

The Requirements Problem for Adaptive Systems

IVAN J. JURETA, Fonds de la Recherche Scientifique – FNRS and University of Namur
 ALEXANDER BORGIDA, Rutgers University
 NEIL A. ERNST, University of British Columbia
 JOHN MYLOPOULOS, University of Trento

Requirements Engineering (RE) focuses on eliciting, modelling, and analysing the requirements and environment of a system-to-be in order to design its specification. The design of the specification, known as the Requirements Problem (RP), is a complex problem solving task, as it involves, for each new system, the discovery and exploration of, and decision making in a new problem space.

A system is adaptive if it can detect deviations between its runtime behaviour and its requirements, specifically situations where its behaviour violates one or more of its requirements. Given such a deviation, an Adaptive System uses feedback mechanisms to analyse these changes and decide, with or without human intervention, how to adjust its behaviour as a result.

We are interested in defining the Requirements Problem for Adaptive Systems (RPAS). In our case, we are looking for a *configurable* specification such that whenever requirements fail to be fulfilled, the system can go through a series of adaptations that change its configuration and eventually restore fulfilment of the requirements.

From a theoretical perspective, this paper formally shows the fundamental differences between standard RE (notably Zave & Jackson [Zave and Jackson 1997]) and RE for Adaptive Systems (see the seminal work by Fickas & Feather [Fickas and Feather 1995], to Letier & van Lamsweerde [Letier and van Lamsweerde 2004], and up to Whittle et al. [Whittle et al. 2010]). The main contribution of this paper is to introduce the RPAS as a new RP class that is specific to Adaptive Systems. We relate the RPAS to RE research on the relaxation of requirements, the evaluation of their partial satisfaction, and the monitoring and control of requirements, all topics of particular interest in research on adaptive systems [de Lemos et al. 2013].

From an engineering perspective, we define a proto-framework for solving RPAS, which illustrates features needed in future frameworks for adaptive software systems.

Categories and Subject Descriptors: D.2.1 [Requirements]: Languages

General Terms: Languages, Theory

Additional Key Words and Phrases: Adaptive systems, requirements engineering, requirements problem, requirements modelling language, requirements problem for adaptive systems

ACM Reference Format:

Ivan J. Jureta, Alexander Borgida, Neil A. Ernst, and John Mylopoulos, 2014. The Requirements Problem for Adaptive Systems. *ACM Trans. Manag. Inform. Syst.* 1, 1, Article 3 (July 2014), 33 pages.
 DOI: <http://dx.doi.org/10.1145/0000000.0000000>

This work was supported in part by ERC advanced grant 267856, titled “Lucretius: Foundations for Software Evolution” <http://www.lucretius.eu>.

Author’s addresses: Ivan J. Jureta, Department of Business Administration, University of Namur, 8 rempart de la Vierge, B-5000 Namur, Belgium, ivan.jureta@unamur.be; Alexander Borgida, Department of Computer Science, Rutgers University, Piscataway, NJ 08855, USA, borgida@cs.rutgers.edu; Neil A. Ernst, Department of Computer Science, University of British Columbia, ICICS/CS Building 201-2366 Main Mall, Vancouver, B.C. V6T 1Z4, Canada, neil.ernst@gmail.com; John Mylopoulos, Information Engineering and Computer Science, University of Trento, 5 Via Sommarive, I-38123 Povo, Italy, jm@disi.unitn.it.

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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 2158-656X/2014/07-ART3 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

1.1. Uncertainty and Risk in Requirements

Requirements Engineering (RE) focuses on eliciting, modelling, and analysing the requirements and environment of a system-to-be in order to design a specification. The design of the specification, usually called the *Requirements Problem* (RP), is a complex problem solving task, as it involves the discovery and exploration of, and decision making in, a large problem space. Unsurprisingly, it has been repeatedly found through experience that it is hard to get the requirements right, as “it is really impossible for a client, even working with a software engineer, to specify completely, precisely, and correctly the exact requirements of a modern software product before trying some versions of the product” [Brooks 1986].

A fundamental difficulty in solving the RP is that requirements often reflect engineers’ (be they software engineers, product designers, or business analysts) interpretations of future expectations and needs of stakeholders.

Uncertainty is therefore an essential property of requirements. So is risk: uncertainty is usually quantified by probability, and risk as the sum of the products of the probability of each undesirable outcome and its corresponding cost.

It is impossible to eliminate all requirements risks at design-time, when the specification of the first, or any subsequent release is produced. It is impossible because the system-to-be is situated in a complex environment of societal rules, commitments, incentives, laws, contracts, organisations, and, to borrow again from Brooks, “[t]hese all change continually, and their changes inexorably force change upon the software product” [Brooks 1986].

1.2. Adaptive Systems as a Response to Requirements Risk

While there is no silver bullet for requirements uncertainty, there are strategies to mitigate risk. Design-time, for example, may involve prototypes and frequent releases, collecting feedback, then taking that feedback into account for the design of each next release.

A risk management strategy that has been steadily gaining attention is the engineering of *Adaptive Systems* [Cheng et al. 2009; de Lemos et al. 2013]. A system is adaptive if it can detect differences between its requirements and runtime performance, and can adjust its behaviour in order to cope with such deviations.

Adaptive Systems engineering, when seen as a risk management strategy is one that spans design-time and run-time. Design-time, because design decisions influence the range of monitored inputs for the system, and the feedback mechanisms it will have. Run-time, because these mechanisms enable the system to react to at least some changes, rather than ignore them.

1.3. Exactly Why and How is Requirements Engineering for Adaptive Systems Non-Standard?

It has been argued that RE for adaptive software systems is non-standard [Cheng et al. 2009; Whittle et al. 2010; de Lemos et al. 2013], because it is impossible to know at design-time all possible changes in requirements and environment conditions, and therefore build into the specification all the best ways for it to respond at run-time. Although some changes can be anticipated at design-time, many will only be identified and understood when, or after they happen.

The observation that we are not very good at predicting the future, can, however, be made independently of Adaptive Systems. It applies for RE for all kinds of systems. The interesting question for research for RE on Adaptive Systems is therefore this:

Exactly why and how is Requirements Engineering for Adaptive Systems different from standard Requirements Engineering?

It is interesting, because an answer will suggest what can be retained, discarded from, or added to standard RE modelling languages, frameworks, and methodologies to address Adaptive Systems RE, that is, to help design Adaptive Systems.

The aim of this paper is to give a formal answer to this question, and thereby help connect existing research, inform future work, and move closer to a deeper classification of research in Adaptive Systems RE. It might also help skeptics decide if/how novel is the problem posed by Adaptive Systems RE.

1.4. Contributions: A New Requirements Problem

To answer the central question in this paper, we start from the *de facto* standard RP from Zave & Jackson [Zave and Jackson 1997], then illustrate and discuss its differences from the *Requirements Problem for Adaptive Systems* (RPAS hereafter).

We then introduce a formalism which is just expressive enough to be useful in solving RPAS, includes concepts and relations already appearing in mainstream RE, and includes features central in RE for Adaptive Systems. These features are:

- Monitoring and control of requirements [Fickas and Feather 1995; Feather et al. 1998; Robinson 2006]), i.e., the idea that satisfaction of requirements should be monitored, and that failure to reach required satisfaction levels should trigger changes in system behavior;
- Probabilistic relaxation of requirements [Letier and van Lamsweerde 2004], i.e., instead of requiring that some requirement be satisfied to the same extent all of the time, probabilistic relaxation weakens such requirements so that they can be satisfied some of the time, to some desired frequency;
- Fuzzy relaxation of requirements [Whittle et al. 2010; Baresi et al. 2010]), where black-and-white satisfaction of a requirement is replaced with satisfaction to some level on a continuous scale defined by a fuzzy membership function;
- Evolution requirements [Souza et al. 2012] that concern how to change the system when some other requirements succeed or fail.

It turns out that extending the simple, but abstract modelling language *Techne* [Jureta et al. 2010] to allow constraints over quantitative (numerical) variables is enough to capture key ideas in monitoring, relaxation, and adaptation, and to show how they contribute to the RPAS.

This paper presents two main contributions:

- The RPAS and its solution concept, called *Configurable Specification*, are formally defined. A Configurable Specification amounts to a set of requirements configurations and evolution requirements for switching between configurations. Each configuration is shown to satisfy all properties required by Zave & Jackson [Zave and Jackson 1997], and in our prior work [Jureta et al. 2008]. These parallels show exactly how the requirements problem and Configurable Specification concepts are different from more traditional RE.
- To introduce the RPAS and Configurable Specification concepts, we define a simple formalism in which we can model (i) probabilistic and fuzzy relaxation, (ii) monitored and controlled variables, (iii) evolution requirements. It is a *proto-framework* because (a) it is impractical in its current form, but is enough to allow us to present and discuss the salient properties of the RPAS and Configurable Specification concepts, and (b) it can serve as a prototype for early formal modelling languages for RE of Adaptive Systems.

1.5. Organisation

This paper is organised as follows. Section 2 gives an informal introduction to, and the definition of the RPAS. Section 3 contrasts RPAS to the the standard RP in RE; there, we use a simple example to discuss how the RPAS departs from the standard RP. Section 4 introduces informally a modelling framework, which we use in the rest of the paper to model instances of the RPAS. Section 5 recalls key ideas in the existing RE research on Adaptive System, and thereby identifies the notions that have to be accommodated in the RPAS, its solution concept, and the modelling framework. Section 6 focuses on the solution concept, the Configurable Specification, and defines its components. We then revisit in Section 7 the question of how exactly RE for Adaptive System differs from standard RE. We summarise conclusions and open issues in Section 8. This paper has electronic appendices, where we formally define the modelling framework, and explain how to use it to capture the key notions from related work on RE for Adaptive System.

2. THE REQUIREMENTS PROBLEM FOR ADAPTIVE SYSTEMS

2.1. Premises for the Definition of the Problem

The Requirements Problem for Adaptive Systems is based on the premise that *the overall aim when designing an Adaptive System, is to make sure that the system satisfies its stakeholders' requirements as best as feasible over time, as its environment conditions change and / or its components fail.*

In order to adapt, the Adaptive System has to gather data about events in its operating environment and about the functioning of its own components. At all times, and on the basis of these observations, the Adaptive System has to estimate the level to which it satisfies the stakeholders' requirements. If the levels of satisfaction are inadequate, the Adaptive System has to make changes to its operational configuration in order to satisfy requirements.

This leads to key observations about the *run-time* of Adaptive System. (i) The level at which requirements are satisfied will vary, due to failures in the system and changes in its environment. (ii) It is necessary to monitor the level of satisfaction, in order to know when the system needs to adapt. (iii) When the system adapts, it may have different ways of adapting, and each of these ways may have a different impact on requirements satisfaction levels. (iv) Whenever it needs to adapt, the system should adapt in a way that optimises levels of requirements satisfaction, relative to the newly observed failure of a component, or of a change in the environment.

The observations about run-time have important implications for the *design-time* of Adaptive Systems. Due to observation (i), it may be too idealistic and impractical to think of requirements as being either satisfied or not, since this may lead to too many failed requirements, too often. It can be more practical, therefore, to define multivalued scales of requirements satisfaction, where failure equates to only some of many values. This is done through the **relaxation of requirements**, where the idea is to replace binary levels of satisfaction with, for example, continuous scales of satisfaction, or by letting the requirement be binary, but tracking the frequency of them being satisfied or failing, then using that frequency as the measure of the degree to which these requirements are satisfied. Observation (ii) suggests that it is necessary, at design-time, to define the levels of requirements satisfaction that trigger adaptation. If the requirement has a binary satisfaction scale, being either satisfied or not, it may be relevant to define the minimal probability of observing its satisfaction; for example, in an ambulance dispatch system, asking for the probability of at least 0.95, that an ambulance arrives to an incident location within 14 minutes of being dispatched to it. This would translate, at run-time, into looking at the frequency of incidents where the

ambulance arrived 14 minutes or more, and triggering adaptation if that frequency is 5% or more of all incidents to which an ambulance was dispatched. If the requirement has a scale with many levels of satisfaction, then a threshold value has to be defined on that scale, such that, when the satisfaction is below threshold, the system needs to adapt. This has led to research on **awareness requirements**, which are used to define these thresholds for triggering adaptation.

At run-time, when awareness requirements are satisfied, feedback loops become active, and the system adapts. Because of observations (iii) and (iv), it is necessary to define at design-time the requirements that the system should satisfy when adapting. These are the so-called **evolution requirements** [Souza et al. 2012], and place constraints on how the system adapts. In the terminology of research on the RE for Adaptive Systems, evolution requirements place constraints on the range of **reconciliation tactics** that the system may choose to apply, when adapting. For example, if the requirement to compute ambulance location fails, one evolution requirement could be to try to satisfy instead the requirement that assistants should compute ambulance location manually. Another evolution requirement could be to try to compute ambulance location from ambulance's onboard record.

2.2. Adaptation as Switching between Configurations

Given the observations and implications discussed in the preceding section, the run-time of an Adaptive System looks like a sequence of time periods of two kinds: stability, when no awareness requirement is violated, and adaptation, which lasts while awareness requirements are violated and the system is moving through unstable states towards another stable period. Evolution requirements define the transitions that take place during adaptation periods.

This leads to the following terminology in this paper. A **configuration** describes the tasks that the system should execute, and the environment conditions within which these tasks can be executed. If these conditions change, the same tasks may no longer be feasible. A configuration describes system responsibilities and its environmental assumptions.

The difference between the notion of configuration and the usual notion of specification, is that when we talk about configurations, we assume that there is a so-called Configurable Specification, and that one configuration amounts to one set of values of all the parameters of that specification. By assigning a value to each configurable parameter of the specification, we are choosing one of its possible configurations.

During adaptation the system **switches** from one configuration to another, and does so because awareness requirements became violated during the last period of stability. When switching from one configuration to another, that is, when changing the parameters of the specification, the system must satisfy the evolution requirements; it has to choose the new set of tasks by taking into account the environment conditions, the requirements to satisfy, and the evolution requirements.

An individual configuration has to satisfy a number of properties; we mention them here, and define them formally in Section 6. It has to be consistent, so that executing tasks does not make it impossible to execute others in the same configuration, or to violate the environment conditions. It must satisfy all mandatory requirements to their threshold levels, as it would otherwise not be acceptable to stakeholders. It has to include only those tasks that are necessary and sufficient to reach thresholds for the satisfaction of the requirements, and no unused tasks. Finally, it has to be Pareto efficient with regards to optional requirements; these requirements are optional in the sense of being nice-to-have, but if they fail, the system still remains acceptable to the stakeholders. If there were two configurations that the system could choose from, to exit an adaptation period, and if their only difference is that one satisfies more desir-

able optional requirements, then the system choosing the other would not be Pareto efficient with regards to optional requirements.

2.3. Configurable Specifications

An individual configuration, or in other words a non-parameterizable specification, cannot be the solution to the RPAS. This is because a single configuration is only applicable in the environment conditions that it describes, and when the system is capable of executing those tasks that it describes. The single configuration says nothing about the system when these environment conditions change, or when some of its capabilities fail, making it unable to execute the tasks that the configuration prescribes.

At run-time, the Adaptive System will be switching between configurations, and should do so in such a way that it satisfies the evolution requirements. The Adaptive System will at run-time be realizing a specific sequence of configurations.

A specific sequence of configurations is not the solution to the RPAS, because it is not possible to predict how the environment conditions will vary over time, and when exactly the system will fail, and which tasks it will fail to execute. In other words, we cannot at design-time predict the best sequence of configurations to follow at run-time.

Moreover, a motivation for enabling adaptivity, is to delegate to the system at least part of the decision-making involved in choosing which configuration to switch to. The solution should recognize that there can be varying degrees of autonomy with which the system will choose the next configuration, when the current one fails.

The discussion above led us to introduce the concept of **Configurable Specification**, as the solution concept for the RPAS.

Definition 2.1. Configurable Specification: A Configurable Specification is a pair $\mathcal{R} = (\mathcal{S}, \mathcal{E})$, where $\mathcal{S} = \{S_1, \dots, S_n\}$ is a set of n configurations, and \mathcal{E} is a set of evolution requirements.

The solution to the RPAS includes no concept, which is specifically useful in describing how exactly adaptation should be done. Evolution requirements impose constraints on this, and if needed, feedback loops can be defined in tasks, in configurations.

2.4. The Requirements Problem for Adaptive Systems

The Requirements Problem for Adaptive Systems is the problem of designing a **Configurable Specification** which can satisfy the given requirements to the highest feasible level, for the predicted failures and environment changes.

Definition 2.2. The Requirements Problem for Adaptive Systems (RPAS): **Given** (i) stakeholders' requirements, (ii) stakeholders' preferences over different levels of requirements satisfaction, and (iii) domain assumptions about the operating environment and its expected changes at run-time, **find** the Configurable Specification such that, if the system-to-be is implemented so as to be capable of realizing the configurations and satisfy the evolution requirements in the Configurable Specification, then it is expected at run-time to satisfy the requirements to their most desirable feasible levels.

The RPAS does not commit to what it means exactly that requirements are satisfied to their most desirable feasible levels; this will depend on the preferences over requirements, and the specifics of scales of satisfaction used to quantify the degree of satisfaction. Also note that configurations might share the same tasks, yet these tasks could be parametrised differently in different configurations.

3. FROM THE STANDARD REQUIREMENTS PROBLEM TO THE REQUIREMENTS PROBLEM FOR ADAPTIVE SYSTEMS

3.1. The Standard Requirements Problem

The standard view in RE, is that the system specification is produced incrementally, starting from requirements and domain knowledge. The latter captures properties of the environment, where the system will run. The goal is to produce a specification which, together with the given domain knowledge, satisfies requirements, and is consistent with both requirements and domain knowledge.

