2010 Seventh IEEE International Conference and Workshops on Engineering of Autonomic and Autonomous Systems

Extracting component-oriented behaviour for self-healing enabling

Marco Bakera*, Christian Wagner*, Tiziana Margaria*, Emil Vassev†, Mike Hinchey‡, Bernhard Steffen§
*Universität Potsdam, Chair Service and Software Engineering, Potsdam, Germany
{margaria,wagner}@cs.uni-potsdam.de
†University College Dublin, Ireland, emil.vassev@lero.ie
‡University of Limerick, Ireland, mike.hinchey@lero.ie
§Technische Universität Dortmund, Chair Programming Systems, Dortmund, Germany
steffen@cs.tu-dortmund.de

***Abstract*—Rich and multifaceted domain specific specification languages like the Autonomic System Specification Language (ASSL) help to design reliable systems with self-healing capabilities. The GEAR game-based Model Checker has been used successfully to investigate in depth properties of the ESA ExoMars Rover. We show here how to enable GEAR's game-based verification techniques for ASSL via systematic model extraction from a behavioral subset of the language, and illustrate it on a description of the Voyager II space mission. This way, we close the gap between the design-time and the run-time techniques provided in the SHADOWS platform for self-healing of concurrency, performance, and functional issues.[1]**

I. Introduction

The SHADOWS project (**S**elf-**h**ealing **A**pproach to **D**esigning **C**o**m**plex Soft**w**are **S**ystems) [19], [20], [21] aims at developing technologies that augment large software systems with a sort of immune response against various issues and contingencies that can occur at design-time or runtime. It targets general issues in the areas of performance, concurrency, and functional problems. Without self-healing protection, these would result in a costly, partial or complete breakdown of the system. By design, the SHADOWS techniques and methodology have wide applicability, reaching well beyond the scope of the case studies and application domains investigated in the project. In fact, the aim is to provide general techniques that augment a given or planned system independently from the system's functionality.

Within SHADOWS, one of the central application domains brought in by the industrial partners and driving the development of the new techniques was avionics and space. They were contributed by our partners IAI - Israel Aerospace Industry, based in Haifa (IL) - and Artisys (based in Brno, CZ), a supplier to international constructors of airplanes and spacecrafts. Focussing on functional healing at design time, we developed a number of enabling techniques to functional self-healing. In particular, we introduced game based model checking of behavioral models in the GEAR tool [1], [15], [2] as a deep diagnosis tool for early realignment between behavioral models and requirements expressed as temporal properties.

[1]A preliminary version of this work was presented at NFM'09[3].

We successfully applied this technique to investigate the recovery behavior of the ExoMars Rover[2], as described in Kapellos [13], and we were able to modify and adapt the display and the illustration of the game and it playing in a way acceptable to engineers.

The weak point was however the lack of a link to an adequate, formal description of the Rover's behavior. We derived our models and properties from the literature (textual descriptions and previous studies) [4], [13], while for a stringent demonstration of the techniques and for a validation of the underlying SHADOWS methodology it would have been advantageous to start from real models.

In this paper we show 1) how we are able to link the behavioral modelling style of our techniques with ASSL [23], a rich domain-specific language for the specification of autonomous systems, equipped with a formal semantics [23], and 2) how we can easily and systematically translate (parts of) the specification of the Voyager's behavior into Service Logic Graphs (SLGs, introduced formally in Section III-A), thus enabling the application of the SHADOWS technologies to the large class of autonomous systems describable in ASSL. The advantage of SLGs over other models is that they are closer to the field engineer's understanding, thus making advanced game-based diagnosis features accessible to non-experts in formal methods and models.

A. The ExoMars Rover Case Study

In the concrete mission example we examined in SHADOWS, the ESA ExoMars Rover is sent on a surface mission on Mars where it has to accomplish several tasks, including the acquisition of subsurface soil samples using a drill. As customary, the mission is organized in a hierarchical three-tier control model which accounts for partial autonomy of the Rover. *Mission* plans are designed and enforced by the ground control center, while finer-grained operational decisions, at the *task* level, are completely autonomous: the Rover has its own planning capabilities, which allows it to transform a task

[2]The ESA ExoMars Rover was studied in the FORMID Project (*FO*rmal *R*obotic *M*ission *I*nspection and *D*ebugging), that aimed at creating a development environment for the verification and analysis of robotic missions [4].

978-0-7695-4004-7/10 $26.00 © 2010 IEEE
DOI 10.1109/EASe.2010.23

assignment into a suitable executable sequence of *actions* in a context-dependent and error-aware way.

We showed in diverse publications about verification [1], [15], [2] how to take advantage of the interactive and exploratory benefits of game-based verification technologies. In the case of problems within highly reactive and concurrent systems – as in the context of autonomous aerospace missions – it is hard to automatically find recovery mechanisms to overcome these problems. Even for human system developers it is non-trivial to completely understand the nature of a problem if mismatches between the behavioral specification and the system implementation occur.

