Taming Uncertainty in Self-Adaptive Software

Naeem Esfahani
Department of Computer Science
George Mason University

nesfaha2@gmu.edu

Ehsan Kouroshfar
Department of Computer Science
George Mason University

ekourosh@gmu.edu

Sam Malek
Department of Computer Science
George Mason University

smalek@gmu.edu

ABSTRACT

Self-adaptation endows a software system with the ability to satisfy certain objectives by automatically modifying its behavior. While many promising approaches for the construction of self-adaptive software systems have been developed, the majority of them ignore the uncertainty underlying the adaptation decisions. This has been one of the key obstacles to wide-spread adoption of self-adaption techniques in risk-averse real-world settings. In this paper, we describe an approach, called *POssIbilistic SElf-aDaptation* (*POISED*), for tackling the challenge posed by uncertainty in making adaptation decisions. POISED builds on possibility theory to assess both the positive and negative consequences of uncertainty. It makes adaptation decisions that result in the best range of potential behavior. We demonstrate POISED's application to the problem of improving a software system's quality of service via runtime reconfiguration of its customizable software components. We have extensively evaluated POISED using a prototype of a robotic software system.

Categories and Subject Descriptors

D.2.10 [**Software Engineering**]: Design – *Methodologies*.

General Terms

Algorithms, Performance, Design.

Keywords

Uncertainty, Self-Adaptation, Software Architecture

1. INTRODUCTION

Self-adaptation is an effective approach in dealing with the changing dynamics of many application domains, such as mobile and pervasive systems. In response to changes in the environment or requirements, a self-adaptive software system modifies itself to satisfy certain objectives [2, 16]. While the benefits of such systems are plenty, their development has shown to be more challenging than traditional software systems [2]. One key culprit is that self-adaptation is subject to *uncertainty* [2].

We distinguish between the external and internal uncertainty. *External uncertainty* arises from the environment or domain in which the software is deployed. For example, external uncertainty for a software system deployed in an unmanned vehicle may include the likelihood of certain weather conditions occurring. Software self-adaptation is one approach in dealing with the

effects of external uncertainty, e.g., in a snow storm the vehicle's navigator component may be replaced with a more conservative navigator to avoid a collision. On the other hand, *internal uncertainty* is rooted in the difficulty of determining the impact of adaptation on the system's quality objectives, e.g., determining the impact of replacing a software component on the system's responsiveness, battery usage, etc. While both sources of uncertainty are of great concern, in this paper we focus our attention on the challenges posed by *internal uncertainty*.

Uncertainty can be observed in every facet of adaptation, albeit at varying degrees. It follows from the fact that the system's user, adaptation logic, and business logic are loosely coupled, introducing numerous sources of uncertainty [3]. A key observation is that while the level of uncertainty could vary, no self-adaptive software system is ever completely free of it.

The research community has made great strides in tackling the complexity of constructing self-adaptive software systems [2, 16]. However, as corroborated by others [2], there is a dearth of applicable techniques for handling uncertainty in this setting. A few researchers have recently begun to address uncertainty issues in requirements specification [1, 28] and resource prediction [23], but no approach that we are aware of has tackled the challenge posed by uncertainty in making adaptation decisions. We believe this has been one of the primary obstacles to wide-spread adoption of self-adaptation in risk-averse domains.

This is precisely the challenge we have aimed to address in this paper. We present a general quantitative approach for tackling the complexity of automatically making adaptation decisions under internal uncertainty, called *POssIbilistic SElf-aDaptation* (*POISED*). Estimates of uncertainty in the elements comprising a self-adaptation problem are incorporated in *possibilistic analysis* of the adaptation choices. Possibilistic analysis is founded on the principles of fuzzy mathematics [29], which provides a sound basis for representing uncertainty, as well as dealing with its negative and positive consequences on the adaptation choices. POISED redefines the conventional definition of optimal adaptation decision to one that has the best range of behavior. In turn, the selected solution has the highest likelihood of satisfying the system's quality objectives, even if due to uncertainty, properties expected of the system are not borne out in practice.

We demonstrate POISED by applying it to the problem of improving a software system's quality of service via runtime reconfiguration of its customizable software components. We have evaluated POISED under numerous circumstances and using a prototype of a robotic software system. The results demonstrate POISED's ability to deal with uncertainty by making adaptation decisions that are superior to those of the conventional approach.

This paper is organized as follows. Section 2 motivates the research. Section 3 provides an overview of POISED. Section 4 formalizes the self-adaptation problem used to describe and

evaluate POISED. Section 5 describes techniques for quantifying uncertainty. Section 6 presents our possibilistic analysis approach, followed by detailed evaluation in Section 7. The paper concludes with a discussion of the related literature and future research.

2. MOTIVATING EXAMPLE

We use a subset of a robotic software system developed in another project [18] to motivate and describe this research. The robotic software is part of a distributed search and rescue system aimed at supporting the government agencies in dealing with emergency crises (e.g., fire, hurricane). Figure 1b provides an abridged view of the robotic system's architecture. The software components comprising the robotic system range from abstractions of the physical entities, such as software controlled sensors and actuators on board the robot, to purely logical functionalities, such as image detection and navigation. The bold path in Figure 1b indicates the *Maneuver* execution scenario, which aims to safely steer the robot. The *Camera* feed is sent to *Obstacle Detector*, which runs an image processing algorithm to identify obstacles. Obstacle information is used by *Navigator* to plan the direction and speed of movement, which are then put into effect by the *Controller*.

The software components comprising this system are *customizable* meaning that they can be configured to operate in different modes of operation. Figure 1a shows some of the available configuration dimensions. For instance, *Power* is a configuration dimension for the *Controller* component. A *Controller* could operate in either *Energy Saving* or *Full Power* mode. A component may have many configuration dimensions.

The configuration of a software component determines its quality attributes (e.g., response time) and resource usage (e.g., memory), which could also impact the properties of the entire system. For instance, given the resource constrained nature of the mobile robots, the configuration decisions of each component have a significant impact on the system's performance as well as its battery life. Such decisions can only be effectively made at
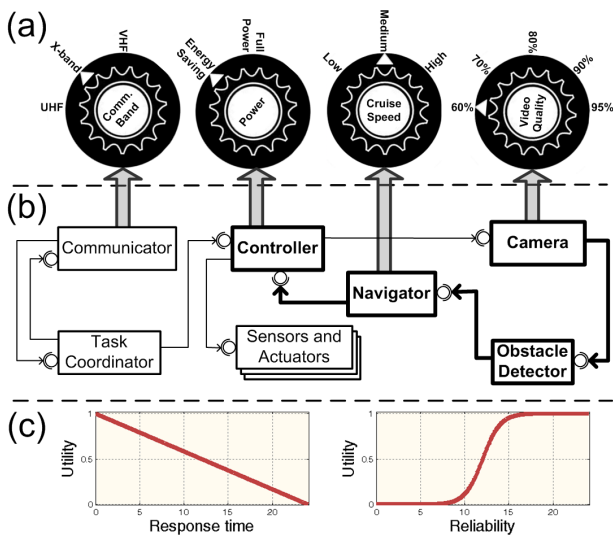

Figure 1. A subset of the robotic software: (a) configuration dimensions and alternatives for components of the robot, (b) software architecture, where the bold path indicates the components participating in the Maneuver execution scenario, and (c) utility functions defined in terms of quality attributes.

runtime, since the system properties (e.g., available bandwidth) are often not known at design-time and may change at runtime.

As shown in Figure 1c, for making runtime decisions, utility functions capturing the user's satisfaction with different levels of quality attribute (e.g., availability) are used. The adaptation logic uses analytical models to estimate the effect of configuration decision on the system's quality attributes, and in turn the resulting utility. For example, given the configuration of the robot's components, an analytical model, such as Queueing Network model [20], may be used to quantify the response time for its *Maneuver* scenario. The objective is to find a configuration that achieves the maximum overall utility.

