

An Architecture for Requirements-Driven Self-reconfiguration

Fabiano Dalpiaz, Paolo Giorgini, and John Mylopoulos

University of Trento - DISI, 38100, Povo, Trento, Italy
{fabiano.dalpiaz, paolo.giorgini, jm}@disi.unitn.it

Abstract. Self-reconfiguration is the capability of a system to autonomously switch from one configuration to a better one in response to failure or context change. There is growing demand for software systems able to self-reconfigure, and specifically systems that can fulfill their requirements in dynamic environments. We propose a conceptual architecture that provides systems with self-reconfiguration capabilities, enacting a model-based adaptation process based on requirements models. We describe the logical view on our architecture for self-reconfiguration, then we detail the main mechanisms to monitor for and diagnose failures. We present a case study where a self-reconfiguring system assists a patient perform daily tasks, such as getting breakfast, within her home. The challenge for the system is to fulfill its mission regardless of the context, also to compensate for failures caused by patient inaction or other omissions in the environment of the system.

1 Introduction

There is growing demand for software systems that can fulfill their requirements in very different operational environments and are able to cope with change and evolution. This calls for a novel paradigm for software design where monitoring, diagnosis and compensation functions are integral components of system architecture. These functions can be exploited at runtime to monitor for and diagnose failure or under-performance, also to compensate through re-configuration to an alternative behavior that can better cope with the situation on-hand. Self-reconfiguration then is an essential functionality for software systems of the future in that it enables them to evolve and adapt to open, dynamic environments so that they can continue to fulfill their intended purpose.

Traditionally, self-reconfiguration mechanisms are embedded in applications and their analysis and reuse are hard. An alternative approach is externalized adaptation [1], where system models are used at runtime by an external component to detect and address problems in the system. This approach – also known as model-based adaptation – consists of monitoring the running software, analyzing the gathered data against the system models, selecting and applying repair strategies in response to violations.

We analyze here the usage of a special type of system models: requirements models. Deviations of system behavior from requirements specifications have been discussed in [2], where the authors suggest an architecture (and a development process) to reconcile requirements with system behavior. Reconciliation is enacted by anticipating deviations

at specification time and solving unpredicted circumstances at runtime. The underlying model is based on the goal-driven requirements engineering approach KAOS [3].

In this paper, we propose a conceptual architecture that, on the basis of requirements models, adds self-reconfiguration capabilities to a system. The architecture is structured as a set of interacting components connected through a Monitor-Diagnose-Compensate (MDC) cycle. Its main application area is systems composed of several interacting systems, such as Socio-Technical Systems [4] (STSs) and Ambient Intelligence (AmI) scenarios. We have chosen to use Tropos [5] goal models as a basis for expressing requirements, for they suit well for modeling social dependencies between stakeholders. We enrich Tropos models adding activation events to trigger goals, context-dependent goal decompositions, fine-grained modeling of tasks by means of timed activity diagrams, time limits within which the system should commit to carry out goals, and domain assumptions that need to be monitored regardless of current goals.

We adopt the BDI paradigm [6] to define how the system is expected to reason and act. The system is running correctly if its behavior is compliant with the BDI model: when a goal is activated, the system commits to it by selecting a plan to achieve it. The architecture we propose monitors system execution and looks for alternatives when detects no progress or inconsistent behaviour.

The closest approach to our work is Wang et al. [7], which proposes a goal-oriented approach for self-reconfiguration. Our architecture differs from hers in the details of the model we use to monitor for failures and violations. These details allow us to support a wider class of failures and changes, also to compensate for them.

This paper is structured as follows: Section 2 presents the baseline of our approach, Section 3 describes our proposed architecture for self-reconfiguration, whereas Section 4 explains how to use it. Section 5 details the main monitoring and diagnosis mechanisms the architecture components use. Section 6 shows how the architecture can be applied to a case study concerning smart-homes. Section 7 presents related work and compares our approach to it. Finally, Section 8 discusses the approach and draws conclusions.

2 Baseline: Requirements Models

A requirements-driven architecture for model-based self-reconfiguration needs a set of models to support full modeling of requirements. A well established framework in Requirements Engineering (RE) is goal-oriented modeling [3], where software requirements are modelled as goals the system should achieve (with assistance from external agents). Among existing frameworks for requirements models, we have chosen Tropos [5], for it allows to describe systems made up of several socially interacting actors depending on each other for the fulfillment of their own goals. Recently, Jurata et al. [8] have revisited the so-called “requirements problem” – what it means to successfully complete RE – showing the need for requirements modeling frameworks richer than existing ones. The core ontology they propose is based on the concepts of goal, softgoal, quality constraint, plan, and domain assumption. Direct consequence of this result is that goal models alone are insufficient to completely express system




Fig. 1. Enriched Tropos goal model used by our architecture

requirements, and in our framework we support some of the suggested ingredients to express requirements.