This important and general conceptualization of the aim in RE is most clearly formulated in Zave & Jackson's seminal paper, "Four dark corners of requirements engineering" [Zave and Jackson 1997]. This view, called ZJ hereafter, is echoed in some of the most influential research in the field, which both preceded and followed the said paper, including, for example, contributions from Boehm et al. [Boehm 1988; Boehm et al. 1995], van Lamsweerde et al [Dardenne et al. 1993; Darimont and van Lamsweerde 1996; van Lamsweerde et al. 1998; van Lamsweerde and Letier 2000; Letier and van Lamsweerde 2004], Mylopoulos et al. [Mylopoulos et al. 1992; Greenspan et al. 1994; Castro et al. 2002], Robinson et al. [Robinson et al. 2003], Nuseibeh et al. [Nuseibeh et al. 1994; Hunter and Nuseibeh 1998], to name some.

More specifically, Zave & Jackson suggested that, in any concrete systems engineering project, RE is successfully completed when the following five conditions are satisfied [Zave and Jackson 1997]:

- (1) *"There is a set R of requirements. Each member of R has been validated (checked informally) as acceptable to the customer, and R as a whole has been validated as expressing all the customer's desires with respect to the software development project.*
- (2) *There is a set K of statements of domain knowledge. Each member of K has been validated (checked informally) as true of the environment.*
- (3) *There is a set S of specifications. The members of S do not constrain the environment; they are not stated in terms of any unshared actions or state components; and they do not refer to the future.*
- (4) *A proof shows that $K, S \vdash R$. This proof ensures that an implementation of S will satisfy the requirements.*
- (5) *There is a proof that S and K are consistent. This ensures that the specification is internally consistent and consistent with the environment."*

The conditions above, according to Zave & Jackson, "establish minimum standards for what information should be represented in a requirements language".

If the satisfaction of these conditions marks the end of RE in any systems engineering project, then we can give the following compact formulation of the standard RP, which we call the ZJ RP hereafter.

Definition 3.1. Zave & Jackson Requirements Problem (ZJ RP): Given a set R of requirements and a set K of domain knowledge, find a specification S such that (i) there is a proof of R from K and S , written $K, S \vdash R$, and (ii) K and S are consistent.

It is useful to observe the following about the ZJ RP. Firstly, it suggests that the basic categories of information to distinguish in the RP are requirements, domain knowledge, and specification. Secondly, it imposes two conditions on the members of these categories. The first condition will be called the *Satisfaction Condition*, and the second the *Consistency Condition* in the rest of the paper.

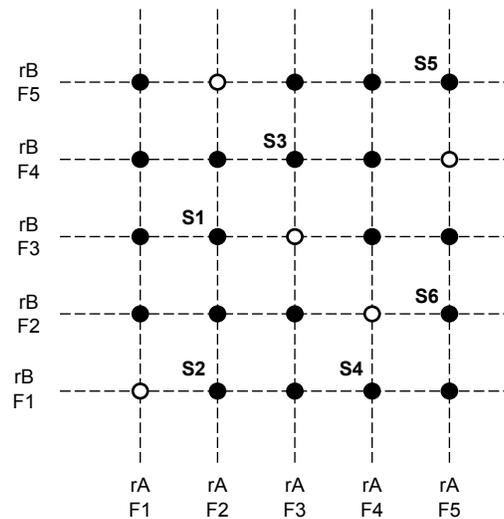


Fig. 1. rA and rB are two functional requirements that both need to be satisfied by a system. $rAF1$ to $rAF5$ are alternative functionalities that satisfy rA , while $rBF1$ to $rBF5$ are alternative functionalities that satisfy rB . Filled circles are combinations of functionalities that satisfy both rA and rB , and are thereby alternative configurations of a system; empty circles denote incompatible combinations of functionalities, and are not configurations.

3.2. Adaptation to Change

Having recalled the *ZJ RP* as the standard *RP* for *RE*, the goal now is to clarify what is meant by adaptation, when talking about *RE* for Adaptive Systems. This will be used in subsequent sections to discuss why and how the *RPAS* differs from the *ZJ RP*.

To clarify what is meant by adaptation, consider a trivial example. In this example, by *functional requirement*, we mean a requirement that is either satisfied or not; it is, if a system can deliver a specific functionality, which is, roughly speaking, that it can do something we can observe. By *quantitative variable requirement*, we mean a requirement that assigns a desirable value or range of values to a variable, which is not binary.

We have only two functional requirements rA and rB to satisfy. We know that we can satisfy rA by implementing one of five different functionalities, denoted $rAF1$ to $rAF5$, and rB by implementing another five different functionalities, denoted $rBF1$ to $rBF5$. For simplicity, let all ten functionalities be different and not related in terms of refinement or parthood, that is, they are neither more detailed variants of one another, nor parts of one another.

With two functional requirements and 5 functionalities satisfying each, there are 25 combinations of the 10 functionalities. But, some functionalities are not compatible. This means that we cannot make a system which includes them both. Some combinations of functionalities therefore do not give a configuration which satisfies both rA and rB .

The part of the example introduced so far can be drawn as in Figure 1. There, filled circles are configurations, that is, combinations of functionalities that satisfy both rA and rB . Empty circles are incompatible combinations of functionalities, and since we cannot make a system that has those functionalities, these empty circles are not configurations.

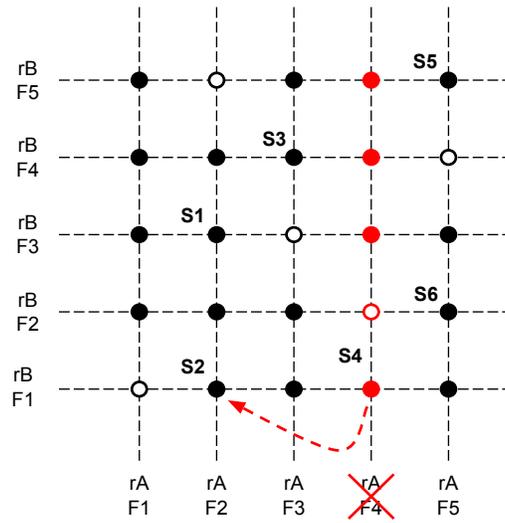


Fig. 2. System runs according to configuration S4. Then, functionality rAF4 fails, and the system activates functionality rAF2 instead, switching thereby from S4 to S2.

If our problem was to find one combination where functionalities are compatible, and which satisfies both rA and rB, then this can be any one of the 20 configurations in Figure 1.

Consider adaptation to the failure of a functionality. If we were to design a system according to one configuration among those in Figure 1, suppose that we chose, for example, S4, so that the system has functionalities rAF4 and rBF1. If rAF4 fails, the system fails. If it were an Adaptive System, then it would be designed so as to switch to another functionality at run-time, for example, from rAF4 to rAF2. Such a change is, in the terminology introduced so far, a switch from one configuration to another, by changing the parameters of the Configurable Specification. This is illustrated in Figure 2.

So far, and including adaptation shown in Figure 2, there is nothing that the ZJ RP fails to capture. Namely, if S in the ZJ RP is not one filled circle, but a set that includes all filled circles, together with conditions for activating each subset S1, S2, and so on, the switching between these subsets requires no changes to the ZJ RP: the problem is still to find S, given R and K, such that S satisfies the Satisfaction Condition and Consistency Condition.

Now, assume that we have two quantitative variable requirements to satisfy, in addition to rA and rB. The first relates to scalability and the second to how the system compares to its competitors. Let Var1 be the variable in the first, and Var2 in the second quantitative variable requirement. Var1 can be “number of users that can simultaneously use the system” and Var2 “number of products available for purchase”.

We prefer large to small values for both Var1 and Var2, and we cannot accept values that are below some threshold. This is drawn in Figure 3, where T1 is the threshold for Var1 and T2 for Var2, so that the shaded area shows all acceptable combinations of Var1 and Var2 values.

If the system were running according to S4, then it satisfies all four requirements, as its values over Var1 and Var2 are above their respective thresholds. If rAF4 fails, the system would need to switch from S4 to either S3 or S5 in order to still satisfy all

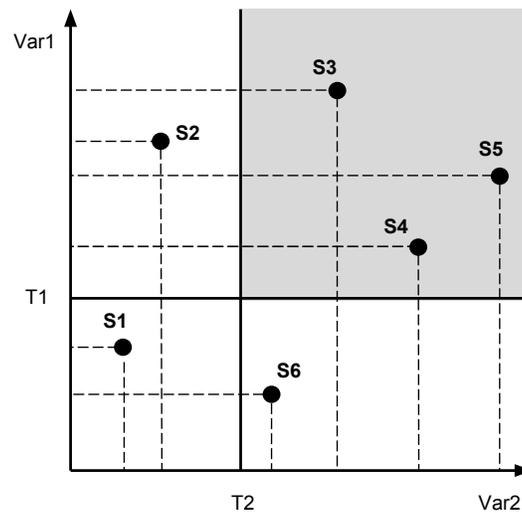


Fig. 3. Variables Var1 and Var2 quantify the level of satisfaction of two quantitative variable requirements. Hypothetical simulations of configurations S1 to S6 yield values show in the figure. T1 is the threshold value for Var1, and T2 for Var2.

four requirements; if it switched to S2, it would satisfy r_A , r_B , and the requirement quantified by Var1, but not the requirement quantified by Var2.

As long as the system can switch from one configuration to *any* other configuration, provided that the latter satisfies all requirements and domain knowledge, then we can capture with the ZJ RP the problem of designing that system's specification.

But this fails to capture two important insights. Firstly, the ZJ RP can be read as the problem of finding any single configuration, which satisfies the Satisfaction Condition and the Consistency Condition. There is no reason why such a specification should result in a system capable of coping with requirements failure or environment changes. In the RPAS, we are looking for a collection of such specifications, each of which applies in different conditions. We need a Configurable Specification, which can be configured into each of these different non-adaptive specifications. Secondly, once there are different configurations that the system can switch to, a mechanism is needed to decide which of the configurations it is best to switch to.

3.3. Why is RE for Adaptive Systems Non-Standard?

To see what optimality means here, and why it makes RPAS non-standard, we continue the example from the previous section.

We said that we prefer higher values of Var1 and Var2. To make this more precise, we need to say which combinations of values we prefer over others. Since the region above the thresholds T1 and T2 is large, it is interesting to indicate (i) the shape of indifference curves in that region, and (ii) the direction where these indifference curves are over more desirable combinations of Var1 and Var2 values.

Figure 4 shows hypothetical indifference curves in the region above thresholds T1 and T2. Each indifference curve is the set of Var1 and Var2 value combinations that are equally preferred. So every specification that has Var1 and Var2 values on the same indifference curve as S3 is equally preferred to S3. The arrow indicates the direction where Var1 and Var2 value combinations are preferred, so that any specification on the

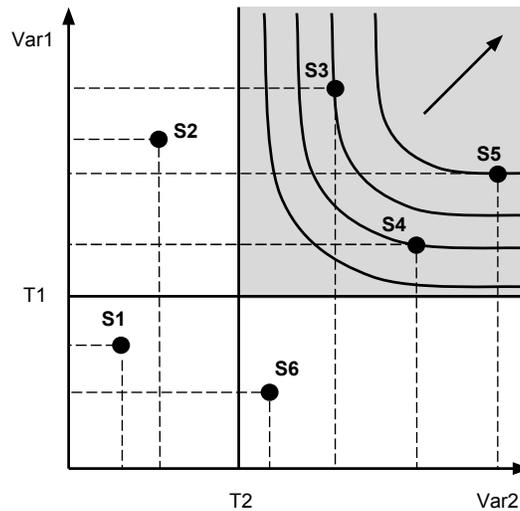


Fig. 4. Hypothetical indifference curves over value combinations of Var1 and Var2. One indifference curve includes all Var1 and Var2 value combinations that are equally desirable. For example, any configuration on the same indifference curve as S3 is equally desirable as S3. The arrow indicates the direction in which value combinations are more desirable. Therefore, S5 is preferred to S3, and S3 is preferred to S4.

indifference curve with S5 is strictly preferred to any specification on the indifference curve with S3.

Having clarified with indifference curves what we mean by preference for higher Var1 and Var2 values, we now go back to Figure 1. There, we had 20 alternative configurations. Moreover, we said that adaptation amounts to moving from one to another of these configurations, based on monitored input of the system, and feedback mechanisms that indicate what configuration to switch to. As discussed earlier, only some of this can be captured in terms of a ZJ RP.

By adding the two quantitative variable requirements, with their Var1 and Var2, we had restricted the set of acceptable configurations to some subset of the 20 shown in Figure 1.

Given several configurations, all above T1 and T2 thresholds, and all satisfying rA and rB, which one should we choose?

The answer is simple: given the indifference curves, we should choose any configuration which is on the most desirable indifference curve. More generally, we want the system to switch, every time it needs to adapt, to the configuration that is the most desirable, among those that are feasible.

This is not to say that we cannot capture also this notion of optimality in ZJ RP. For example, we can have as an additional member of R in ZJ RP, the proposition that there should be no feasible configuration which is on a more desirable indifference curve than the chosen configuration. We can, therefore see RPAS as a subclass of the ZJ RP, although doing so seems rather odd, for there were no considerations of configurations, adaptation, preference, uncertainty, or optimality in defining the ZJ RP.

RE for Adaptive Systems is non-standard precisely because it is not clear how exactly notions key to Adaptive Systems, such as configurations, adaptation, preference, uncertainty, optimality, and so on, fit in the ZJ RP. Clarifying this requires that we define the RPAS, and thereby answer exactly *how* RE for Adaptive Systems differs from standard RE.

4. BASIC FRAMEWORK

The aim in this section is to introduce a framework needed to model the kinds of examples used in the previous section. We need it, in order to make our discussion of RPAS more precise. We introduce the framework informally in this section. It is formalised in the electronic appendices of this paper.

The framework distinguishes different types of requirements, and allows relations between requirements. We discuss how these types and relations are used to compare alternative requirements configurations. To illustrate the framework, we use the London Ambulance Service (LAS) case study [Anonymous 1993].

4.1. Classification of Requirements

The term *quantitative variable* refers to any variable which takes a number as its value. The term *propositional variable* is used as in classical propositional logic – it is a variable that refers to a proposition and gets either True or False as its value.

A requirement is either one of the following three types:

- (1) the assignment of True to a propositional variable; for example, “ q_1 should be true”, where q_1 refers to the proposition “Ambulance shall arrive at an incident location within 14 minutes”,
- (2) the assignment of a value to a quantitative variable; for example, “ v should be exactly 50”, where v is the number of ambulances per 1.000.000 inhabitants,
- (3) the restriction of the range of values of a quantitative variable; for example, “average value of $t_{1,c}$ should be at most 15 seconds”, where $t_{1,c}$ is the time the caller waits for a control assistant.

In item 1 above, we talk about a **functional requirement**, in items 2 and 3 of a **quantitative variable requirement**.

Figures 5–6 show some of the functional and quantitative variable requirements for LAS. Annotations of functional requirements, visible in the two figures, indicate their types according to the two additional classification dimensions, discussed next.

The functional and quantitative distinction gives one classification dimension for requirements. The second classification dimension distinguishes requirements according to the role they play in the RP. For that second dimension, we use the types introduced in our prior work on a core ontology for RE [Jureta et al. 2008]; this means that a requirement is:

k: a **domain assumption** if it is assumed to hold for the domain. Consider two examples:

- In Figure 5, $\mathbf{k}(r_3)^M$ is “Callers report imprecise incident location”.
- Let μ be a function over a quantitative variable t_x , where t_x is the time to establish if two or more calls refer to the same incident location, and $\mu(t_x)$ the degree of satisfaction of LAS stakeholders for a given value of t_x . (Another way to see μ is that it is a utility function.) The domain assumption $\mathbf{k}(\mu(t_x = 30 \text{ sec}) = 0.6)$ says that the degree of satisfaction with value $t_x = 30 \text{ sec}$ is 0.6 on some scale.

g: a **goal** is a propositional variable that describes a desired state (by the stakeholders); for example, in Figure 5, the goal $\mathbf{g}(p_4)$ indicates that it is desired that “Fill incident report” is True. Goals are always functional requirements. Quality constraint, below, is analogous to a goal, but for quantitative variable requirements.

q: a **quality constraint** if it is desired that a quantitative variable gets the value (in the range of values that) the requirement assigns to it. For example, the quality constraint $\mathbf{q}(t_x = 60 \text{ sec})$ says that t_x should have the value 60 seconds, while $\mathbf{q}(t_x \leq 60 \text{ sec})$ would say that t_x should have the value of at most 60 seconds.