B. The NASA Voyager Mission Case Study

The NASA Voyager Mission started in 1977 and was designed for exploration of the outer planets of the Solar System. As the twin spacecraft Voyager I and Voyager II flew, they took pictures of planets and their satellites in 800x800 pixel resolution, then radiotransmitting them to Earth. Voyager II has two on-board television cameras - one for wide-angle images and one for narrow-angle images - that record images in black and white. Each camera is equipped with a set of colour filters, which help images to be reconstructed as fully-colored ones. Voyager II uses radar-like microwave frequencies to send the stream of pixels toward Earth. The signal suffers on this distance a 20 billion times attenuation [6].

In Vassev and Hinchey [24], the mission is specified as an autonomic system composed of the Voyager II spacecraft and four antennas on Earth, all specified as distinct autonomic elements. This paper bases on this specification and on those results on the behavior of the system.

In the rest of this paper, we briefly sketch ASSL (Sect. II) and then how to map the ASSL specification with our models (Sect. III), and illustrate it on the model for the NASA Voyager mission. We then discuss verification issues (Sect. IV) and how this model generation technique enables the use of SHADOWS self-healing techniques in a smooth fashion that combines design- and runtime (Sect. V. We then discuss some related work (Sect. VI), and finally conclude (Sect. VII).

II. ASSL

The Autonomic System Specification Language (ASSL) is a framework that provides a multi-tier structure for specifying and validating autonomic systems and targets the generation of an operational prototyping model for any valid ASSL specification [23]. ASSL provides a multi-tier specification model that tackles autonomic systems (ASs) as composed of autonomic elements (AEs) interacting over interaction protocols (ASIP and AEIP). We concentrate here on the behavioral aspects of the AS and AE description, since they are the part of ASSL that finds direct counterpart in the GEAR behavioral models.

- The AS tier - provides a general and global AS perspective. It defines the general system rules in terms of *service-level objectives (SLO)* and *self-management policies*, *architecture topology*, and global *actions*, *events*, and *metrics* applied in these rules. It is similar to the mission and task level of the ExoMars description.
- the AE tier - provides a unit-level perspective, It defines interacting sets of individual autonomic elements (AEs) with their own behavior. This tier is composed of AE rules (*SLO* and *self-management policies*), an *AE interaction protocol (AEIP)*, *AE actions*, *AE events*, and *AE metrics*. It is similar to the Action level of the ExoMars description.

A. How the Voyager takes pictures

When a space picture must be taken and sent to Earth, the Voyager exhibits autonomous-specific behavior. The spacecraft must detect on the fly interesting objects and take their pictures. This reveals a sort of autonomic event-driven behavior that can be easily specified with ASSL at the three main tiers - AS (autonomic system) tier, ASIP (autonomic system specification protocol) tier, and AE (autonomic element) tier [23].

The Voyager II spacecraft and the antennas on Earth are specified at both AS and AE tiers as autonomic elements that follow their autonomic behavior encoded as a self-management policy called `IMAGE_PROCESSING`. ASSL specifies self-management policies with special ASSL constructs - fluents[3] and mappings [23]. Whereas the former are special ASSL constructs used to denote specific system states, the latter simply map fluents to ASSL actions (actions to be performed when the system gets into a fluent).

B. AS Tier Specification.

The `IMAGE_PROCESSING` self-management policy is specified at the AS tier to process images from four antennas on Earth located in Australia, Japan, California, and Spain. In fact, we consider this specification as forming the autonomic image-processing behavior of the Voyager Mission base on Earth.

As shown in Figure 1, the policy is specified with four policy *fluents* - one per antenna. Fluents denote specific system states. They are *initiated* by events prompted when an image has been received and *terminated* by events prompted when the received image has been processed. Further, all the four fluents are *mapped* to an ASSL *action*: that is to be performed when the system enters in one of the fluents. Figure 2 shows the specification of the events that initiate and terminate the fluent presented by Figure 1. Note that the first event is prompted to occur in the system when a special message has been received. In addition, a `processImage` action (see [25] for this action's specification) is specified to process images from all four antennas.

At the autonomic system interaction protocol (ASIP) tier, the image messages (one per antenna), a communication channel that is used to communicate these messages, and communication functions to send and receive these messages over that communication channel to the Earth are specified [25].