We adopt an enriched version of Tropos, which contains additional information to make it suitable for runtime usage: (i) activation events define when goals are triggered; (ii) commitment conditions express a time limit within which an agent should commit to its goal; (iii) contexts express when certain alternatives are applicable (like in Ali et al. [9]); (iv) preconditions define tasks applicability. In Fig. 1, two agents (*Patient* and *Supermarket*) interact by means of a dependency for goal *Provide Grocery*. The top-level goal of patient – *Have lunch* – is activated when it's 12AM, and the patient should commit to its achievement within one hour since activation. Two alternatives are available to achieve the goal, that is *Prepare lunch* and *Get lunch prepared*. In this scenario, the former option is applicable only in context *c1*, that is when patient is autonomous, whereas the latter option is applicable when the patient is not autonomous (*c2*). Goal *Prepare lunch* is and-decomposed to sub-goals *Get needed ingredients* and *Cook lunch*. The former goal is a leaf-level one, and there are two tasks that are alternative means to achieve it (means-end): *Take ingredients from cupboard* and *Order food by phone*. The latter task requires a dependency for goal *Provide grocery* on agent *supermarket*.

A shared language to express information about domain is clearly needed. This language is used to formally express contexts, preconditions, domain assumptions, and any relation between domain and requirements. We exploit an object diagram (as in [9]), where context entities are objects, their properties are attributes, and relations between entities are association links. For instance, the precondition for task *Order food by phone* (*Patient.house.hasPhone = true*) can be expressed in an object model with classes *Patient* and *House*, where *Patient* is linked to *House* by an aggregation called *house*, and *House* has a boolean attribute *hasPhone*. Domain assumptions are rules that should hold regardless of current goals. For example, a domain assumption for our small example is that each patient has exactly one house. Finally, we use a fine grained definition of tasks, in which each task is a workflow of monitorable activities to be carried out within time constraints. The completion of each activity is associated to the detection of an associated event, which is expressed over the context model. We provide further details about this formalism (timed activity diagram) in Section 5.

3 System Architecture

In this section we propose our conceptual architecture for structuring systems able to self-reconfigure. We present the architecture logical view in Fig. 2, exploiting an UML 2.0 component diagram to show the components and the connections among them. Component diagrams depict not only the structure of a system, but also the data flow between components (through provided and required interfaces).




Fig. 2. Logical view on the proposed architecture for self-reconfiguration

3.1 External Components

Our architecture supports systems characterized by decentralized and non-monolithic structure, such as Socio-Technical Systems (STSs) and Ambient Intelligence (AmI) scenarios, which require quick and effective reconfiguration in response to context change or failure. A set of external components interacts with the self-reconfiguration component, providing inputs and enacting reconfigurations.

The component *Context sensor* represents any system providing up-to-date information about the context where the system is running. In AmI settings, sensors are spread throughout the environment and collect data such as temperature, light level, noise, presence. Also desktop applications have several context sensors that provide useful values such as free memory, CPU utilization, mainboard temperature, and list of active processes. The component context sensor provides changes in the context through the interface *Events*. Possible events are changes in the light level, detection of humans in front of the door, identification of loud noise in the bathroom.

Monitored system is the system the self-reconfiguration component assists, that is the stakeholder whose requirements are monitored to diagnose and compensate failures. This system need not necessarily be software or hardware, but can be – and often is – a human or an organization. Examples of monitored systems are anti-virus software, patients living in smart-homes, firemen in crisis management settings. This component provides all available information concerning the current status of the system through the interface *Log*, and requires from the interface *System pushes* advice on what should be done (which *goals*) and how it should act (which *tasks*). A patient can be reminded to take her medicine by sending an SMS to her mobile phone (system pushes interface).

Support system represents any system connected to the monitored system by requirements level links by goal, task, or resource dependencies from the monitored system. For example, anti-virus software may depend on update sites for the resource “updated virus definition file”, while patients may depend on social workers for the goal “prepare breakfast”. The provided interface *Interaction log* contains information about the status of dependencies with the monitored system; the required interface *Task assignments* provides the tasks or goals for which the monitored system depends on support systems. If the patient should prepare breakfast but did not commit to it, the self-reconfiguration component can order breakfast from a catering service (the support system), enacting a dependency from the patient to the catering service for goal “prepare breakfast”.

Context actuator identifies any actuator in the environment which can receive commands to act on the context. Examples of actuators in AmI scenarios are sirens, door openers, automated windows, and remote light switches. The component gets from the required interface *Actuations* the commands to enact.