The above approach is rather myopic, since it does not consider the uncertainty of information used in making adaptation decisions. Consider that almost every facet of the approach outlined above faces some form of uncertainty:

- *Uncertainty in System Parameters*: The monitoring data obtained from a running system rarely corresponds to a single value, but rather a distribution of values obtained over the observation period. For instance, a sensor monitoring the available network bandwidth may return a slightly different number every time a sample is collected. This variation could be either due to actual changes in the bandwidth or the error (noise) in the employed probes.
- *Uncertainty in Analytical Models*: Analytical models often make simplifying assumptions, and thus provide only estimates of the system's behavior. For instance, an analytical model quantifying the system's response time may account for the dominant factors, such as execution time of components, and ignore others, such as the transmission delay difference between TCP and UDP. Response time estimates provisioned by such a formulation are not only error-prone, but also the magnitude of error varies depending on the circumstances.
- *Uncertainty in User Preferences*: Eliciting user's preferences in terms of utility functions, such as those depicted in Figure 1c, is a well-known challenge [2]. Often users have difficulty expressing their preferences and thus the overall accuracy of the utility functions remains subjective, making the analysis based on them prone to uncertainty.

The uncertainty in these elements challenges the system's ability in making decisions that bring about the intended effects.

3. APPROACH

In this section, we first provide an intuitive description of our approach. Afterwards, we provide a brief background on the applicable mathematical techniques for realizing our approach, and justify our choices.

3.1 Underlying Insights

Figure 2a shows the typical behavior of a self-adaptive system that does not incorporate uncertainty in its analysis. We abstractly refer to this as the *traditional approach*. The system is initially executing with utility U_1 prior to time T_1. At time T_1, due to either an internal or external change, the system's utility drops to U_2. By time T_2, the self-adaptation logic detects this drop in utility, finds and effects an optimal configuration, which is conventionally defined as the one achieving the maximum utility. As shown in Figure 2a, this corresponds to U_3, which represents the *expected* utility of the best configuration for the system. In practice, however, the *actual* utility of the system may vary between the
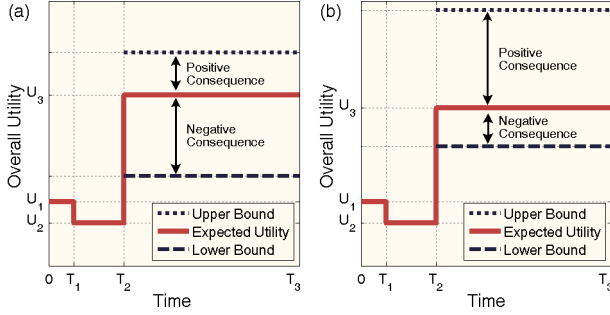
Figure 2. Impact of uncertainty on the utility of a self-adaptive system: traditional (part a) versus POISED's (part b) selection of optimal configuration.

two dashed lines, representing the likely positive and negative consequences of uncertainty. By not accounting for uncertainty, the approach is vulnerable to gross overestimation of the utility.

The centerpiece of POISED is the reconceptualization of what is traditionally considered as the optimal solution, such that the uncertainty is incorporated into the analysis. We illustrate the insights underlying POISED using Figure 2b. Similar to the scenario of Figure 2a, a new configuration is effected at time T_2, except POISED concurrently follows three objectives in selecting the new configuration: (1) maximizes U_3, which represents the most likely utility for the system under uncertainty; (2) maximizes the *positive consequence* of uncertainty, which represents the likelihood of the solution being better than U_3; and (3) minimizes the *negative consequence* of uncertainty, which represents the likelihood of the solution being worse than U_3.

The details of our approach, including how the likelihood of each objective is calculated, are described in Sections 5 and 6. Nevertheless, we can make a general observation. As depicted in Figure 2, concurrent satisfaction of the three objectives may result in a smaller value of *expected* utility (i.e., U_3) in POISED compared to that of the traditional approach. But since the information used to estimate the expected utility is uncertain, *expected* utility is not guaranteed to occur in practice. We argue the true quality of a solution is determined by the <u>range</u> of possible utility. As depicted here and evaluated in Section 7, POISED's objective is to find solutions with a better range.

3.2 Analysis Technique

The above discussion assumes one could quantify the range of utility under uncertainty. There are two general techniques to estimating uncertainty: *probability theory* and *possibility theory*. Probability theory is concerned with the analysis of random phenomena and forms the foundation of statistics. Possibility theory is founded on the concept of fuzzy set [30]. In a fuzzy set, the elements have a degree of membership. Degree of membership is a value between zero and one: a value of zero indicates the element is certainly not a member of the set, a value of one indicates the element is certainly a member of the set, and a value in between indicates the extent of certainty that the element is a member of the set. In possibility theory, the concept of *possibility* is defined as the degree of membership, which plays a similar role as that of probability in statistics. The *most optimistic* and *most pessimistic* values have a degree of membership of zero, while the *most possible* value has a degree of membership of one. The distinction between the two concepts is that *Probability*

theory deals with the statistical characteristic of data, while *possibility theory* focuses on the meaning of data [30].

For an illustration of this difference, see Figures 3 and 4 in which the uncertainty in available network bandwidth is modeled using possibility and probability distributions, respectively. The details of how such functions can be obtained are described in Section 5. Possibility distribution models a *fuzzy variable*, while probability distribution models a *random variable*. Possibility distribution in Figure 3 returns the *degree of membership*, while probability distribution in Figure 4 returns the *probability density*.

While uncertainty can be represented using both approaches, the ability to make adaptation decisions is significantly impacted by the representation. The operations research technique for making decisions under probability theory is called *stochastic programming* [26]. When all the probability distributions follow normal distribution, by applying the central limit theorem [13], the problem can be solved effectively [14]. Central limit theorem allows for the definition of commonly used algebraic operations on the normally distributed variables. However, as further demonstrated in this paper, uncertainty in self-adaptive software systems rarely follows a normal distribution, which severely limits the application of stochastic programming.

An alternative approach that does not suffer from these limitations is *possibilistic programming*, which is founded on possibility theory. Possibilistic programming is widely used for making decisions under uncertainty in many fields of engineering, including control theory, robotics, and artificial intelligence. Advantages of possibilistic programming are twofold: generality and efficiency. Unlike stochastic programming, where simple algebraic operations require special considerations (e.g., central limit theorem), and not even always possible, in possibilistic programming, fuzzy variables can be simply operated on using traditional algebraic operators. This is true even if the distributions are complex and unique. Possibilistic programming problems can be solved much more efficiently than stochastic programming problems [14], making them desirable in many practical engineering problems, including self-adaptive systems. For these reasons we have adopted possibilistic programming as the technique for achieving POISED's objectives.

While the characteristics of our problem made possibility theory a better fit, probability theory has shown to be suitable in many areas of software engineering. Thus, the conclusions made here are only with respect to POISED's objectives. A more comprehensive analysis of the differences between the two theories can be found in the extensive body of literature [5, 7].

4. SELF-ADAPTATION PROBLEM

In this section, we provide a formal specification of the self-adaptation problem introduced in Section 2, which we use to demonstrate and evaluate POISED. However, note that the underlying concepts and techniques in POISED are generally applicable and not tied to the problem described here.

4.1 Configuration

A system like the one depicted in Figure 1b consists of several software components, which we denote as set C. Each component $c \in C$ may have several configuration *dimensions*, which we denote as set D_c. Configuration dimensions correspond to the knobs depicted in Figure 1a. Each configuration dimension $d \in D_c$ may have A_d configuration *alternatives*. For example, *Video Quality* dimension of the *Camera* component in Figure 1 is

comprised of the following alternatives: *60%, 70%, 80%, 90%,* and *95%*. Configuration alternatives within the same dimension are mutually exclusive, e.g., *Video Quality* could be exactly in one of the 5 possible alternatives at any point in time.

We define the configuration space of component $c \in C$ as the Cartesian product of all the available configuration alternatives for that component:

$$ConfSpace_c \stackrel{\text{def}}{=} \otimes_{d \in D_c} \otimes_{a \in A_d} dom(x_{c,d,a})$$

Where x represents a decision variable with a binary domain and indicates whether an alternative has been selected or not.

