve better quality results, even when using different underlying MOEA. PLATO, on the other hand, is further constrained by its nature of weighted sum objective functions, preventing it from finding some of the better trade-off points. Although FUSION applies an exact optimization algorithm which should obtain the optimal solution, the fact of the given large search space has prevented it from reaching such an optimality in reasonable time, henceforth forcing it to terminate when the time threshold is hit. Moreover, since the state-of-the-art frameworks do not consider all dependency, their ratio of valid solutions ranges from 8.39% to 64.71% only. For the case of *SOA-SAS*, as shown in Figure 34 and Table 3, we see that FEMOSAA with all three MOEAs yield significantly better results ($p < 0.05$) than PLATO on both quality attribute. Through spending the highest cost, DUSE achieves higher value than that of FEMOSAA on throughput under MOEA/D-STM, but the comparison is statistically insignificant. Further, FEMOSAA is superior to

¹⁰DUSE is actually the same as FEMOSAA-N with NSGA-II which we have evaluated in the previous sections.

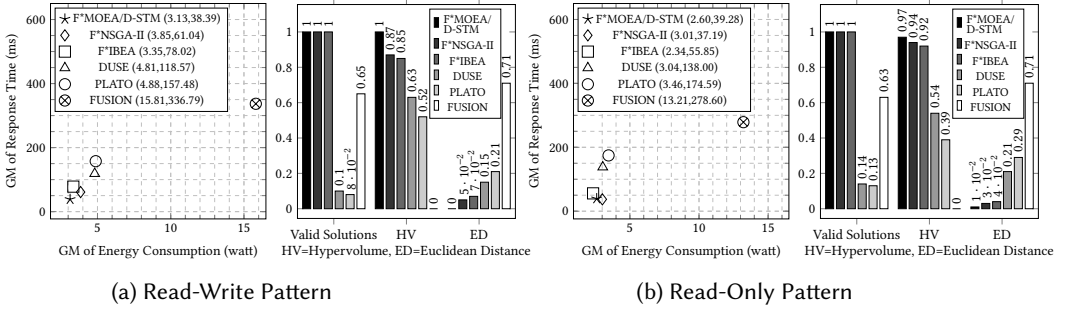


Fig. 33. Comparing FEMOSAA (under different MOEAs) with state-of-the-art frameworks for RUBiS-SAS.. (GM denotes Geometric Mean. F*MOEA/D-STM, F*NSGA-II and F*IBEA denote FEMOSAA with MOEA/D-STM, NSGA-II and IBEA, respectively.)

Table 2. Wilcoxon Signed-Rank Test Results Between FEMOSAA with Different MOEAs and State-of-the-art Frameworks for RUBiS-SAS. (The Significant Statistics, i.e., $p < 0.05$, are Highlighted and Shown in Bold)

(a) Read-Write Pattern

	<i>p</i> value (effect size) for Response Time (ms)			<i>p</i> value (effect size) for Energy Consumption (watt)		
	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA
DUSE	.004 (medium)	.043 (small)	.327 (small)	<.001 (large)	.032 (medium)	<.001 (large)
PLATO	<.001 (large)	.001 (large)	.015 (medium)	<.001 (large)	<.001 (large)	<.001 (large)
FUSION	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)

(b) Read-Only Pattern

	<i>p</i> value (effect size) for Response Time (ms)			<i>p</i> value (effect size) for Energy Consumption (watt)		
	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA
DUSE	.004 (medium)	<.001 (large)	.029 (medium)	.246 (small)	.624 (trivial)	.005 (medium)
PLATO	<.001 (large)	<.001 (large)	<.001 (large)	.010 (medium)	.303 (small)	.001 (large)
FUSION	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)

and more balanced than DUSE overall as evidenced by the greatly improved, statistically significant results on rest of the cases, as well as better HV and ED values. We can also observe that FUSION has very competitive result on cost, which is the best for this quality. However, this is due to the fact that the two quality attributes are strongly conflicting, which causes the exact search in FUSION to trap at a local area of the search space within the given time. This eventually resulted in an unwisely strong bias towards the improvement of cost, as evident by the even larger degradation on throughput in contrast to FEMOSAA, as well as the poor HV and ED values. FUSION leads to only 3% valid solutions as it considers categorical dependencies only while most of the cross-branched dependencies in SOA-SAS are numeric.

In conclusion, those results suggest that:

Answering RQ4—In contrast to the state-of-the-art frameworks (DUSE, PLATO and FUSION), FEMOSAA achieves better optimized quality for SAS with statistically significant results and non-trivial effect sizes on at least one quality attribute, even when using different underlying MOEAs.

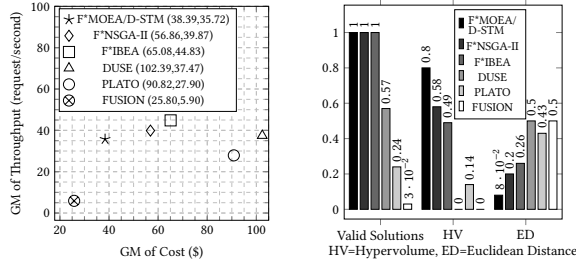


Fig. 34. Comparing FEMOSAA (under different MOEAs) with state-of-the-art frameworks for SOA-SAS. (GM denotes Geometric Mean. F*MOEA/D-STM, F*NSGA-II and F*IBEA denote FEMOSAA with MOEA/D-STM, NSGA-II and IBEA, respectively.)