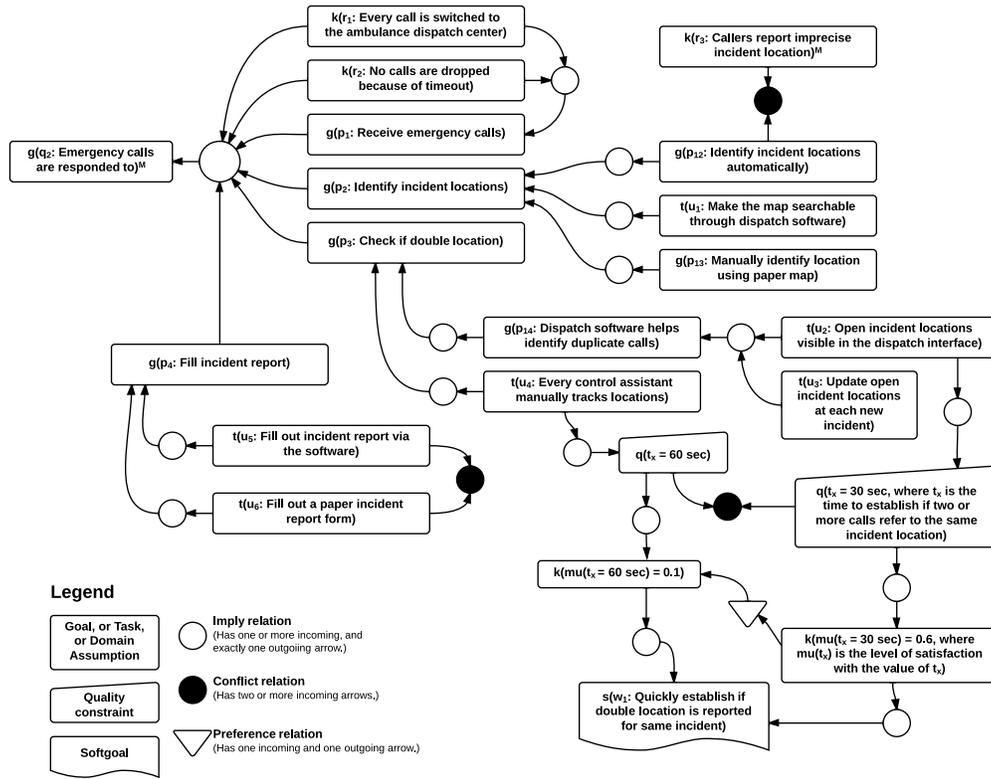


Fig. 5. One part of requirements for LAS.

s: a **softgoal** refers to a desirable range of values for a variable where the range is only vaguely specified. Consider “Quickly establish if double location is reported for same incident”, where “quickly” makes the statement vague, leading us to consider this a softgoal, denoted $s(\tilde{w}_1)$, where \tilde{w}_1 refers to the given vague statement.

t: a **task** is an action/function that should be executed. In Figure 5, for example, a task is $t(u_1)$, where u_1 refers to the proposition “Make the map searchable through dispatch software”. When a task is over a quantitative variable, it says that it is necessary to ensure that the given variable gets the assigned value, and that it is known how to do so (otherwise, it would have not been a task, but a quality constraint). For example, if v is the number of new control assistants hired every year, there can be a task $t(v = 3)$ to indicate simply that 3 control assistants should be hired every year, and it is not necessary to model the details of how it will be ensured that 3 control assistants are hired every year.

The third, and final classification dimension for requirements concerns if they *must* be satisfied (i.e., obtain the value True if functional, or some specified value if quantitative), or can be allowed to fail. A requirement can be **mandatory**, **optional**, or **neither**. $g(q_2)^M$ in Figure 5 is a mandatory goal, which means that it must be satisfied. $g(p_{15})^O$ in Figure 6 is an optional goal, meaning that it is more desirable that it is satisfied, than not, but we will still accept a system which fails to satisfy $g(p_{15})^O$.

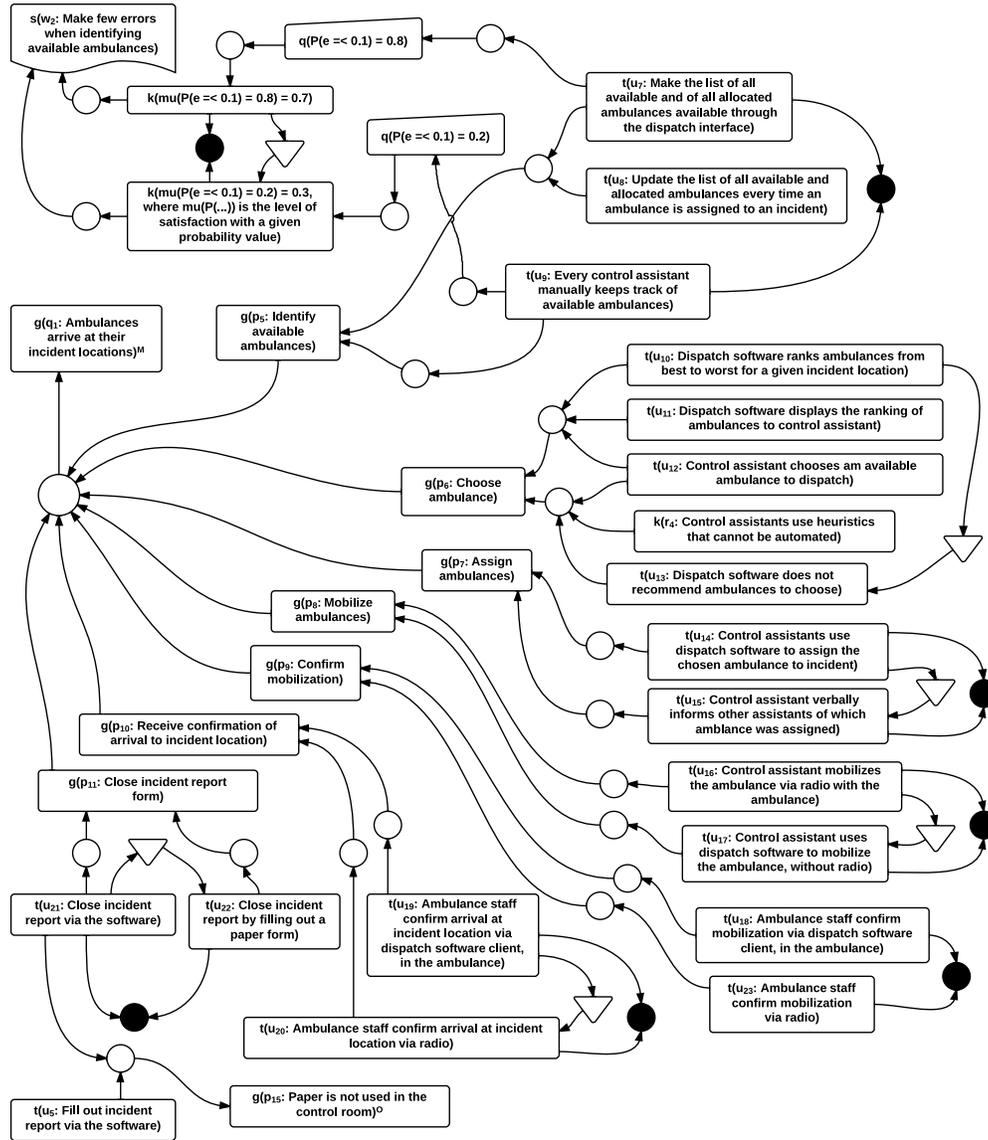


Fig. 6. Another part of requirements for LAS.

4.2. Relations between Requirements

We will use three kinds of relations between requirements:

- The imply relation drawn with empty circles in Figures 5 and 6 represents non-material **implication** [Jureta et al. 2010], understood as a conditional, if-then relation, so that the line from $t(u_1)$ to $g(p_2)$ abbreviates the formula $i(t(u_1) \rightarrow g(p_2))$, i.e., a relation according to which, if we can deduce $t(u_1)$, then we can also deduce $g(p_2)$. When the left-hand side of the implication includes only tasks and domain assumptions, we say that it *operationalises* the requirement on the right-hand side.

In this example, $\mathbf{t}(u_1)$ operationalises $\mathbf{g}(p_2)$. The implication relation is understood as an *operationalisation* relation if and only if it relates tasks and domain assumptions to goals or quality constraints; otherwise, it is read as a *refinement* relation: in Figure 5, $\mathbf{i}(\mathbf{k}(r_1) \wedge \mathbf{k}(r_2) \wedge \mathbf{g}(p_1) \wedge \mathbf{g}(p_2) \wedge \mathbf{g}(p_3) \wedge \mathbf{g}(p_4) \rightarrow \mathbf{g}(q_2)^{\mathbf{M}})$ says that the domain assumptions $\mathbf{k}(r_1)$ and $\mathbf{k}(r_2)$, and the goals $\mathbf{g}(p_1)$ to $\mathbf{g}(p_4)$ refine $\mathbf{g}(q_2)^{\mathbf{M}}$.

- Logical inconsistency is represented by the **conflict** relation; for example, we assume in Figure 5 that $\mathbf{t}(u_5)$ cannot be carried out together with $\mathbf{t}(u_6)$. The conflict relation between these tasks, drawn in Figure 5, abbreviates the assumption $\mathbf{i}(\mathbf{t}(u_2) \wedge \mathbf{t}(u_4) \rightarrow \perp)$: if we can deduce both, then we will also conclude logical inconsistency.
- Finally, we can have a **preference** relation between requirements, to indicate their relative desirability. For example, we will write $\mathbf{t}(u_{10}) \succ \mathbf{t}(u_{13})$ to say that making sure $\mathbf{t}(u_{10})$ gets the value True is strictly more desirable than making sure of the same for $\mathbf{t}(u_{13})$. We draw this as in Figure 6.

Preferences are not written in $\mathbf{i}()$, as they differ from implication and conflict: the latter two are used when doing deduction from requirements, while preferences are not, and are instead used only for the comparison of requirements.

The idea of using only the implication and conflict relations in deductions fits the observation from Robinson et al. [Robinson et al. 2003], that there are essentially two kinds of relations used in RE: positive and negative. Negative relations exist between requirements that cannot be satisfied together at all, or perhaps they can, but not all to the desired extent. Typical kinds of negative relations are conflict [van Lamsweerde et al. 1998] and obstruction [van Lamsweerde and Letier 2000]. Positive relations exist between requirements which can be satisfied together, and are relations found to be useful in designing or comparing solution to requirements problems. Typical examples are goal refinement [Darimont and van Lamsweerde 1996], means-ends [Yu and Mylopoulos 1994], and operationalisation [Dardenne et al. 1993] relations. Appendix C in the electronic appendices discusses how implication and conflict can be used to define oft-cited positive and negative relations in RE research, and thereby illustrates the versatility of the relatively simple requirements ontology we outlined using our classification dimensions on requirements and the relations above.

4.3. Comparison of Configurations

Figures 5–6 gave alternative operationalisations of the two mandatory top-level goals $\mathbf{g}(q_1)^{\mathbf{M}}$ and $\mathbf{g}(q_2)^{\mathbf{M}}$.

Each combination of operationalisations, which results in the satisfaction of all mandatory goals in those Figures is also called a **configuration**.

There are many alternative configurations that can be identified in Figures 5 and 6. Of these, three are shown in Figures 7–9. They are referred to, respectively, by S_1 , S_2 , and S_3 ; Table I shows the domain assumptions and tasks included in each.

Given several configurations, we want to be able to compare them in order to choose the most desirable one. Figures 7, 8, and 9 each show one configuration. Each configuration satisfies the two mandatory top-level goals, $\mathbf{g}(q_1)^{\mathbf{M}}$ and $\mathbf{g}(q_2)^{\mathbf{M}}$. One difference between S_2 and both S_1 and S_3 , is that in S_2 , the dispatch software does not keep track of open incident locations automatically, be it because of malfunction, or otherwise, requiring the control assistants to manually keep track of incident locations. Another difference of S_2 from both S_1 and S_3 is that we could switch to S_2 if it was no longer possible to use dispatch software to keep track of which ambulances are available, that is, if $\mathbf{t}(u_7)$ and/or $\mathbf{t}(u_8)$ fail, we would need to switch to $\mathbf{t}(u_9)$. Notice that switching to $\mathbf{t}(u_9)$ does not necessarily mean always switching to S_2 , as there can be other configurations, not shown in this paper, that include $\mathbf{t}(u_9)$ and are different from S_2 .

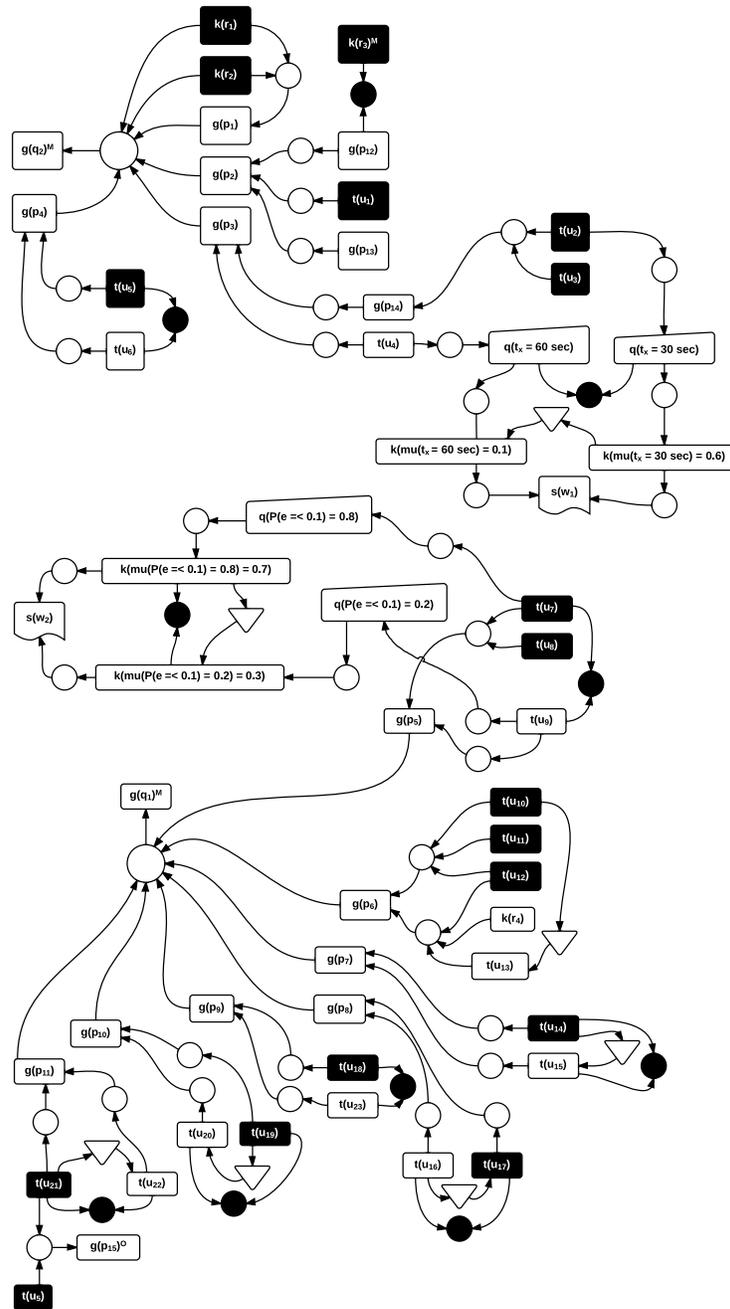


Fig. 7. One configuration, denoted S_1 for LAS requirements from Figures 5 and 6. Highlighted domain assumptions and tasks operationalise the two mandatory top-level goals, $g(q_1)^M$ and $g(q_2)^M$.

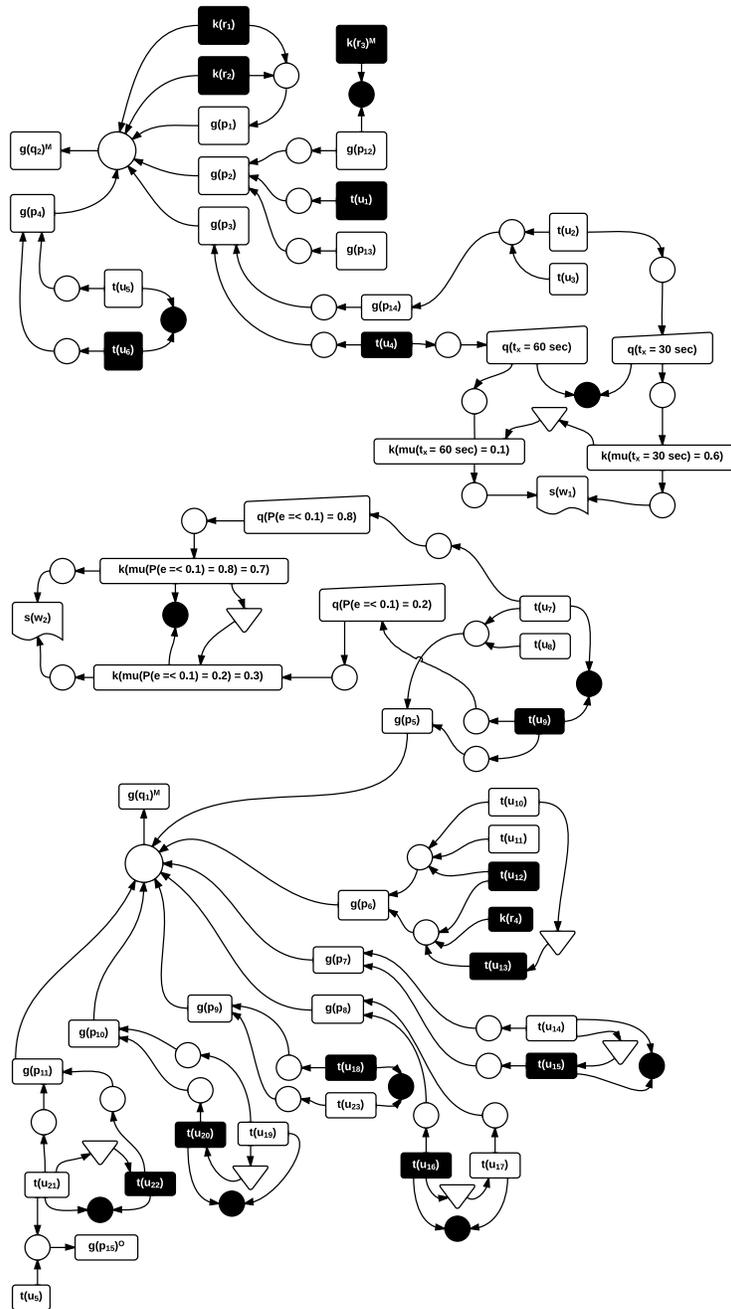


Fig. 8. Another configuration, denoted S_2 , for LAS requirements from Figures 5 and 6. Highlighted domain assumptions and tasks operationalise the two mandatory top-level goals, $g(q_1)^M$ and $g(q_2)^M$.