3.2 Self-reconfiguration Component

The self-reconfiguration capabilities of our architecture are provided by the component *self-reconfiguration*. We identified three major sub-components in the reconfiguration process, each enacting a phase in a Monitor-Diagnose-Compensate cycle. *Monitor* is in charge of collecting, filtering, and normalizing events and logs; *Diagnoser* identifies failures and discovers root causes; *Reconfigurator* selects, plans and deploys compensation actions in response to failures.

The monitoring phase starts with the *Event normalizer* gathering the current status of the monitored system, of the context, and of the interaction with support systems. These events are normalized according to a shared context model, e.g. defining transformation schemes using XSLT [10]. The event normalizer provides the translated data through the interface *Normalized events*. This interface is required by three different components, each handling a specific type of events. *Dependency monitor* computes

the status of existing dependencies and exposes it through the provided *Dependencies status* interface. *Context sensor* is in charge of updating the interface *Current context*, processing the normalized events related to changes in the context. For instance, if the house door is closed (*door.status = closed*) and we just received an event such as *open(door, time_i)*, the status of the door will change to open (*door.status = open*). The component *Task execution monitor* handles events concerning the execution of tasks and provides the interface *Task execution status*. For example, if the patient is executing the task “Open door” and event *pressed(patient, button, time_j)* is received, the status of task “Open door” will turn to success.

The diagnosis phase – responsibility of the component *Diagnoser* – is essentially a verification of the current status against requirements models. Models specify what should happen and hold: which goals should / can / cannot be achieved, which tasks can / cannot be executed, the domain assumptions that should not be violated. The richer the requirements models are, the more accurate the diagnosis will be. In contrast, the granularity of detected events is bounded by technological and feasibility concerns, and also increases the overhead introduced by the architecture. Detecting if a patient is sitting on a sofa is reasonably realizable (e.g., using pressure sensors), while detecting if she is handling a knife the wrong way is far more complex.

Contextual goal model manager analyzes the goal model to identify goals and tasks that should / can / cannot be achieved, and provides this output through the interface *Goals / Tasks applicability*. The component *Domain assumption verifier* checks a list of domain assumptions against the current context, and exposes identified violations through the provided interface *Violated domain assumptions*.

Dependency diagnoser computes problems in established dependencies. Dependencies fail not only if the dependee cannot achieve the goal or perform the task (e.g., the nurse cannot support the patient because she’s busy with another patient), but also when changes in the context modify goal applicability and the dependency is not possible anymore (e.g., the patient exits her house and thus cannot depend on a catering service anymore). *Task execution diagnoser* is needed to verify whether the current execution status of tasks is compliant with task applicability. For example, if the patient is preparing breakfast but already had breakfast, something is going wrong and this failure should be diagnosed. This component provides the interface *Failed tasks / goals*. *Goal commitment diagnoser* is in charge of detecting those goals that should be achieved but for whose fulfillment no action has been taken. In our framework, each top-level goal has a commitment time, a timeout within which a commitment to achieve the goal should be taken (i.e., an adequate task should begin). For instance, the patient should have breakfast within two hours since waking up. This component provides the interface *Uncommitted goals*.

The component *Failure diagnoser* requires the interfaces containing the identified failures and *Tolerance policies* provided by component *Policy manager*. The policy manager – handling policies set by system administrators – specifies when failures do not lead to reconfiguration actions. For example, lack of commitment for washing dishes can be tolerated if the patient’s vital signs are good (she may wash dishes after next meal). Diagnoses to be compensated are exposed through *Failure diagnosis*.

The reconfiguration phase – carried out by component *Reconfigurator* – should define compensation / reconfiguration strategies in response to any kind of failure. Its effectiveness depends on several factors: number of tasks that can be automated, available compensation strategies, extent to which the monitor system accepts suggestions and reminders. In our architecture we propose general mechanisms, but the actual success of compensation strategies is scenario-dependent and difficult to assess in a general way. Suppose a patient suddenly feels bad: if she lives in a smart-home provided with a door opener, the door can be automatically opened to the rescue team; otherwise, the rescue team should wait for somebody to bring the door keys.

The component *Prioritize diagnosis* selects a subset of failures according to their priority level and provides them through the interface *Selected Diagnosis*. Common criteria to define priority are failure severity, urgency of taking a countermeasure, time passed since failure diagnosis. Selected diagnoses are then taken as input by the component *Reaction strategy selector*, which is in charge of choosing a reaction to compensate the failure. This component acts as a planner: given a failure, it looks for appropriate reconfigurations, and selects one of them. Three different types of reconfigurations are supported by our architecture, each manifested in a specific interface. *Task reassignment reconfigurations* contains reconfigurations that involve the automated enactment dependencies on support systems. For example, if the patient didn't have breakfast and the commitment time for the goal is expired, the system could automatically call the catering service. *Push system reconfigurations* includes strategies that push the monitored system to achieve its goals (reminding goals or suggesting tasks). A push strategy for the patient that forgot to have breakfast is sending an SMS to her mobile phone. *Actuate reconfigurations* consists of compensations that will be enacted by context actuators. For instance, if the patient feels bad, the door can be automatically opened by activating the door opener. Three components use the interfaces provided by reaction strategy selector: *Task assigner*, *System pushing*, and *Actuator manager*. Their role is to enact the reconfigurations that have been selected, and each component provides a specific interface.