A system may have several execution *scenarios* denoted as list $S \stackrel{\text{def}}{=} \langle s_1, s_2, \ldots, s_n \rangle$, where each scenario $s_i \subseteq C$. A scenario represents a high-level functional capability involving the services provided by software components. For example, the components bolded in Figure 1b form a scenario dealing with the robot maneuvering a terrain. Same components can potentially form different scenarios, thus we represent execution scenarios as a list.

Configuration space of each scenario is the Cartesian product of the configuration space of all the components that participate in it:

$$ConfSpace_s \stackrel{\text{def}}{=} \otimes_{c \in s} ConfSpace_c$$

We use $ConfSpace$ with no subscript to denote the set of all possible configurations of components in a system (i.e., $s = C$):

$$ConfSpace \stackrel{\text{def}}{=} \otimes_{c \in C} ConfSpace_c$$

4.2 Configuration Constraint
Not all configurations for a software system are valid. There may be some constraints among the system's configuration dimensions. For example, one may not be able to configure the robot in a way that *Cruise Speed* is set to *High* and *Power* is set to *Energy Saving*. Another class of constraints are among the alternatives in a given configuration dimension, i.e., the fact that configuration alternatives are mutually exclusive (recall Section 4.1). We represent these constraints as follows:

$$ConfConstraints: ConfSpace \rightarrow \{0,1\}$$

Where given a configuration, $ConfConstraints$ returns 0, if at least one constraint is violated, and 1, otherwise.

4.3 Quality Attribute
We use Q to denote the set of quality attributes, which are quantifiable non-functional properties of interest (e.g., response time). A quality attribute may take either discrete or continuous values, e.g., response time may take continuous values bigger than 0, while security may take an enumeration of discrete values.

The quality attributes of an execution scenario are determined by the configuration of components participating in that scenario. For example, the response time of the *Maneuver* scenario depicted in Figure 1b is affected by the configuration of its four components. Given a configuration of a scenario, a quality attribute is estimated via an analytical formula (model). These analytical formulas are used by the adaptation logic for making decisions. We represent an analytical formula estimating the configuration's impact on quality attribute $q \in Q$ of execution scenario $s \in S$ as:

$$\widetilde{QE_{s,q}}: ConfSpace_s \rightarrow dom(q)$$

The tilde is the conventional notation for representing uncertainty. Description of analytical models for estimating quality attributes is beyond the scope of this paper. Numerous previous studies (e.g., [6, 9, 12, 21]) have developed analytical approaches for estimating quality attributes in terms of the system's architectural configuration. Regardless of the approach, the analytical models provide only estimates, and thus represent a source of uncertainty.

4.4 Resource
We use set R to denote the different computing resources utilized by the software system. For each resource $r \in R$, we use $\widetilde{Capacity_r}$ to represent the maximum available resource. While in some cases the available resource is a known constant (e.g., physical memory), in others the available resource may fluctuate (e.g., network bandwidth), and thus introduce uncertainty.

The configuration of a system determines its resource usage. For example, consider that in the robotics system, when the *Power* dimension of the *Controller* component is configured to operate at *Full Power*, the system's battery consumption increases. We represent an analytical formula estimating the configuration's impact on the system's resource $r \in R$ as follows:

$$\widetilde{RE_r}: ConfSpace \rightarrow dom(r)$$

Numerous previous studies (e.g., [23-25]) have developed resource usage models that can be used in such setting. While sophisticated models may reduce the inaccuracy, they are not ever completely free of it, challenging the confidence with which adaptation decisions are made.

4.5 User Preference
Similar to the previous research [9, 23, 24, 27], we use utility functions to represent the user's preferences for changes in the quality attributes. A utility function representing the user's satisfaction with quality attribute $q \in Q$ of an execution scenario $s \in S$ is represented as:

$$\widetilde{UP_{s,q}}: ran(\widetilde{QE_{s,q}}) \rightarrow [0,1]$$

A higher value indicates more user satisfaction with the system. Given a vector $\overrightarrow{cnf} \in ConfSpace$, we define the *overall utility* $\widetilde{U}$ to be the cumulative satisfaction of all the user preferences:

$$\widetilde{U} \stackrel{\text{def}}{=} \sum_{\forall s \in S} \sum_{\forall q \in Q} \widetilde{UP_{s,q}} \left(\widetilde{QE_{s,q}} \left(\overrightarrow{cnf_s} \right) \right) \qquad \text{Eq. 1}$$

Where $\overrightarrow{cnf_s}$ is defined as the *projection* of $\overrightarrow{cnf}$ from $ConfSpace$ onto $ConfSpace_s$: $\overrightarrow{cnf_s} = proj_{\overrightarrow{I_s}}(\overrightarrow{cnf})$, and $\overrightarrow{I_s}$ is the *identity* vector for $ConfSpace_s$. In other words, since not every component participates in every execution scenario of interest s, the above projection removes the unnecessary elements from vector $\overrightarrow{cnf}$ to derive $\overrightarrow{cnf_s}$. The assumption in the formulation of Eq. 1 is that if the user has not specified a preference for a quality attribute of an execution scenario, the corresponding utility function returns only zero.

4.6 Optimization Problem
The objective of the adaptation logic is to find a configuration that maximizes the system's overall utility:

$$argmax_{(\overrightarrow{cnf} \in ConfSpace)} \widetilde{U} \qquad \text{Eq. 2}$$

The solution maximizing the above objective should satisfy two constraints. First, ensure that the solution satisfies the configuration constraints:

$$ConfConstraints(\overrightarrow{cnf}) = 1 \qquad \text{Eq. 3}$$

Second, ensure that the resource usage does not exceed the available resources:

$$\forall r \in R, \quad \widetilde{RE_r}(\overrightarrow{cnf}) \leq \widetilde{Capacity_r} \qquad \text{Eq. 4}$$

For simplicity, the above formulation assumes the capacity of all resources is uncertain. But as you may recall from Section 4.4, this may not always be the case.

5. QUANTIFYING UNCERTAINTY

Before uncertainty can be dealt with in the analysis, it needs to be quantified for a given configuration of the system. POISED's accuracy depends on the ability to: (1) identify the sources of uncertainty, and (2) estimate the level of uncertainty. To put it boldly, our approach addresses "known unknowns", not "unknown unknowns". However, even if the two conditions are partially satisfied (i.e., only some sources of uncertainty are identified and estimated), POISED produces better results than the traditional approach by incorporating the known uncertainties.

In this section, we first describe two ways of estimating uncertainty: eliciting it from the stakeholders (e.g., user, engineer), and observing it in the system. These techniques are not intended to be exhaustive, or even generally applicable, but rather concrete examples to illustrate the feasibility of our work. We then describe how uncertainty in the individual elements can be combined to quantify the overall uncertainty for the system.

5.1 Eliciting Uncertainty from Stakeholders

Stakeholders often provide inputs for different facets of a self-adaptive system. One of the most crucial and commonly elicited inputs is the user's quality preferences. It is commonly agreed that eliciting user's preferences in terms of complex utility functions is challenging [2]. The specification of such utility functions is highly subjective and inevitably prone to uncertainty. Engineers may also provide inputs for certain software properties that cannot be easily monitored. For instance, the maximum memory consumed by a software component is a property that may be available from the component's source code, but not easily obtainable through runtime monitoring. Similarly, engineers may provide inputs for certain systems parameters, such as the available network bandwidth, based on a combination of past experiences, hardware specifications, similar systems, etc.

For the inputs provided by the stakeholders, it is also reasonable to ask them to estimate the range of uncertainty based on the expected level of variation in the input. For illustration, Figure 3 shows how an engineer may estimate the range of uncertainty in the network bandwidth in the form of a *triangular* possibility distribution. Possibility distribution can be modeled in different ways (e.g., Gaussian, Triangular) [32], but for simplicity we use only triangular distribution in this paper. The horizontal axis marks the network bandwidth, while the vertical axis marks the possibility (i.e., degree of membership). This distribution indicates that the range of feasible values for network bandwidth is anywhere between the *most pessimistic*, denoted with $Capacity_{bw}^{p}$, and the *most optimistic*, denoted with $Capacity_{bw}^{o}$. In a triangular distribution, as we reach the boundaries, the possibility of getting the expected value drops to 0. The *most possible* value is $Capacity_{bw}^{m}$, which always has the value of 1.