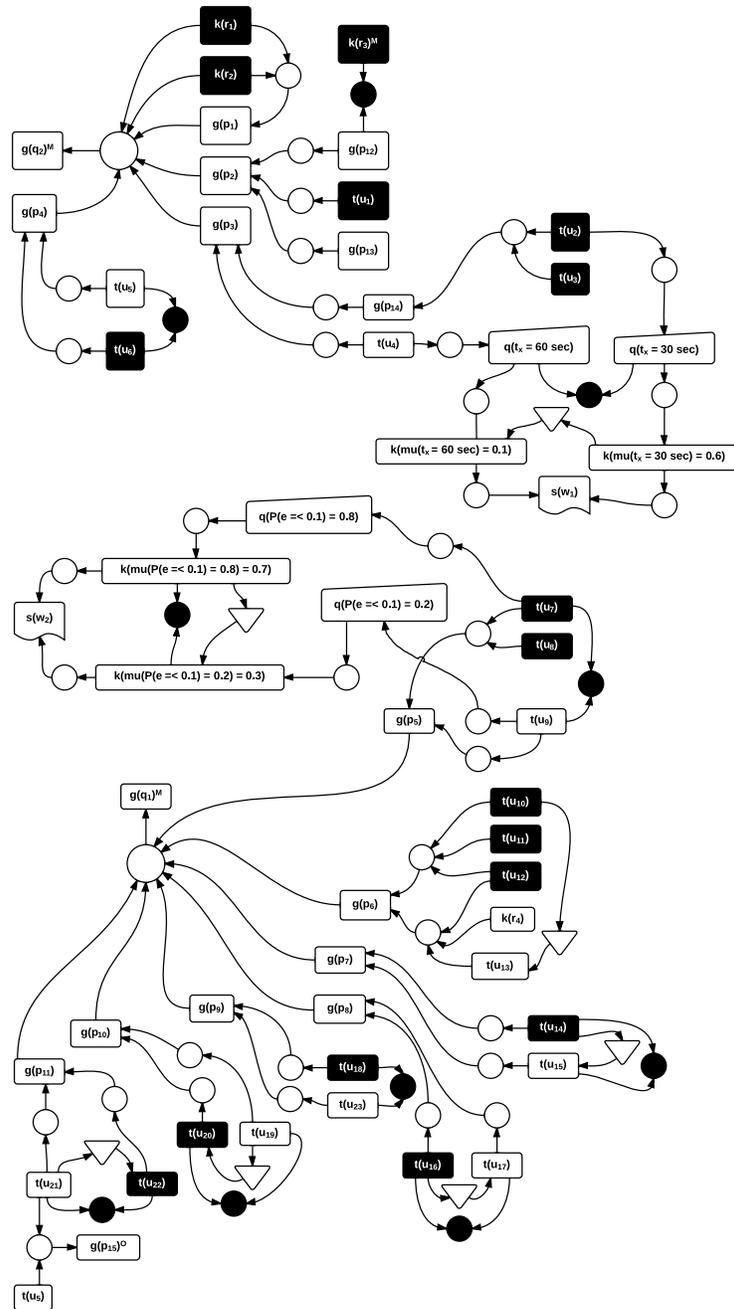


Fig. 9. A third configuration, denoted S_3 , for LAS requirements from Figures 5 and 6. Highlighted domain assumptions and tasks operationalise the two mandatory top-level goals, $g(q_1)^M$ and $g(q_2)^M$.

Table I. Content of each configuration shown in Figures 7–9

		S_1	S_2	S_3
Every call is switched to the ambulance dispatch center	$\mathbf{k}(r_1)$	x	x	x
No calls are dropped because of timeout	$\mathbf{k}(r_2)$	x	x	x
Callers report imprecise incident location	$\mathbf{k}(r_3)^{\mathbf{M}}$	x	x	x
Control assistant uses heuristics that cannot be automated	$\mathbf{k}(r_4)$		x	
Make the map searchable through dispatch software	$\mathbf{t}(u_1)$	x	x	x
Open incident locations visible in the dispatch interface	$\mathbf{t}(u_2)$	x		x
Update open incident locations at each new incident	$\mathbf{t}(u_3)$	x		x
Every control assistant manually tracks locations	$\mathbf{t}(u_4)$		x	
Fill out incident report via the software	$\mathbf{t}(u_5)$	x		
Fill out a paper incident report form	$\mathbf{t}(u_6)$		x	x
Make the list of all available and of all allocated ambulances available through the dispatch interface	$\mathbf{t}(u_7)$	x		x
Update the list of all available and allocated ambulances every time an ambulance is assigned to an incident	$\mathbf{t}(u_8)$	x		x
Every control assistant manually keeps track of available ambulances	$\mathbf{t}(u_9)$		x	
Dispatch software ranks ambulances from best to worst for a given incident location	$\mathbf{t}(u_{10})$	x		x
Dispatch software displays the ranking of ambulances to control assistant	$\mathbf{t}(u_{11})$	x		x
Control assistant chooses available ambulances to dispatch	$\mathbf{t}(u_{12})$	x	x	x
Dispatch software does not recommend ambulances to choose	$\mathbf{t}(u_{13})$		x	
Control assistants use dispatch software to assign the chosen ambulance to incident	$\mathbf{t}(u_{14})$	x		x
Control assistant verbally informs other assistants of which ambulances are assigned	$\mathbf{t}(u_{15})$		x	
Control assistant mobilizes the ambulance via radio with the ambulance	$\mathbf{t}(u_{16})$		x	x
Control assistant uses dispatch software to mobilize the ambulance, without radio	$\mathbf{t}(u_{17})$	x		
Ambulance staff confirm mobilization via dispatch software client, in the ambulance	$\mathbf{t}(u_{18})$	x	x	x
Ambulance staff confirm arrival at incident location via dispatch software client, in the ambulance	$\mathbf{t}(u_{19})$	x		
Ambulance staff confirm arrival at incident location via radio	$\mathbf{t}(u_{20})$		x	x
Close incident report via the software	$\mathbf{t}(u_{21})$	x		
Close incident report by filling out a paper form	$\mathbf{t}(u_{22})$		x	x
Ambulance staff confirm mobilization via radio	$\mathbf{t}(u_{23})$			

Both Figures 5 and 6 include information for the comparison of configurations:

- There are softgoals, such as $\mathbf{s}(\tilde{w}_1)$, which come with a satisfaction function μ over the quantitative variable t_x . Such softgoals are used to compare configurations as follows: different configurations will result in different values of t_x , and different values $\mu(t_x)$. According to Figures 7–9, task $\mathbf{t}(u_2)$ in configurations S_1 and S_3 satisfies the quality constraint $\mathbf{q}(t_x = 30 \text{ sec})$, while the task $\mathbf{t}(u_4)$ in configuration S_2 satisfies the quality constraint $\mathbf{q}(t_x = 60 \text{ sec})$. Using μ , we then know that satisfaction is higher when $\mathbf{q}(t_x = 30 \text{ sec})$ is satisfied, than when $\mathbf{q}(t_x = 60 \text{ sec})$ is satisfied. Finally, because the softgoal says that we prefer to quickly establish double location, the softgoal gives the preference $\mathbf{k}(\mu(t_x) = 0.6)^{\mathbf{M}} \succ \mathbf{k}(\mu(t_x) = 0.1)^{\mathbf{M}}$, indicating, for this softgoal only, that there are better configurations than S_2 . We will later revisit how the softgoals, in general, give us such preferences.

- Preference relations are not only generated via softgoals, as in the case above. The preference $\mathbf{t}(u_{14}) \succ \mathbf{t}(u_{15})$ is independent of softgoals, and will be satisfied if we choose a configuration that includes $\mathbf{t}(u_{14})$.
- Different configurations may satisfy different optional goals. The optional goal $\mathbf{g}(p_{15})^\circ$ is only satisfied by configuration S_1 in Figure 7.

5. KEY IDEAS IN ADAPTIVE SYSTEMS RE

This section reviews key ideas in RE for Adaptive System. The language introduced in the previous section, together with the key ideas presented in this section, form the basis for defining the RPAS in Section 6.

5.1. Monitoring and Reconciliation Tactics

Adaptive behavior entails monitoring of requirements, and control over a collection of behaviors, each of which can satisfy requirements. Fickas & Feather [Fickas and Feather 1995] suggested that for this we must start with a specification of all alternative behaviors to be considered in a requirements model, that is, alternative refinements and operationalisations of all requirements. The model is then a source of *monitored* variables, the observed values of which will trigger behaviors, called “reconciliation tactics”, that are intended to change the values of *control* variables, which can affect how one operationalises a requirement. This forms a *control loop*, through which the system observes and reacts by switching from one configuration to another, at runtime.

For illustration, consider the goal $\mathbf{g}(p_3$: Check if double location) for LAS in Figure 5, which has two operationalisations: one way to satisfy $\mathbf{g}(p_3)$ is by executing both tasks $\mathbf{t}(u_2)$ and $\mathbf{t}(u_3)$, the other is via $\mathbf{t}(u_4)$. A reconciliation tactic could be, for example, “if $\mathbf{t}(u_2)$ fails, then operationalise $\mathbf{g}(p_3)$ by $\mathbf{t}(u_4)$ ”.

Feather et al. [Feather et al. 1998] combined the requirements monitoring mechanisms from Fickas & Feather with KAOS [Dardenne et al. 1993] to make an RE framework, where requirements can be converted into constraints on events, and events are monitored at runtime. Robinson [Robinson 2006] suggested a requirements monitoring platform that can be combined with a requirements modelling language via translation rules between the expressions of the latter and those accepted by the former, then suggested how to use Object Constraint Language as the modelling language for the platform [Robinson 2008].

5.2. Computed Reconciliation Tactics

An important characteristic of a reconciliation tactic (and same applies to adaptation goals [Baresi et al. 2010]), is that it is defined in relation to *local* requirements, i.e., it says what to do in relation to a single requirement which is not being satisfied. There are two problems with defining a reconciliation tactic by looking mainly at local requirements:

- It is not clear that following any or every reconciliation tactic will make sure that the new system configuration is consistent with the requirements. It is not clear what happens if the reconciliation tactic works locally, but activates tasks which are inconsistent with some other already active part of the system.
- One local change may seem fine by itself, and may result in a new and consistent configuration, but this new configuration may be a dominated one: i.e., maybe instead of satisfying one reconciliation tactic and obtaining some quality level, we could have used several adaptations, which result in a configuration that ensures higher quality. The local approach to the definition of reconciliation tactics comes with the risk of adapting to a sub-optimal or inconsistent configuration.

We therefore contend that instead of defining reconciliation tactics by considering their impact only on a small set of requirements, *reconciliation tactics should be defined so as to ensure that the configuration obtained after adaptation via reconciliation tactics is desirable in its entirety in terms of various criteria used to compare alternative requirements configurations* (including, for example, preferences over requirements, probabilities of requirements satisfaction, and so on). This has two consequences on the formulation of the problem and solution concepts in Adaptive Systems RE:

- (1) We do not define reconciliation tactics *while writing the requirements model*, or in other words at design-time, but **compute** reconciliation tactics after we have identified at least some configurations from a set of alternative configurations, all of which satisfy some set of properties (for example, are consistent) and satisfy some preferences. Instead, we define evolution requirements, which the computed reconciliation tactics have to satisfy. Consider Figures 7–9, where each configuration satisfies the mandatory goals from in Figures 5–6 for the LAS system: for example, if the system is configured according to S_1 , and $\mathbf{t}(u_5)$ fails, then it may be more desirable to adapt to S_3 than to S_2 , so that reconciliation tactics can be identified directly by comparing S_1 to S_3 , and defined as what tasks to add and delete in order to switch from $\mathbf{t}(u_{21})$ to $\mathbf{t}(u_{22})$, from $\mathbf{t}(u_{19})$ to $\mathbf{t}(u_{20})$, from $\mathbf{t}(u_{17})$ to $\mathbf{t}(u_{16})$, and from $\mathbf{t}(u_5)$ to $\mathbf{t}(u_6)$.
- (2) The local/global distinction has an impact on how we *select* configurations: it is likely that the system *can* adapt from one to *one among several alternative configurations*; for example, if $\mathbf{t}(u_{17})$ fails in configuration S_1 , then we can consider both S_2 and S_3 , which use the alternative $\mathbf{t}(u_{16})$. Commitment to one of these configurations should take into account the evolution requirements and the consequences of adapting to the configuration: for example, perhaps configuration S_2 is more desirable than S_3 if only these two are compared, but we may know that if we select S_2 , then in the long term it is not sustainable, in order to keep satisfying some quality constraint; or, we may have an evolution requirement which says that, if $\mathbf{t}(u_{17})$ fails, then configuration to adapt to, must exclude $\mathbf{t}(u_{23})$.

5.3. Probabilistic and Fuzzy Relaxation of Requirements

Adaptive behavior by definition means that the system cannot satisfy all the specific requirements and meet the desired quality levels, all of the time – it may do worse, but it may also do better.

Relaxation of requirements has been suggested to cover both cases: a requirement can be relaxed to allow the system to fail it, but we restrict how often it can do so (via probabilistic relaxation), or to avoid overly constraining the system when it may deliver higher quality than expected (via fuzzy relaxation). Relaxation of requirements is closely related to the evaluation of their satisfaction, as one aim of relaxation is to quantify the degree of satisfaction of a requirement. Two broad approaches to relaxation and satisfaction evaluation have been suggested based on *objective criteria* (where measurement is in terms of quantifiable phenomena in the environment or the system) or on *subjective criteria* (where a number is given to reflect personal impression of satisfaction, regardless of measurable phenomena in the domain). Souza et al. [Souza et al. 2012] called awareness requirements those requirements that one obtains, after relaxing some original requirements.

An important approach to relaxation, based on objective criteria was suggested by Letier & van Lamsweerde [Letier and van Lamsweerde 2004], who extended the goal models of KAOS to allow the specification of measures on system behaviors, and probability functions over the values of these measures. Using such a formalism, they can state, for example, that in the LAS system, 95% of ambulances should reach the des-

ignated place of incident in at most 14 minutes after an incident is reported. Values of measures and probabilities are computed for every requirements configuration, all of which are inputs to a decision-making process that aims at selecting a single configuration.

Approaches based on subjective criteria were revived recently for RE of Adaptive Systems. Whittle et al. [Whittle et al. 2010] associate fuzzy membership functions with functional requirements, in order to allow and quantify departures from the ideal satisfaction of such requirements. The fuzzy membership function can be interpreted as a satisfaction function, and they use modalities to relax requirements: for example, to relax the task $t(u_{18})$ in Figure 6, we would write [As early as possible] $t(u_{18})$. This means that we would add a function μ which returns a level of satisfaction when given a duration between time “now” and the time when $t(u_{18})$ is executed, whereby μ “has its maximum value at 0 (i.e., at the current point in time) [and] tails off gradually ad infinitum (i.e., it has a triangular membership graph that is asymptotic)” [Whittle et al. 2010].

Baresi et al. [Baresi et al. 2010] also use fuzzy membership functions in relaxing requirements, and so are able to treat functional requirements as quantitative ones. They introduce degrees of satisfaction between the binary 0 (for violated) and 1 (satisfied); thus, adjusted requirements are associated to “adaptive” requirements that relate ranges of values of the fuzzy membership function to trigger switching between configurations. For example, if there is a quality constraint $q(v < 6 \text{ hrs})$ — which would be written as a goal $\mathcal{G}(v < 6 \text{ hrs})$ in their notation, as they allow quantitative variables in goals — then relaxing it would amount to replacing it with $q(v <_f 6 \text{ hrs})$ (in their notation, $\mathcal{G}(v <_f 6 \text{ hrs})$), where $<_f$ is a fuzzy operator. Their interpretation of $v <_f 6 \text{ hrs}$ is that there is a fuzzy membership function μ which returns the level of satisfaction as a function of v , and the shape of μ is predefined (for $<_f$ in $\mathcal{G}(v <_f 6 \text{ hrs})$, it is positive and constant until $v = 6$, then decreases up to the satisfaction value 0 for some $v > 6$).

6. CONFIGURABLE SPECIFICATIONS, CONFIGURATIONS, AND EVOLUTION REQUIREMENTS

The RPAS and the corresponding Configurable Specification concept are needed, because monitoring, control, evolution requirements, and probabilistic and fuzzy relaxation are not formulated in the standard RP. These features affect the properties that solutions need to satisfy, and influence the criteria used to compare alternative solutions. The aim of this section is to present and discuss the formal properties of the Configurable Specification concept, as the solution concept for RPAS. In Section 7, we compare RPAS and Configurable Specification concepts to Zave & Jackson [Zave and Jackson 1997] and Techne [Jureta et al. 2010].

The requirements problem is a *design* and *decision-making* problem. By design, we mean that alternative solutions to the problem are not available, but need to be created: for example, in Figures 5–6, we start from two broad goals, $\mathbf{g}(q_1)^M$ and $\mathbf{g}(q_2)^M$, so that we need to elicit further requirements and information, refine goals, identify conflicts, define operationalisations, all of which fall under design.¹ Decision-making requires the identification of criteria for the comparison of alternative solutions, and the application of a decision rule to rank alternative solutions from the most desirable to the least desirable.

¹Note that “design” is typically used for steps of systems development that come after RE, such as architectural design. The term “design” is used here in a broader sense, to encompass all activities (e.g., elicitation, validation) which are necessary to make alternative options on which we then need to decide.

To provide definitions, we start with the notion of a requirements database Δ – a set of requirements and relations between them. For example, the requirements database for LAS in this paper includes all requirements and relations in Figures 5 and 6.

The novelty of the RPAS and Configurable Specification concepts lies in the properties of requirements configurations, evolution requirements, and in how Configurable Specifications can be ranked. We discuss these in turn below.

6.1. Properties of Requirements Configurations

Let us first look at the configurations appearing in a Configurable Specification.