4 Creating the Architecture for an Existing System

We describe how the architecture can be used in practice to add self-reconfiguration capabilities to an existing distributed socio-technical system. The required input is a set of interacting sub-systems – sensors and effectors – that compose the distributed system. The following steps should be carried out: (i) define context model (ii) define requirements models; (iii) establish traceability links for monitoring; (iv) select tolerance policies for diagnosis; and (v) choose reconfiguration and compensation mechanisms.

Steps (i) and (ii) output the models we presented in Section 2, that is the context model, Tropos goal models, timed activity diagrams for tasks, and domain assumptions. Step (iii) defines what to monitor for at runtime, by connecting requirements to code. Traceability is ensured by associating events – produced by sensors – to *activities* that are part of a task, to *task preconditions*, to *contexts*, and to *activation conditions* for top-level goals. Events should also be normalized according to the context model defined in step (i).

Step (iv) is carried out to specify tolerance policies for failures. Indeed, some failures have to be addressed through reconfiguration, whereas some others can be tolerated. In step (v) the reaction mechanisms enacting self-reconfiguration are defined. Two sub-steps should be carried out: (i) definition of a *compensation plan* to revert the effects of the failed strategies, and (ii) identification of a *reconfiguration strategy* to retry goal achievement. Both steps exploit the actuation capabilities of the distributed system, i.e. reconfigurations consist of giving commands to effectors (execute a task, enact a dependency, issue a reminder).

5 Monitoring and Diagnosis Mechanisms

We detail now monitoring and diagnosis mechanisms included in our architecture. Efficient and sound algorithms need to be defined for successfully diagnosing problems in the running system. Failures are identified by comparing *monitored behavior* of the system to *expected and allowed behaviors*. Failures occur when (a) monitored behavior is not allowed or (b) expected behavior has not occurred.

Table 1 defines expected and allowed goals and tasks. We use first-order logic rules for clarity, but our prototype implementation is based on disjunctive Datalog [11]. We suppose that each goal instance differs from other instances of the same goal class for

Table 1. First-order logic rules to define expected and allowed goals and tasks

$\begin{array}{l} \text{goal}(G) \wedge \text{goal_parameters}(G, P) \wedge \neg \text{done}(G, P) \wedge \text{activation_evt}(G, P, T) \\ \wedge T \leq \text{current_time} \wedge \nexists G_p \text{ s.t.} \\ \quad \text{goal}(G_p) \wedge \text{decomposed}(G_p, G) \end{array}$ <hr style="width: 50%; margin: 0 auto;"/> $\text{should_do}(G, P)$	(i)
$\text{should_do}(G, P)$ <hr style="width: 50%; margin: 0 auto;"/> $\text{can_do}(G, P)$	(ii)
$\begin{array}{l} \text{goal}(G) \wedge \text{goal_parameters}(G, P) \wedge \neg \text{done}(G, P) \\ \wedge \exists G_p \text{ s.t.} \\ \quad \text{goal}(G_p) \wedge \text{goal_parameters}(G_p, P_p) \wedge \text{decomposed}(G_p, G, Dec) \\ \quad \wedge \text{can_do}(G_p, P_p) \wedge \text{context_cond}(Dec) \\ \quad \wedge \forall p \in P \text{ s.t. } (\exists p_p \in P_p \text{ s.t. } \text{name}(p, n) \wedge \text{name}(p_p, n)), \\ \quad \quad \text{value}(p, v) \wedge \text{value}(p_p, v) \end{array}$ <hr style="width: 50%; margin: 0 auto;"/> $\text{can_do}(G, P)$	(iii)
$\begin{array}{l} \text{task}(T) \wedge \text{task_parameters}(T, P) \wedge \text{pre_cond}(T, P) \wedge \neg \text{done}(T, P) \\ \wedge \exists G \text{ s.t.} \\ \quad \text{goal}(G) \wedge \text{goal_parameters}(G, P_p) \wedge \text{means_end}(G, T, Dec) \\ \quad \wedge \text{context_cond}(Dec) \wedge \text{can_do}(G, P_p) \\ \quad \wedge \forall p \in P \text{ s.t. } (\exists p_p \in P_p \text{ s.t. } \text{name}(p, n) \wedge \text{name}(p_p, n)), \\ \quad \quad \text{value}(p, v) \wedge \text{value}(p_p, v) \end{array}$ <hr style="width: 50%; margin: 0 auto;"/> $\text{can_do}(T, P)$	(iv)

actual parameters; for example, a patient's goal *Have breakfast* can be repeated every day, but with different values for the parameter *day*.

