Self-Explanation in Adaptive Systems Based on Runtime Goal-Based Models

Kris Welsh[1(✉)], Nelly Bencomo[2], Pete Sawyer[3], and Jon Whittle[3]

[1] School of Computing, University of Kent, Canterbury, UK
k.welsh@kent.ac.uk
[2] School of Engineering and Applied Science, Aston University, Birmingham, UK
nelly@acm.org
[3] School of Computing and Communications, Lancaster University, Lancaster, UK
{sawyer,whittle}@comp.lancs.ac.uk

Abstract. The behaviour of self adaptive systems can be emergent, which means that the system's behaviour may be seen as unexpected by its customers and its developers. Therefore, a self-adaptive system needs to garner confidence in its customers and it also needs to resolve any surprise on the part of the developer during testing and maintenance. We believe that these two functions can only be achieved if a self-adaptive system is also capable of self-explanation. We argue a self-adaptive system's behaviour needs to be explained in terms of satisfaction of its requirements. Since self-adaptive system requirements may themselves be emergent, we propose the use of goal-based requirements models at runtime to offer self-explanation of how a system is meeting its requirements. We demonstrate the analysis of run-time requirements models to yield a self-explanation codified in a domain specific language, and discuss possible future work.

Keywords: Self-explanation · Self-adaptive · Goals · Claims

1 Introduction

Self-adaptive systems are able to adjust their behaviour according to changes in their operating environment. Uncertainty in the operating environment may cause the behaviour of self-adaptive systems to be emergent. A system whose behaviour cannot be accurately predicted poses serious problems in terms of assurance and acceptance. A lack of intelligibility may cause users to stop using a self-adaptive system [1–3]. Because its behaviour is emergent, a self-adaptive system needs to garner confidence in its stakeholders, and allow developers to

This paper is an extended version of the paper "Towards Requirements Aware Systems: Run-time Resolution of Design-time Assumptions" by Bencomo, Welsh, Sawyer and Whittle, 17th IEEE International Conference International Conference on Engineering of Complex Computer Systems ICECCS 2012, Paris, France, July 2012.

© Springer-Verlag Berlin Heidelberg 2014
R. Kowalczyk and N.T. Nguyen (Eds.): TCCI XVI, LNCS 8780, pp. 122–145, 2014.
DOI: 10.1007/978-3-662-44871-7_5

understand observed behaviour [3]. We believe that these two functions can only be achieved if a self-adaptive system is also capable of self-explanation.

We argue that a self-adaptive system's behaviour is best explained in terms of the satisfaction of its requirements. Observing the degree to which a system satisfies its requirements is well-discussed in requirements monitoring literature [4], and addresses questions of *what* the system is doing. The ability of a self-adaptive system to select alternative configurations based on environmental triggers raises questions on *how* the system is doing it, with more useful explanations offering clues to *why* the system is behaving as observed. Readily-understandable explanations are challenging to produce, with several key challenges preventing developers from creating such functionality. These challenges are discussed in the following paragraphs.

Firstly, an ability to explain behaviour relies upon an ability to monitor, introspect and reason about the system's current and past behaviour. There has been significant research interest in providing support for requirements monitoring [4,5], and in the specific area of self-adaptive systems, advances have also been made towards better support for introspection by adaptive middleware [6,7] and other frameworks [8,9]. However, work seeking to combine these two capabilities with reasoning still needs more research effort. The new and broader research area of requirements-aware systems covers similar interests [3,10,11].

Secondly, explanations need to be created at a sufficiently high level as to be understandable by a variety of interested stakeholders (e.g. end-users, but also by maintainers and support personnel). Ideally, users should interact with the system at a level of abstraction that is meaningful to them. This requires that the system is able to trace backwards and forwards between abstractions at the user's level and abstractions used by the systems at lower levels (e.g. components, component configurations, etc.). Furthermore, a trace of relevant events in the history of the adaptations the system has gone through should be kept by the system.

Thirdly, for self-explanations to be trustable, a self-adaptive system should be able to trace down from goals towards code to keep a synchronized link between requirements and architecture during execution. This trace needs to consider the dynamic changes that will affect requirements and the architecture of the system at runtime and keep a causal connection between the two.

Finally, a self-adaptive system should be able to reproduce a trace history of the adaptations it has performed in a way that is meaningful to support self-explanation.

In [12], we described our view of requirements-aware systems. In our work, representations of assumptions are made explicit using the concept of claims in goal models at design time. Using what we call claim refinement models (CRMs), we have defined the semantics for claims in terms of their impact on alternative strategies that can be used to pursue the goals of the system. The impact is calculated in terms of satisfaction and trade-off of the system's non-functional requirements (modeled as softgoals). Crucially, at runtime, when the executing system monitors that a given claim does not hold anymore, the system may adapt

to an alternative goal realization strategy that may be more suitable for the new contextual conditions. Importantly, our approach tackles uncertainty, i.e. the new goal realization strategy may imply a new configuration of components that was not necessarily foreseen at design time. With the potential for unforeseen behavior, self-explanation capabilities are crucial. In this paper we build on the approach described in [12] to address the challenges posed by self-explanation described above.

The rest of the paper is organized as follows: In Sect. 2 we present the motivation of the paper using a simple but yet useful discussion. In Sect. 3, we discuss our initial progress towards a mechanism by which self-explanation can be achieved. In Sect. 4, we apply this means of providing self-explanation to a short case study. In Sect. 5, we propose a simple domain-specific language in which to convey self-explanations generated using our technique. Section 6 describes relevant related work. Section 7 concludes the paper and discusses future work.

2 Motivating Example

Consider the example of a robotic vacuum cleaner for domestic apartments, which uses self-adaptation to balance two conflicting non-functional requirements: to avoid causing a danger to people within the apartment (*avoid tripping hazard*) and to be economical to run (*minimise energy costs*). The cleaner supports two modes of operation: clean at night and clean when empty. Cleaning at night will likely yield lower energy costs, but could cause the occupants to trip should they awake and move about the apartment. Cleaning when the apartment is empty eliminates this hazard, but if the apartment is only empty during daytime this will come at a cost of increased energy costs. A standard goal model, showing the different ways in which the robot can clean the apartment, and each method's impact on the two competing NFRs (which can be modelled as softgoals) would be deadlocked, with no clear favourable goal operationalisation strategy. We have previously discussed [13] the use of claims, which were first proposed in the Non-functional Requirements (NFR) Framework [14], to model an assumption made to break the deadlock in a goal model. In this case, we can make an assumption that the tripping hazard is unlikely to cause an accident. We illustrate this using an i* [15] Strategic Rationale (SR) model, which models how an agent achieves its goals, and allows alternative goal satisfaction strategies to be compared in terms of their impact on softgoals. The model in Fig. 1 shows a claim "No Tripping Hazard" breaking the deadlock that would otherwise occur.

In Fig. 1, the vacuum cleaner's "Clean Apartment" goal may be satisfied either by the "Clean at night" task, or the "Clean when empty" task. Cleaning at night *helps* satisfy the "Minimise energy costs" softgoal, but *hurts* the "Avoid tripping hazard" softgoal, as represented by the *contribution links* attached to the task. The "Clean when empty" task makes the inverse contributions to each of the softgoals.