Definition 6.1. A **configuration** in a Configurable Specification is a set S of domain assumptions and tasks in Δ such that S satisfies the Consistency, Functional threshold achievement, Quantitative threshold achievement, Conformity, Dominance, and Minimality conditions.

6.1.1. Consistency. Every requirements configuration S must be consistent, i.e., $S \not\vdash_{\triangleright} \perp$, where \vdash_{\triangleright} is the consequence relation defined in Section A.2. The Adaptive System should work to satisfy a consistent set of requirements. In Figures 7–9, every configuration is such that there are no conflicts between its members. Remark that in a Configurable Specification, $\bigcup S$ can, and will often be inconsistent, as alternative configurations need not be consistent together. If we adapt from S_1 to S_2 , note that $S_1 \cup S_2 \vdash_{\triangleright} \perp$, since there are conflicts (e.g., between $\mathbf{t}(u_{16})$ and $\mathbf{t}(u_{17})$).

6.1.2. Functional Threshold Achievement. Every configuration S should operationalise every mandatory goal in Δ — a set written as $\Delta_{\mathbf{g}}^{\mathbf{m}}$. An operationalization defines tasks and domain assumptions with which we can fulfill a goal. For example, in Figure 5, if the task $\mathbf{t}(u_1)$ is satisfied, then $\mathbf{g}(p_2)$ is satisfied, so that the former operationalises the latter. Since a requirement can be operationalised by alternative sets of tasks and domain assumptions, we have a function $\text{Op}(\varphi)$ (defined in Section A), which given a mandatory goal $\varphi \in \Delta_{\mathbf{g}}^{\mathbf{m}}$, returns the set of all operationalisations of φ . For example, $\text{Op}(\mathbf{g}(p_7))$ includes two operationalisations of $\mathbf{g}(p_7)$ in Figure 6. This property is restricted to *functional* requirements, for which satisfaction is binary, and thus to goals, since a goal is either satisfied or not. The property requires *threshold* satisfaction, since it requires only that all *mandatory* goals be satisfied by a configuration.

A consistent set of requirements must meet both the functional and quantitative thresholds, for it otherwise fails to satisfy the requirements which must be satisfied.

6.1.3. Quantitative Threshold Achievement. Every configuration S should operationalise every mandatory quality constraint in Δ — a set written as $\Delta_{\mathbf{q}}^{\mathbf{m}}$. Once again, the operationalisation function $\text{Op}(\varphi)$ is used to return all sets of operationalisations of a quality constraint $\varphi = \mathbf{q}(\alpha)$, whereby some set Π is such an operationalisation iff it includes domain assumptions and tasks in which there are assignments of values to all variables in α , and these assignments are such that the constraint in α is satisfied.

Quality constraints identify desirable values of quantitative variables. For example, quality constraints may thus place constraints on values of measures of some behavior of the system-to-be. In LAS, and in relation to the response to emergency calls and the mobilization of ambulances, the time that these activities take is a crucial part of quality of service. We can use the following variables, where c identifies a call and e a unique incident:

$t_{1,c}$: Time the caller waits for a control assistant;
 $t_{2,c}$: Time to identify incident location;
 $t_{3,c}$: Time to fill out incident report;
 $t_{4,e}$: Time to mobilize an ambulance;
 $t_{5,e}$: Time for the mobilized ambulance to arrive and confirm arrival at incident location.

Instead of defining bounds on all of these variables, a government standard may not be as specific, and may impose a maximal time over a subset of activities that need to be executed between the placement of the emergency call and the confirmation of arrival of an ambulance to location; for example:

$\mathbf{q}(t_{6,c} \leq 3\text{min}; t_{6,c}$ is the duration between the switching of the call c to the dispatch center to the mobilization of the ambulance to the incident location)

where it is clear that the value of $t_{6,c}$ depends on $t_{1,c}$, $t_{2,c}$, $t_{3,c}$ and $t_{4,e}$. If we assume that $t_{6,c} = t_{1,c} + t_{2,c} + t_{3,c} + t_{4,e}$, then we add this to the requirements database as a domain assumption over quantitative variables, $\mathbf{k}(t_{6,c} = t_{1,c} + t_{2,c} + t_{3,c} + t_{4,e})$. The domain assumption says that there is a *quantitative refinement* relation between variables, and it specifies the functional relation between the refined $t_{6,c}$ and the variables refining it. In contrast to the refinement relation, quantitative refinement is not over requirements, but variables in quantitative variable requirements.

While it may be useful to set a precise bound on the value of an aggregate variable, as $t_{6,c}$ above, it may be more interesting to set bounds relative to the values of other variables. For example, if we want $t_{2,c}$ to be at most 110% of its average value over the past three months, then we write

$$\mathbf{q}(t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1}));$$

$\mathbf{k}(v_{1,m} = n(m)^{-1} \sum_{i=1}^{n(m)} \sum_{c=1}^{c(i)} t_{2,c})$, where m is the month identifier, $n(m)$ the number of days in m , c is the call identifier on a given day, $c(i)$ is the total number of calls received on day i of month m ;

where the domain assumption is a quantitative refinement of $v_{1,m}$, the average time, over a month, that it takes to identify the location.

The quality constraints and quantitative refinements need to be related to goals, tasks and domain assumptions in order to determine if the former are satisfied. If there is a running system, the values of $t_{2,c}$ are recorded, and it is straightforward to check if the constraint $t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1})$ is satisfied.

Before the system is in operation, we can simulate values of $t_{2,c}$, by assuming that $t_{2,c}$ is a random variable which has some probability distribution, so that we would have $\mathbf{k}(t_{2,c} \sim \mathcal{N}(60\text{sec}, 45\text{sec}^2))$ if we assume that $t_{2,c}$ follows a normal distribution with mean 60sec and variance 45sec^2 when the task $\mathbf{t}(u_1)$ is satisfied, which we model by a refinement $\mathbf{k}(\mathbf{t}(u_1) \rightarrow \mathbf{k}(t_{2,c} \sim \mathcal{N}(60\text{sec}, 45\text{sec}^2)))$. This assumption may be based on data from a pilot study, from expert opinion, or from data on systems which also satisfy $\mathbf{t}(u_1)$ and are already in operation.

6.1.4. Conformity. *Every configuration S must include all strict domain assumptions and all mandatory tasks, i.e., $\Delta_{\mathbf{k}}^{\mathbf{m}} \cup \Delta_{\mathbf{t}}^{\mathbf{m}} \subseteq S$.* This property asks that all strict domain assumptions are not violated and all mandatory tasks are executed in every configuration.

The Functional and Quantitative threshold achievement and Conformity properties ensure that a configuration satisfies all that *must* be satisfied.

6.1.5. Dominance. *Every configuration S must be maximal w.r.t. optional requirements, i.e., $\exists S' \subseteq \Delta$ such that (i) both S and S' satisfy the Consistency, Functional threshold*

achievement, Quantitative threshold achievement, and Conformity conditions, (ii) $S \subset S'$, and (iii) $S' \setminus S$ contains optional **k** or **t** requirements.

This condition formalises the idea of optional requirements, which are desirable to satisfy, but can be violated. A configuration should include as many such defeasible domain assumptions and as many optional tasks, up to the point at which adding any further defeasible domain assumptions and/or optional tasks violates the Consistency, Functional and Quantitative threshold achievement, and Conformity properties. The Dominance property ensures that every configuration is Pareto efficient with regards to optional requirements, as this condition makes it impossible to add optional domain assumptions and tasks to any S and still ensure that S is a configuration.

6.1.6. Minimality. A set S satisfying all the above properties must be minimal in order to qualify as a configuration. Minimality requires that a configuration includes only the domain assumptions and tasks which are needed to satisfy exactly the Consistency, Functional threshold achievement, Quantitative threshold achievement, Conformity, and Dominance properties. Minimality requires that a configuration includes only the domain assumptions and tasks which are needed to satisfy exactly the Consistency, Functional threshold achievement, Quantitative threshold achievement, Conformity, and Dominance properties.

All three configurations S_1 , S_2 , and S_3 in Figures 7–9 satisfy all six properties.

In terms of our discussion of RE for Adaptive System in Section 3, we have defined above the properties that are satisfied by every individual configuration shown in Figure 1. We now consider the switching between configurations, via evolution requirements.

6.2. Evolution Requirements

An evolution requirement specifies the differences that are allowed between the previous configuration and the next configuration to switch to. It does so by describing the presence/absence of certain requirements in pairs of related configurations, and is triggered by the failure of monitored assumptions, tasks, or quality constraints. Any requirement monitored for failure, is an **awareness requirement**.

In Figure 7, if task $\mathbf{t}(u_{21})$ fails, we might want to replace $\{\mathbf{t}(u_{19}), \mathbf{t}(u_{17}), \mathbf{t}(u_5)\}$ with $\{\mathbf{t}(u_{22}), \mathbf{t}(u_{20}), \mathbf{t}(u_{16}), \mathbf{t}(u_6)\}$. In general, we will therefore want to specify the conditions monitored for failure, and the changes (additions, removals) that must be seen in the successor configuration when failure occurs. The above evolution requirement can then tell us that if we are running the configuration S_1 in Figure 7, and the system fails to satisfy $\mathbf{t}(u_{21})$, then it could adapt from S_1 to S_3 . Note that an evolution requirement $\langle T, A, D \rangle$ might then be thought of as a simple operator in AI planning, where T is trigger condition, while A and D are add/delete lists.

Definition 6.2. An **evolution requirement** is an operator of the form $\langle T, A, D \rangle$, where T is a set of requirements present in the initial configuration (which are monitored to fail) while A and D are, respectively, sets of requirements that must be present and absent in the final configuration, if the operator applies.

Note therefore that if operator $\langle T, A, D \rangle$ applied in going from configuration S_i to S_{i+1} , then $T \cup D \subseteq S_i$ and $A \cap S_i = \text{must hold}$, and analogously with S_{i+1} . Note also that it is quite acceptable that multiple operators apply at the same time, when connecting S_i to S_{i+1} , and that (as with the so-called ramification problem in AI), there may be additional changes needed in S_{i+1} in order to make it a valid configuration.

In the above sense, requirements are monitored, while evolution requirements act as a control mechanism, which guides adaptation by changing the target requirements configuration that the Adaptive System needs to satisfy.

6.3. Comparison of Configurations in a Configurable Specification

Given several configurations, the configuration which ranks best, is the optimal configuration. Ranking of configurations requires the identification of comparison criteria, and the application of a decision rule which establishes a ranking of configurations in a Configurable Specification, on the basis of criteria.

Comparison criteria are either optional requirements or preferences.

Preferences can be individually added to Δ , or can be obtained through the relaxation of requirements, or from softgoals. Preferences are illustrated in Figures 5 and 6, where, for example, $\mathbf{t}(u_{16})$ is strictly preferred to $\mathbf{t}(u_{17})$, which indicates that, for this criterion only, a configuration having the former task is strictly more desirable than the configuration having the latter task.

6.3.1. Preferences through Probabilistic Relaxation. Continuing the example from Section 6.1.3, the upper bound on $t_{2,c}$ in $\mathbf{q}(t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1}))$ may still be too idealistic, as callers may provide information of very different quality about the incident location.

Probabilistic relaxation of a quality constraint is done in two steps:

- (1) The variable constrained in \mathbf{q} is redefined as a random variable, and an assumption is made on the probability distribution of that random variable. Adding $\mathbf{k}(t_{2,c} \sim \mathcal{N}(60sec, 45sec^2))$ to Δ makes of $t_{2,c}$ a random variable which follows a normal distribution.
- (2) The quality constraint to be relaxed is removed from Δ , and a new quality constraint is added. The new constraint specifies a bound not on the value of the now random variable, but on the probability that its value is in some range. Since we made $t_{2,c}$ into a random variable in the first step, we now replace $\mathbf{q}(t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1}))$ with $\mathbf{q}(P(t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1})) \geq 0.90)$, that is, we now require that the minimal probability should be 0.90 for $t_{2,c}$ to be at most 110% of its three-month average.

Random variables and quality constraints thereon are important for decision-making, as we use them to set desired probability levels over random variables. In more informal terms, this means that we can write quality constraints and domain assumptions which reflect, respectively, the confidence that we desire to reach in relation to system behaviors, and confidence that we actually have.

Figure 6 indicates that $\mathbf{t}(u_7)$ operationalises the quality constraint $\mathbf{q}(P(e \leq 0.10) = 0.8)$, while $\mathbf{t}(u_8)$ operationalises $\mathbf{q}(P(e \leq 0.10) = 0.2)$. The two quality constraints suggest that the two tasks result in different probabilities of having less than 10% of erroneous identifications of available ambulances. Figure 6 further indicates the level of satisfaction with each of the two probability levels.

If we needed to do probabilistic relaxation of a goal, rather than of a quality constraint, the process above is changed as follows. Given a goal, say $\mathbf{g}(p_9)$ in Figure 5, to do probabilistic relaxation of it, we remove it from the requirements, and add a quality constraint $\mathbf{q}(p'_9 \geq 90\%)$, where p'_9 is the percentage of confirmed mobilizations over all mobilizations, and 90% a threshold value. We then proceed to relax this quality constraint in the same way as above.

This same approach applies when doing probabilistic relaxation of a domain assumption which is over a propositional variable, such as $\mathbf{k}(r_3)^M$.

When we do probabilistic relaxation of an optional requirement, the new requirement obtained by relaxation is also optional; if it was mandatory, the new requirement would also remain mandatory.

6.3.2. *Preferences through Fuzzy Relaxation.* Fuzzy relaxation of a quality constraint involves two steps:

- (1) The quality constraint is removed, and a fuzzy membership function is defined over the variable from the removed quality constraint. In the second step, a softgoal is defined over the variable from the relaxed requirement. Consider again $\mathbf{q}(t_{2,c} \leq (1.1/3)(v_{1,m-3} + v_{1,m-2} + v_{1,m-1}))$ and we now apply fuzzy relaxation. We start by removing this quality constraint from Δ . A fuzzy membership function over $t_{2,c}$, denote it $\mu(t_{2,c})$, depends on how stakeholders evaluate, in terms of desirability, the various values of $t_{2,c}$. For example, we could use $\mu(t_{2,c}) = e^{-t_{2,c}}$, so that the higher the value of $t_{2,c}$, the lower $\mu(t_{2,c})$ is, which reflect the idea that the more time it takes to identify an incident location, the more the stakeholders are dissatisfied, whereby their satisfaction increases as $t_{2,c}$ approaches 0. If we adopt μ as defined, we have removed the quality constraint on $t_{2,c}$ and we quantify satisfaction as a function of $t_{2,c}$.
- (2) To finish with the fuzzy relaxation of the quality constraint on $t_{2,c}$, a softgoal needs to be added to the requirements database over values of $t_{2,c}$. Softgoals have had various definitions in RE, but there seems to be agreement on their two properties: (i) they are used for the comparison of alternatives, and (ii) they are vague, as they refer to some desirable values of variables, even though it may not be clear which exact values, or of which variables. For example, *Response time should be low* is a typical softgoal, where the variable is suggested, but it is not clear which specific range of its values qualify as *low*; in the softgoal *High safety*, not only is it not clear how to measure *safety*, it is also not apparent when *safety* is *high*. When used in fuzzy relaxation, a softgoal is defined over a known variable: in our example, it is reasonable to prefer lower over higher values of $t_{2,c}$, and we consequently have the softgoal

```

s( LET  $x_1 \stackrel{\text{def}}{=} \text{VAL}(S_1, t_{2,c})$  AND  $x_2 \stackrel{\text{def}}{=} \text{VAL}(S_2, t_{2,c})$ ;
  IF  $\mu(x_1) > \mu(x_2)$  THEN ADD
  { $\mathbf{k}(t_{2,c} = x_1), \mathbf{k}(t_{2,c} = x_2), \mathbf{k}(t_{2,c} = x_1) \succ \mathbf{k}(t_{2,c} = x_2)$ }
  TO  $\Delta$ ;)

```

where $\text{VAL}(S_1, t_{2,c})$ returns the value of $t_{2,c}$ in configuration S_1 . Although the formulation above seems very different from saying “Low $t_{2,c}$ ”, this is precisely what it does, as long as $\mu(t_{2,c})$ decreases when $t_{2,c}$ increases. \succ denotes the strict preference relation, and $\mathbf{k}(t_{2,c} = x_1) \succ \mathbf{k}(t_{2,c} = x_2)$ says that $\mathbf{k}(t_{2,c} = x_1)$ is strictly more desirable than $\mathbf{k}(t_{2,c} = x_2)$.

The important point is that the softgoal specifies a macro, and the macro generates domain assumptions and preference relations. For any two configurations S_1 and S_2 , the macro compares the satisfaction (returned by satisfaction function μ) with values that $t_{2,c}$ has in each of those two configurations. Depending on the comparison between these values, the macro adds two domain assumptions and a preference relation to Δ . The reason we add them to Δ and not individual configurations is because we can add one of these domain assumptions only if doing so does not violate the conditions that a configuration must satisfy (i.e., adding one of these domain assumptions may make a configuration inconsistent, and thus no longer a configuration at all).

The macro above ensures that if we compare two configurations, S_1 and S_2 , $t_{2,c}$ obtains the value x_1 in S_1 and the value x_2 in S_2 , then we will prefer *over this criterion (independently of other criteria)* the configuration in which $t_{2,s}$ obtains the value which results in higher satisfaction. The macro thus conveys the idea that, whenever we are given two values of $t_{2,c}$, we prefer the one that we are more satisfied with.

Fuzzy relaxation of a quality constraint over a variable v thus works by (i) removing the quality constraint on v , (ii) adding a fuzzy membership function $\mu(v)$ on v , (iii) interpreting $\mu(v)$ as the level of satisfaction with the value v , and (iv) adding a softgoal macro which generates preference relations that reflect the shape of μ .