In the example of Figure 3, the engineer sets the most possible value to be the expected network bandwidth, which may be based on the engineer's past experiences, similar systems, etc. The most pessimistic value is set to 0, representing network failure, and the most optimistic value to the ideal network bandwidth, as advertised by the network provider. By connecting the most pessimistic and optimistic points with the most possible point, we arrive at a triangular possibility distribution representing the uncertainty in the network bandwidth. A similar approach could be used for eliciting and quantifying uncertainty with the other types of stakeholder provided inputs, such as utility functions.

5.2 Measuring Uncertainty via Monitoring

A self-adaptive software system often relies on monitoring to reason about changes in the execution condition. Dynamically fluctuating system parameters are one type of phenomena in this setting that are monitored. For instance, consider that while capacity of certain resources (e.g., available physical memory) may be known prior to system's deployment, other resources (e.g., available battery charge) may change at runtime, and thus would need to be sampled during the system's execution. On top of the uncertainty created by the fluctuations in the monitored phenomenon, monitoring is also impacted by the error in the sensors used for data collection. Such an error is in particular unavoidable with digital sensors that take samples of a continuous physical phenomenon. Even if the observed phenomenon is constant, the collected data may vary, due to noisy sensors.

System parameters may not be the only elements that need to be monitored. In [9] we showed that analytical models used for making self-adaptation decisions are abstractions of the system and by definition make simplifying assumptions, which if not held at runtime may make the estimates inaccurate. The inaccuracy of an analytical model could be determined by comparing its estimates (i.e., $\widetilde{QE}$, $\widetilde{RE}$) against the actual behavior of the system. This can be achieved either prior to system's deployment by benchmarking the system, or through runtime observation.

The uncertainty corresponding to any monitored phenomenon can be estimated as a probability distribution. For example, if the engineer is not able to use the technique from Section 5.1 to manually specify the uncertainty in network bandwidth, an alternative approach is to estimate it by monitoring the variations and constructing the equivalent probability distribution. Figure 4 shows the probability distribution corresponding to the data collected for variations in the network bandwidth (i.e., $Capacity_{bw}$). There are numerous approaches for deriving a *probability density function* [13] that represents the probability distribution of collected data, including Q-Q plot [13] that could be used to estimate well-known (e.g., *normal*, *beta*) distributions, and Quantile-Regression [15] that could be used to estimate arbitrary complex distributions. In our experiments, we found Q-Q plot to be sufficient, as monitoring data often follows one of the well-known distributions. For the parameters that only have

Figure 3. Triangular possibility distribution.

Figure 4. Probability distribution.

monitoring noise, we obtained approximate *normal* distribution. However, for most parameters, the distributions were skewed and best represented using *beta* distribution. As you may recall from Section 3.2, the fact that the majority of distributions were not normal was one of the primary motivations that led us to use possibilistic programming, as opposed to stochastic programming.

The majority of existing self-adaptive systems (e.g., [6, 9, 21, 24]) ignore the probability distribution of the collected data, and simply use the mean value in their analysis. By basing the analysis on the mean behavior, they essentially ignore the statistical characteristics of the data, and thus the underlying uncertainty.

5.3 Quantifying the Overall Uncertainty

From the estimates of uncertainty in the elements comprising a self-adaptive software system, we can estimate the overall uncertainty in the system's ability to satisfy its objectives (i.e., uncertainty in the overall utility). As mentioned in Section 3, for efficient and effective analysis of uncertainty, we have adopted the possibilistic model of uncertainty in POISED. However, to quantify the overall uncertainty under the possibility theory, and as a fuzzy variable, we also need the uncertainty associated with each of the elements to be expressed as a fuzzy variable.

As you may recall from Section 5.1, the approach for estimating uncertainty in the inputs provided by the stakeholder already produces fuzzy variables. Recall that a possibility distribution, such as the one depicted in Figure 3, defines a fuzzy variable. On the other hand, we need to transform the probability distribution representing the uncertainty in monitored elements to the equivalent possibility distribution.

We demonstrate this transformation via the network bandwidth example. From the probability distribution of Figure 4, we can derive the corresponding possibility distribution of Figure 3 as follows: (1) calculate the *confidence interval* [13] and *mode* [13] of probability distribution (*mode* is the value that occurs most frequently in a distribution), (2) set the most pessimistic value equal to the *low confidence limit* of the probability distribution, (3) set the most optimistic value equal to the *high confidence limit* of the probability distribution, (4) set the most possible value equal to the mode of probability distribution, (5) connect the most pessimistic and optimistic points with the most possible point to arrive at the triangular possibility distribution of the collected data. This approach could be used to derive the possibility distribution for all of the monitored elements.

The possibility distribution derived in this way is an approximation of the probability distribution. This is because the most pessimistic and optimistic values are determined based on the confidence limits, which are not the absolute pessimistic and optimistic values for a probability distribution. While in theory the two may not be identical, in practice, selecting a large confidence level (e.g., 99%) results in negligible difference.

Using the possibility distributions quantifying the uncertainty in the individual elements of our problem, we quantify the overall

uncertainty in system's ability to satisfy its overall utility (see Eq. 1) via a possibility distribution specified as follows:

$$
\begin{cases}
U^p \stackrel{\text{def}}{=} \sum_{\forall s \in S} \sum_{\forall q \in Q} UP^p_{s,q}\left(QE^p_{s,q}\left(\overrightarrow{cnf_s}\right)\right) \\
U^m \stackrel{\text{def}}{=} \sum_{\forall s \in S} \sum_{\forall q \in Q} UP^m_{s,q}\left(QE^m_{s,q}\left(\overrightarrow{cnf_s}\right)\right) \\
U^o \stackrel{\text{def}}{=} \sum_{\forall s \in S} \sum_{\forall q \in Q} UP^o_{s,q}\left(QE^o_{s,q}\left(\overrightarrow{cnf_s}\right)\right)
\end{cases}
$$

Where $\overrightarrow{cnf} \in ConfSpace$, and U^p, U^m, and U^o denote the most pessimistic, possible, and optimistic values for the overall utility of the system, respectively. The insight is that the most pessimistic overall utility U^p occurs when all of the quality estimates QE and user preferences UP are also pessimistic. Similar insight holds for U^m and U^o. For simplicity, this formulation assumes the utility functions are monotonic, which is often the case with user-specified preferences [27]. If that is not the case, by taking the derivative of the functions we could find their extremums ; we elide the details for brevity.

6. POSSIBILISTIC ANALYSIS
We now describe our approach for formulating the problem of Section 4 as a Possibilistic Linear Programming (PLP) [14] problem, and solving it using conventional solvers [22].

6.1 Possibilistic Formulation of the Problem
As you may recall from Section 3, POISED manages the uncertainty by finding a solution that has the best range of overall utility, where the range depends on the level of uncertainty in the system. Figure 5a depicts the intuition behind our approach, which involves pursuing three concurrent objectives: (1) select a configuration that maximizes $z_m \stackrel{\text{def}}{=} U^m$, (2) minimize negative consequence of uncertainty $z_p \stackrel{\text{def}}{=} |U^m - U^p|$, and (3) maximize positive consequence of uncertainty $z_o \stackrel{\text{def}}{=} |U^o - U^m|$. In essence, the objective is to find a configuration that has the highest likelihood of satisfying the user's preferences given the level of uncertainty. We rewrite our objective (Eq. 2) as a PLP as follows:

$$
\begin{cases}
argmin_{(\overrightarrow{cnf} \in ConfSpace)} z_p \\
argmax_{(\overrightarrow{cnf} \in ConfSpace)} z_m \qquad \text{Eq. 5}\\
argmax_{(\overrightarrow{cnf} \in ConfSpace)} z_o
\end{cases}
$$

Figure 5b shows the two possible scenarios that could occur as PLP searches for the optimal solution among three hypothetical configurations. Each configuration corresponds to a possibility distribution in Figure 5b, which is unlike the traditional approach, where each configuration would be a point in the utility space. The first scenario occurs when a configuration is inferior to others with respect to all objectives. For instance, in Figure 5b, *Conf_a* is inferior to *Conf_b* and *Conf_c* with respect to all three objectives. That is *Conf_a* has a larger z_p, and a smaller z_m and z_o. The second scenario occurs when there are trade-offs. For instance, *Conf_b* and *Conf_c* present a trade-off, as *Conf_b* is superior to *Conf_c* with respect to z_o and z_p, and inferior with respect to z_m. Section 6.2 describes the process through which such trade-offs can be resolved.