The "No tripping hazard" claim *breaks* the negative contribution made to the "Avoid tripping hazard" softgoal by the "Clean at night" task, which means

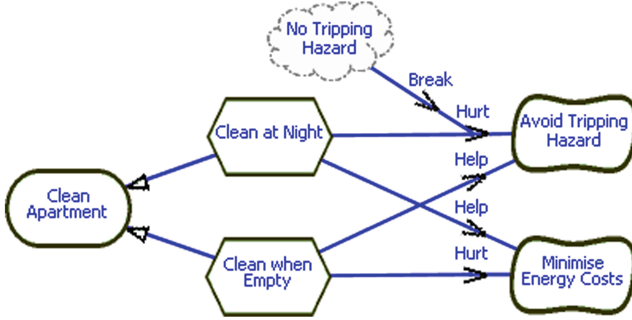

Fig. 1. Goal model of a robot vacuum cleaner from [16]

that this contribution should be lent less credence, or disregarded completely when deciding between the competing goal operationalisation strategies. With this assumption made, the decision to clean at night follows naturally.

Although assuming that the tripping hazard doesn't pose any real risk makes for a convenient way to break the deadlock in the goal model, the assumption is mere conjecture and would prove difficult to verify at design time. Thus, the robot vacuum cleaner is provided with a means of verifying the assumption at runtime, using monitoring. The broad nature of the "No tripping hazard" claim makes it more difficult to identify a suitable monitoring mechanism, so we use a *claim refinement model* (CRM) to decompose the claim hierarchically into its underlying assumptions, until some more precise, and crucially monitorable, assumptions are identified. We consider a claim refinement model to be sufficiently complete when all leaf claims are either: monitorable, axiomatic or considered an unmitigatable risk. In the latter case, the claim marks the edge of the contextual envelope in which the system is capable of tailoring itself to suit.

In this example, our "No tripping hazard" claim can be decomposed into the CRM shown in Fig. 2. There are four sub-claims organized in two ANDed branches (claims may also be OR-ed). Together, the branches illustrate the rationale for why the root claim should hold. In this case, "No tripping hazard" holds because there is no-one in the room in which the vacuum cleaner is working AND no external impact has been detected by the vacuum cleaner. The leaf claims of the CRM, "Light level [remains] constant" and "No shock detected" are directly monitorable via events or statistical data collected by the system. We refer to claims that are possible to directly monitor and verify at runtime as *monitorables*. If a monitorable turns out to be false, for example, if the vacuum's inertial sensor detects an external shock, then claim falsification propagates upwards towards the root. Thus, in this case, the impact event would falsify the "No tripping hazard" claim by propagation. Similarly, a sudden increase in the light level would indicate that a light has been switched on by a woken occupant, and the "No Tripping Hazard" claim would again be falsified by propagation.

With a means of run time verification for the deadlock-breaking "No tripping hazard" assumption having being found, the robot vacuum cleaner can be

Fig. 2. Claim refinement model for robot vacuum cleaner

specified as using a clean at night strategy unless a shock is detected or a light is switched on, in which case the robot should self-adapt to use the "Clean when empty" strategy.

However, after it has been in operation for some time, the owners of the robot vacuum cleaner find that it is costing more to run than expected. A self-explanation capability would mean that the vacuum cleaner could explain that it is required to avoid causing a tripping hazard, and that it has been unable to clean at night because the occupants frequently wake up and turn the lights on. In this scenario, the explanation would help the users to understand the system's behaviour, and help to pinpoint the reason the system is not behaving in the manner they would have imagined. The customer understands the reasoning, but is still dissatisfied because the operating costs are unacceptable. They submit a change request to the developer for the vacuum cleaner to be modified so that it only adopts the clean when empty strategy if two consecutive nights' cleaning have been interrupted.

In isolation, this change request may seem unimportant to the developers, especially if the change request is scant on background information justifying it. To contextualize the request, they interrogate the vacuum cleaner to determine its history of operation, with special attention to its history of self-adaptation and the events sensed in its environment that triggered adaptations. They discover the light detection event is being triggered more frequently than expected, and understand by consultation of the requirements model that this is interpreted as invalidation of the assumption that underpins prioritization of energy cost minimization.

The developers realize that running costs are high but note also that the customer does move around the apartment at night. They modify the vacuum cleaner's software to adopt a new strategy; they relax [17] the clean apartment goal by accepting that the clean apartment goal may be satisfied at a later time. The user change request is accepted; when interrupted, the robot tries to clean the following night before resorting to the clean when empty strategy.

In this simple example, the information contained within the explanation offered by the system could be obtained by analysis of standard debugging output or logs, and by deduction. However, these sources of information are low-level

artefacts of particular code execution paths, and such analysis is performed by the system's developers, who will need time to perform the analysis. The potential for a self-adaptive system to adopt an unexpected configuration, or adopt an expected configuration in unexpected circumstances, means that there is a need for users to be able to understand what the system is doing, and why.

Our interest lies in reconciling a higher-level trace of the system's behaviour with its requirements, to establish whether the system's behaviour is appropriate, or better optimal, and whether the requirements themselves are correct. Although an explanation in terms of requirements may still prove too complex for some users to be able to understand a system's operation in some circumstances, the higher-level explanation may allow non-developer support personel to resolve queries without requiring developer input.

3 Self-Explanation Through Run-Time Requirements Models

Andersson *et al.* propose a means of characterising the change a self-adaptive system is designed to tolerate. Changes can be *foreseen*, *foreseeable* or *unforeseen*, as explained in [18]. We ignore here systems dealing with unforeseen change, which are more properly a topic for artificial intelligence research and pose a different order of challenge both for self-adaptation and self-explanation.

Much of our previous work has concerned requirements modeling for systems dealing with *foreseen* change [13,16,19]. Where change is foreseen, the set of contexts that the system may encounter are known at design time. Here, a self-adaptive system can be defined as a set of pre-determined system configurations that define the system's behaviour in response to changes of environmental context. Thus, there is little or no uncertainty about the nature of the system's environment and, if it is developed to high quality standards, satisfaction of the systems requirements should be deterministic.

More recently [12], we have started to address systems dealing with change that is, in [18]'s terms, merely *foreseeable*. Here, the key challenge is uncertainty, where at design time some features of the problem domain are unknown, perhaps even unknow*able*. Crucially, and in contrast to unforeseeable change, the fact of this uncertainty can be recognized, offering the possibility of mitigating it by resolving the uncertainty at runtime. The uncertainty associated with foreseeable change typically forces the developers to make assumptions in order to define the means to achieve the system's requirements. Thus, for example, a particular environmental context may be assumed to have particular characteristics and the system's behaviour defined accordingly. If the context turns out to have different characteristics, the system may behave in a way that is inappropriate. This has led us to exploit the concept of markers of uncertainty. Markers of uncertainty serve as an explicit marker of an unknown that forces the developer to make an assumption. We implement markers of uncertainty using claims as described in the previous section. A benefit of using claims to represent design-time assumptions is that the uncertainty is bounded and thus the risk of the

system behaving in an inappropriate may be mitigated by monitoring, claim and goal evaluation, and adaptation.

Our solution uses i* goal and claim refinement models, as depicted in Figs. 1 and 2. As described in the previous section, claim monitoring may permit assumptions to be verified during operation. Where a claim turns out to be false, the corresponding portion of the goal model can be re-evaluated at run-time with the claim removed, or its effect on the model weakened. If, as a consequence of this, the original goal operationalisation strategy no longer evaluates as the optimal solution, an alternative goal operationalisation strategy can be substituted dynamically, using the system's adaptation mechanism. We have applied our work to the domain of wireless sensor networks where our run-time models are supported by advanced adaptive middleware and domain-specific component models [6]. The overview of the approach is shown in Fig. 3. The overview is explained in terms of the development process (the box @design time) and the run-time components (the box @runtime). The module *Self-Explanation* is part of the module *Runtime Reasoner*, which is responsible for the transformation of the run-time goal models, as will be explained further in the next section.