Rule (i) defines when a top-level goal should be achieved. This happens if G is a goal with parameters set P , the goal instance has not been achieved so far, the activation event has occurred before the current time, and G is a top-level goal (there is no other goal G_p and/or-decomposed to G). Rule (ii) is a general axiom saying that whenever a goal instance should be achieved, it is also allowed. Rules (iii) and (iv) define when tasks and decomposed goals are allowed, respectively. A goal instance G with parameters set P can be achieved if it has not been done so far and exists an achievable goal G_p with parameters P_p that is decomposed to G , the context condition on the decomposition is true, and the actual parameters of G are compatible with the actual parameters of G_p . A similar condition holds for means-end tasks, with two main differences: tasks are also characterized by a precondition – which should hold to make the task executable – and are connected to goals through means-end (rather than by and/or decomposition).

Expected and allowed goals and tasks identified by rules (i-iv) are used by Algorithm 1 to diagnose goals and tasks failures, comparing the monitored behavior to expected and allowed behaviors. The parameters of COMPUTEFAILURES are the monitored system, the examined goal instance, and the set of failures (initially empty). The algorithm is invoked for each top-level goal of the monitored system and explores the goal tree recursively; all parameters are passed by reference.

Algorithm 1 starts by setting the status of goal g to *uncommitted*, and the variable *means_end* to false. Lines 3-10 define the recursive structure of the algorithm. If the goal is and/or decomposed (line 3), the set G contains all the sub-goals of g (line 4), and the function COMPUTEFAILURES is recursively called for each sub-goal (lines 5-6). If the status of all the sub-goals is *success*, also the status of g is set to *success* (lines 7-8). If the goal is means-end decomposed (lines 9-10), G contains the set of tasks that are means to achieve the end g , and *means_end* is set to true.

If g is still *uncommitted* (line 11) each sub-goal (or means-end decomposed task) is examined (lines 12-39). If g is and-decomposed (lines 13-23), two sub-cases are possible: (a) if the sub-goal g_i is not allowed but its status is different from *uncommitted*, and the status of g is still *uncommitted*, the status of g is set to *fail* and the cycle is broken, for the worst case – failure – has been detected (lines 14-17); (b) if the status of g_i is *fail* (lines 18-23), the status of g is set to *fail* in turn, and the cycle is broken (lines 19-21); if g_i is in progress, the status of g is set to *in_progress*.

If g is or-decomposed or means-end (lines 24-39), it succeeds if at least one sub-goal (or task) succeeds. If g is means-end decomposed, the algorithm calls the function MONITORSTATUS, which diagnoses the execution status of a task (lines 24-25). If g_i is not allowed and its status is different from *uncommitted*, g_i is added to the set of failures (line 27), and if g is not committed its status is set to *fail* (lines 28-29). If g_i is allowed or its status is *uncommitted* (line 30), three sub-cases are possible: (a) if the status of g_i is *success*, the status of g is set to *success* and the cycle is terminated (lines 31-33); (b) if g_i is in progress, the status of g is set to *in_progress* and the loop is continued (lines 34-35); (c) if the status of g_i is *fail*, g_i is added to the set of failures, and if g is still *uncommitted* its status is set to *fail*. If g is a top-level goal that should be achieved, its status is *uncommitted*, and the commitment condition is true, then the status of g is set

to fail because no commitment has been taken (lines 40-41). If the status of g is fail, it is added to the list of failures (lines 42-43).

Algorithm 1. Identification of goal and task failures.

```

COMPUDEFAILURES( $s$  : System,  $g$  : Goal,  $F$  : Failure [])
1   $g.status \leftarrow$  uncommitted
2   $means\_end \leftarrow$  false
3  if  $\exists g_1 \in s.goals$  s.t.  $decomposed(g, g_1, dec)$ 
4    then  $G \leftarrow \{g_i \in s.goals$  s.t.  $decomposed(g, g_i, dec)\}$ 
5     for each  $g_i$  in  $G$ 
6     do COMPUDEFAILURES( $s, g_i, F$ )
7     if  $\forall g_i$  in  $G$ ,  $g_i.status = success$ 
8     then  $g.status \leftarrow$  success
9  else  $G \leftarrow \{t \in s.tasks$  s.t.  $means\_end(g, t, dec)\}$ 
10   $means\_end \leftarrow$  true
11  if  $g.status = uncommitted$ 
12  then for each  $g_i$  in  $G$ 
13    do if  $and\_decomposed(g, g_i, dec)$ 
14      then if  $g_i.can\_do = false$  and  $g_i.status \neq uncommitted$ 
15        and  $g.status = uncommitted$ 
16        then  $g.status \leftarrow$  fail
17         break
18      else switch
19        case  $g_i.status = fail$  :
20           $g.status \leftarrow$  fail
21          break
22        case  $g_i.status = in\_progress$  :
23           $g.status \leftarrow$  in\_progress
24      else if  $means\_end = true$ 
25        then  $g_i.status \leftarrow$  MONITORSTATUS( $g_i, g$ )
26        if  $g_i.can\_do = false$  and  $g_i.status \neq uncommitted$ 
27        then  $F \leftarrow F \cup g_i$ 
28         if  $g.status = uncommitted$ 
29         then  $g.status \leftarrow$  fail
30        else switch
31          case  $g_i.status = success$  :
32             $g.status \leftarrow$  success
33            break
34          case  $g_i.status = in\_progress$  :
35             $g.status \leftarrow$  in\_progress
36          case  $g_i.status = fail$  :
37             $F \leftarrow F \cup g_i$ 
38            if  $g.status = uncommitted$ 
39            then  $g.status \leftarrow$  fail
40  if  $g.should\_do = true$  and  $g.status = uncommitted$  and  $g.comm\_cond = true$ 
41  then  $g.status \leftarrow$  fail
42  if  $g.status = fail$ 
43  then  $F \leftarrow F \cup g$ 