The next step is the formulation of constraints with uncertainty, which in our problem corresponds to Eq. 4, ensuring the resource usage does not exceed the available capacity. In Section 5, we described how the range of uncertainty in resource estimate and capacity can be quantified in the form of a triangular possibility distribution. Since we are dealing with a constraint, we would like to assess it under the worst case.

Figure 5. Possibilistic Linear Programming: (a) insight, and (b) trade-off between objectives.

Using the possibility distributions, we can reformulate Eq. 4 as follows:

$$\forall r \in R, \quad RE_r^p\left(\overrightarrow{cnf}\right) \leq Capacity_r^p \qquad \text{Eq. 6}$$

RE_r^p is the maximum resource usage, while $Capacity_r^p$ is the minimum resource capacity. While both RE_r^p and $Capacity_r^p$ represent the most pessimistic points in the corresponding possibility distributions, the two are semantically inverse of one another. Finally, as you may recall from Section 4.4, the capacity of certain resources may be a crisp value (e.g., available physical memory), in which case $Capacity_r^p = Capacity_r$.

6.2 Solving the Possibilistic Problem

The PLP problem is an instance of a multi-objective problem, and to solve it using commonly available linear programming solvers, we first need to transform it to an equivalent single-objective problem. This is necessary to allow us to reason about the objective trade-offs, such as those depicted in Figure 5b. Intuitively the transformation process entails (1) normalizing the objective functions, (2) combining the objective functions, and (3) if necessary, specifying priorities among the objectives. In this section, we describe the details of these three steps.

6.2.1 Normalizing the Objectives

Since the three PLP objectives are defined differently in terms of U, their range may not be the same. Therefore, to avoid one objective to dominate the other ones as we combine the three objectives, we first have to normalize them. We use normalizing linear membership function [17], which is a function μ that maps each objective z to a value between 0 and 1:

$$j \in \{p, m, o\}, \quad \mu_{z_j}: dom(z_j) \longrightarrow [0,1]$$

This allows us to have objective functions with the same range. However, for defining each function μ, we first need to determine the two extremums for each objective function z: given a configuration, the extremum maximizing the objective is called *Positive Ideal Solution* (*PIS*), and the one minimizing the objective is called *Negative Ideal Solution* (*NIS*). Note that the definitions of *NIS* and *PIS* are reversed when we are dealing with a minimization objective (i.e., z_p in Eq. 5).

We can obtain these values by performing the following six single objective optimizations:

$$z_p^{PIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmin} z_p, \quad z_p^{NIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmax} z_p$$

$$z_m^{PIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmax} z_m, \quad z_m^{NIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmin} z_m$$

$$z_o^{PIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmax} z_o, \quad z_o^{NIS} \overset{def}{=} \underset{(\overrightarrow{cnf} \in ConfSapce)}{argmin} z_o$$

We specify μ to return 1 for the *PIS* value, 0 for the *NIS* value, and proportionally linear between the two extremums:

$$\mu_{z_p} \begin{cases} 1 & z_p < z_p^{PIS} \\ \dfrac{z_p^{NIS} - z_p}{z_p^{NIS} - z_p^{PIS}} & z_p^{PIS} \leq z_p \\ & z_p \leq z_p^{NIS} \\ 0 & z_p > z_p^{NIS} \end{cases} \quad \mu_{z_m} \begin{cases} 1 & z_m > z_m^{PIS} \\ \dfrac{z_m - z_m^{NIS}}{z_m^{PIS} - z_m^{NIS}} & z_m^{NIS} \leq z_m \\ & z_m \leq z_m^{PIS} \\ 0 & z_m < z_m^{NIS} \end{cases}$$

Function μ_{z_o} is specified similar to μ_{z_m}. Figure 6 shows two instances of μ_{z_p} and μ_{z_m} that normalize the possible outputs of z_p and z_m, respectively. As the definitions of *NIS* and *PIS* are reversed for z_p, the normalizing function μ_{z_p} is decreasing, while


Figure 6. The normalizing linear membership functions for (a) z_p, and (b) z_m.

μ_{z_m} and μ_{z_o} are increasing. This is to transform the minimization objective (i.e., z_p in Eq. 5) into a maximization one.

6.2.2 Combining the Objectives

Given the maximization objectives with the same range, we use a well-known technique[1] to specify the single-objective problem equivalent to PLP problem of Eq. 5 as follows:

$$argmax_{(\overrightarrow{cnf} \in ConfSapce)} \varphi$$

Subject to: $\quad j \in \{p, m, o\}, \mu_{z_j} \geq \varphi$

In the above formulation, we have reformulated the objective to maximize the auxiliary decision variable $\varphi \in [0,1]$ representing the overall satisfaction with the three normalized objectives μ_{z_p}, μ_{z_m}, and μ_{z_o}. In other words, the auxiliary objective is to find a configuration that maximizes φ, which is constrained by the three normalized objectives, resulting first in their optimization.

6.2.3 Specifying Priorities

In the above formulation, the three objectives (expressed as constraints) have the same importance. But in certain domains, some of the objectives may have a higher priority. For instance, in a mission critical system, minimizing z_p may take precedence over maximizing z_o, since a solution capable of providing certain guarantees in the worst case scenario would be desirable. This may not be necessarily the case in other domains that are willing to tolerate higher risks with the potential of higher utility.

We achieve this by assigning weights w_p, w_m, and w_o to objectives μ_{z_p}, μ_{z_m}, and μ_{z_o}, respectively. The weights specify the importance of each objective, and $w_p + w_m + w_o = 1$. Thus, the final complete optimization problem, including the constraints (Eq. 3 and 6), can be formulated as follows:

$$argmax_{(\overrightarrow{cnf} \in ConfSapce)} \varphi$$

Subject to: $\quad j \in \{p, m, o\}, \quad (1 - w_j)\mu_{z_j} \geq \varphi$

$$ConfConstraints\left(\overrightarrow{cnf}\right) = 1$$

$$\forall r \in R, \quad RE_r^p\left(\overrightarrow{cnf}\right) \leq Capacity_r^p$$

To give a normalized objective μ_{z_j} more priority over others, we make the corresponding constraint (i.e., $\mu_{z_j} \geq \varphi$) more restrictive than others. For that reason, in the above formulation, we multiply each normalized objective μ_{z_j} with $1 - w_j$, where as the value of w_j increases, the related constraint becomes more restrictive.

7. EVALUATION

We have evaluated POISED on an extended version of the robotic software system that was developed in a previous project [18].

---

[1] *Canonical description of the mathematical technique for translating a maximization multi-objective problem to a maximization single-objective problem* (see [17, 31]): the problem of maximizing a set of objective functions $F = \{\mu_1, \mu_2, \cdots, \mu_n\}$, where all functions have the same range, is equivalent to the maximization of auxiliary variable φ, where $\forall \mu_i \in F, \mu_i \geq \varphi$.