In the context this paper, the key feature of foreseeable change is that it may result in behaviour that is emergent. Emergent behaviour may surprise stakeholders who may require the behaviour to be explained in order to build and maintain their confidence in the system. Our thesis is that the same run-time requirements models that we employ to handle unforeseen change can also be employed as the basis of a self-explanation capability. Partially based on [20], a useful self-explanation of an adaptation needs to include:

1. Details of any change in priority, or the proposed degree of satisfaction of a system (soft)goal.
2. Details of the adaptation performed by the system.
3. The history of the adaptation, and the related events that triggered it.

In the next section, we illustrate how a self-explanation of a system's behaviour that contains this information may be provided using our run-time requirements models solution for the GridStix wireless sensor network.

4 Case Study

To demonstrate self-explanation in the context of a system which adapts to contexts not fully foreseen, we present the GridStix flood prediction system [21]. We have previously discussed this system in the context of requirements modelling [13], and have recently been exploring run-time uses of these requirements models. In [12], we discuss systems using run-time goal-based models to guide adaptation to circumstances where assumptions on which the originally prescribed configuration(s) rely no longer hold. In this paper, we show how claims and run-time requirements models that have been implemented for GridStix support self-explanation.

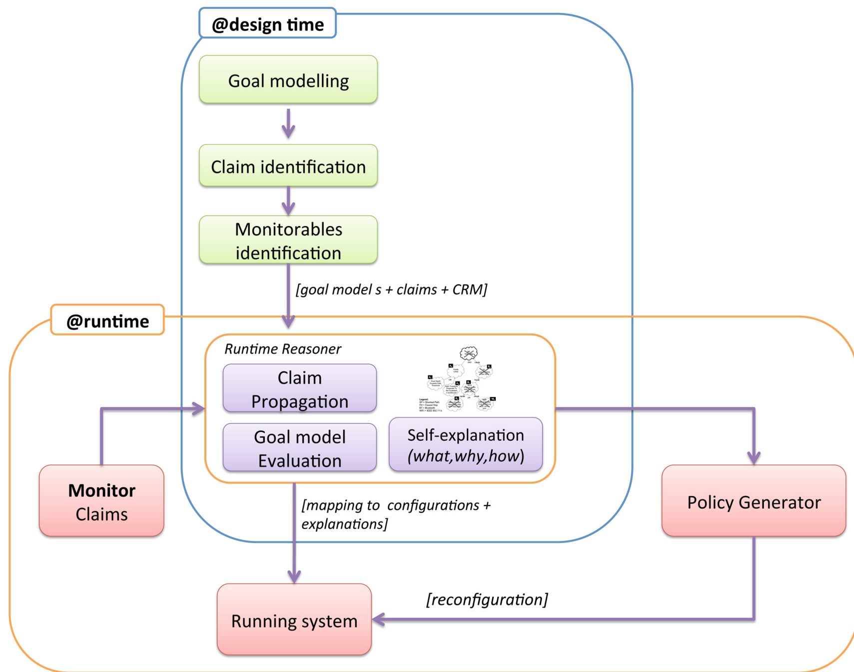

Fig. 3. Overview of the approach

The GridStix system is a wireless sensor network (WSN) for detecting and predicting flooding, versions of which were deployed on the river Ribble in North-West England and on the River Dee in North Wales. GridStix comprises a number of nodes (14 on the Ribble installation), each of which are equipped with sensors for detecting water depth and flow rate. The captured sensor data is processed by a stochastic model of the river to predict future river state. A feature of this algorithm is that it is distributed and lightweight enough to be executable by the GridStix nodes. Incremental results are cascaded from the most up-stream node down to the gateway node and from there via a GSM link to Lancaster University. Its accuracy is a function of the number of nodes contributing data.

GridStix is deployed in relatively remote, inaccessible locations with no mains power available, requiring that GridStix nodes rely on batteries and solar panels for power. As a result, energy conservation is a key non-functional requirement. GridStix uses an ad-hoc overlay network in which nodes can communicate using Bluetooth or WiFi, configured as either a shortest-path or fewest-hop spanning tree.

To help test feasibility and derive requirements for GridStix, empirical data was collected from experiments with a laboratory-based prototype. Data was collected to measure (among other metrics) resilience and power consumption [6],

as illustrated by the graphs in Fig. 4. Here, resilience is a measure of network fragmentation; the more nodes become isolated from the gateway (uplink) node, the less resilient is the network. If too many nodes become isolated from the gateway node, it becomes impossible for the system to offer an accurate flood prediction. Power consumption measures per-hop power consumed during the transmission of 1 KB of data from each node to the gateway. The graph *Physical Network Resilience* in Fig. 4 shows that the greater range of WiFi meant that data from each node could be routed to the gateway by a larger number of paths with WiFi than using Bluetooth, while the graph *Physical Network Power Consumption* in Fig. 4 shows that the additional resilience comes at the cost of higher power consumption.

Similarly, the graph *Spanning Tree Resilience* shows that, for a small number of nodes (nodes B, H and I), the number of routes to the gateway affected by node failure is much higher when using a shortest-path (SP) spanning tree algorithm than when using a fewest-hop (FH) spanning tree. This means that fewer nodes are likely to become isolated from the gateway node when GridStix is configured to use its FH spanning tree. The graph *Spanning Tree Power Consumption* shows that for the nodes furthest from the gateway node (nodes L, M, N and O) the

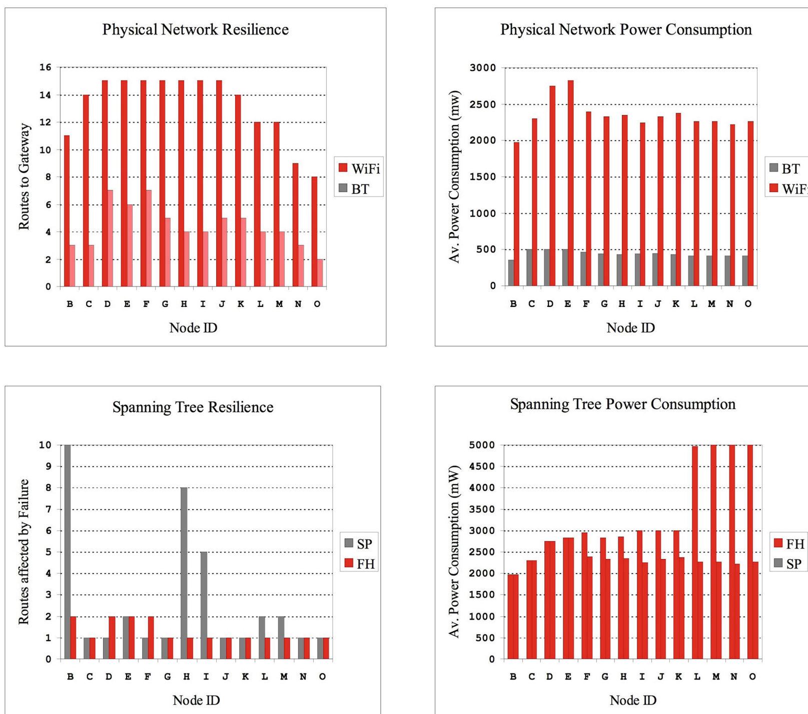

Fig. 4. Laboratory performance data (reproduced from [6])

power consumed in transmitting the data is significantly higher for a FH than SP spanning tree.