6.3.3. Preferences based on Softgoals. Not all softgoals are macros used in fuzzy relaxation. A softgoal can be introduced in Δ even if we do not know the exact variable it refers to, or the exact range of values that it makes desirable. Very broad softgoals are allowed, such as

$\mathbf{s}(\tilde{p}_{17}$: Ambulances should arrive quickly at incidents)

where we write the tilde over p because the content of the softgoal is not a propositional variable as in a goal or task, since it is not clear from the statement referred to by \tilde{p}_{17} how to measure the time of arriving to incident scenes (for example, does it begin at call reception, or at ambulance mobilization?) and it is not clear when a time to arrive at an incident counts as *quick*. There are two ways to **approximate** such softgoals: by refinement or by satisfaction functions, and both can result in preferences being added to Δ .

When a softgoal is refined, it is treated just as any other requirement. For example, we may define the following refinement of $\mathbf{s}(\tilde{p}_{17})$

$\mathbf{q}(t_{7,e} \leq 15\text{min})$
 $\mathbf{k}(\mathbf{q}(t_{7,e} \leq 15\text{min}) \rightarrow \mathbf{s}(\tilde{p}_{17}))$
 $\mathbf{k}(t_{7,e} = t_{1,c} + t_{2,c} + t_{3,c} + t_{4,e} + t_{5,e})$

where we refined the softgoal with the quality constraint over the quantitative variable $t_{7,e}$ (i.e., if we satisfy $\mathbf{q}(t_{7,e} \leq 15\text{min})$ then the if-then domain assumption above, with implication, tells us to deduce $\mathbf{s}(\tilde{p}_{17})$), which is itself a function of times introduced before. We could find another refinement of $\mathbf{s}(\tilde{p}_{17})$, and have a preference between the two.

A softgoal can be approximated using a satisfaction function, and by proceeding in a similar way to fuzzy relaxation. Given $\mathbf{s}(\tilde{p}_{17})$, we decide which variable it refers to, and we assume it is $t_{7,e}$, such that $\mathbf{k}(\mathbf{q}(t_{7,e} \leq 15\text{min}) \rightarrow \mathbf{s}(\tilde{p}_{17}))$. Secondly, we remove the softgoal $\mathbf{s}(\tilde{p}_{17})$. Thirdly, we define a function $\mu_1(t_{7,e})$, and we interpret the value given by $\mu_1(t_{7,e})$ as the level of satisfaction with the value $t_{7,e}$. Fourthly, we add a softgoal macro over $t_{7,e}$:

$\mathbf{s}(\text{LET } x_1 \stackrel{\text{def}}{=} \text{VAL}(S_1, t_{7,e}) \text{ AND } x_2 \stackrel{\text{def}}{=} \text{VAL}(S_2, t_{7,e});$
 $\text{IF } \mu_1(x_1) > \mu_1(x_2) \text{ THEN ADD}$
 $\{\mathbf{k}(t_{7,e} = x_1), \mathbf{k}(t_{7,e} = x_2), \mathbf{k}(t_{7,e} = x_1) \succ \mathbf{k}(t_{7,e} = x_2)\}$
 $\text{TO } \Delta;)$

This case is also illustrated with the softgoal $\mathbf{s}(\tilde{w}_1)$ and function $\mu(t_x)$ in Figure 5.

6.3.4. Use of Criteria for Ranking. Given a set of criteria and Configurable Specifications, two kinds of *decision rules* are needed: (i) to establish a ranking of configurations, from the most desirable to the least desirable, in a Configurable Specification; (ii) to establish a ranking of configuration sequences that can be generated with a Configurable Specification. A configuration sequence is generated from the configurations and evolution requirements in a Configurable Specification.

RPAS, Configurable Specification, and configuration concepts are too general to prescribe a particular decision rule. This also reflects the idea that no universal decision rule can be given: every decision rule gives one or more criteria priority over others, so that domain- and/or project-independent decision rules should be much more interesting. Note that the decision rule to apply will depend on how tradeoffs are resolved,

through, among others, stakeholder negotiations, these concerns remain outside the scope of RPAS definition.

Various decision rules can be used to rank configurations. Examples are:

- **R1:** Rank highest the configuration S if S maximizes some quantitative variable v , and rank other configurations in a descending order by value of v that each achieves: i.e., S s.t., $\text{MAX}(v)$ iff $\forall S', S'$ is a configuration and $\text{max}(\text{VAL}(S, v)) \geq \text{max}(\text{VAL}(S', v))$, where max returns the largest constant in a given set. The variable v may be quantifying utility, for example, and its value may be obtained through the aggregation of values of quantitative variables appearing in each configuration; we do not address how exactly the aggregation is done, but contributions on multi-criteria decision-making may be relevant.
- **R2:** Same as R1, but rank highest the configuration which minimizes v , i.e., rank from $\text{MIN}(v)$ to configuration which results in the highest value of v .
- **R3:** Same as R1, augmented for the following condition: the chosen configuration should satisfy the maximal number of optional and preferred requirements, where in a preference relation $\phi \succ \psi$ or $\phi \succeq \psi$, ϕ is the preferred requirement.

R1 and R2 ignore preferences and optional requirements which are independent of v , and both give highest priority to the value of a chosen variable. In that case, the decision rule requires us to solve a single-objective optimization problem. The third rule is significantly different, as it has two high priorities, namely the value of the chosen quantitative variable and the number of optional and preferred requirements. The application of the third rule above requires the resolution of a multi-objective optimization problem.

The decision-rules that ranks configuration sequences can concern the maximization or minimization of a quantitative variable over a sequence of configurations from Configurable Specification. We may ask, with such a decision rule, that optimization satisfies constraints on efficiency of evolution. For example, the decision rule may be as follows:

- **R4:** Choose configuration sequence X if (i) it maximizes the sum of values of the quantitative variable v over all configurations in X , under the constraints that (ii) v must never fall below value q in X , and (ii) no configuration in X differs from the configuration which immediately precedes it by more than w requirements.

R4 is a decision rule pattern, where a concrete rule is obtained by setting q and w to constants. R4 illustrates that constraints need not be limited on individual configurations, but also on the sequence of configurations, as it limits the number of changes that can be performed, which may reflect the necessity to conserve resources during evolution.

The choice of a Configurable Specification does not impose one sequence of configurations that the system should follow: adaptive behavior means that, for example, if we adopt rule R3 above, the system will start from the configuration in which the domain assumptions are consistent with the conditions in the environment of the system. It will then adapt according to the satisfaction/failure of monitored requirements. It may be relevant to dynamically (at runtime) compute new configurations, the discussion of which we leave as an open issue.

7. HOW EXACTLY IS RE FOR ADAPTIVE SYSTEM NON-STANDARD?

Monitoring and control of requirements are not ideas that came with RE for Adaptive Systems. They have been around at least since mid-1990s [Fickas and Feather 1995]. In a sense, any change of requirements that is intended to broaden the range of stimuli (desirable or undesirable) that a system can observe and react to, is a matter of

requirements adaptation. And this is hardly a new issue in RE. Relaxation, be it probabilistic or fuzzy, is also not a new topic; the former has been discussed at least since 2004 [Letier and van Lamsweerde 2004], the latter at least since 1996 [Liu and Yen 1996] in relation to imprecise requirements (i.e., softgoals in this paper).

A novelty in Adaptive Systems RE is that in order to provide methodological support to solve RPAS, it is necessary to design concepts, ontologies, formalisms, and other tools for dealing with both anticipated change in the system environment, and unanticipated change.

Design of configurations and Configurable Specifications is a response to anticipated changes. Fuzzy and probabilistic relaxation are responses only to some forms of unanticipated change, and are insufficient: they still assume that all changes will produce stimuli that are observable using sensors engineered into the system, and that the values of measurements through these sensors and on those stimuli will fall within ranges covered by relaxed requirements.

We have identified no mechanism in this paper, which deals with change that fails these properties, i.e., which involves stimuli for which there are no sensors, or where measurements result in values outside the scales defined through requirements relaxation.

Novelty, then, is in the necessity to have RE frameworks that are capable of both the specification of, and reasoning about monitoring and control considerations, and of probabilistic and fuzzy relaxation. And if the aim is to also address a broader range of unanticipated changes, then there is a need for conceptual tools that are unclear to us at the moment. This also answers the question of what to retain, discard from or add to existing RE modelling languages, frameworks, methodologies to address Adaptive Systems RE: such artifacts and methodologies are likely to need means for requirements relaxation and the specification of evolution requirements, to be relevant for Adaptive Systems RE.

If we take the framework introduced in the paper as one that has means for doing RE while taking monitoring and control aspects, and fuzzy and probabilistic relaxation into account, then we can give more precise observations on how RPAS and Configurable Specification differ from the ZJ RP and standard RE.

There is in the ZJ RP no information to define comparison criteria. We offered a revised definition of the requirements problem [Jureta et al. 2008] (JMF hereafter), which has a richer ontology (to account for concepts that have established themselves in RE, such as, for example, goals and softgoals), allows inconsistencies between alternative configurations (each of which satisfies ZJ conditions), introduces preferences to allow the comparison of configurations in terms of desirability, and we argued that logical consequence in early RE is more adequately described by a nonmonotonic and paraconsistent consequence relation \vdash_{\neg} [Jureta et al. 2010] than \vdash from classical logic. These revisions make it hard to synthesize the requirements problem as $K, S \vdash R$, since consistency and achievement become only a *subset* of the conditions that a solution should satisfy. Note that alternative solutions are indistinguishable in ZJ RP because of the absence of comparison criteria.

Both the ZJ RP and JMF formulations of the requirements problem place considerable emphasis on properties which can be established from propositional variables alone. With ZJ RP, the aim was to design one solution which is consistent and achieves requirements. With JMF, the aim was to design several configurations, all of which are consistent and achieve a threshold (i.e., all mandatory requirements), then proceed to decision-making, that is, compare configurations in terms of which preferred and optional requirements they satisfy, then choose one of them.

In RPAS, the aim is to choose Configurable Specifications, whereby the configurations in a Configurable Specification all are consistent and achieve requirements (as

in ZJ RP), but also distinguish mandatory from optional requirements and can be compared in terms of preference (as in JMF). By satisfying Consistency and Functional threshold achievement, each configuration in a Configurable Specification satisfies the ZJ RP conditions. By satisfying Conformity in addition, every configuration satisfies all properties that a candidate solution should satisfy in JMF. From there on, the departures of the configuration, evolution requirement, Configurable Specification, and RPAS concepts from ZJ RP and JMF are:

- (1) A configuration satisfies Quantitative threshold achievement, Dominance, and Minimality, in addition to Consistency, Functional threshold achievement, and Conformity.
- (2) Decision rules for the ranking of configurations and of Configurable Specifications can be formulated as an optimization problem, where the aim is to optimize one or more quantitative variables. This was not of interest in ZJ RP, and was very limited in JMF.
- (3) There are new kinds of preferences and tradeoffs to consider. While a configuration may give an optimal value of a variable in one Configurable Specification, we may prefer another Configurable Specification, which fails to include that configuration, but ensures some suboptimal, but acceptable value of the variable over several configurations in the Configurable Specification. Instead of focusing on the desirability of individual solutions, it is necessary to look at the desirability of Configurable Specifications.

As the problem of designing, finding, and ranking Configurable Specifications, RPAS reflects the consequences of adaptation behavior: a Configurable Specification is a set of configurations, which need not have a predefined sequence, but evolution requirements, while the system needs to be designed so as to be able to monitor requirements, and satisfy evolution requirements by controlling its behavior, whereby it switches between the highest ranking, yet feasible configurations.

8. CONCLUSIONS AND OPEN ISSUES

We have presented two results. Firstly, by defining the configuration, evolution requirement, and Configurable Specification concepts, we have suggested a precise definition of the RP for Adaptive Systems. We have related these concepts to the key notions from RE of Adaptive Systems, namely, monitoring, control, evolution requirements, probabilistic and fuzzy relaxation. This led us to argue that there are fundamental differences between Zave & Jackson's and our conceptions of the requirements problem, and the RPAS, its solution concepts, and the decision rules used to rank solutions. Secondly, we have used a simple modelling framework throughout the paper, based on *Techné* [Jureta et al. 2010], and which serves as a proto-framework, being an illustration of features needed in future requirements modelling languages relevant to RE for Adaptive Systems.

There are a number of interesting open issues. The shift to configurations and Configurable Specifications suggests that research into extended planning and single- and multi-objective optimization may be a source for further advances in RE frameworks and tool support. It is necessary to look into how the RPAS relates to mixed-integer programming, and mixed-variable programming, how to make adaptation rules on the basis of Configurable Specifications, and what decision rules may be relevant for decision-making in RE. New frameworks for solving RPAS need to be connected to frameworks for controlling the evolution of services, for example, which are capable of addressing structural, behavioral, and quality of service changes [Andrikopoulos et al. 2012]. Work is needed on the integration of probabilistic and fuzzy relaxation, monitor-

ing and control within *practical* modelling languages, for early and late requirements phases.

Our focus in this paper was on the modelling and reasoning aids for the design of the technological side of evolving socio-technical systems. Our conceptualisation of the RPAS did not pay particular attention to potentially complicated interactions between creating Adaptive Systems and the organisational, economic, social, or political aspects of environment in which they are created and/or will run. Arriving at the current RPAS and its solution concept, while relating it to a considerable body of existing work has been itself a challenge. This obliged us to leave outside the scope of our work the interplay between the possibilities, or affordances [Zammuto et al. 2007], that RPAS, Configurable Specification, and the resulting Adaptive Systems do and will offer to organisations, the interdependencies between solving RPAS and influencing or designing the human organisation that will use it, and the effect of understanding these interdependencies on the very formulation of the RPAS, its solution concept, and conceptual tools for finding solutions. This paper will hopefully inform these future efforts.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

REFERENCES

- V. Andrikopoulos, S. Benbernou, and M. P. Papazoglou. 2012. On the Evolution of Services. *IEEE Trans. Software Eng.* 38, 3 (2012), 609–628.
- Anonymous. 1993. *Report of the Inquiry Into The London Ambulance Service*. Technical Report. The Communications Directorate, South West Thames Regional Authority.
- L. Baresi, L. Pasquale, and P. Spoletini. 2010. Fuzzy Goals for Requirements-driven Adaptation. In *IEEE Int. Req. Eng. Conf.*
- B. W. Boehm. 1988. A Spiral Model of Software Development and Enhancement. *IEEE Computer* 21, 5 (1988), 61–72.
- B. W. Boehm, P. K. Bose, E. Horowitz, and M. J. Lee. 1995. Software Requirements Negotiation and Renegotiation Aids: A Theory-W Based Spiral Approach. In *ICSE*, Dewayne E. Perry, Ross Jeffrey, and David Notkin (Eds.). ACM, 243–253.
- F. P. Brooks. 1986. No Silver Bullet - Essence and Accidents of Software Engineering (Invited Paper). In *FIP Congress*. 1069–1076.
- J. Castro, M. Kolp, and J. Mylopoulos. 2002. Towards requirements-driven information systems engineering: the Tropos project. *Inf. Syst.* 27, 6 (2002), 365–389.
- B. H. C. Cheng, R. de Lemos, H. Giese, P. Inverardi, and J. Magee (Eds.). 2009. *Software Engineering for Self-Adaptive Systems [outcome of a Dagstuhl Seminar]*. Lecture Notes in Computer Science, Vol. 5525. Springer.
- A. Dardenne, A. van Lamsweerde, and S. Fickas. 1993. Goal-Directed Requirements Acquisition. *Sci. Comput. Program.* 20, 1-2 (1993), 3–50.
- R. Darimont and A. van Lamsweerde. 1996. Formal Refinement Patterns for Goal-Driven Requirements Elaboration. In *SIGSOFT FSE*.
- R. de Lemos, H. Giese, H. A. Müller, and M. Shaw (Eds.). 2013. *Software Engineering for Self-Adaptive Systems II - International Seminar, Dagstuhl Castle, Germany, October 24-29, 2010 Revised Selected and Invited Papers*. Lecture Notes in Computer Science, Vol. 7475. Springer.
- N. A. Ernst, A. Borgida, I. J. Jureta, and J. Mylopoulos. 2013. Agile Requirements Engineering via Paraconsistent Reasoning. *Information Systems* In press (2013).
- M. S. Feather, S. Fickas, A. Van Lamsweerde, and C. Ponsard. 1998. Reconciling System Requirements and Runtime Behavior. In *IWSSD*. IEEE Computer Society, Washington, DC, USA, 50.
- S. Fickas and M. S. Feather. 1995. Requirements monitoring in dynamic environments. In *IEEE Int. Req. Eng. Conf.* IEEE Computer Society, 140–147.
- S. Greenspan, J. Mylopoulos, and A. Borgida. 1994. On formal requirements modeling languages: RML revisited. In *Proc. 16th Int. Conf. Software Eng.* 135–147.
- A. Hunter and B. Nuseibeh. 1998. Managing Inconsistent Specifications: Reasoning, Analysis, and Action. *ACM Trans. Softw. Eng. Methodol.* 7, 4 (1998), 335–367.