Table 3. Wilcoxon Signed-Rank Test Results Between FEMOSAA with Different MOEAs and State-of-the-art Frameworks for SOA-SAS. (The Significant Statistics, i.e., $p < 0.05$, are Highlighted and Shown in Bold)

	<i>p</i> value (effect size) for Throughput (request/second)			<i>p</i> value (effect size) for Cost (\$)		
	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA	FEMOSAA with MOEA/D-STM	FEMOSAA with NSGA-II	FEMOSAA with IBEA
DUSE	.451 (trivial)	.300 (small)	.007 (medium)	<.001 (large)	<.001 (large)	<.001 (large)
PLATO	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)
FUSION	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)	<.001 (large)

Table 4. Comparing Mean Running Time For Producing An Adaptation Solution.

	<i>RUBiS-SAS</i> with Read-Write Pattern	<i>RUBiS-SAS</i> with Read-Only Pattern	SOA-SAS
FEMOSAA with MOEA/D-STM	1.96s	1.88s	0.30s
FEMOSAA-N with MOEA/D-STM	2.04s	2.5s	0.43s
FEMOSAA-0/1 with MOEA/D-STM	22.02s	22.94s	1.06s
FEMOSAA with NSGA-II	0.9s	0.86s	0.08s
FEMOSAA-N with NSGA-II	0.94s	3.53s	0.04s
FEMOSAA-0/1 with NSGA-II	10.65s	10.48s	0.24s
FEMOSAA with IBEA	1.09s	1.12s	0.20s
FEMOSAA-N with IBEA	1.13s	1.12s	0.14s
FEMOSAA-0/1 with IBEA	20.90s	20.59s	1.45s
DUSE	0.94s	3.53s	0.04s
PLATO	0.99s	0.79s	0.02s
FUSION	37.98s	55.09s	54.54s

7.6 Runtime Overhead

Since we are interested in optimizing SAS at runtime, the running overhead of the MOEAs guided by FEMOSAA is an important aspect to evaluate. In Table 4, we compare the runtime overhead of FEMOSAA, FEMOSAA-N, FEMOSAA-0/1, DUSE, PLATO and FUSION under *RUBiS-SAS*'s two different workload patterns and *SOA-SAS*. For FEMOSAA and its variants, we examine them using all the studied MOEAs. We report on the mean overhead over all the timesteps. As we can see, under all MOEAs and both subject SAS, FEMOSAA (less than 2s) yields much smaller overhead than FEMOSAA-0/1 (up to 22.94s) as the former encodes the elitist features only, which fundamentally reduces the search space and also simplifies the process of generating new solutions (individuals) in MOEAs. Interestingly, we note that for *RUBiS-SAS*, FEMOSAA (0.9s to 1.96s) has slightly smaller overhead than FEMOSAA-N (0.94s to 3.53s), which is quite surprising as we expected that the former should introduce slightly bigger overhead as it exploits additional processes in the reproduction operators and the selection of final solution for adaptation. These results could be attributed to two reasons: (i) the extra efforts spent in dependency aware operators and knee selection are negligible, and (ii) the dependency aware operators tend to produce solutions that can affect the running time

of MOEAs, for example, in NSGA-II, if the number of solutions in most of the higher ranked fronts is smaller, then the calculation of their crowding distances would yield less running time.

In contrast to the state-of-the-art frameworks, FEMOSAA achieves competitive results on runtime overhead, the actual time taken by FEMOSAA depends on the underlying MOEA though. Notably, FUSION has the biggest overhead as exact algorithm fails to scale with large search space and the optimization runs are often forcibly returned as they hit our predefined threshold of 40s.

To summary, the results reveal that:

Answering RQ5—While the running time of FEMOSAA depends on the underlying MOEA, FEMOSAA has very competitive runtime overhead in contrast to the state-of-the-art frameworks. In addition, the extra efforts spent in dependency aware operators and knee selection are negligible; sometimes, they can even slightly speedup the running time of MOEAs.

7.7 Discussion

7.7.1 FEMOSAA Benefits and Applicability. The most notable benefit of FEMOSAA is that it advances the synergy between software engineering for SAS and evolutionary computation. Without in-depth expertise on evolutionary algorithms in general, software engineers are granted the ability to influence the behaviors of a MOEA in a way that they are familiar with, i.e., the feature model design of a SAS. On the other hand, such a design serves as strong domain knowledge that can extend the MOEA and guide the evolutionary search behavior, i.e., in form of elitist chromosome and dependency aware operators, to produce better solutions, .