In other words, GridStix was predicted to be relatively resilient to node failure when configured to use WiFi and a fewest hop spanning tree, but at the cost of high power consumption.

Resilience and power consumption were two of GridStix's important non functional requirements. However, as shown in the experiments it is hard to optimize for both, meaning that one would have to be prioritized over the other. However, a feature of self-adaptive systems is that the extent to which any NFR must be satisficed (sufficiently satisfied) tends to be context-dependent, and this was the case with GridStix. Goal-based models, and specifically softgoals, support reasoning about tradeoff decisions that are aimed at achieving optimal goal satisfaction.

For the purposes of GridStix, expert environmental scientists had partitioned river behaviour into three distinct operating conditions (*domains*); *quiescent*, *high flow* and *flood*. Quiescence was predicted to be the most common domain over time and so, with the need for the nodes to retain enough power to react when the river state changed, energy efficiency was the priority. When in the flood and high flow domains, by contrast, resilience was prioritized to better tolerate any node loss that could impair the accuracy of GridStix's flood predictions. Thus, a particular GridStix configuration was specified for each domain, with (what was predicted to be) adaptation from one configuration to another specified to happen when the river was observed to change from one domain to another. These domain changes were based on sound knowledge and were therefore *foreseen*, meaning that we knew that the river's state would change and could specify the behaviour required of GridStix for each domain. Figure 5 shows the goal model for the flooding domain (which we call S3). The figure shows the claims "Bluetooth too risky for S3", "SP too risky for S3" and "Single node image processing not accurate enough for S3". Each claim records an assumption about a design-time choice of goal operationalization, made because of uncertainty about the relative performance of alternative operationalizations in the field. The tasks (goal operationalisation strategies) chosen are in white (i.e. WiFi, and FH). Note that for simplicity reasons the single-node and multi-node image processing shown in the figure is not part of the explanation. However, similar conclusions can be made if we take into account these operationalizations and their effect on the NFRs, therefore the *calculate flow rate* goal should be ignored in the figure.

The configurations that were specified at design-time for each domain were based on the performance of the alternative communication technologies and spanning tree configurations observed in the laboratory experiments described above. However, we were aware that the lab results might prove imperfect predictors of how GridStix performed in the field. The initial River Ribble deployment confirmed that the effects of radio signal absorption by the river banks, rain, trees, etc., had a significant affect on performance [21]. To make GridStix more tolerant of these effects, it was augmented with claims to monitor

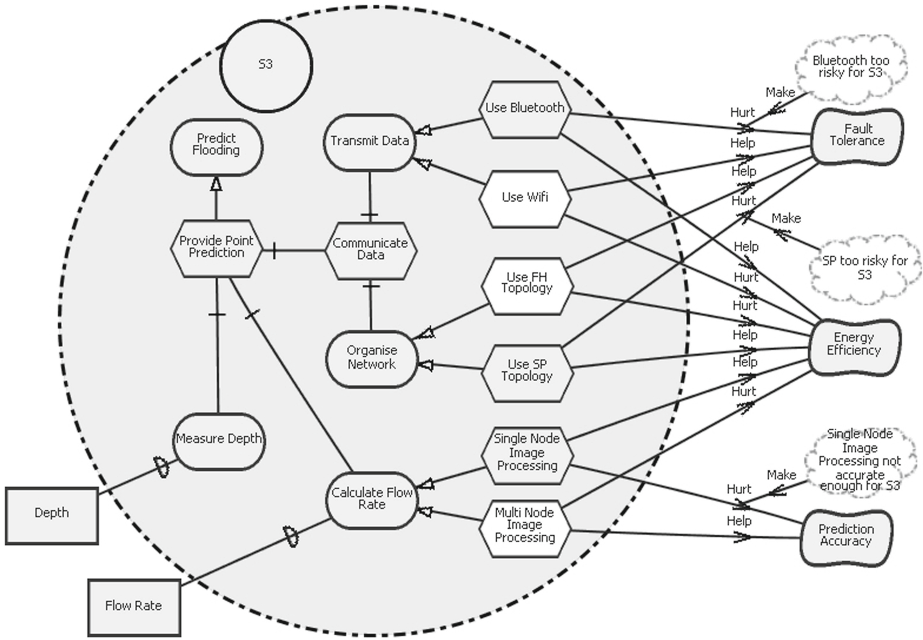

Fig. 5. Gridstix goal models for the flooding state of the river

the design-time assumptions, and to adapt to an alternative configuration if monitoring suggested that the alternative configuration could perform better. This was an important change because it meant that, in addition to the changes foreseen by knowledge of the different river domains, change as a consequence of operational experience was also *foreseeable*. When using claim monitoring, GridStix can decide by itself to adapt to a new configuration under some circumstances that were not predefined at design time. Thus, whereas GridStix's *adaptive* behaviour had been deterministic (even if its adequacy as a WSN had not been), its adaptive behaviour was now non-deterministic. Such non-deterministic behavior could cause "surprise" to an operator of the system, and therefore a self-explanation capability is appropriate.

A portion of the claim refinement models used by the GridStix *flood* and *high flow* domains is presented in Fig. 5. There is one top-level claim shown (in bold). This represents assumptions derived from the laboratory experiments that Bluetooth communication technology is too risky. In other words, the assumption is that if GridStix was configured to use Bluetooth, network resilience would likely be poor; implicitly poor*er* than if WiFi was used instead. The associated claim refinement model represents derivation of the means to sustain the claim and results in (using the labels in Fig. 6 as shorthand for the subclaims):

$$BT_Too_Risky \Leftrightarrow (A_0 \Leftrightarrow (A_1 \vee A_2)) \wedge (B_0 \Leftrightarrow (B_1 \Leftrightarrow (B_2 \vee \neg B_3)))$$

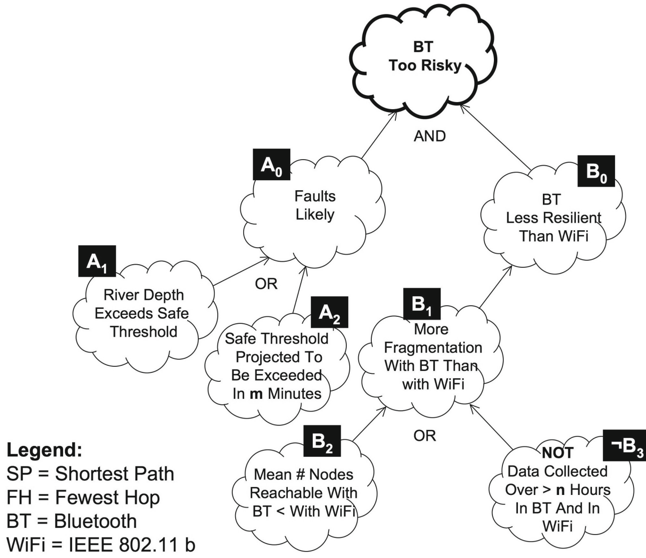

Fig. 6. GridStix claim refinement model justifying choice of WiFi for Inter-Node Communication

Thus, our root assumption, that using Bluetooth will lead to greater fragmentation than using WiFi (the *BT Too Risky* claim), will be disproved if any of the leaf (*monitorable*) subclaims is negated. In other words, Bluetooth is not likely to fragment the network if the river depth is below the safe threshold level or, at the current rate of change it will not exceed the safe level anytime soon. Similarly, Bluetooth is unlikely to lead to excessive fragmentation of the network if the rate of fragmentation when using Bluetooth is no higher than when using WiFi or, if there is data that contradicts this, there is too little data to make the contradiction statistically sound.