```

Due to space limitations, the algorithms for diagnosing task failures are only sketched here, but they are fully described and discussed in [12]. We define tasks as workflows of activities each occurring within well-defined time limits (we refer to this formalism using the term “timed activity diagram”). Successful completion of an activity a is associated to the happening of an event e ($\text{happens}(e) \rightarrow \text{success}(a)$). Activities can be connected sequentially, in parallel (by fork and join nodes), and conditionally (through branch and merge nodes). A graphical example of this formalism is given in Fig. 4. The allowed branch of a decision point is defined by a branch condition. At any instant, a well-formed model has exactly one allowed branch. `MONITORSTATUS` is invoked by Algorithm 1 to check the status of a particular task; its parameters are a task t and a goal g linked through means-end. This algorithm returns task failure if the first activity of the task has happened but beyond the commitment time for goal g , whereas returns uncommitted if the first activity hasn’t happened so far but the commitment time for g has not expired. If no failure or uncommitment has been identified, the algorithm retrieves the next node in the timed activity diagram that defines the task, calling the recursive function `CHECKNODE`.

The behavior of Algorithm `CHECKNODE` depends on node type. If the node is an *activity*, we check if the associated event happened within the time limit. If it happened beyond time limit we return failure, if it happened before time limit we recursively check the next node. If the event hasn’t happened so far: (a) if time limit has expired we return failure; (b) if the activity is still within its time limit we return `in_progress`. If a *fork* is found, `CHECKNODE` is recursively called for all the forks. If any fork failed a fail value is returned. If all the forks joined a recursive check is performed on the node that follows the join, with time limit starting from the last joined fork. Otherwise, the algorithm returns `in_progress`. If a *branch* is met, all the branches should be checked. We return failure if any branch that was disallowed happened or an activity in an allowed branch failed. We return `in_progress` if allowed activities occurred within time limits and no disallowed activity happened. If none of these conditions hold, we check the node following the merge construct and return its status.

6 Case Study: Smart Homes

We show now a promising application for our architecture, emphasizing how requirements models are used to define and check allowed and expected behaviors, and how the architecture performs the MDC cycle. Our case study concerns smart homes: a patient lives in a smart home, a socio-technical system supporting the patient in everyday activities (such as eating, sleeping, taking medicine, being entertained, visiting doctor). Both smart home and patient are equipped with AmI devices that gather data (e.g., patient’s health status, temperature in the house) and enact compensations (e.g., open the door). The partial goal model in Fig. 3 represents the requirements of the patient; due to space limitations, we present here only the top-level goal “Have breakfast”.

Goal $g1$ is activated when the patient wakes up (activation event); a commitment to achieve $g1$ should be taken (either by the patient or by other agents) within two hours since goal activation. Four different contexts characterize the scenario: in $c1$ the patient is autonomous, in $c2$ the patient is not autonomous, in $c3$ the patient is at home, in $c4$




Fig. 3. Contextual goal model describing the patient health care scenario

the patient is not at home. If the patient is autonomous ($c1$ holds) $g1$ is decomposed into the subtree of goal “Eat alone” ($g2$); if $c2$ holds $g1$ is decomposed into the subtree of goal “Get eating assistance” ($g22$). In the former case, $c3$ enables the subtree of goal “Eat at home” ($g3$), whereas $c4$ enables the subtree of goal “Eat outside” ($g7$). When eating at home, the patient has to prepare food ($g4$), eat breakfast ($g5$), and clean up ($g6$). Goal $g4$ is means-end to two alternative tasks: “Prepare autonomously” ($p1$) and “Order catering food” ($p2$). The latter task requires interaction with actor “Catering service”, which should fulfill goal “Provide food” to execute $p2$. The other subtrees of Fig. 3 are structured in a similar way, thus we don’t detail them here.