The robotic software used in our experiments was comprised of 12 software components/connectors, and 50 configuration alternatives. POISED treats components and connectors the same, as they are both configurable. The self-adaptation logic was tasked with satisfying 5 user preferences in terms of utility expressed as sigmoid functions; therefore, the maximum achievable overall utility (recall Eq. 1) was 5. We used an implementation of the robotic software running on top of Prism-MW [19], which is a middleware platform with extensive support for runtime monitoring and adaptation. Interested reader can find additional details about the robotic software system in [18].

For the experiments we setup a controlled environment that allows us to create and measure the effect of uncertainty in the system. For that purpose, we used XTEAM [8], an architectural modeling, analysis, and simulation environment that has been integrated with Prism-MW [19]. Through this integration, the XTEAM models are kept in sync with the software running on top of Prism-MW, and vice versa. XTEAM can also be used to control the execution of the software running on Prism-MW. We used XTEAM to simulate uncertainty by controlling the extent of random changes in the system (e.g., available network bandwidth, memory consumption of configuration alternatives). However, neither the robotic software nor POISED was controlled, which allowed them to behave as they would in practice.

The analytical models used in our experiments were derived using the reinforcement learning technique developed in our recent work [9]. The adaptation logic was realized as a three step model interpreter engine that: (1) generates the PLP from the runtime model of the system, (2) solves it using the conventional linear programming solvers (e.g., [22]), and (3) changes the runtime model using the XTEAM's API [8], which automatically effects the changes to the software running on Prism-MW [19].

7.1 Quality Trade-Offs

We compared the quality of solutions selected by POISED with the traditional approach in 10 different experiments. In all of the experiments we applied both approaches on the same system. However, the components used in each experiment were instrumented to have different levels of uncertainty in their response time. As you may recall from Section 3, the traditional approach is representative of the majority of existing literature that ignore the uncertainty in the adaptation decisions (i.e., base the analysis on purely crisp values obtained by calculating the mean behavior of the system properties).

We performed two types of comparison: (1) For each experiment, we compared the expected quality of solutions (configurations) selected by each approach. We refer to these results as *expected*, since they are based on the calculated consequences of uncertainty in the solutions selected by POISED. (2) We then executed the software system in the selected configuration, and observed the actual quality of solution. We refer to these results as *actual*, since they are based on the data collected after the solution was put into effect.

We show the expected results in Figure 7a. The triangular possibility distribution values correspond to the solution selected by each approach. We observe a similar pattern to what was hypothesized in Figure 2. While POISED's solution may have a slightly lower *mode* [13] compared to that of the traditional approach, the overall range is always better—POISED's most pessimistic and optimistic points are higher than that of traditional approach. This is expected, since traditional approach aims to maximize the mean behavior of the system, while POISED aims to maximize the range of behavior.

We complemented the expected ranges of the system's behavior with the actual results obtained in 30 different executions of the system under each configuration. The results are shown in Figure 7b. For a fair comparison, in each experiment, we used XTEAM to fix the application workload, as well as the range of uncertainty in the execution context (e.g., network bandwidth). By "fixing the range of uncertainty" we mean controlling the range of random behavior within each source of uncertainty. Thus, different executions still resulted in different observed behaviors. We can see that the observed utilities are very closely correlated to the corresponding possibility distribution in Figure 7. The results show that in comparison to the traditional approach, POSIED is more likely to select a solution with better overall utility.

As the level of uncertainty in the system increases, it is more likely for POISED to select solutions that are better than the traditional approach. For instance, we instrumented the level of uncertainty in E4 to be higher than the level of uncertainty in E5. This can be gleaned from Figure 7a, as the utility range for the solutions selected by both POISED and traditional approach in E4 are larger than E5. Consequently, from Figure 7b we can see that the solution selected by POISED in E4 achieves more improvement over the traditional than in E5. The level of uncertainty, however, is not the only factor, as the structure of the problem also affects the amount of improvement over traditional.

For a meaningful comparison, in these experiments, we did not specify stringent resource constraints, which as shown next could significantly influence the outcome of both approaches.

7.2 Violation of Resource Constraints

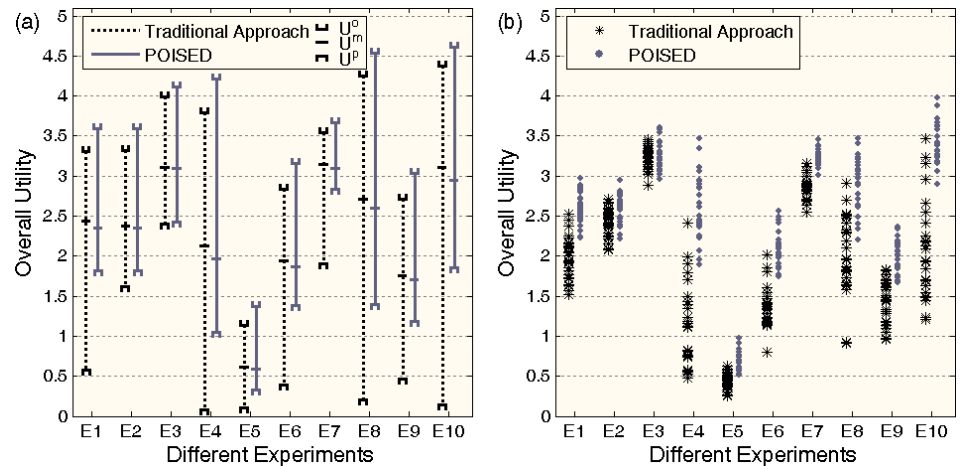We evaluated POISED's ability to satisfy the resource constraints

Figure 7. Comparison of POISED with traditional approach in 10 different experiments: (a) possibility distribution for the selected configuration, (b) 30 actual observations for each selected configuration.
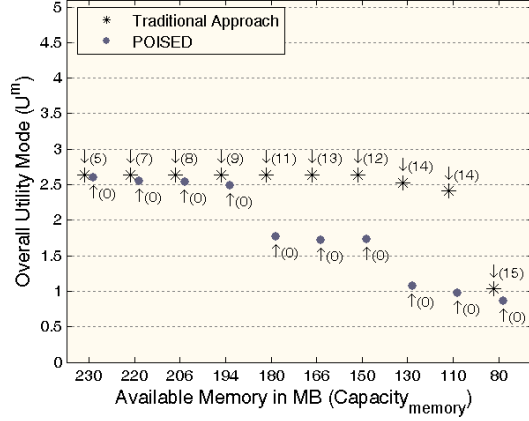
Figure 8. Impact of uncertainty on the overall utility and resource constraints. Number of violations is parenthesized.

Figure 9. Impact of weights on the selected configuration.

under uncertainty, and compared its results against the traditional approach. We ran both approaches on the same adaptation problem but with varying levels of uncertainty in the available memory. The overall utility mode corresponding to the solution selected by each approach is shown in Figure 8. The robotic software system corresponding to each selected configuration was then executed 30 times. We instrumented our controlled environment to change the memory limit used by the adaptation logic (i.e., Capacity variable from Section 4.4). Parenthesized annotations in Figure 8 show the number of times a memory violation was observed in the actual executions of the system.

We can make several observations from these results. POISED incorporates the uncertainty in the resource usage estimates, and aims to satisfy the worst case (most pessimistic) formulation of resource constraints. Therefore, as the available memory decreases, POISED continues to select solutions that do not violate the memory constraint, but naturally have a lower utility compared to those of traditional approach. On the other hand, since the traditional approach ignores the underlying uncertainty in the estimates, as the available memory decreases, the likelihood of selecting configurations that would violate the memory constraint increases. This pattern persists until the available memory decreases to 80MB, which is less than the mean of the memory usage estimate for configurations with high utility. The traditional approach is thus forced to select configurations with a relatively low utility. But even then, since it does not consider uncertainty, 15 executions violate the memory constraint.

7.3 Effect of Weights

In previous experiments, we placed the same weight on all objectives (i.e., $w_p = w_m = w_o = \frac{1}{3}$). But as you may recall from Section 6.2.3, this may not always be the case. We evaluated POISED's sensitivity to these weights on an instance of the robotic software. For a meaningful comparison, with the exception of weights, all other attributes of the system were fixed, including the range of uncertainty. Figure 9 shows the overall utility for the experiments. The solid bar shows the possibility distribution corresponding to the selected configuration under each weight assignment. The dots depict the observed utility for 30 executions of the software in the selected configuration.