Because the River Ribble deployment of GridStix has been decommissioned, we used a simulator to observe the system's behaviour when experimenting with claim monitoring. The simulator has been developed using the collected data of the several months GridStix was deployed with the advantage that we can run experiments when needed. The simulator handles factors such as: power usage by batteries of nodes and according to whether the nodes were configured to use WiFi or Bluetooth, fewest hops or shortest path; whether the nodes were idling or performing computationally intensive tasks; and power replenishment from solar panels depending on time of day, amount of sunlight received or how cloudy the weather is, among others. Using a simulator constructed for GridStix, we ran an experiment to compare the longevity of the claim-augmented version of GridStix with the original. Longevity in this context means the length of time during which a sufficient number of nodes were connected to allow a meaningful result to be returned by the gateway. The simulator includes randomization

to simulate jitter and packet loss. We complemented this with random node failures to simulate those actually observed. We ran the simulator with a profile of river behaviour over a fixed period comprising a sequence of flow rate and depth values that simulated the river in every mood from quiescent to flood. We varied a single variable; the amount of sunlight received by the nodes' solar panels, using percentage of cloud cover during daylight hours as a proxy. The experiment was run three times and the results averaged to account for the randomization elements.

The experiments suggest no significant benefit from claim augmentation when cloud cover is above approximately 40 %. Once cloud cover drops below 40 %, however, the augmented version has significantly greater longevity. For example, at 30 % cloud cover, instead of failing after approximately 180 h of operation, GridStix survives for approximately 250 h.

The increase in GridStix's longevity under some conditions appears to correlate with a particular self-adaptation being performed in these simulations. In those simulations where the claim-augmented version of the system outperformed the original version, GridStix substituted the use of WiFi for communication whilst the river was flooding, as originally specified, for the use of Bluetooth. The history of the monitoring data shows that over the defined minimum period of accumulating data, network fragmentation was no less during that period when using Bluetooth than when using WiFi. The effect of this on the claim refinement model (Fig. 6) in which the falsified monitorable claims $B_2 \vee \neg B_3$ ... became ... $\neg B_2 \wedge B_3$ and propagated up the hierarchy to falsify the top-level *BT Too Risky* claim that justified the original (design-time) choice of WiFi over Bluetooth. This in turn triggered the run-time re-evaluation of the goal model, revealing that the operationalization of the *Transmit Data* goal now favoured the use of Bluetooth rather than WIFi because Bluetooth's net impact on power consumption and resilience had become more +ve (positive) than that of WiFi. The goal model was thus changed to select Bluetooth as *Transmit Data*'s operationalization which in turn triggered the GridStix middleware to adopt a new component configuration, dynamically binding the Bluetooth component in place of the WiFi component (Fig. 7).

Discussion

Revisiting the challenges presented in the introduction of this paper we conclude that our approach:

1. Offers suitable monitoring capabilities for self-explanation through the use of claim monitoring. As for reasoning capability, our claim refinement models allow a change of configuration to be traced back to the monitored, and falsified, assumption that caused it.
2. Can offer self-explanation and discourse at the level of requirements. Self-adaptation (that maybe misunderstood by operators) can be explained in terms of goals, operationalisations and assumptions. This level of abstraction

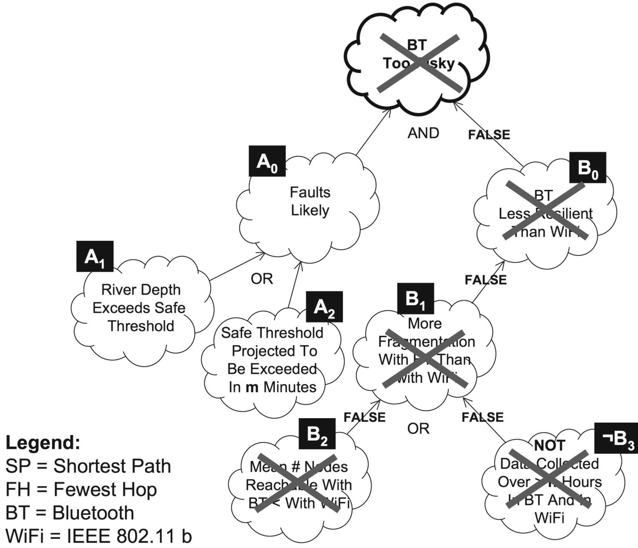

Fig. 7. Falsified claim propagation

is closer to natural language than architectural or code-level descriptions, offering more understandable explanations.

3. Provides the required link between the requirements and the architecture. The currently active configuration, at an architecture level, is linked to goal operationalisations, to the goals they achieve and their expected impact on system NFRs. (e.g. a proper link between the architectural use of BlueTooth and the "Communicate Data" goal it achieves, and the effect on system NFRs such as power consumption).
4. Finally, allows the system to be able to reproduce a trace history (i.e. a sequence of steps) of monitored events, assumption falsifications and resultant reconfigurations to explain the reasons behind a self-adaptation being carried out.

5 A Domain-Specific Language for Self-Explanation

An obvious goal for self-explanation research would be to offer natural language explanations, understandable by a system's users and other stakeholders. However, at present we feel that offering explanations in a simply-structured domain-specific language (DSL) is a more feasible goal. We propose a self-explanation DSL designed to convey explanations of self-adaptations in terms of requirements met, softgoals satisficed and claims invalidated, and supporting the identification of the run-time monitoring data that triggered a self-adaptation. Although explanations expressed in our DSL can be trivially transformed into

text approaching natural language, knowledge and understanding of the system's requirements models is required to fully understand the explanation. Thus, our self-explanation DSL is targetted at support personel rather than end users. To extend our approach to come closer to the goal of offering user-understandable natural language self-explanations, some means of explaining the requirements models themselves would have to be devised.

As discussed in the previous section, self-adaptation in systems using our run-time requirements models can be either planned or unplanned. Planned self-adaptation fits the definition of foreseen adaptive behaviour in [18]. For unplanned self-adaptation, the acts of identifying an assumption, codifying it in a claim and devising a monitor for it imply a degree of foresight on the part of the developers that the system may need to tolerate conditions under which the assumption doesn't hold. However, the exact nature of the circumstances under which a claim (or some combination of claims) would be invalidated remain ambiguous, meaning that unplanned adaptation is most analagous to [18]'s foreseeable adaptive behaviour.

Planned self-adaptation takes place between a pair of pre-specified configurations, in response to a trigger conditions identified during the RE process. Planned self-adaptation requires the least detailed self-explanation, with the system needing to explain that the trigger conditions for a specific planned adaptation were met, and that the system has reconfigured itself to a specific, pre-specified configuration. This simpler self-explanation offers little more than standard logging output, and as such our DSL is designed to convey explanations of unplanned self-adaptation.

Unplanned self-adaptation takes place when a claim is invalidated, and the system's run-time requirements models are re-evaluated. Prior to an unplanned self-adaptation, the system may have adopted one of its pre-specified configurations, or another unplanned adaptation may have occurred, leaving the system in a configuration not foreseen at design time. A suitable self-explanation of an unplanned self adaptation needs to identify:

1. The configuration change that has been made.
2. The goal model change that resulted in the configuration change.
3. The softgoal(s) expected to be better satisfied in the new configuration.
4. The monitoring data that triggered the goal model change.