Our requirements models characterize each task with a precondition that, if false, inhibits the execution of the task. If a task is executed but its precondition is false, a failure occurs (see rule (iv) in Table 1). A possible precondition for task “Prepare autonomously” is that there are both bread and milk in the house; a precondition for task “Order catering food” is that the house is provided with a landline phone.

Fig. 4 is a timed activity diagram for task $p1$. The task starts as goal “Prepare food” is activated. When the patient enters the kitchen ($a1$), there is evidence that she is going




Fig. 4. Timed activity diagram for monitoring the task “Prepare autonomously”

to prepare food. If this doesn’t happen within 45 minutes from goal activation the task fails. After a_1 , a fork node creates two parallel execution processes. In the first fork, the patient should open the fridge (a_2) and put the milk on stove (a_4); in the second fork, the bread cupboard should be opened (a_3) and bread has to be put on the table (a_5). The forks are then merged, and the next activity is to turn on the stove (a_6) within a minute since the last completed activity. Sequentially, the task requires the stove to be turned off within 5 minutes (a_7) and the milk to be poured into the cup (a_8).

We conclude this section with a description of a possible reconfiguration process. Let’s suppose that patient Mike wakes up at 8.00 am. Mike is autonomous (c_1) and at home (c_3); goal g_1 is expected, and the subtree of g_3 is the only allowed one (see rules (i) and (ii) in Table 1). At 8.20 am Mike enters the kitchen: checking the activity diagram for p_1 against this event changes the status of the goal g_4 to *in_progress*. In turn, this status is propagated bottom-up till g_1 (see Algorithm 1). At 8.25 Mike hasn’t neither opened the fridge nor opened the bread cupboard. This violates the specification of p_1 (Fig. 4). The reconfiguration strategy selector component selects to push the system, and the system pushing component sends an SMS message to remind Mike to have breakfast. The strategy succeeds, for Mike opens both the fridge (a_2) and the bread cupboard (a_3), then he puts bread on table (a_5). These events are compliant with the task specification of Fig. 4, thus the task is evaluated as *in_progress*. Anyhow, Mike does not put milk on stove (a_4) within one minute since a_2 , therefore a new failure is diagnosed. The compensation to address this failure is to automate p_2 , and the task assigner component selects a catering service. In an alternative scenario Mike exits house (the context c_4 is true, c_3 is false). This would change the tasks that can happen: the subtree of g_7 becomes the only possible one, and this influences all its sub-goals (rule (iii) in Table 1) and the tasks linked to leaf-level goals (rule (iv) in Table 1).

7 Related Work

Self-adaptive software has been introduced by Oreizy et al. [13] as an architectural approach to support systems that modify their behaviour in response to changes in the operating environment. This class of systems performs self-reconfiguration according to the criteria specified at development time, such as under what conditions reconfiguring, open/closed adaptation, degree of autonomy. The building units for self-adaptive software should be components and connectors. Compared to our work, the solution proposed in [13] is generic and flexible to many reconfiguration criteria, whereas ours is focused on particular types of models, that is requirements models.

Rainbow [1] is an architecture-based framework that enables self-adaptation on the basis of (i) an externalized approach and (ii) software architecture models. The authors of Rainbow consider architecture models as the most suitable level to abstract away unnecessary details of the system, and their usage both design- and at run-time promotes the reuse of adaptation mechanisms. Our proposal shares many features with Rainbow; the main difference is that we use higher level models to support the ultimate goal of any software system, that is to meet its requirements. A drawback of our choice is that establishing traceability links between requirements and code is more complex.

Sykes et al. [14] propose a three-layer architecture for self-managed software [15] that combines the notion of goal with software components. This approach is based on a sense-plan-act architecture made up of three layers: *goal management* layer defines system goals, *change management* layer executes plans and assembles a configuration of software components, *component* layer handles reactive control concerns of the components. Our proposal exploits a more elaborate goal representation framework, and differs in planning (predefined plans instead of plan composition) and reconfiguration.

Wang's architecture for self-repairing software [7] uses one goal model as a software requirements model, and exploits SAT solvers to check the current execution log against the model to diagnose task failures. We propose more expressive goal models, accurate specification of tasks based on timed activity diagrams, allow for specifying multiple contexts, and support dependencies on other actors / systems.

Feather et al. [2] address system behaviour deviations from requirements specifications; they introduce an architecture (and a development process) to reconcile requirements with behaviour. This process is enacted by jointly anticipating deviations at specification time and solving unpredicted situations at runtime, and examine the latter option using the requirements monitoring framework FLEA [16]. FLEA is used in conjunction with the goal-driven specification methodology KAOS [3]. Our architecture differs in the usage of different requirements models (Tropos rather than KAOS), support to a wider set of failures ([2] is focused on obstacle analysis), and applicability to scenarios composed of multiple interacting actors (such as AmI ones).