FEMOSAA exploits MOEA which particularly fits for cases where the relative weights between objectives are unknown or it is too difficult to quantify them. In this perspective, the benefit of FEMOSAA is that it does not require one to specify weights, which could be labor intensive. Indeed, we acknowledge that there are scenarios where the relative importance between conflicting objectives is explicitly known, and their weights can be precisely specified. In those cases, the concept of MOEA and knee selection employed by FEMOSAA might be less sensible. However, it is possible for FEMOSAA to work with single-objective evolutionary algorithms using aggregation of the objectives while deactivating the knee selection, in which case the SAS can still benefit from the power of the elitist chromosome representation and the dependency aware operators.

Another point which is worth mentioning is that MOEA does not guarantee optimal solutions; however, it is very efficient in producing good approximation to complex and non-linear problems that would be otherwise unsolvable by exact optimization. Thus, we would not recommend to use FEMOSAA on SAS that are simple, small in the search space and can be handled by exact search that leads to an optimal solution.

7.7.2 FEMOSAA Running Time. Indeed, in contrast to FEMOSAA, the classic rule and policy based decision making approach is very fast when adapting a SAS. However, its effectiveness of adaptation quality relies on several important factors, including:

- (1) The full knowledge of every possible condition that the SAS may encounter.
- (2) Some theoretical assumptions on the SAS and the environment that underpins the rules and policies.
- (3) The manual reasoning of the optimal adaptation decision under a given condition (i.e., the mapping between a condition and an adaptation decision).

For SAS works in highly dynamic and uncertain environments (e.g., in cloud computing and software defined networking), rule and policy based approaches would fail due to the fact that they

heavily rely on human knowledge and there are emergent conditions that have not been accounted for, given the requirements of points (1) and (2). Further, SAS often has a large variability, i.e., a large number of alternatives as the two SAS we studied. This makes point (3) in the rule and policy based approaches unrealistic; in fact, an exact search optimization would also fail under the problem with a large search space, as such a problem itself is intractable. In contrast, search-based software engineering techniques, particularly evolutionary algorithms, offer a promising way to our SAS optimization problem. This is because being a metaheuristic, evolutionary algorithms are dynamic in nature that are able to perform optimization without in-depth knowledge and assumptions of the problem in hand (e.g., the property of the SAS and the environment). In addition, the notion of natural evolution and population permits its ability to find approximately optimal solutions even for intractable problems.

As we illustrated in Section 7.6, FEMOSAA achieves a runtime overhead of seconds under the two SAS studied. This may cause delay in adaptation under extreme scenarios where the transition needs to be completed in microseconds. However, for other cases where the requirements of transition time can be relaxed, the better quality of adaptation that is generated by the underlying evolutionary algorithm has made the cost of its runtime overhead negligible, especially considering that the SAS optimization problem would otherwise be unsolvable with rule and policy based approaches. In fact, as we have shown throughout Section 7, the proposed elitist chromosome representation and dependency aware operators in FEMOSAA have enabled the underlying evolutionary algorithm to reach a better adaptation quality with even smaller runtime overhead.

7.7.3 Threats to Validity. Some threats to validity of FEMOSAA can be discussed as the following:

Threats to construct validity is concerned with whether the used metrics can indeed reflect what we intend to measure. In this work, our experiments have selected wider range of quality attributes, i.e., response time, energy consumption, throughput and cost. Those quality attributes and their metrics are the most commonly assessed quality aspects for SAS from the literature [39][49][23][42]. Further, we assessed the aggregated results of different quality objectives using HV and ED, which are widely applied metrics to measure the quality of solutions for multi-objective optimization problems [50][37]. Threats to construct validity could be also related to the stochastic nature of the considered MOEAs in experiments which can influence the measurements, especially for those that do not account for dependency and knee in the optimization. Indeed, to draw a meaningful conclusion of the measurements for stochastic search-based optimization, repeated runs are necessary as suggested in [4]. We have mitigated such a bias by following the design introduced in [4], including: conducting 102 optimization runs¹¹ for FEMOSAA and the others; exploiting a statistical test to verify the significance of comparisons and reporting the effect size for confirming the measured results are not occurred due to chance. Further, in the *RUBiS-SAS* case, FEMOSAA has been evaluated and measured following the realistic FIFA98 workload trace.

Threats to internal validity is related to the values of parameters for the MOEAs. The setup in this work has been carefully tailored such that it produces good trade-off between the quality of optimization and the overhead. However, these values might vary depending on the characteristics in the context, e.g., the given feature model, the types of SAS and the environmental conditions, which itself could be a topic for future research. Threats to internal validity can be also arisen from the accidental bugs in the experimental implementations, which is always possible when writing any kind of software, especially for SAS which is naturally more complex than others. However, we have tried to mitigate this unavoidable phenomenon by (i) exploiting well-established open-source

¹¹The number of runs is an implied result of the client emulator's setup for *RUBiS-SAS* and *SOA-SAS*

framework whenever possible (e.g., jMetal [21]); (ii) following the exact guidance given in some of the compared work; and (iii) debugging through formal software testing procedure.