Explaining the details of the configuration change is a relatively simple matter, that could be achieved with more traditional logging or monitoring techniques. In terms of our run-time requirements models, the information required for the explanation is the task in the goal model whose selection to satisfy a parent goal is no longer supported by analysis of its contribution links and claims influencing them, and the task selected to replace it. Likewise, the monitoring data that triggered the goal model change is captured during the model transformation process, and could equally be captured through logging or monitoring.

The source of variability in our run-time requirements goal models is the invalidation of claims. Thus, the goal model change that resulted in an unplanned

self-adaptation will be the invalidation of a specific claim, or combination of claims. In cases where claims are invalidated by propagation, the explanation should allow an indirectly invalidated claim's invalidation to be traced back to the original claim, whose invalidation triggered the propagation.

By default, the transformation to the run-time requirements models applied as a result of claim invalidation is an inversion of the claim's contribution link, by which it is connected to the model. Other transformations, such as the removal of claims or removing a malfunctioning goal operationalisation strategy, are supported. However, these other transformations either require further claims to be added to the model, or risk deadlock in the requirements model after claim invalidation, and we also prefer to invert claim contribution links where possible. For simplicity, we prefer to use a single root claim in the goal model to represent a specific area of uncertainty surrounding a variation point.

The softgoal(s) that are expected to be better satisfied as a result of the configuration change can be identified through model analysis. If the default (contribution inversion) transformation is used, and our model construction guideline of using a single claim per variation point has been followed, the promoted softgoal or softgoals are readily identifiable. Table 1 shows details, for each combination of the value of the contribution link connected to a claim (the "claim contribution") and the polarity of the contribution link to which the claim's contribution is attached (the "attached contribution"); which softgoal(s) can be expected to be better satisfied as a result of the self-adaptation.

In Table 1, the "Complementary" softgoals are defined as the softgoals connected to the task the attached contribution link belongs (and thus connects) to *with the same polarity*. Conversely, "Competing" softgoals are defined as those connected to the task the attached contribution link belongs with the *opposing polarity*.

An unplanned self-adaptation may also be triggered in response indicating that a claim that was previously invalidated does, in fact, appear to be valid once more. This would typically occur as a result of developers having missed some operating context encountered by the system only temporarily. We refer to this as claim revalidation. For this class of unplanned self-adaptation, the softgoal(s) expected to be better satisfied as a result of the configuration change are different to those in Table 1. Table 2 details the softgoals affected by claim revalidation, using the same columns and terminology as Table 1.

Table 1. Softgoals promoted by claim invalidation, for different model structures

Claim contribution	Attached contribution polarity	Softgoals affected
Make	Positive	Competing
Break	Positive	Complementary
Make	Negative	Complementary
Break	Negative	Competing

Table 2. Softgoals promoted by claim revalidation, for different model Structures

Claim contribution	Attached contribution polarity	Softgoals affected
Make	Positive	Complementary
Break	Positive	Competing
Make	Negative	Competing
Break	Negative	Complementary

Even if the explanation of an unplanned adaptation contains all of the information discussed so far, it may still prove difficult to understand the circumstances surrounding, and the reason for, an unplanned adaptation in cases where the system has previously performed another unplanned adaptation. In such circumstances, the configuration of the system prior to the adaptation under query may not be one of the configurations specified for the system at design time. Therefore, explanations of an unplanned self-adaptation must include the information of any previous unplanned self-adaptations performed by the system.

To summarise, our DSL is designed to convey self-explanations in terms of the satisficement of softgoals, reconfigurations, model transformations, claims invalidated, and monitoring events. Monitoring events are fired by monitors upon the collection and analysis of data indicating that a claim does not hold. Claims are invalidated as a result of monitoring data, and should the root claim in a claim refinement model (i.e. a claim upon which a decision rests) become invalidated then a model transformation is performed. A model transformation is a change to the run-time requirements model which may, if analysis of the modified run-time model indicates it is necessary, lead to a reconfiguration being performed by the system to better satisfice a softgoal. These concepts are all at the level of abstraction used in our requirements models, and thus this is the level of abstraction used by our explanations.

In Sect. 3, we proposed that a meaningful self explanation of a self-adaptive system's current behaviour needs to include:

1. Details of any change in priority, or the proposed degree of satisfaction of a system (soft)goal.
2. Details of the adaptation performed by the system.
3. The history of the adaptation, and the related events that triggered it.

In these terms, an explanation of a planned self-adaptation consists of details of the change in context identified by the system, the self-adaptation performed, and details of previous self-adaptations that have been performed. For unplanned self-adaptations, an explanation consists of details of the softgoals that are to be better satisfied by the adaptation, details of the configuration change performed, and details of the goal model changes that were made (claims invalidated) triggering the self-adaptation, along with the monitoring data that caused the goal model changes. Further history is also provided by including details of previous unplanned self-adaptations.

Our self-explanation DSL targets the more complex unplanned self-adaptations performed by a system, and presents the information discussed in tuples of:

{*Softgoals Satisficed, Change Performed, Claim Invalidated, Cause*}

A tuple is required for every change to the run-time requirements models performed by the running system. As a result, not all values in the tuple may be populated for every model transformation (e.g. the invalidation of an intermediate claim in a run-time Claim Refinement Model) as some information (e.g. the softgoals better satisficed by a self-adaptation) will not yet be available. The contents of the first three values in each tuple are as discussed so far in this section. Acceptable values for the "Cause" value in the tuple, however, vary depending on the type of change being made to the run-time requirements model. For a claim being invalidated directly by its own monitoring data, the cause is the monitoring data indicating the claim's invalidity. For claims invalidated by propagation, the cause identifies the claim invalidation propagated from, along with an indication of the Claim Refinement Model semantics dictating the claim be invalidated by propagation.

Analysis performed on a collection of tuples for an unplanned adaptation can yield a textual explanation, in terms of requirements, softgoals and claims, that approaches natural language. For the robot vacumm cleaner example discussed in Sect. 2, an explanation of the unplanned adaptation to clean the apartment when empty as opposed to cleaning at night would be explained as:

Adapted to "Clean When Empty" instead of "Clean at Night" to better satisfice "Avoid Tripping Hazard", due to the (invalidated) "No Tripping Hazard" claim. The "No Tripping Hazard" claim was invalidated because its supporting "No Foot Impact" claim was invalidated. The "No Foot Impact" claim was invalidated because monitoring data indicating its invalidity (FootShockEvent) was received.

Of course, our self-explanation DSL is tightly-coupled to our run-time requirements modelling approach, and to the use of claim invalidation as a source of run-time requirements model variability. However, our ability to generate self-explanations using our DSL from a running system, and to interpret them into a usable textual explanation serves to demonstrate the approach's feasibility. We also note that, at present, our method for generating output in the self-explanation DSL from a running system equipped with run-time requirements models using our run-time requirements models and reasoning tools depends on the (previously considered optional) one-claim-per-variation point modelling guideline being followed, and supports only one of three model transformations supported by the our run-time reasoner. We consider the use of more complex claim hierarchies, and particularly the "remove claim on invalidation" run-time model transformation with our self-explanation generator and DSL as interesting areas for future exploration.