[3]ASSL adopts some AI-planning terminology: a fluent is comparable to a state variable in our transition system view.

```
FLUENT inProcessingImage_AntSpain {
  INITIATED_BY { EVENTS.imageAntSpainReceived }
  TERMINATED_BY { EVENTS.imageAntSpainProcessed }}

MAPPING {
  CONDITIONS { inProcessingImage_AntSpain}
  DO_ACTIONS { ACTIONS.processImage("Antenna_Spain") }}
```

Fig. 1. An IMAGE_PROCESSING Fluent

```
EVENT imageAntSpainReceived {
  ACTIVATION { RECEIVED {
    ASIP.MESSAGES.msgImageAntSpain } }}

EVENT imageAntSpainProcessed { }
```

Fig. 2. AS-tier Events

C. ASIP Tier Specification.

They concern the autonomic system interaction protocol (ASIP) [23], which is used by the four antennas when communicating with the Voyager Mission base on Earth. Here, at this tier we specified four image messages (one per antenna), a communication channel that is used to communicate these messages, and communication functions to send and receive these messages over that communication channel [25].

D. AE Tier Specification.

At this tier, we have five autonomic elements: the Voyager II spacecraft and the four antennas on Earth. For each, an own part of the IMAGE_PROCESSING self-management policy is specified.

a) AE Voyager.: The spacecraft's IMAGE_PROCESSING self-management policy (see Figure 3) uses two fluents. The inTakingPicture fluent is initiated by a timeToTakePicture event and terminated by a pictureTaken event. This event also initiates the inProcessingPicturePixels fluent, which is terminated by the pictureProcessed event. The fluents are mapped to the actions takePicture and processPicture respectively. Metrics are used e.g. to count all the detected interesting objects which the Voyager AE takes pictures of.

b) AE Antenna.: Also the four antennas receiving signals from the Voyager II spacecraft are specified as autonomic elements. Their IMAGE_PROCESSING self-management policy uses pairs of fluents inStartingImageSession - inCollectingImagePixels, one for each colour filter. These sets of fluents determine the states of the antenna AEs when an image-receiving session is starting and when an antenna AE is collecting the image pixels.

Since the Voyager AE processes the images by applying different filters and sends each filtered image separately, we have distinct fluents for each colour and antenna. This allows an antenna AE to process a collection of multiple filtered

```
AESELF_MANAGEMENT {
  OTHER_POLICIES {
    POLICY IMAGE_PROCESSING {
      FLUENT inTakingPicture {
        INITIATED_BY { EVENTS.timeToTakePicture }
        TERMINATED_BY { EVENTS.pictureTaken }
      }
      FLUENT inProcessingPicturePixels {
        INITIATED_BY { EVENTS.pictureTaken }
        TERMINATED_BY { EVENTS.pictureProcessed }
      }
      MAPPING {
        CONDITIONS { inTakingPicture }
        DO_ACTIONS { ACTIONS.takePicture }
      }
      MAPPING {
        CONDITIONS { inProcessingPicturePixels }
        DO_ACTIONS { ACTIONS.processPicture }
      }
    }
  }
} // AESELF_MANAGEMENT


FLUENT inStartingGreenImageSession {
  INITIATED_BY { EVENTS.
            greenImageSessionIsAboutToStart }
  TERMINATED_BY { EVENTS.
            imageSessionStartedGreen }
}

FLUENT inCollectingImagePixelsBlue {
  INITIATED_BY { EVENTS.imageSessionStartedBlue }
  TERMINATED_BY { EVENTS.imageSessionEndedBlue }
}


EVENT greenImageSessionIsAboutToStart {
  ACTIVATION { SENT { AES.Voyager.
    AEIP.MESSAGES.msgGreenSessionBeginAus } }
}
EVENT imageSessionStartedBlue {
  ACTIVATION { RECEIVED { AES.Voyager.
    AEIP.MESSAGES.msgBlueSessionBeginAus } }
}
```

Fig. 3. AE antenna self-management policies, fluents, events

images simultaneously.[4] It is the Voyager AE that notifies an antenna that an image-sending session begins and ends. Figure 3 shows two of the IMAGE_PROCESSING fluents. They are further mapped to AE actions that collect the image pixels per filtered image (see [25]).

In Figure 3 we see how two of the events initiate the AE Antenna fluents. The greenImageSessionIsAboutToStart event is prompted (triggered) when the Voyager's msgGreen-SessionBeginSpn message has been sent and the image-SessionStartedBlue event is prompted when the Voyager's

[4]Note that according to the ASSL formal semantics, a fluent cannot be re-initiated while it is initiated, thus preventing the same fluent be initiated simultaneously twice or more times [23].

```
FLUENT inStartingGreenImageSession {
  INITIATED_BY { EVENTS.greenImageSessionIsAboutToStart }
  TERMINATED_BY { EVENTS.imageSessionStartedGreen }
}
FLUENT inCollectingImagePixelsBlue {
  INITIATED_BY { EVENTS.imageSessionStartedBlue }
  TERMINATED_BY { EVENTS.imageSessionEndedBlue }
}
```

Fig. 4. AE Antenna Fluents

`msgBlueSessionBeginSpn` message has been received by the antenna.