Threats to external validity is linked to the benchmark and scenario that are used in the experiments. To improve generalization of the experimental evaluations, we have also evaluated FEMOSAA with three widely applied but distinct MOEAs under two running SAS and different workload patterns, which can diversify the runtime behaviors of the SAS. In particular, one of the SAS (i.e., *RUBiS-SAS*) contains a stack of real-world software, which help to emulate more realistic scenario of the running SAS. Though our experimentation on the cases approximate real and industrial scale, it is difficult to claim complete generality; such claim would require much large number of independent domain-specific cases and needs to be performed by independent future adopters. In future work, we plan to evaluate FEMOSAA in other extreme context, e.g., in mobile environment where computational resource is rather limited. However, it is worth noting that the efforts for building the experiments and deploying the running SAS are non-trivial, which could involve potential expenses.

8 RELATED WORK

Search-based optimization has been widely applied for SAS, either constitutes as a general framework or being tailored to a specific domain, e.g., service systems or cloud systems. In this section, we provide an overview of the most notable and relevant research in this area while examining them in the light of FEMOSAA.

8.1 Evolutionary Optimization for SAS

A common way of optimizing multiple objectives is to simply aggregate different objectives (e.g., weighted sum) such that they can be resolved in a single objective function. Hence, there exists some work that leverages single-objective evolutionary algorithms for optimizing SAS based on an objective aggregation. PLATO [42] and VALKYRIE [27] are two examples. However, as opposed to the automatic transposition approach in FEMOSAA, they have relied on manual encoding of the SAS into chromosome representation and they do not consider dependency between features. Further, it is well-known that the relative weights are difficult to be tailored by the engineers; and a single aggregation could restrict the search, causing limitation when searching for good solutions spread over the search space.

Exploiting MOEAs to handle the trade-off between conflicting objectives is an emerging trend for optimizing SAS at runtime. Among others, DUSE [1] uses NSGA-II to produce a set of non-dominated solutions for SAS. Similarly, NSGA-II and other EAs, as well as other meta-heuristics algorithms, have also been applied in [28][22][51][12][15], which focus on general SAS and the specific SAS for cloud and service systems. Nevertheless, their encoding of the SAS has been manual and no dependency between features are considered. Consequently, a significant number of function evaluations would be wasted to search for invalid solutions, which as we have shown, can degrade the quality of solutions found. Furthermore, they rely on the non-dominated set, from which any solution can be used for adaptation, which can entail imbalanced trade-off. On contrary, FEMOSAA finds a knee solution that is generally preferable. Moreover, FEMOSAA relies on automatic approach where the elitist features are identified and the dependency is extracted to guide MOEA.

The closely related and most recent work is probably the one produced by Pascual et al. [41], where they proposed a framework for self-adaptive mobile system, which uses the feature model with MOEA while considering feature dependencies in the optimization. However, FEMOSAA is different from their work in various aspects:

- Pascual et al.'s work relies on a binary representation of all features and only categorical features are considered; while FEMOSAA encodes elitist features into the chromosome only to form a polyadic representation which reduces the number of genes, and thereby considerably shrinking the search space.
- Instead of modifying the reproduction operators, the work by Pascual et al. still allow the operators to explore and generate invalid solutions but fixes those solutions using a random repair strategy, i.e., each gene, which violates the categorical dependency, is fixed. However, such a fix is not guided by the dependency chain, thus there is no guarantee that the fixed gene would not cause additional violations (if there is a chain of dependency), and thereby it cannot ensure that a valid solution is always produced. Indeed, in their work, the process is repeated if the previous fix has not resolved all the violations, and the repair stops when a maximum number of repeats has been reached. On the contrary, FEMOSAA extracts dependencies and directly injects them to both the mutation and crossover operators, which are explicitly guided by the dependency chains and the related value trees; this fundamentally prevents invalid solutions from being explored.
- The nature of binary representation in Pascual et al.'s work implies that it is difficult for them to handle numeric dependencies, which are covered by FEMOSAA.
- FEMOSAA considers knee selection for adaptation while Pascual et al.'s work selects any non-dominated solution for adaptation, which could be highly imbalanced.

8.2 Evolutionary Optimization in Software Product Line Engineering

The feature model is widely used in Software Product Line (SPL) for modeling variability, which is similar to our use of feature model on SAS. In particular, MOEA has been applied to SPL, see [45][46][31] for example. However, unlike SPL, where the objectives are highly concerned with designing software products that expose features' richness, diversity and their known defects, *etc*, our focus in SAS is on optimizing non-functional quality attributes, seeking adaptation decisions that can better respond to dynamics and uncertainty in the environment with limited or no human intervention. In addition, SPL assume design time while the SAS mainly focus on runtime.

Although given the differences mentioned above, some SPL work exhibits resemblance to FEMOSAA: they both aim to transpose the design of the feature model into the context of MOEA. Therefore, here we compare FEMOSAA with those SPL preventative approaches, where we specifically look at how the feature model is transposed into MOEA and their strategies to handle feature dependencies. This consideration is particularly essential to explain why it is insufficient to directly apply the transposition used in the existing studies of the SPL domain, and what are the improvements we have made in this aspect.