Robinson's ReqMon [17] is a requirements monitoring framework for specific usage in enterprise systems. ReqMon integrates techniques from requirements analysis (KAOS) and software execution monitoring. Although ReqMon's architecture covers all the reconfiguration process, accurate exploration is provided only for the monitoring and analysis phases. Our approach has broader applicability, can diagnose a larger set of failure types, and supports more reconfigurations mechanisms; on the contrary, ReqMon is particularly suitable for enterprise systems.

8 Discussion and Conclusion

We have proposed a novel architecture for self-configuring systems founded on principles adopted from Goal-Oriented Requirements Engineering, externalized adaptation, and BDI paradigm. Our approach adds self-reconfiguration capabilities to a wide variety of systems, among which Ambient Intelligence scenarios and Socio-Technical Systems. The architecture is a model-based one, with requirements models used to specify what can, should, and should not happen. We have detailed the main mechanisms for

monitoring and diagnosis, which describe how requirements models are checked against monitored information. We also introduced a case study – smart homes – to show how a realization of the architecture works in practice.

Several aspects will be addressed in future work. Firstly, we need a complete implementation of our architecture, as well as further experimentation on the smart-home case study. We also want to extend our framework so that it deals with a broader class of monitored phenomena, including attacks and failures caused by false domain assumptions. Finally, we propose to introduce mechanisms through which a system can extend its variability space through collaboration with external agents.

Acknowledgements

This work has been partially funded by EU Commission, through the SERENITY project, and by MIUR, through the MEnSA project (PRIN 2006).

References

1. Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B., Steenkiste, P.: Rainbow: architecture-based self-adaptation with reusable infrastructure. *Computer* 37(10), 46–54 (2004)
2. Feather, M., Fickas, S., Van Lamsweerde, A., Ponsard, C.: Reconciling system requirements and runtime behavior. In: *IWSSD 1998*, pp. 50–59 (1998)
3. Dardenne, A., Lamsweerde, A., Fickas, S.: Goal-directed requirements acquisition. In: *Science of computer Programming*, pp. 3–50 (1993)
4. Emery, F.: *Characteristics of socio-technical systems*. Tavistock, London (1959)
5. Bresciani, P., Perini, A., Giorgini, P., Giunchiglia, F., Mylopoulos, J.: Tropos: An agent-oriented software development methodology. *JAAMAS* 8(3), 203–236 (2004)
6. Rao, A., Georgeff, M.: An abstract architecture for rational agents. In: *KR&R 1992*, pp. 439–449 (1992)
7. Wang, Y., McIlraith, S., Yu, Y., Mylopoulos, J.: An automated approach to monitoring and diagnosing requirements. In: *ASE 2007*, pp. 293–302 (2007)
8. Jureta, I.J., Mylopoulos, J., Faulkner, S.: Revisiting the core ontology and problem in requirements engineering. In: *RE 2008* (2008)
9. Ali, R., Dalpiaz, F., Giorgini, P.: Location-based software modeling and analysis: Tropos-based approach. In: Li, Q., Spaccapietra, S., Yu, E., Olivé, A. (eds.) *ER 2008*. LNCS, vol. 5231, pp. 169–182. Springer, Heidelberg (2008)
10. Clark, J., et al.: Xsl transformations (xslt) version 1.0. W3C Recommendation 16(11) (1999)
11. Eiter, T., Gottlob, G., Mannila, H.: Disjunctive datalog. *ACM Transactions on Database Systems* 22(3), 364–418 (1997)
12. Dalpiaz, F., Giorgini, P., Mylopoulos, J.: An architecture for requirements-driven self-reconfiguration. Technical Report DISI-09-010, DISI, University of Trento (2009)
13. Oreizy, P., Medvidovic, N., Taylor, R.N.: Architecture-based runtime software evolution. In: *ICSE 1998*, pp. 177–186 (1998)
14. Sykes, D., Heaven, W., Magee, J., Kramer, J.: From goals to components: a combined approach to self-management. In: *SEAMS 2008*, pp. 1–8. ACM Press, New York (2008)
15. Kramer, J., Magee, J.: Self-managed systems: an architectural challenge. In: *ICSE 2007*, pp. 259–268. IEEE Computer Society, Washington (2007)
16. Fickas, S., Feather, M.S.: Requirements monitoring in dynamic environments. In: *RE 1995*, pp. 140–147. IEEE Computer Society Press, Washington (1995)
17. Robinson, W.N.: A requirements monitoring framework for enterprise systems. *Requirements Engineering* 11(1), 17–41 (2006)