6 Related Work

Although, to our knowledge, we are the first to discuss self-explanation in the context of self-adaptive systems; the desire for systems to produce output at a higher level to improve understanding is long-standing. For example, there has been significant research into Natural Language Generation by the Artificial Intelligence and Computational Linguistics communities.

In [22], Duggan and Bent present an algorithm, designed to infer the type of variables during compilation of programs written in implicitly typed languages such as ML or Haskell, where explicit variable type declarations are not used. The algorithm infers variable type by analysis of variable usage, annotating the program's syntax tree as it progresses. Inference is performed using a set of rules; for example a variable to which the addition operator is applied, with a right hand operand of 1, is an integer. A variable whose type is determined to be integer through this example rule would have the following explanation annotated to the program's syntax tree: +(x,1) gives x: int. Explanations can become considerably more complex when a variable's type is dependent on that of one or more other variables, however the base format remains the same. In this work, the explanation is used by the algorithm itself to allow explanation fragments previously generated to guide later type inferences, but the authors consider the approach potentially useful in providing debugging support.

Similarly, in [23], Van Baalen *et. al.* retrofit a domain specific code generator with explanatory capability for use at NASA. In this work, the explanation covers the relationship between a specification, domain theory and synthesised code. The explanation is relatively low-level, designed to allow developers to prove correctness, given NASA's obvious need for high-assurance software.

In [24], Huang and Fiedler discuss the PROVERB text planner, which verbalises mathematical (natural deduction) proofs. The planner uses a three-stage approach, with the first stage responsible for hierarchically decomposing complex proofs into a series of subproofs, the second stage identifies possible opportunities to "paraphrase" (or rather combine proof elements into larger, useful sentences) with the third stage actually generating the textual output. A more general overview of the state of the art in automated theorem provers, including discussion of the usefulness of their output, is offered in [25].

Although this work shows a research interest in providing high-level output to ease human understanding, our focus is not on programs providing natural language output, but in providing explanations of observed behaviour. Furthermore, the explanations offered by [22,23] are aimed at developers and mathematicians, respectively. The self-explanations we advocate are at a higher-level of abstraction, aimed at users and support personnel.

Debugging mechanisms, even those considered high-level [26,27], are focussed on data structures and code rather than on requirements, goals and operationalisation strategies. More closely-related work can be found in the field of requirements monitoring [4,5], from which we derive our claim monitoring. [28] proposes "awareness requirements", which are requirements that refer to the success or failure of other requirements. The authors state that awareness requirements

may refer to goals, tasks, quality constraints and domain assumptions. Claim monitoring in our work is similar to domain assumption awareness requirements in [28], but their focus is on the mapping from requirements models to feedback loops, with no run-time representation of the awareness requirements.

The claim reasoning we use to demonstrate the utility of run-time requirements models in offering self-explanation is based on a combination of two previous streams of our work. In [13], we discuss the use of claims to highlight assumptions made during self-adaptive system specification, with a view to them being revisited in light of later requirement changes. In [10], we make the case for the run-time use of requirements models, with the ability to rectify deficiencies in requirements satisfaction using self-adaptation being a key motivator. Although we use reasoning of run-time claim refinement models to offer a limited form of self-explanation, the type of self-adaptive system claim reasoning proves most useful for are those with a limited number of potential goal operisational strategies, or where self-adaptation is being used to balance a set of conflicting non-functional requirements.

Approaches such as RELAX [17] and FLAGS [29] adopt fuzziness in requirements to allow self-adaptation to prioritise and optimise their satisfaction. In these approaches, a run-time requirements model could be used to record the (re)prioritisations that take place, and to allow explanations of adaptations in the context of which requirements were compromised and which favoured. Approaches such as [30], which use KAOS [31] goal models, could benefit from run-time analysis of obstacle models to offer self-explanation in terms of which obstacles have been detected in the operating environment, and which goal operationalisation strategies have been adopted to overcome them.

When tackling uncertainty, the ideas discussed in [32] are also related. As we do, the authors of [32] argue that uncertainty plays an important role in any software based system that needs adapt continuously to meet the goals. They argue that the focus of managing uncertain information should be on the rationale used to come to a decision. We emphasize the importance of being able to explain this rationale. In their case, the decision may be taken either during design or requirements (i.e. before execution). In our case, we go further because the self-adaptive system is able to make decisions at runtime as well. Finally, we believe our work is relevant to the implementation of dynamic traceability needed when dealing with self-adaptive systems where little work has yet been done. The authors of [33] discuss traceability in the presence of uncertainty. Similar to our work on claims, the authors of [33] propose to attach supplementary information to traceability links. This additional information describes the confidence and the rationale for its creation. The authors take into account the fact that the rationale that supports design decisions is often based on assumptions and beliefs. However, in contrast to our work, their work focuses on the case of software product lines and their evolution during software life cycle rather than on runtime adaptation

7 Conclusions and Future Work

This paper has argued that self-adaptive systems with the potential to behave in a manner not prescribed at design-time require self-explanation to allow emergent behaviour to be diagnosed, understood and explained. Self-explanation is important because it provides a means to increase confidence in, and resolve queries about, the behaviour of a self-adaptive system by its users. Self-explanation can also aid developers in understanding the behaviour of a self-adaptive system by tracing observed run-time behaviour (the *what*) to design-time assumptions, instrospect the strategy chosen (the *how*) and the extent to which they proved to be valid in operation (the *why*).

As already described in [12,16] we have developed an approach to creating self-adaptive systems capable of tailoring their behaviour to an operating environment not fully foreseen at design-time, using run-time requirements models. These systems are capable, indeed likely, to exhibit emergent behaviour. In this paper we show how self-explanation of such behaviour might be generated from the systems' adaptive reasoning machinery. The particular run-time requirements models used by our approach are in-memory representations of i* Strategic Rationale and NFR framework Claim Refinement Models, which are notably high-level in their nature. Our hypothesis is that these dynamic models, interpreted through the history of observed behaviour and adaptation events can provide a plausible means of explaining why the observed behaviour came about. This contrasts with the use of low-level reconfigurations and executed code paths used in standard debugging tools which are difficult to interpret in terms of systems' requirements, even for expert developers working on systems that don't have the added complexity of a self-adaptive capability.

We have demonstrated how a system equipped with our run-time requirements models for self-adaptation may also analyse these models to support self-explanation. We have introduced a simple self-explanation domain-specific language, and have shown that analysis by a running system of an explanation conveyed in our self-explanation DSL can yield a near natural language explanation in terms of requirements, claims and monitors.

There are several ways in which our approach can be improved. Currently, the claim reasoning and model transformation based adaptation mechanism discussed in this paper applies where goals are achieved by selecting from a finite number of goal operationalisation strategies defined a-priori but selected dynamically. Our approach is able to improve the flexibility of an executing system facing unforeseen situations, but the potential operationlization strategies, and the goals they achieve are defined and analysed at design time. Where new goal operationalisation strategies may themselves be emergent (e.g. through dynamic service discovery), further research is needed. This is one of the topics we are investigating in the FP7 CONNECT project[1]. Specifically, we are studying ways in which a new goal operationalisation strategy can be conceived at runtime.

[1] http://connect-forever.eu//

One of the challenges explored is that of updating goal models during execution to keep the required causal link between architecture and requirements.

Looking to more distant research prospects, two clear goals of self-explanation research are to provide natural language explanations that are within the capabilities of users to understand, and to allow users themselves to instruct the system to perform a self-adaptation by interacting with the system at this same level. The ability for a running system to use run-time requirements models to trace between high-level concepts such as goals and claims and lower-level system configurations, as demonstrated in this paper, indicates that this may be possible with future work.