Among others, Sayyad et al. [45][46] exploit NSGA-II and IBEA for finding optimal design of product line using binary encoding of the features into chromosome representation. In contrast, FEMOSAA does not rely on the lengthy binary representation and it encodes a polyadic chromosome using elitists feature, which as we have shown, leading to better optimization results and running time. Hierons et al. [31] proposed an extended MOEA encoding method for the feature model, in which the basic encoding is still binary. Similar to our motivation, Hierons et al. [31] seek simplification by eliminating features which are the root of a *OR* group as their variability can be represented by their children, e.g., *Cache Mode* in Figure 4. However, FEMOSAA goes one step ahead by discarding the binary chromosome representation; this is achieved through selecting the elitists features that cannot be removed without affecting the variability while minimizing the length of encoding, which does not only eliminating any features whose variability can be represented by their children (not restricted to the features that are root of *OR* group as in [31]), e.g., *Cache* and

Cache Mode in Figure 4, but also eliminating those whose variability can be represented by their parent, e.g., *CPU*'s children in Figure 4. Unlike their work, we do not simply remove all mandatory features; we retain those with a *XOR* group of children as they would help us to considerably simplify the chromosome representation even more (e.g., *CPU* in Figure 4). Further, they have ignored numeric features which are common in SAS. The numeric features, if not specifically handled, would lead to high computational overhead under binary encoding.

Another fact that distinguishes FEMOSAA and existing SPL work on the transposition process is the way in which the dependencies are handled. Hierons et al. [31] and Sayyad et al. [45][46] simply formulate dependency compliance as an additional objective to be optimized. Their formulation stems from the fact that it is too complex to explicitly handle dependency during the evolution due to the inherited difficulty of the binary encoding. The basic assumption is that, with a large number of generations, the MOEA would eventually discovery many valid solutions. While this schema may be sensible for design time optimization problems, it is ill-suited for SAS where the optimizations occur at runtime, because the extra dimension of dependency objective would impose too much additional difficulty of the problem and, at the same time, still fail to guarantee valid solutions. With the binary chromosome representation, Henard et al. [30] aim to overcome this by combing MOEA with an off-the-shelf CSP solver, in which the solvers act similarly to our dependency aware operators. However, those solvers are general and have not been specifically tailored to the needs of problem, which are often computationally expensive and they would not guarantee valid solutions either. FEMOSAA, on the other hand, extracts the dependencies with respect to the elitist chromosome representation at design time, and those dependencies are injected into the operators of MOEA to fundamentally prevent the exploration of invalid solutions during runtime evolution. Further, FEMOSAA handles complex dependencies related to the numeric features.

8.3 Non-evolutionary Optimization for SAS

Beside evolutionary algorithms, exact algorithms are also utilized for optimizing SAS [23][24][9], as they guarantee to find the optimal solution and can easily work with single objective. Here, we aim to consider the most noticeable non-evolutionary optimization approaches for SAS which do not tie to the special characteristics of a specific application domain. FUSION [23] is a general framework that optimizes SAS using exact algorithm. At design time, it also applies the feature model to represent the design of a SAS. However, unlike FEMOSAA, it does not consider numeric features and they formulate the problem as Integer Optimization problem using binary encoding of all the features. Eshafani et al. [24] also follow similar idea of FUSION but they additionally use Fuzzy logic to constitute the objective functions. MOSES [9] is a framework that is designed for self-adaptive service systems, where the optimization is formulated as Linear Programming problem that is convex and can be solved exactly.

One fundamental issue with the non-evolutionary and exact algorithm is that they fail to scale when the search space is large, which is often the case for modern SAS. Furthermore, exact approaches tend to be highly sensitive to the nature of the problem, e.g., whether they are convex or concave. This could impose extra difficulty because the analysis of problem nature for SAS is very difficult, if not impossible, due to the dynamic and uncertain nature of the context. Additionally, exploiting exact algorithms often need to work on an objective aggregation, which limits their applicability and capability.

In contrast, FEMOSAA relies on MOEA, which is a type of stochastic evolutionary optimization algorithms that is specifically designed to handle multi-objective optimization problems. It is known that MOEA can efficiently find optimal (or near-optimal) solutions for problems with large search space and is capable to reveal fine-grained trade-off surface without the need of objective

aggregation. In addition, MOEA is problem agnostic and thereby it is less sensitive to the nature of a given problem.

8.4 Other Approaches for Self-Adaptation at Runtime

Advanced control theory has also been used for SAS decision making because of their low latency. Among others, Filieri et al. [25] propose a multi-objective controller where each objective would have independent sensors and actuators; the reasoning relies on aggregation of objectives, however. Simplex [47] is also another recent control theoretic method for SAS, where simplex optimization algorithm is used in conjunction with the updates of controller's gains.

Reinforcement Learning (RL) [8] is another thread that regards the SAS optimization problem as a learning problem. However, in RL, there is no explicit optimization process due to the absence of clear model, and therefore, the adaptation decision is often tailored in a trial-and-error manner which could impose an expensive exploration phase.

Without an explicit search behavior, both control theoretic and learning based approaches lack the ability to perform exploration without affecting the SAS. In addition, they do not consider dependency constraints and they are difficult to be adopted for effectively reasoning about trade-offs at runtime. In contrast, FEMOSAA exploits MOEA and the feature model, aiming to explicitly handle multi-objective optimization while considering dependencies.