In the two experiments with high w_p, POISED selects a conservative solution, i.e., puts more emphasis on minimizing the
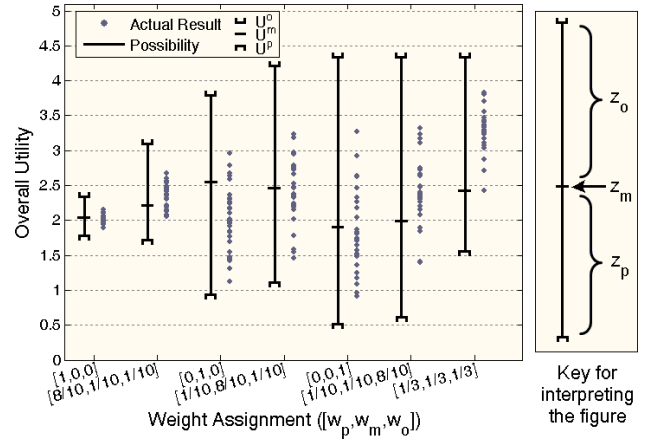
negative consequence of uncertainty (z_p from Section 6.1). In the two experiments with high w_o, POISED selects a risky solution, i.e., puts more emphasis on maximizing the positive consequence of uncertainty (z_o from Section 6.1). Both approaches come at the cost of achieving mediocre overall utility mode (z_m from Section 6.1). In the two experiments with high w_m, POISED selects a solution with the best utility mode (z_m from Section 6.1), while ignoring the negative and positive consequences of uncertainty. In the last experiment, with a balanced assignment of weights, the solution achieves neither the best U^m, nor does it provide guarantees on the consequences of uncertainty. But since all of the objectives are given the same weight, it achieves a balanced set of trade-offs. However, as mentioned earlier, we can envision situations in which placing emphasis on one of the objectives may be more appropriate, which POISED allows for naturally.

7.4 Sensitivity to Uncertainty Estimates

We performed a set of experiments to evaluate how sensitive is POISED to the accuracy of uncertainty estimates. Figure 10 shows the results. For all of the experiments we used XTEAM to fix the range of uncertainty in the system parameters, as well as the workload. We changed the accuracy of uncertainty estimates used in our analysis. To that end, we simply changed the confidence level used for transforming the probability distribution corresponding to the monitored data to the equivalent possibility distribution (recall Section 5.3). As one decreases the confidence
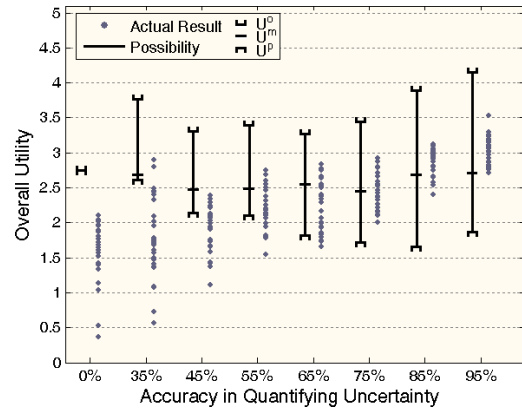

Figure 10. Impact of the accuracy of uncertainty estimates on the quality of POISED solutions.

level in a probability distribution, such as the one depicted in Figure 4, *low confidence* and *high confidence limits* converge to *mode*, resulting in underestimation of the range of uncertainty.

The confidence levels shown on the horizontal axis of Figure 10 denote the accuracy of uncertainty estimates. As we decrease the confidence level from 95% to 0%, thereby making the uncertainty estimates less accurate, POISED selects configurations for which overall utilities are not borne out in practice. More specifically, since by decreasing the confidence level we underestimate the uncertainty, the actual results are lower than the expected utility. Overestimating uncertainty would have the opposite effect.

Finally, in the experiment with 0% accuracy, the most pessimistic, possible, and optimistic points are overlapping. By not considering the uncertainty, POISED is behaving the same as the traditional approach. Comparing the results of experiment with 0% accuracy to others, corroborates our assertions in Section 5 that even with partially accurate estimates of uncertainty, POISED selects solutions that are better than the traditional approach.

7.5 Performance Trade-Offs

Table 1 shows the results of benchmarks comparing the execution time of traditional approach with that of POISED. It took POISED longer to compute the optimal solution, which is not surprising, since as you may recall from Section 6.2, POISED requires 6 additional optimizations to calculate *PIS* and *NIS* values pairs for the three objective functions. While it takes longer to execute POISED, we found its performance to be reasonable for our problem. It took 4.6 seconds to find the optimal configuration in the largest problem, consisting of 100 components and 20 different configuration alternatives, for a total of $20^{100} = 1.2 \times 10^{130}$ possible combinations.

8. RELATED WORK

The literature in this area of research is extensive. In lieu of enumerating all of the related studies, we refer the reader to [2, 16] for a comprehensive analysis of the state-of-the-art in self-adaptation. We focus our discussion here to those works that are of utmost relevance. The challenge posed by uncertainty in the construction of dependable self-adaptive software system is an established concept [2, 10]. A few works [1, 3, 4, 9, 23, 28] have aimed to tackle the different facets of this challenge as follows.

Whittle et al. [28] introduced RELAX, a formal requirements specification language that relies on Fuzzy Branching Temporal Logic to specify the uncertain requirements in self-adaptive systems. In a subsequent publication [1], Cheng et al. extended RELAX with goal modeling to specify the uncertainty in the objectives. Our approach is complementary to their work, as both RELAX and POISED are based on fuzzy mathematics, except RELAX targets requirements specification phase, while POISED targets decision making phase at runtime.

Chuang and Chan [4] presented a QoS management framework that uses a hierarchical fuzzy control model. Their work aims at making it easier for the users to specify their QoS requirements, which are then translated into fuzzy rules. Unlike POISED, their objective is QoS rule satisfaction, and does not target the challenge of making adaptation decisions under uncertainty.

Dynamic configuration of resource-aware services was studied by Poladian et al. [24], where they showed how to select an appropriate set of services to carry out a user task, and allocate resources among those services at runtime. Subsequently, the

work was extended to make anticipatory decisions [23], and considered the inaccuracy of future resource usage predictions. Unlike POISED, their approach neither aims to satisfy a utility range, nor employs possibilistic analysis in incorporating the effect of uncertainty in decisions.

Cheng and Garlan [3] described three specific sources of uncertainty (*problem-state identification*, *strategy selection*, and *strategy outcome*) in self-adaptation and provided high-level guidelines for mitigating them in Rainbow [11]. In this paper, we have presented a novel approach for tackling the challenge of *strategy outcome*, i.e., the impact of uncertainty on the selected solution, and techniques to deal with it.

Finally, in our recent work [9], we presented FUSION, a learning based approach to engineering self-adaptive systems. Instead of relying on static analytical models that are subject to wrong assumptions, FUSION uses machine learning to self-tune the adaptive behavior of the system to unanticipated changes, but does not address making adaptation decisions under uncertainty.

9. CONCLUSION AND FUTURE WORK

We presented a novel quantitative approach, called POISED, for making adaptation decisions under uncertainty. Unlike any other related work, POISED adopts a possibilistic method to assess the positive and negative consequences of uncertainty in its analysis. The centerpiece of our work is the reconceptualization of what is typically considered to be the optimal solution as one that has the best range of possible behavior. POISED's analysis can be made as risk-averse as desired via a set of knobs (weights). While POISED is a general approach that can be applied to many types of adaptation problems, it was described and extensively evaluated in the context of a self-adaptation problem aimed at improving a system's quality attributes via runtime reconfiguration of its customizable software components.