- I. J. Jureta, A. Borgida, N. A. Ernst, and J. Mylopoulos. 2010. Techne: Towards a New Generation of Requirements Modeling Languages with Goals, Preferences, and Inconsistency Handling. In *IEEE Int. Req. Eng. Conf.*
- I. J. Jureta, J. Mylopoulos, and S. Faulkner. 2008. Revisiting the Core Ontology and Problem in Requirements Engineering. In *IEEE Int. Req. Eng. Conf.* IEEE Computer Society.
- E. Letier and A. van Lamsweerde. 2004. Reasoning about partial goal satisfaction for requirements and design engineering. In *SIGSOFT FSE*, Richard N. Taylor and Matthew B. Dwyer (Eds.). ACM, 53–62.
- X. F. Liu and J. Yen. 1996. An analytic framework for specifying and analyzing imprecise requirements. In *Proceedings of the 18th international conference on Software engineering*. IEEE Computer Society, 60–69.
- J. Mylopoulos, L. Chung, and B. Nixon. 1992. Representing and Using Nonfunctional Requirements: A Process-Oriented Approach. *IEEE Trans. Softw. Eng.* 18, 6 (1992), 483–497. DOI: <http://dx.doi.org/10.1109/32.142871>
- B. Nuseibeh, J. Kramer, and A. Finkelstein. 1994. A Framework for Expressing the Relationships Between Multiple Views in Requirements Specification. *IEEE Trans. Software Eng.* 20, 10 (1994), 760–773.
- N. A. Qureshi, I. J. Jureta, and A. Perini. 2011. Requirements Engineering for Self-Adaptive Systems: Core Ontology and Problem Statement. In *Conference on Advanced Information Systems Engineering*. London.
- W. N. Robinson. 2006. A requirements monitoring framework for enterprise systems. *Requir. Eng.* 11, 1 (2006), 17–41.
- W. N. Robinson. 2008. Extended OCL for Goal Monitoring. *ECEASST 9* (2008).
- W. N. Robinson, S. D. Pawlowski, and V. Volkov. 2003. Requirements interaction management. *ACM Comput. Surv.* 35, 2 (2003), 132–190.
- V. E. S. Souza, A. Lapouchnian, K. Angelopoulos, and J. Mylopoulos. 2012. Requirements-driven software evolution. *Computer Science - Research and Development* (2012), 1–19. DOI: <http://dx.doi.org/10.1007/s00450-012-0232-2>
- V. E. S. Souza, A. Lapouchnian, W. N. Robinson, and J. Mylopoulos. 2013. Awareness Requirements. In *Software Engineering for Self-Adaptive Systems II*, Rogério Lemos, Holger Giese, Hausi A. Müller, and Mary Shaw (Eds.). Lecture Notes in Computer Science, Vol. 7475. Springer, 133–161.
- A. van Lamsweerde, R. Darimont, and E. Letier. 1998. Managing Conflicts in Goal-Driven Requirements Engineering. *IEEE Trans. Software Eng.* 24, 11 (1998), 908–926.
- A. van Lamsweerde and E. Letier. 2000. Handling Obstacles in Goal-Oriented Requirements Engineering. *IEEE Trans. Software Eng.* 26, 10 (2000), 978–1005.
- J. Whittle, P. Sawyer, N. Bencomo, B. H. C. Cheng, and J.-M. Bruel. 2010. RELAX: a language to address uncertainty in self-adaptive systems requirements. *Requir. Eng.* 15, 2 (2010), 177–196.
- E. S. K. Yu and J. Mylopoulos. 1994. Understanding "Why" in Software Process Modelling, Analysis, and Design. In *Proc. 16th Int. Conf. Software Eng.* 159–168.
- Raymond F Zammuto, Terri L Griffith, Ann Majchrzak, Deborah J Dougherty, and Samer Faraj. 2007. Information technology and the changing fabric of organization. *Organization Science* 18, 5 (2007), 749–762.
- P. Zave and M. Jackson. 1997. Four Dark Corners of Requirements Engineering. *ACM T. Softw. Eng. Methodol.* 6, 1 (1997), 1–30.

Received April 2013; revised December 2013; accepted April 2014

Online Appendix to: The Requirements Problem for Adaptive Systems

IVAN J. JURETA, Fonds de la Recherche Scientifique – FNRS and University of Namur
ALEXANDER BORGIDA, Rutgers University
NEIL A. ERNST, University of British Columbia
JOHN MYLOPOULOS, University of Trento

A. FORMALISATION

The aim in this section is to define the modelling language used throughout the paper. The modelling language is called *Techné 2* (T2 hereafter), as it extends *Techné 1* [Ernst et al. 2013] by adding quantitative variable requirements, quality constraints, softgoals, and preferences.

A.1. Language

The set of requirements \mathcal{L} is the union of four disjoint sets of formulas: simple requirements on propositional variables \mathcal{L}^P , simple requirements on quantitative variables \mathcal{L}^N , simple softgoals \mathcal{L}^S , and complex requirements \mathcal{L}^C .

Every **simple functional requirement**, i.e., every $a \in \mathcal{L}^P$, satisfies the following BNF specification:

$$a ::= \mathbf{k}(p) \mid \mathbf{g}(p) \mid \mathbf{t}(p) \quad (\text{A.1})$$

where p is a symbol for a propositional variable. A requirement over a propositional variable can be a domain assumption (**k**), a goal (**g**), or a task (**t**). The informal reading of the members of \mathcal{L}^P is

- k**(p): it is believed that p is satisfied;
- g**(p): it is desired that p be satisfied;
- t**(p): the execution of the task makes p satisfied.

Every **simple quantitative variable requirement**, i.e., every $b \in \mathcal{L}^N$, satisfies the following BNF specification:

$$x ::= n \mid v \mid x + x \mid x - x \mid x \cdot x \mid x/x \mid x^x \quad (\text{A.2})$$

$$y ::= x > x \mid x < x \mid x = x \mid x \geq x \mid x \leq x \mid x \neq x \mid x \sim pdf \quad (\text{A.3})$$

$$b ::= \mathbf{k}(y) \mid \mathbf{q}(y) \mid \mathbf{t}(y) \quad (\text{A.4})$$

where n is a real number, v is a quantitative variable, and pdf is a probability density function. The informal reading is:

- k**(y): it is believed that the condition y is satisfied;
- q**(y): it is desired that the condition y be satisfied;
- t**(y): the execution of the task makes y satisfied.

Simple softgoals are kept outside \mathcal{L}^P and \mathcal{L}^N because the natural language statement in a softgoal (e.g., *High usability*, *High performance*), denoted \tilde{p} , refers to some desirable values of variables in such a way that it is not clear which exact values, or how the values of these variables can be observed or measured. Every $c \in \mathcal{L}^S$ satisfies:

$$c ::= \mathbf{s}(\tilde{p}). \quad (\text{A.5})$$

A \tilde{p} is called the content of a softgoal, and is a vague proposition. By being vague, \tilde{p} is neither a propositional variable, nor a quantitative variable, nor a condition on quantitative variables. The symbol used is intentionally different from those inside domain assumptions, goals, quality constraints, and tasks.

A **complex requirement** relates simple requirements and/or softgoals, or is a mandatory or optional requirement. Every $h \in \mathcal{L}^C$ satisfies the BNF specification

$$d ::= a \mid b \mid c \quad (\text{A.6})$$

$$e ::= \bigwedge_{i=1}^{m \geq 1} d_i \rightarrow d \mid \bigwedge_{i=1}^{m \geq 2} d_i \rightarrow \perp \quad (\text{A.7})$$

$$f ::= \mathbf{i}(e) \quad (\text{A.8})$$

$$g ::= d^M \mid d^O \mid f \quad (\text{A.9})$$

As a convention, we consider every complex requirement with an implication connective as information that defines a configuration. This is because we consider such statements as not being given in an RPAS, but part of what we add during the design of the solution.

We use lowercase letters of the Greek alphabet to denote generic requirements, members of $\mathcal{L} = \mathcal{L}^P \cup \mathcal{L}^N \cup \mathcal{L}^S \cup \mathcal{L}^C$. Uppercase Greek letters denote sets of requirements. Preferences over simple requirements and softgoals are kept in the separate set.

The set of **preferences** \mathcal{P} includes all preference relations over simple requirements and/or softgoals; every $w \in \mathcal{P}$ conforms to the specification

$$w ::= d \succeq d \mid d \succ d \mid d \approx d. \quad (\text{A.10})$$

Note that this disallows preferences between complex requirements; and preferences themselves cannot be marked mandatory or optional, nor can they participate in a refinement, or in a conflict.

Definition A.1. The language is the tuple $(P, \mathbb{R}, V, \tilde{P}, \mathcal{L}, \mathcal{P})$, where \mathcal{L} the set of requirements which satisfy the BNF specification in Equations A.1–A.9 and \mathcal{P} is the set of preferences which satisfy the BNF specification in Equation A.10, whereby the propositional variables in requirements and preferences come from the set P , numbers from the set \mathbb{R} of real numbers, quantitative variables from V , and vague propositions from \tilde{P} .

Definition A.2. Any $\Delta \subseteq \mathcal{L} \cup \mathcal{P}$ is a requirements database.

A.2. Consequence Relation

We use the following consequence relation, defined by extending the Techne consequence relation to allow deduction over statements with quantitative variables.

Definition A.3. The **consequence relation** $\vdash_{\tilde{\gamma}}$ is such that, for $\Pi \subseteq \mathcal{L}$, $\phi \in \mathcal{L}$ and $x \in \{\phi, \perp\}$:

- $\Pi \vdash_{\tilde{\gamma}} \phi$ if $\phi \in \Pi$, or
- $\Pi \vdash_{\tilde{\gamma}} x$ if for every i such that $1 \leq i \leq n$, $\Pi \vdash_{\tilde{\gamma}} \phi_i$ and $\mathbf{i}(\bigwedge_{i=1}^n \phi_i \rightarrow x) \in \Pi$, or
- If every y_1, \dots, y_n is a statement about real values and variables that conforms to Eq. A.3, and if $\{y_1, \dots, y_{n-1}\} \vdash y_n$ in real arithmetic, then $\mathbf{x}_1(y_1), \dots, \mathbf{x}_{n-1}(y_{n-1}) \vdash_{\tilde{\gamma}} \mathbf{x}_n(y_n)$, where $\mathbf{x}_1, \dots, \mathbf{x}_{n-1} \in \{\mathbf{k}, \mathbf{t}\}$ and $\mathbf{x}_n \in \{\mathbf{k}, \mathbf{t}, \mathbf{q}\}$.

The consequence relation $\vdash_{\tilde{\gamma}}$ is sound with regards to standard entailment \vdash in classical propositional logic, but is incomplete in two ways: it only considers deducing posi-

tive atomic facts, and no ordinary proofs based on arguing by contradiction go through, thus being paraconsistent.

A.3. Operationalisation Function

The purpose of the operationalisation function is, given a requirement ϕ and a set of requirements X , to return, if it exists, every set $Y \subseteq X$ which satisfies the following conditions. If ϕ is a requirement over a propositional variable, then Y includes only those domain assumptions and tasks which are necessary to deduce ϕ . If ϕ is a requirement over a quantitative variable, then Y includes only those domain assumptions and tasks which ensure that the condition in ϕ is satisfied, when all of the variables in that condition are replaced by the values that they are assigned by the requirements in Y .

To define the operationalisation function, we define below the Select function and one useful set. We assume that $\mathcal{O} = \{\mathbf{k}, \mathbf{g}, \mathbf{q}, \mathbf{s}, \mathbf{t}, \mathbf{i}\}$ and we denote by $\wp(X)$ the powerset of X .

Definition A.4. The select function

$$\text{Select} : \mathcal{O} \times \{\text{"empty"}, \mathbf{o}, \mathbf{m}\} \times \wp(\mathcal{L}) \longrightarrow \wp(\mathcal{L}) \quad (\text{A.11})$$

is defined as follows, for $\mathbf{x} \in \mathcal{O}$, $\mathbf{y} \in \{\text{"empty"}, \mathbf{o}, \mathbf{m}\}$, and $\Pi \subseteq \mathcal{L}$:

$$\text{Select}(\mathbf{x}, \mathbf{y}, \Pi) \stackrel{\text{def}}{=} \{\phi \mid \mathbf{x}(\phi)^{\mathbf{y}} \in \Pi\}. \quad (\text{A.12})$$

To simplify notation, we use the following abbreviation:

$$\text{Select}(\mathbf{x}, \mathbf{y}, \Pi) \equiv \Pi_{\mathbf{x}}^{\mathbf{y}}. \quad (\text{A.13})$$

Given a set of expressions, the Select function returns its subset, in which all members have the labels $\mathbf{x} \in \mathcal{O}$ and $\mathbf{y} \in \{\text{"empty"}, \mathbf{o}, \mathbf{m}\}$.

Definition A.5. Useful set: Let $\Pi \subseteq \mathcal{L}$. Then

$$\text{CON}(\Pi) \stackrel{\text{def}}{=} \{\Phi \subseteq \Pi \mid \Phi \not\vdash_{\tau} \perp \text{ and } \Delta_{\mathbf{i}} \cup \bigcup_{\mathbf{x} \in \mathcal{O}} \Delta_{\mathbf{x}}^{\mathbf{m}} \subseteq \Phi\} \quad (\text{A.14})$$

$\text{CON}(\Pi)$ is the set of all consistent subsets of Π , whereby every one of these subsets must include all mandatory requirements from Δ .

The reason why every $\phi \in \text{CON}(\Pi)$ must include all mandatory requirements is that we require the set of mandatory requirements to be consistent and, as we will see below, to be included in all operationalisations, because an operationalisation must not be inconsistent with mandatory requirements.

Definition A.6. The operationalisation function

$$\text{Op} : \bigcup_{\forall \mathbf{x}, \mathbf{y}} \Delta_{\mathbf{x}}^{\mathbf{y}} \longrightarrow \wp \left(\wp \left(\bigcup_{\forall \mathbf{z}, \mathbf{w}} \Delta_{\mathbf{z}}^{\mathbf{w}} \right) \right) \quad (\text{A.15})$$

for $\mathbf{y}, \mathbf{w} \in \{\text{"empty"}, \mathbf{o}, \mathbf{m}\}$, $\mathbf{x} \in \{\mathbf{g}, \mathbf{q}, \mathbf{s}\}$, and $\mathbf{z} \in \{\mathbf{k}, \mathbf{t}, \mathbf{i}\}$, is defined as follows:

$$\begin{aligned} \text{Op}(\phi) \stackrel{\text{def}}{=} & \{\Pi \in \text{CON}(\Delta) \mid \Pi \vdash_{\tau} \phi \\ & \text{and } \Pi \setminus \{\phi\} \subseteq \bigcup_{\forall \mathbf{z}, \mathbf{w}} \Delta_{\mathbf{z}}^{\mathbf{w}}, \\ & \text{and } \nexists \Phi \subset \Pi, \Phi \vdash_{\tau} \phi\}. \end{aligned} \quad (\text{A.16})$$

Every member of the set $\text{Op}(\phi)$ is a minimal consistent set of tasks and domain assumptions that is sufficient to operationalise the goal, quality constraint, or softgoal

ϕ . Informally, $\text{Op}(\phi)$ tells us all ways in (i.e., subsets of) Δ of satisfying a goal, quality constraint, or softgoal, which are consistent with all mandatory requirements.

A.4. Softgoal Macros

Softgoal macros automatically change the requirements database when it already contains some specific requirements.

Macros rely on the special *monitoring function* VAL . Given a set of requirements $X \subseteq \Delta$, and a quantitative variable v , if there are tasks and domain assumptions in X which assign a value $n \in \mathbb{R}$ to v (e.g., a task such that $\mathbf{t}(v = n)$ where n is a constant), then $n \in \text{VAL}(X, v)$.

Definition A.7. The **monitoring function**

$$\text{VAL} : \wp(\Delta) \times V \longrightarrow \wp(\mathbb{R}) \quad (\text{A.17})$$

is defined as follows, for $X \subseteq \Delta$, $v \in V$, $\mathbf{x} \in \{\mathbf{k}, \mathbf{q}, \mathbf{t}\}$, and $\mathbf{y} \in \{\text{"empty"}, \mathbf{o}, \mathbf{m}\}$:

$$\text{VAL}(X, v) = \{n \in \mathbb{R} \mid \exists Y \subseteq X \text{ s.t. } Y \in \text{Op}(\mathbf{x}(v = n)^{\mathbf{y}})\} \quad (\text{A.18})$$

$\text{VAL}(X, v)$ is a set of all constants that are assigned through operationalisations in X to v .

The softgoal macro applies to any quantitative variable v for which $\text{VAL}(\Delta, v) \neq \emptyset$ and for which there is a satisfaction function $\mu(v)$. This is the case when there is in Δ an operationalisation of a requirement which assigns a value to v and a domain assumption which defines a variable v' as the output of the satisfaction function on v , i.e., there exists $\mathbf{k}(v' = \mu(v)) \in \Delta$. The softgoal macro automatically adds preference relations between sets of requirements which operationalise requirements that assign constants to v , whereby the preferences are defined to reflect the values returned by μ .

Definition A.8. For every quantitative variable $v \in V$ such that $\text{VAL}(\Delta, v) \neq \emptyset$ and there exists $\mathbf{k}(v' = \mu(v)) \in \Delta$, where μ is informally interpreted as a satisfaction function, there is the following **softgoal macro**:

```

s(
   $\forall x_1, x_2 \in \text{VAL}(\Delta, v), x_1 \neq x_2,$ 
  IF  $\mu(x_1) = \mu(x_2)$  THEN ADD
     $\{\mathbf{k}(v = x_1), \mathbf{k}(v = x_2),$ 
      $\mathbf{k}(v = x_1) \approx \mathbf{k}(v = x_2)\}$ 
  TO  $\Delta;$ 
  IF  $\mu(x_1) > \mu(x_2)$  THEN ADD
     $\{\mathbf{k}(v = x_1), \mathbf{k}(v = x_2),$ 
      $\mathbf{k}(v = x_1) > \mathbf{k}(v = x_2)\}$ 
  TO  $\Delta;$ 
)
```

Two remarks are in order. Firstly, note that the preferences are added to reflect the values not of v but of the satisfaction level $\mu(v)$ on v . Secondly, note the difference between the definition of the softgoal macro above, and the macro we used in Section 6.3.2: there, the macro was simpler because we assumed that we already knew the configurations.