9 CONCLUSION

This paper presents FEMOSAA, a novel framework that systematically and automatically synergizes the feature model of SAS and a MOEA, to optimize the SAS at runtime. At design time, FEMOSAA finds the elitist features, including categorical and numeric ones, to create a polyadic chromosome representation; it then extracts the dependencies between the genes, which are then used to extend the underlying MOEA for runtime optimization. The feature model serves as the engineers' domain knowledge that can reduce the search space (fundamentally and virtually) and guide the search, henceforth increasing the chance for finding better solutions. Further, FEMOSAA finds the knee solutions that achieve a balanced trade-off. By extensively comparing FEMOSAA with its variants and the state-of-the-art frameworks on two complex real-world SAS, using three widely applied MOEAs and under two workload patterns for optimizing various conflicting objectives, we show that FEMOSAA produces statistically better and more balanced results for trade-off with reasonable overhead. In particular, the most notable observations of FEMOSAA are that:

- Applying the elitist chromosome representation to encode the problem into MOEA helps to produce better quality and smaller runtime overhead for optimizing SAS, but such improvement tends to become marginal when the number of effected solution is small, e.g., the workload is heavy.
- The dependency aware operators can properly guide the search, finding solutions with better convergence and diversity, leading to better quality of SAS optimization. However, applying dependency aware operators without ensuring the balance of the selected adaptation solution might obscure its effectiveness.
- The knee selection helps to find more balanced solution for adaptation. However, applying knee selection without ensuring the quality of searched valid solution can obscure its effectiveness.
- In contrast to the state-of-the-arts framework from the literature, FEMOSAA, with all the three studied MOEAs, produces statistically and practically better quality for optimizing SAS at runtime.

- Overall, FEMOSAA has very competitive runtime overhead in contrast to the state-of-the-art frameworks. Further, the extra efforts spent in dependency aware operators and knee selection are negligible; sometimes, they can even slightly speedup the running time of MOEAs.

Our work impacts and advances the synergy between software engineering for SAS and evolutionary computation, combining the strengths from both fields. Particularly, with FEMOSAA, software engineers can exploit MOEAs to tackle SAS optimization without prior extensive expertise of the MOEA. On the other hand, automatic transposition of the feature model into MOEA's context can improve MOEA, and make the domain knowledge systematic and comprehensible for MOEA researchers, who in turn can design more effective algorithms for SAS. In contrast to many SBSE work, our deeper synergy takes one step further by automatically and dynamically extracting the domain information of SAS to extend the internal structure of MOEA. In future work, we plan to apply FEMOSAA in other domains of SAS and extend it for managing more conflicting objectives.

REFERENCES

- [1] ANDRADE, S. S., AND D. A. MACDO, R. J. A search-based approach for architectural design of feedback control concerns in self-adaptive systems. In *2013 IEEE 7th International Conference on Self-Adaptive and Self-Organizing Systems* (Sept 2013), pp. 61–70.
- [2] APACHE SOFTWARE FOUNDATION. Apache tomcat. <http://tomcat.apache.org/>.
- [3] APACHE SOFTWARE FOUNDATION. Ehcache. <http://www.ehcache.org/>.
- [4] ARCURI, A., AND BRIAND, L. A practical guide for using statistical tests to assess randomized algorithms in software engineering. In *Proceedings of the 2011 IEEE 33rd International Conference on Software Engineering (ICSE)* (2011), pp. 1–10.
- [5] ARLITT, M., AND JIN, T. A workload characterization study of the 1998 world cup web site. *Netwrk. Mag. of Global Internetwkg.* 14, 3 (May 2000), 30–37.
- [6] BENAVIDES, D., SEGURA, S., AND RUIZ-CORTS, A. Automated analysis of feature models 20 years later: A literature review. *Information Systems* 35, 6 (2010), 615–636.
- [7] BOURDON, A., NOUREDDINE, A., ROUYVOY, R., AND SEINTURIER, L. PowerAPI: A Software Library to Monitor the Energy Consumed at the Process-Level. *ERCIM News* 92 (Jan. 2013), 43–44.
- [8] BU, X., RAO, J., AND XU, C.-Z. Coordinated self-configuration of virtual machines and appliances using a model-free learning approach. *IEEE Transactions on Parallel and Distributed Systems* 24, 4 (2013), 681–690.
- [9] CARDELLINI, V., CASALICCHIO, E., GRASSI, V., IANNUCCI, S., PRESTI, F. L., AND MIRANDOLA, R. Moses: A framework for qos driven runtime adaptation of service-oriented systems. *IEEE Transactions on Software Engineering* 38, 5 (Sept 2012), 1138–1159.
- [10] CHEN, T., AND BAHSOON, R. Self-adaptive and sensitivity-aware qos modeling for the cloud. In *Proceedings of the 8th International Symposium on Software Engineering for Adaptive and Self-Managing Systems* (2013), pp. 43–52.
- [11] CHEN, T., AND BAHSOON, R. Self-adaptive and online qos modeling for cloud-based software services. *IEEE Transactions on Software Engineering* 43, 5 (2017), 453–475.
- [12] CHEN, T., AND BAHSOON, R. Self-adaptive trade-off decision making for autoscaling cloud-based services. *IEEE Transactions on Services Computing* 10, 4 (2017), 618–632.
- [13] CHEN, T., BAHSOON, R., WANG, S., AND YAO, X. To adapt or not to adapt? technical debt and learning driven self-adaptation for managing runtime performance. In *Proceedings of the 9th ACM/SPEC International Conference on Performance Engineering (ICPE)* (2018), ACM.
- [14] CHEN, T., BAHSOON, R., AND YAO, X. Online qos modeling in the cloud: A hybrid and adaptive multi-learners approach. In *Proceedings of the IEEE/ACM 7th International Conference on Utility and Cloud Computing* (2014), pp. 327–336.
- [15] CHEN, T., LI, M., AND YAO, X. On the effects of seeding strategies: A case for search-based multi-objective service composition. In *Proceedings of the 20th international Genetic and Evolutionary Computation Conference (GECCO)* (2018), ACM.
- [16] CZARNECKI, K., HELSEN, S., AND EISENECKER, U. Staged configuration using feature models. In *SPLC* (2004), vol. 3154, Springer, pp. 266–283.
- [17] DE LEMOS, R., GIESE, H., MÜLLER, H. A., SHAW, M., ANDERSSON, J., LITOU, M., SCHMERL, B., TAMURA, G., VILLEGAS, N. M., VOGEL, T., WEYNS, D., BARESI, L., BECKER, B., BENCOMO, N., BRUN, Y., CUKIC, B., DESMARAI, R., DUSTDAR, S., ENGELS, G., GEIHS, K., GÖSCHKA, K. M., GORLA, A., GRASSI, V., INVERARDI, P., KARSAI, G., KRAMER, J., LOPES, A., MAGEE, J., MALEK, S., MANKOVSKII, S., MIRANDOLA, R., MYLOPOULOS, J., NIERSTRASZ, O., PEZZÈ, M., PREHOFER, C., SCHÄFER,