Our focus so far has been on the *internal uncertainty*, which is the uncertainty associated with adaptation decisions aimed at satisfying the system's internal quality objectives. In future, we plan to investigate applicability of POISED to *external uncertainty*, which is the uncertainty associated with decisions aimed at satisfying the domain objectives. We also believe POISED could complement the existing efforts aimed at alleviating uncertainty in other facets of self-adaptation. We envision an integration of RELAX [1, 28] with POISED to be a fruitful avenue of future work, as it would allow the traceability of uncertainty from the system's requirements specification to its execution. Proactively adaptive systems, such as those described in [6, 23], face another form of uncertainty—the inaccuracy of future predictions. Investigating the synergy between POISED and such approaches is another interesting avenue of future work.

10. ACKNOWLEDGMENTS

Table 1. Execution time of POISED versus traditional approach.

Problem		Exec. Time (ms)	
# of Comp	# of Conf	Tradit-ional	POISED
8	4	2	30
10	5	6	51
18	7	10	70
25	8	20	180
37	9	28	298
50	10	30	370
62	13	60	630
75	15	130	1520
88	17	290	3740
100	20	400	4600

11. REFERENCES

[1] Cheng, B.H., Sawyer, P., Bencomo, N. and Whittle, J. 2009. A Goal-Based Modeling Approach to Develop Requirements of an Adaptive System with Environmental Uncertainty. *Int'l Conf. on Model Driven Engineering Languages and Systems* (Denver, Colorado, Oct. 2009), 468-483.

[2] Cheng, B. et al. 2009. Software Engineering for Self-Adaptive Systems: A Research Roadmap. *Software Engineering for Self-Adaptive Systems, LNCS Hot Topics*. 1-26.

[3] Cheng, S.W. and Garlan, D. 2007. Handling uncertainty in autonomic systems. *Int'l Wrkshp. on Living with Uncertainty* (Atlanta, Georgia, Nov. 2007).

[4] Chuang, S.-N. and Chan, A.T.S. 2008. Dynamic QoS Adaptation for Mobile Middleware. *IEEE Transactions on Software Engineering*. 34, 6 (Dec. 2008), 738-752.

[5] Coletti, G. and Scozzafava, R. 2004. Conditional probability, fuzzy sets, and possibility: a unifying view. *Fuzzy Sets and Systems*. 144, 1 (May. 2004), 227-249.

[6] Cooray, D., Malek, S., Roshandel, R. and Kilgore, D. 2010. RESISTing Reliability Degradation through Proactive Reconfiguration. *Int'l Conf. on Automated Software Engineering* (Antwerp, Belgium, Sep. 2010).

[7] Dubois, D. and Prade, H. 2001. Possibility Theory, Probability Theory and Multiple-Valued Logics: A Clarification. *Annals of Mathematics and Artificial Intelligence*. 32, (Aug. 2001), 35–66.

[8] Edwards, G., Malek, S. and Medvidovic, N. 2007. Scenario-Driven Dynamic Analysis of Distributed Architectures. *Int'l Conf. on Fundamental Approaches to Software Engineering* (Braga, Portugal, Mar. 2007), 125-139.

[9] Elkhodary, A., Esfahani, N. and Malek, S. 2010. FUSION: A Framework for Engineering Self-Tuning Self-Adaptive Software Systems. *Int'l Symp. on the Foundations of Software Engineering* (Santa Fe, New Mexico, Nov. 2010), 7-16.

[10] Garlan, D. 2010. Software Engineering in an Uncertain World. *FSE/SDP Wrkshp. on the Future of Software Engineering Research* (Santa Fe, New Mexico, Nov. 2010).

[11] Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B. and Steenkiste, P. 2004. Rainbow: Architecture-Based Self-Adaptation with Reusable Infrastructure. *IEEE Computer*. 37, 10 (Oct. 2004), 46-54.

[12] Ghezzi, C. and Tamburrelli, G. 2009. Predicting Performance Properties for Open Systems with KAMI. *Int'l Conf on the Quality of Software Architectures* (East Stroudsburg, Pensylvania, Jun. 2009), 70–85.

[13] Gibbons, J.D. and Chakraborti, S. 2003. *Nonparametric Statistical Inference (4th Edition)*. CRC Press.

[14] Inuiguchi, M. and Ramík, J. 2000. Possibilistic linear programming: a brief review of fuzzy mathematical programming and a comparison with stochastic programming in portfolio selection problem. *Fuzzy Sets Syst.* 111, 1 (Apr. 2000), 3-28.

[15] Koenker, R. 2005. *Quantile regression*. Cambridge University Press.

[16] Kramer, J. and Magee, J. 2007. Self-Managed Systems: an Architectural Challenge. *Int'l Conf. on Software Engineering* (Minneapolis, Minnesota, May. 2007), 259-268.

[17] Lai, Y.-J. and Hwang, C.-L. 1992. A new approach to some possibilistic linear programming problems. *Fuzzy Sets Syst.* 49, 2 (Jul. 1992), 121-133.

[18] Malek, S., Edwards, G., Brun, Y., Tajalli, H., Garcia, J., Krka, I., Medvidovic, N., Mikic-Rakic, M. and Sukhatme, G.S. 2010. An architecture-driven software mobility framework. *Journal of Systems and Software*. 83, 6 (Jun. 2010), 972-989.

[19] Malek, S., Mikic-Rakic, M. and Medvidovic, N. 2005. A Style-Aware Architectural Middleware for Resource-Constrained, Distributed Systems. *IEEE Trans. Softw. Eng.* 31, 3 (Mar. 2005), 256-272.

[20] Menasce, D.A., Dowdy, L.W. and Almeida, V.A.F. 2004. *Performance by Design: Computer Capacity Planning By Example*. Prentice Hall PTR.

[21] Menascé, D.A., Sousa, J.P., Malek, S. and Gomaa, H. 2010. QoS Architectural Patterns for Self-Architecting Software Systems. *Int'l Conf. on Autonomic Computing* (Washington, DC, Jun. 2010).

[22] NEOS Server for Optimization: *http://www-neos.mcs.anl.gov/*. Accessed: 2010-08-17.

[23] Poladian, V., Garlan, D., Shaw, M., Satyanarayanan, M., Schmerl, B. and Sousa, J. 2007. Leveraging Resource Prediction for Anticipatory Dynamic Configuration. *Int'l Conf. on Self-Adaptive and Self-Organizing Systems* (Boston, Massachusetts, Jul. 2007), 214-223.

[24] Poladian, V., Sousa, J.P., Garlan, D. and Shaw, M. 2004. Dynamic Configuration of Resource-Aware Services. *Int'l Conf. on Software Engineering* (Scotland, UK, May. 2004).

[25] Seo, C., Malek, S. and Medvidovic, N. 2008. Component-Level Energy Consumption Estimation for Distributed Java-Based Software Systems. *Int'l Symp. on Component Based Software Engineering* (Karlsruhe, Germany, Oct. 2008).

[26] Stochastic Programming Community Home Page: *http://www.stoprog.org/index.html?spintroduction.html*. Accessed: 2011-02-03.

[27] Walsh, W.E., Tesauro, G., Kephart, J.O. and Das, R. 2004. Utility Functions in Autonomic Systems. *Int'l Conf. on Autonomic Computing* (New York, New York, May. 2004).

[28] Whittle, J., Sawyer, P., Bencomo, N., Cheng, B.H.C. and Bruel, J.-M. 2009. RELAX: Incorporating Uncertainty into the Specification of Self-Adaptive Systems. *Int'l Requirements Engineering Conf.* (Atlanta, Georgia, Sep. 2009), 79-88.

[29] Zadeh, L.A. 1965. Fuzzy sets. *Information and control*. 8, 3 (Jun. 1965), 338–353.

[30] Zadeh, L.A. 1999. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets Syst.* 100, (Jun. 1999), 9-34.

[31] Zimmermann, H.-J. 1978. Fuzzy programming and linear programming with several objective functions. *Fuzzy Sets and Systems*. 1, 1 (Jan. 1978), 45-55.

[32] Zimmermann, H.-J. 2001. *Fuzzy Set Theory and its Applications (4th Edition)*. Springer.