A.5. Configuration Concept

Definition A.9. A **configuration** S , defined from the requirements database Δ , is a set

$$S \subseteq \bigcup_{\forall \mathbf{x}, \mathbf{y}} \Delta_{\mathbf{x}}^{\mathbf{y}}, \text{ for } \mathbf{x} \in \{\mathbf{k}, \mathbf{t}, \mathbf{i}\}, \mathbf{y} \in \{\text{"empty"}, \mathbf{o}, \mathbf{m}\} \quad (\text{A.19})$$

of domain assumptions, tasks, imply and conflict relations, which satisfies the following properties:

- (1) Consistency: $S \not\vdash \perp$;
- (2) Functional threshold achievement: $\forall \phi \in \Delta_{\mathbf{z}}^{\mathbf{m}}, \mathbf{z} \in \{\mathbf{g}, \mathbf{s}\}$, there exists $\Pi \in \text{Op}(\phi)$ such that $\Pi \subseteq S$;
- (3) Quantitative threshold achievement: $\forall \phi \in \Delta_{\mathbf{q}}^{\mathbf{m}}, \exists \Gamma \in \text{Op}(\phi)$ such that $\Gamma \subseteq S$;
- (4) Conformity: $\Delta_{\mathbf{k}}^{\mathbf{m}} \cup \Delta_{\mathbf{t}}^{\mathbf{m}} \subseteq S$;
- (5) Dominance: $\nexists S'$ s.t. both S and S' satisfy conditions 1–4 above, $S \subset S'$, and $\exists \Delta_{\mathbf{x}}^{\circ} = S' \setminus S$, s.t. $\Delta_{\mathbf{x}}^{\circ} \neq \emptyset$ and $\mathbf{x} \in \{\mathbf{k}, \mathbf{t}\}$;
- (6) Minimality: $\nexists S'$ such that S' satisfies the conditions 1–5 above and $S' \subset S$.

The Threshold achievement property requires that a configuration satisfies all mandatory goals, quality constraints, and softgoals. Satisfaction depends on the presence of operationalisations in S , for each of the mandatory goals, quality constraints, and softgoals.

The Conformity property asks that all strict domain assumptions are not violated and all mandatory tasks are executed. The Achievement and Conformity properties ensure that the configuration satisfies all that *must* be satisfied.

According to the Dominance property, every configuration will be maximal with regards to optional requirements. This property formalises the idea of the optional relation, as holding on requirements which are desirable to satisfy, but can be violated. A configuration should include as many defeasible domain assumptions and as many optional tasks, up to the point at which adding any further defeasible domain assumptions and/or optional tasks violates the Consistency, Functional and Quantitative threshold achievement, Conformity, or Minimality properties. The Dominance condition ensures that every configuration is Pareto efficient with regards to optional requirements, as this condition makes it impossible to add optional domain assumptions and tasks to any S and still ensure that S is a configuration.

The Minimality property requires that a configuration includes only the domain assumptions and tasks which are needed to satisfy exactly the Consistency, Threshold achievement, Conformity, and Dominance properties.

B. RELATING THE FORMALISATION TO REQUIREMENTS ENGINEERING FOR ADAPTIVE SYSTEMS

We have illustrated throughout the main paper how the contributions here depart from prior efforts in the understanding of the requirements problem and solution concepts in mainstream RE and in Adaptive Systems RE.

The Configurable Specification concept is inspired by our discussion of the role of contexts, time, and resources in the requirements problem for Adaptive Systems [Qureshi et al. 2011]. We have used *Techne* [Jureta et al. 2010] here as a starting point and extended its syntax to allow requirements which place constraints on quantitative variables. This has resulted in a more general treatment of operationalisation, as we have the operationalisation of both functional and quantitative variable requirements. As in the case of *Techne*, the formalism here does not provide a visual syntax, and is thus not itself a model language for early requirements in the same sense that, for example, KAOS goal models [Dardenne et al. 1993] and *i** actor models [Yu and Mylopoulos 1994] are. Two limitations of *Techne* are overcome here. Firstly, we can make explicit the constraints on quantitative variables. Secondly, *Techne* now distinguishes between stable facts, in the form of mandatory domain assumptions, and defeasible information, in the form of optional domain assumptions.

In the rest of this section, we look at how the language presented in this paper can be used to model information that was recognized as crucial in the research into the relaxation of requirements [Letier and van Lamsweerde 2004; Whittle et al. 2010; Baresi et al. 2010], the evaluation of their partial satisfaction [Letier and van Lamsweerde

2004], and the monitoring and control of requirements [Feather et al. 1998; Robinson 2008]. The aim is not to suggest the language here as a general modelling language, but to show that it covers the main ideas presented in that related research.

B.1. Fuzzy Relaxation

Baresi et al. [Baresi et al. 2010] associate a fuzzy operator with a predefined membership function.

For example, if in our framework there is a quality constraint $\mathbf{q}(v < 6\text{hrs})$ (a goal $G(v < 6\text{hrs})$ for them, as they allow quantitative variables in goals), then relaxing it would amount to replace it with $\mathbf{q}(v <_f 6\text{hrs})$ (in their notation, $\mathcal{G}(v <_f 6\text{hrs})$), where $<_f$ is a fuzzy operator. Their interpretation of $v <_f 6\text{hrs}$ is that there is a fuzzy membership function μ which returns the level of satisfaction as a function of v , and the shape of μ is predefined (for $<_f$, it is positive and constant until $v = 6$, then decreases up to the satisfaction value 0 for some $v > 6$). The operator $<_f$ can be defined in our language by reproducing, in a domain assumption, a function which has the form of the fuzzy membership function for $<_f$, as defined by Baresi et al. We can define as follows a macro which takes a fuzzy goal of the form $\mathcal{G}(v <_f n)$, with $n \in \mathbb{R}$, and transforms it into requirements that can be added to our requirements database Δ :

```

 $\forall \mathcal{G}(v <_f n)$  WHERE  $v \in V$  AND  $n \in \mathbb{R}$  ADD  $\{\mathbf{k}(v' =$ 
 $\mu(v))\}$  TO  $\Delta$ , AND APPLY THE SOFTGOAL MACRO ON
 $\Delta$  and  $v$ ,
```

where $v' \in V$, the set of quantitative variables, and μ is a function which is defined according to the function pattern defined by Baresi et al. for the fuzzy operator $<_f$. It is straightforward to define similar macros for all other fuzzy operators defined by Baresi et al.

Baresi et al. also define binary connectives, such as fuzzy conjunction. These can be defined here as well, as functions of those variables which are defined using fuzzy membership functions. Each binary fuzzy operator gives one function specified in a domain assumption and using an approximation relation. In the formalism from Baresi et al., one way to define fuzzy conjunction \wedge_f between two variables v_1 and v_2 , each of which has an accompanying fuzzy membership function $\mu(v_1)$ and $\mu(v_2)$, is as follows: $\mu(v_1 \wedge_f v_2) = \mu(v_1) \cdot \mu(v_2)$. In our language, $\mu(v_1)$ and $\mu(v_2)$ give satisfaction levels. The fuzzy conjunction connective between two quality constraints, respectively over variables v_1 and v_2 is introduced in Δ as the domain assumption $\mathbf{k}(v_3 = \mu(v_1) \cdot \mu(v_2))$, where v_3 is the joint level of satisfaction over variables v_1 and v_2 .

B.2. Probabilistic Relaxation

To handle idealistic requirements, Letier & van Lamsweerde [Letier and van Lamsweerde 2004] suggest the association of probability estimates to constraints on quantitative variables. This is what we allow in the T2 language, and this has been illustrated earlier (see Section 6.3.1 in the main paper).

B.3. Monitoring & Control Variables

The importance of monitoring and control of requirements is highlighted in Fickas & Feather [Fickas and Feather 1995], Feather et al. [Feather et al. 1998], and Robinson [Robinson 2006]. It is on the basis of the ontology for requirements that we identify controlled variables: if a variable appears in a task, it is a controlled variable. Any variable can be monitored, but all variables in goals, quality constraints, and domain assumptions must be monitored.

B.4. Adaptation Rules

Adaptation rules have been called reconciliation tactics [Feather et al. 1998], adaptive goals [Baresi et al. 2010], adaptivity mechanisms [Souza et al. 2013]. We have not discussed how they ought to be specified or actually implemented, but we have suggested how they can be identified. The Configurable Specification adopted for a system gives a set of configurations and evolution requirements. It consequently indicates what changes between two configurations, i.e., which requirements are dropped, which become optional, which others become mandatory, and so on. Differences between every two consecutive configurations in the Configurable Specification specify the effect that adaptation rules should have, regardless of how they are specified or implemented, while evolution requirements specify constraints that the adaptation rules must not violate.

B.5. Limitations

While it may be possible to specify time and temporal constraints by having a quantitative variable, the values of which are defined by a clock, and to map every configuration to a temporal interval, this would clearly not be a convenient way to model temporal constraints on Configurable Specifications. More appropriate in this respect may be requirements modelling languages built on top of linear temporal logic [Dardenne et al. 1993] or branching temporal logic [Whittle et al. 2010]. It is, however, not clear how requirements models built with such languages relate to Configurable Specifications, that is, are these requirements models describing constraints on a single configuration, or on parts of Configurable Specifications.

Techne has previously been critiqued for ignoring the notion of agent, and the fact that requirements belong to individuals which may thus be modeled as agents. We continue to ignore them here, mainly because they can be introduced in a straightforward manner, yet would complicate notation and presentation, without adding much to the main purpose of this paper.

To add agents, consider first why they need to be added. If the aim is to help figure out who needs to negotiate requirements conflicts, then assume there is a set of identifiers for agents, and add a function which returns, for every requirement, one or more agents who agree on it. If the aim is to assign responsibilities of tasks to agents, then assume a set of identifiers for these agents, and define a function which maps every task to one or more agents responsible for its execution. Either of the two uses of agents has no influence on the structure of the RPAS, other than suggesting that RE does involve negotiation and the assignment of responsibilities.

C. RELATIONS DEFINED USING IMPLICATION AND CONFLICT

We discuss in this section how oft-cited relations between requirements can be defined using our requirements types and the implication and conflict relations. Our aim is to illustrate the versatility of the language we used in the paper, despite its simplicity. We start by introducing the *minimal consistent inference* relation (MCI).

Definition C.1. Let $\Delta \subseteq \mathcal{L}$ and suppose $Impl(\Delta)$ returns all implications in Δ . A proposition q in a T2 requirements database Δ will be said to be in the MCI relation to propositions $\{p_1, \dots, p_n\} \subseteq \Delta$, $n \geq 1$ if and only if:

- (1) $(p_1 \wedge \dots \wedge p_n \rightarrow q) \in Impl(\Delta)$;
- (2) $\nexists \Pi. \Pi \subset \{p_1, \dots, p_n\}$ and $\Pi \cup Impl(\Delta) \not\vdash q$;
- (3) $\nexists \gamma. \gamma \in Impl(\Delta)$, $\{\gamma\} \cup \{\bigwedge_{i=1}^n p_i\} \vdash \perp$.

The first condition requires that there be an implication between the propositions. The second, minimality condition requires that there be no subset of the premises from

which the consequence can be deduced via \vdash_{τ} . The third condition requires that the premises be consistent in one step.

We use MCI to define the concept of *argument* in T2. An argument puts together the premises and the conclusion which stand in an MCI relation.

Definition C.2. The pair (Π, q) is an argument in a requirements theory Δ if and only if:

- (1) $\Pi \subseteq \Delta$,
- (2) q is in the MCI relation to all propositions in $\Pi \setminus Impl(\Delta)$.

By restricting the types of requirements in premises and the conclusion of arguments, we can define a taxonomy of relations, and thereby illustrate that T2 can capture oft-cited relations in requirements modelling languages.

MCI is the root of the relations taxonomy, and is specialized onto the Refinement, Realization, and Conflict relation. Each of these is defined by restricting the types allowed in the premises and the conclusion, as follows:

- Refinement is a relation that concludes a requirement, and not all of its premises are domain assumptions and/or tasks;
- Realization, or operationalisation, is a relation that concludes a requirement, and premises are only domain assumptions and/or tasks;
- Conflict is a relation that concludes \perp .

It is then straightforward to observe the following:

- The goal refinement relation from Darimont and van Lamsweerde [Darimont and van Lamsweerde 1996] can be defined in T2 as a specialization of the Refinement relation. Namely, Goal Refinement in T2 is the Refinement relation, in which all premises and the conclusion are of type Goal. Recall that Darimont & van Lamsweerde defined goal refinement as the relation between a goal being refined and subgoals which refine it, the latter having to satisfy three conditions: (i) be sufficient to deduce the refined goal, (ii) be minimal, and (iii) be consistent. All three conditions are satisfied here, since Goal Refinement is a specialization of Refinement, which is in turn a specialization of MCI.
- Yu and Mylopoulos [Yu and Mylopoulos 1994] introduced task decomposition in i^* (i -star). It is similar to goal refinement, with two differences: (i) the requirement being refined/decomposed must be a task, and (ii) it can be refined by any combination of goals and tasks.² Task Decomposition can be defined in T2 as a specialization of the Realization relation, in which all premises and the conclusion are all of type Task.
- The goal operationalisation relation in KAOS [Dardenne et al. 1993] is similar to the means-ends relation in i^* . The idea of both is that tasks should be executed in order to satisfy goals. Operationalisation in KAOS stands between goals and constraints, whereby a constraint is operational, in the sense that it is formulated in terms of objects and actions available to the agents in/of the system. Means-ends rather emphasizes the role of goals as reasons why tasks are executed, i.e., a task exists in a requirements database because it is a means to a goal. The Goal Operationalisation relation, which captures the idea of both Goal operationalisation and means-ends, is a specialization of the Realization relation, in which all premises are tasks or domain assumptions, and the conclusion is a goal.

²There is no concept in i^* [Yu and Mylopoulos 1994] which corresponds to Domain assumption, so it is not allowed here to have Domain assumptions in a decomposition of a task.

Table II. Common conflict relations translated to T2.

<i>Relation in Robinson and colleagues' survey [Robinson et al. 2003]</i>	<i>Corresponding relation in T2</i>
<i>Process-level deviation</i> : deviation of the actual process of developing the system from the pre-defined process.	Either a Type-B conflict, in which the blocker is a domain assumption stating the deviation between the planned and actual development process, or a Type-C conflict where one of the domain assumptions states that deviation.
<i>Instance-level deviation</i> : an instance of an implemented class violates a requirement.	A Type-B conflict, which has one blocking domain assumption. That domain assumption names the instance responsible for the violation, and the alternative in the Type-B is the requirement violated by that instance.
<i>Terminology clash</i> (also <i>Structure clash</i>): a member of the semantic domain is being referred to using more than one symbol/expression.	None of the conflict relations captures terminology clashes. The terminology clash is an error in the use of a requirements modelling language.
<i>Designation clash</i> : a symbol/expression refers to two or more different members of the semantic domain.	As for the terminology clash, a designation clash is an error in the use of the formalism, and cannot be captured in the formalism.
<i>Conflict</i> : a set of requirements is logically inconsistent.	Conflict relation.
<i>Divergence</i> (also <i>Obstruction</i>): A set of requirements is logically inconsistent when a certain sequence of events can occur.	A Type-B conflict, where the blocked requirements are an alternative and the blocker is a domain assumption describing the problematic sequence of events.
<i>Competition</i> : A kind of divergence where particular instances of a requirement can cause a divergence.	A Type-B conflict, which has one blocking domain assumption. That domain assumption names the instance responsible for the violation, and the alternative in the Type-B is the requirement violated by that instance.

To see the kinds of conflict we can capture using our Conflict relation, we first need the concept of Alternative.

Definition C.3. Given $\Pi \subseteq \Delta$, $\Phi \subset \Pi$ is an **alternative** in Π if and only if:

- (1) There is an argument (Π, \perp) ;
- (2) $\Phi \not\vdash \perp$;
- (3) $\forall \Psi. \Psi \subseteq \Pi$ if $\Psi \not\vdash \perp$ then $\Phi \not\subseteq \Psi$; i.e., Φ is a maximally consistent subset of Π ;
- (4) Φ does not include only implications and/or Domain assumptions. In other words, Φ includes at least one goal and/or task.

The set of all alternatives in Π is denoted $\text{Alt}(\Pi)$.

Using the Alternative concept, we specialize the Conflict relation as follows:

- Type-A Conflict relation is the Conflict relation between premises in the argument (Π, \perp) if and only if there are at least two alternatives in Π , i.e., $|\text{Alt}(\Pi)| \geq 2$.
- Type-B Conflict relation is the Conflict relation between premises in the argument (Π, \perp) if and only if $|\text{Alt}(\Pi)| = 1$. Informally, a Type-B conflict involves domain assumptions which are blocking the satisfaction of goals or the execution of tasks. If instances in Π are in Type-B Conflict, then we have $\text{Block}(\Pi) \stackrel{\text{def}}{=} \Pi - \text{Alt}(\Pi)$, and we call $\text{Block}(\Pi)$ the set of blockers.
- Type-C Conflict relation is the Conflict relation between premises in the argument (Π, \perp) if and only if $|\text{Alt}(\Pi)| = 0$. Type-C Conflict involves a minimally inconsistent set which includes only Domain assumptions.

We can also relate Conflict in T2 with notions of conflict in other requirements modelling languages.

In KAOS, the Conflict relation is also a minimally inconsistent set of requirements. Obstruction and Divergence are two relations, also in KAOS, which involve domain assumptions that block, in the sense discussed above, the satisfaction of a goal or the execution of a task.

Table II summarizes the translation of conflict relations identified in Robinson, Pawlowski & Volkov's survey [Robinson et al. 2003] into T2. Conflicts listed in that table cannot obtain in T2 definitions that are as convenient as, for example, Type-A, Type-B, or Type-C Conflict relations. In a propositional formalism such as T2, there is no elegant way to formally talk about instances of classes, and their deviations: a proposition stating the deviation of an instance will be different from a proposition stating normal behavior of other instances, but there is no relation which would say that the two propositions talk about instances of the same class.