- W., SCHLICHTING, R., SMITH, D. B., SOUSA, J. P., TAHVILDARI, L., WONG, K., AND WUTTKE, J. *Software Engineering for Self-Adaptive Systems: A Second Research Roadmap*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 1–32.
- [18] DEB, K., AND JAIN, H. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Transactions on Evolutionary Computation* 18, 4 (2014), 577–601.
- [19] DEB, K., AND KALYANMOY, D. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [20] DEB, K., PRATAP, A., AGARWAL, S., AND MEYARIVAN, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Transactions on Evolutionary Computing* 6, 2 (April 2002), 182–197.
- [21] DURILLO, J. J., AND NEBRO, A. J. jMetal: A java framework for multi-objective optimization. *Advances in Engineering Software* 42, 10 (2011), 760–771.
- [22] EL KATEB, D., FOUQUET, F., NAIN, G., MEIRA, J. A., ACKERMAN, M., AND LE TRAON, Y. Generic cloud platform multi-objective optimization leveraging models@run.time. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing* (New York, NY, USA, 2014), SAC '14, ACM, pp. 343–350.
- [23] ESFAHANI, N., ELKHODARY, A., AND MALEK, S. A learning-based framework for engineering feature-oriented self-adaptive software systems. *IEEE Transactions on Software Engineering* 39, 11 (Nov 2013), 1467–1493.
- [24] ESFAHANI, N., KOUROSHFAR, E., AND MALEK, S. Taming uncertainty in self-adaptive software. In *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering* (New York, NY, USA, 2011), ESEC/FSE '11, ACM, pp. 234–244.
- [25] FILIERI, A., HOFFMANN, H., AND MAGGIO, M. Automated multi-objective control for self-adaptive software design. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering* (2015), ACM, pp. 13–24.
- [26] FITTKAU, F., FREY, S., AND HASSELBRING, W. Cdosim: Simulating cloud deployment options for software migration support. In *2012 IEEE 6th International Workshop on the Maintenance and Evolution of Service-Oriented and Cloud-Based Systems (MESOCA)* (Sept 2012), pp. 37–46.
- [27] FREDERICKS, E. M. Automatically hardening a self-adaptive system against uncertainty. In *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems* (New York, NY, USA, 2016), SEAMS '16, ACM, pp. 16–27.
- [28] FREY, S., FITTKAU, F., AND HASSELBRING, W. Search-based genetic optimization for deployment and reconfiguration of software in the cloud. In *Proceedings of the 2013 International Conference on Software Engineering* (Piscataway, NJ, USA, 2013), ICSE '13, IEEE Press, pp. 512–521.
- [29] HARMAN, M., BURKE, E., CLARK, J. A., AND YAO, X. Dynamic adaptive search based software engineering. In *Proceedings of the 2012 ACM-IEEE International Symposium on Empirical Software Engineering and Measurement* (Sept 2012), pp. 1–8.
- [30] HENARD, C., PAPADAKIS, M., HARMAN, M., AND LE TRAON, Y. Combining multi-objective search and constraint solving for configuring large software product lines. In *Proceedings of the 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering (ICSE)* (2015), pp. 517–528.
- [31] HIERONS, R. M., LI, M., LIU, X., SEGURA, S., AND ZHENG, W. SIP: Optimal product selection from feature models using many-objective evolutionary optimization. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 25, 2 (2016), 17.
- [32] KAMPENES, V. B., DYB, T., HANNAY, J. E., AND SJBERG, D. I. A systematic review of effect size in software engineering experiments. *Information and Software Technology* 49, 11fi?12 (2007), 1073 – 1086.
- [33] KANG, K., COHEN, S., HESS, J., NOVAK, W., AND PETERSON, A. Feature-oriented domain analysis (foda) feasibility study. Tech. Rep. CMU/SEI-90-TR-021, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, 1990.
- [34] KNOWLES, J. D., AND CORNE, D. Approximating the nondominated front using the pareto archived evolution strategy. *Evolutionary Computation* 8, 2 (2000), 149–172.
- [35] LI, K., DEB, K., ZHANG, Q., AND KWONG, S. An evolutionary many-objective optimization algorithm based on dominance and decomposition. *IEEE Trans. Evolutionary Computation* 19, 5 (2015), 694–716.
- [36] LI, K., ZHANG, Q., KWONG, S., LI, M., AND WANG, R. Stable matching-based selection in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation* 18, 6 (2014), 909–923.
- [37] LI, M., CHEN, T., AND YAO, X. A critical review of a practical guide to select quality indicators for assessing pareto-based search algorithms in search-based software engineering: Essay on quality indicator selection for sbse. In *Proceedings of the 40th international conference on software engineering, NIER track* (2018), IEEE/ACM.
- [38] MIETTINEN, K. *Nonlinear Multiobjective Optimization*, vol. 12. Kluwer Academic Publishers, 1999.
- [39] MINGAY, S. Green it: the new industry shock wave. *Gartner RAS Research Note G 153703* (2007), 2007.
- [40] ORACLE CORPORATION. Mysql. <https://www.mysql.com/>.
- [41] PASCUAL, G. G., LOPEZ-HERREJON, R. E., PINTO, M., FUENTES, L., AND EGYED, A. Applying multiobjective evolutionary algorithms to dynamic software product lines for reconfiguring mobile applications. *Journal of Systems and Software* 103 (2015), 392 – 411.