References

1. Muir, B.M.: Trust in automation: Part I. theoretical issues in the study of trust and human intervention in automated systems. Ergonomics **37**(11), 1905–1922 (1994)
2. Gillieis, A.C., Hart, A.: Using kbs ideas in image processing - a case study in human computer interaction. In: Research and Development in Expert Systems V: Proceedings of Expert Systems '88, pp. 258–268 (1988)
3. Sawyer, P., Bencomo, N., Whittle, J., Letier, E., Finkelstein, A.: Requirements-aware systems: a research agenda for re for self-adaptive systems. In: IEEE International Conference on Requirements Engineering, pp. 95–103 (2010)
4. Robinson, W.: A requirements monitoring framework for enterprise systems. Requir. Eng. **11**(1), 17–41 (2005)
5. Fickas, S., Feather, M.: Requirements monitoring in dynamic environments. In: Second IEEE International Symposium on Requirements Engineering (RE'95) (1995)
6. Grace, P., Hughes, D., Porter, B., Blair, G., Coulson, G., Taiani, F.: Experiences with open overlays: a middleware approach to network heterogeneity. In: Submitted to Eurosys 2008, Glasgow, UK (2008)
7. Coulson, G., Blair, G., Grace, P., Joolia, A., Lee, K., Ueyama, J., Sivaharan, T.: A generic component model for building systems software. ACM Trans. Comput. Syst. **26**(1), 1–42 (2008)
8. Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B., Steenkiste, P.: Rainbow: Architecture-based self-adaptation with reusable infrastructure. IEEE Comput. **37**(10), 46–54 (2004)
9. Georgas, J.C., van der Hoek, A., Taylor, R.N.: Using architectural models at runtime to manage and visualize the adaptation process. In: Bencomo, N., Blair, G.S., France, R. (eds.) Models@run.time, Special Issue. IEEE Computer (2009)
10. Bencomo, N., Whittle, J., Sawyer, P., Finkelstein, A., Letier, E.: Requirements reflection: requirements as runtime entities. In: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering, ICSE '10, vol. 2, pp. 199–202. ACM, New York (2010)
11. Souza, V.E.S., Lapouchnian, A., Robinson, W.N., Mylopoulos, J.: Awareness requirements for adaptive systems. In: ICSE Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2011, Waikiki, Honolulu, HI, USA, 23–24 May 2011, pp. 60–69 (2011)
12. Welsh, K., Sawyer, P., Bencomo, N.: Towards requirements aware systems: Runtime resolution of design-time assumptions. In: Proceedings of the 26th IEEE/ACM International Conference on Automated Software Engineering (2011)

13. Welsh, K., Sawyer, P.: Requirements tracing to support change in dynamically adaptive systems. In: Glinz, M., Heymans, P. (eds.) REFSQ 2009. LNCS, vol. 5512, pp. 59–73. Springer, Heidelberg (2009)
14. Chung, L., Nixon, B.A., Yu, E., Mylopoulos, J.: Non-Functional Requirements in Software Engineering, vol. 5. Springer, Heidelberg (1999)
15. Yu, E.S.K.: Towards modeling and reasoning support for early-phase requirements engineering. In: RE '97: Proceedings of the 3rd IEEE International Symposium on Requirements Engineering (RE'97), Washington, DC, USA (1997)
16. Welsh, K., Sawyer, P.: Understanding the scope of uncertainty in dynamically adaptive systems. In: Wieringa, R., Persson, A. (eds.) REFSQ 2010. LNCS, vol. 6182, pp. 2–16. Springer, Heidelberg (2010)
17. Whittle, J., Sawyer, P., Bencomo, N., Cheng, B.H.C., Bruel, J.M.: Relax: a language to address uncertainty in self-adaptive systems requirement. Requir. Eng. **15**(2), 177–196 (2010)
18. Andersson, J., de Lemos, R., Malek, S., Weyns, D.: Modeling dimensions of self-adaptive software systems. In: Cheng, B.H.C., de Lemos, R., Giese, H., Inverardi, P., Magee, J. (eds.) Self-Adaptive Systems. LNCS, vol. 5525, pp. 27–47. Springer, Heidelberg (2009)
19. Goldsby, H.J., Sawyer, P., Bencomo, N., Hughes, D., Cheng, B.H.: Goal-based modeling of dynamically adaptive system requirements. In: 15th Annual IEEE International Conference on the Engineering of Computer Based Systems (ECBS) (2008)
20. Lim, B.Y., Dey, A.K., Avrahami, D.: Why and why not explanations improve the intelligibility of context-aware intelligent systems. In: Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI '09, pp. 2119–2128. ACM, New York (2009)
21. Hughes, D., Greenwood, P., Coulson, G., Blair, G., Pappenberger, F., Smith, P., Beven, K.: Gridstix: Supporting flood prediction using embedded hardware and next generation grid middleware. In: 4th International Workshop on Mobile Distributed Computing (MDC'06), Niagara Falls, USA (2006)
22. Duggan, D., Bent, F.: Explaining type inference. Sci. Comput. Program. **27**, 37–83 (1995)
23. Van Baalen, J., Robinson, P., Lowry, M., Pressburger, T.: Explaining synthesized software. In: Proceedings of the 13th IEEE International Conference on Automated Software Engineering, ASE '98, pp. 240-249. IEEE Computer Society, Washington DC, USA (1998)
24. Huang, X., Huang, X., Fiedler, A., Fiedler, A.: Proof verbalization as an application of NLG. In: Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI), pp. 965–970. Morgan Kaufmann (1997)
25. Bundy, A.: Automated theorem provers: a practical tool for the working mathematician? Ann. Math. Artif. Intell. **61**, 3–14 (2011). doi:10.1007/s10472-011-9248-8
26. Golan, M., Hanson, D.R.: Duel - a very high-level debugging language. In: USENIX Winter, pp. 107–118 (1993)
27. Yang, J., Soffa, M.L., Selavo, L., Whitehouse, K.: Clairvoyant: a comprehensive source-level debugger for wireless sensor networks. In: Proceedings of the 5th International Conference on Embedded Networked Sensor Systems, SenSys '07, pp. 189–203. ACM, New York (2007)
28. Silva Souza, V.E., Lapouchnian, A., Robinson, W.N., Mylopoulos, J.: Awareness requirements for adaptive systems. Technical report, University of Trento (2010)

29. Baresi, L., Pasquale, L., Spoletini, P.: Fuzzy goals for requirements-driven adaptation. In: 2010 18th IEEE International Requirements Engineering Conference (RE), pp. 125–134 (2010)
30. Nakagawa, H., Ohsuga, A., Honiden, S.: Constructing self-adaptive systems using a kaos model. In: Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, 2008, SASOW 2008, pp. 132–137 (2008)
31. Dardenne, A., van Lamsweerde, A., Fickas, S.: Goal-directed requirements acquisition. Sci. Comput. Program. **20**, 3–50 (1993)
32. Lehman, M.M., Ramil, J.F.: Software evolution: background, theory, practice. Inf. Process. Lett. **88**, 33–44 (2003)
33. Anquetil, N., Grammel, B., Galvao, I., Noppen, J., Shakil, S., Arboleda, H., Rashid, A., Garcia, A.: Traceability for model driven, software product line engineering. In: ECMDA (2008)