- [42] RAMIREZ, A. J., KNOESTER, D. B., CHENG, B. H. C., AND MCKINLEY, P. K. Plato: a genetic algorithm approach to run-time reconfiguration in autonomic computing systems. *Cluster Computing* 14, 3 (2011), 229–244.
- [43] RICE UNIVERSITY. Rice university bidding systems. <http://rubis.ow2.org/>.
- [44] ROY, N., DUBEY, A., GOKHALE, A., AND DOWDY, L. A capacity planning process for performance assurance of component-based distributed systems. In *Proceedings of the 2Nd ACM/SPEC International Conference on Performance Engineering* (New York, NY, USA, 2011), ICPE '11, ACM, pp. 259–270.
- [45] SAYYAD, A. S., INGRAM, J., MENZIES, T., AND AMMAR, H. Scalable product line configuration: A straw to break the camel's back. In *Proceedings of the 2013 IEEE/ACM 28th International Conference on Automated Software Engineering (ASE)* (2013), pp. 465–474.
- [46] SAYYAD, A. S., MENZIES, T., AND AMMAR, H. On the value of user preferences in search-based software engineering: a case study in software product lines. In *Proceedings of the 2013 International Conference on Software Engineering* (2013), IEEE Press, pp. 492–501.
- [47] SHEVTSOV, S., AND WEYNS, D. Keep it simple: Satisfying multiple goals with guarantees in control-based self-adaptive systems. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering* (2016), ACM, pp. 229–241.
- [48] UNIVERSITY OF CAMBRIDGE COMPUTER LABORATORY. Xen: a virtual machine monitor. <http://www.xenproject.org/>.
- [49] WADA, H., SUZUKI, J., YAMANO, Y., AND OBA, K. E³: A multiobjective optimization framework for sla-aware service composition. *IEEE Transactions on Services Computing* 5, 3 (2012), 358–372.
- [50] WANG, S., ALI, S., YUE, T., LI, Y., AND LIAAEN, M. A practical guide to select quality indicators for assessing pareto-based search algorithms in search-based software engineering. In *Proceedings of the 38th International Conference on Software Engineering* (New York, NY, USA, 2016), ICSE '16, ACM, pp. 631–642.
- [51] YUSOH, Z. I. M., AND TANG, M. Composite saas placement and resource optimization in cloud computing using evolutionary algorithms. In *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on* (June 2012), pp. 590–597.
- [52] ZHANG, Q., AND LI, H. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation* 11, 6 (2007), 712–731.
- [53] ZITZLER, E., AND KÜNZLI, S. Indicator-based selection in multiobjective search. In *Proceedings of the PPSN VIII: the 8th International conference on Parallel Problem Solving from Nature* (2004), pp. 832–842.
- [54] ZITZLER, E., LAUMANN, M., AND THIELE, L. SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In *Evolutionary Methods for Design, Optimisation, and Control* (2002), pp. 95–100.
- [55] ZITZLER, E., AND THIELE, L. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation* 3, 4 (Nov 1999), 257–271.