III. Mapping ASSL to GEAR Models

ASSL specifications describe all the different aspects of an autonomic system in one comprehensive document. This is practical, but by nature in realistic cases it becomes very complex, the complexity to a good extent due to the many cross references between the specification elements. A trace through the specified autonomic system may request jumping between different aspects (e.g. from messages $\rightarrow$ events $\rightarrow$ fluents $\rightarrow$ mappings $\rightarrow$ actions) and "pages". Another submission to this workshop proposes mapping ASSL specifications to LTS, in order to verify LTL properties [26] and with focus on concerns of state space explosion. Here, we address a different mapping, that privileges intuition of the graphical models, expression of constraints in any mu-calculus derivative, and a deep support to diagnosis by means of reverse model checking and games. Our models are Service Logic Graphs (SLG).

A. Behavioral models: Service Logic Graphs

To complement the original textual view, and in perspective to visualize and reify certain aspects of the SOS semantics of ASSL, we map selected behavioral elements of the specification to GEAR's behavioral models. These can be visualized as Service Logic Graphs (SLG) in the jABC framework [18], [22] (of which GEAR is the model checking plugin) and analyzed, guiding the user through the processes and workflows of the specified autonomic system. These same models are directly amenable to model checking.

SLGs themselves are composed of reusable building blocks that are called *Service Independent Building Blocks* (SIBs) [10], [11], and may represent both a single atomic service or a whole subgraph (i.e. another SLG). Thus SLGs can be hierarchical, which grants a high reusability not only of the building blocks, but also of the models themselves, within larger systems. SLGs formally stem from the concept of Kripke Transition Systems [16].

Kripke Transition System: A Kripke Transition System K is defined as a tuple $(S, \mathsf{Act}, \rightarrow, I)$ over a set of atomic propositions AP, disjoint from Act, where

- S are the states of the model,
- Act is a set of actions,
- $\rightarrow \subseteq S \times \mathsf{Act} \times S$ are the possible transitions in the model, and
- a labelling interpretation function $I : S \rightarrow 2^{\mathsf{AP}}$ equips states with atomic propositions.

A KTS is best-suited for verification tasks that focus on transitions of the system as being the edges. On the contrary, one can think of an SLG as being the engineer's view on the system that focuses on the *actions* of the system as being the nodes.

Service Logic Graph: A *Service Logic Graph (SLG)* is defined as a tuple $(S, \mathsf{Act}, \rightarrow, I)$ over a set of atomic propositions AP, disjoint from Act, where

- S represents the occurrences of the Service Independent building blocks (SIBs), which are the actions or functions in the graph
- Act is the set of possible branching conditions, to be determined upon execution of the preceding SIB,
- $Trans = (s, a, s)$ is a set of transitions where $s, s \in S$ and $a \in Act$, and
- a labelling interpretation function $I : S \rightarrow 2^{\mathsf{AP}}$ equips SIB occurrences with atomic propositions.

The structural match, KTS and SLGs are both graph structures with labeled branches and nodes that are enriched with atomic propositional properties, sufficies to adopt the established model checking technologies for the SLGs.[5].

In mapping the elements of the ASSL specification to a graphical representation in the behavioral model we focus on those constructs that describe behavioral and self-* aspects: these are the central elements which will be most frequently used to specify autonomic systems. We currently cover

- the AS tier: Service-Level Objectives, Self-Management Policies, Actions and Events.
- the AE tier: Service-Level Objectives, Self-Management Policies, Actions and Events, and additionally behavioral models and outcomes.

Architectural, communication, and quantitative aspects will be dealt in future work.

B. Mapping ASSL Elements

From the point of view of model generation, AS and AE specifications are structurally similar wrt. events, self management policies, and actions, but differ in the scoping: while the AS specification has a global scope, the AE specification is only valid for the local element. Due to the similarities, we focus in the description on the autonomic element (AE) tier. The AS tier is captured similarly, by means of hierarchy (where single nodes of the AS-level KTS are expandable to AE-level models).

We refer to Figure 6, showing a specification fragment for the Voyager (right) and the corresponding section of the behavioral model (left). In the textual specification (right), we have two events, one fluent with a mapping, and one action. Dashed arrows illustrate a trace of an event within the specs. Arrows indicate the correspondence between elements of the

[5]In fact both system representation styles can easily be translated into each other by adequately mapping edges to nodes and vice versa

```
EVENT greenImageSessionIsAboutToStart {
  ACTIVATION { SENT { AES.Voyager.AEIP.MESSAGES.msgGreenSessionBeginAus } }
}
EVENT imageSessionStartedBlue {
  ACTIVATION { RECEIVED { AES.Voyager.AEIP.MESSAGES.msgBlueSessionBeginAus } }
}
```

Fig. 5. AE Antenna Fluents

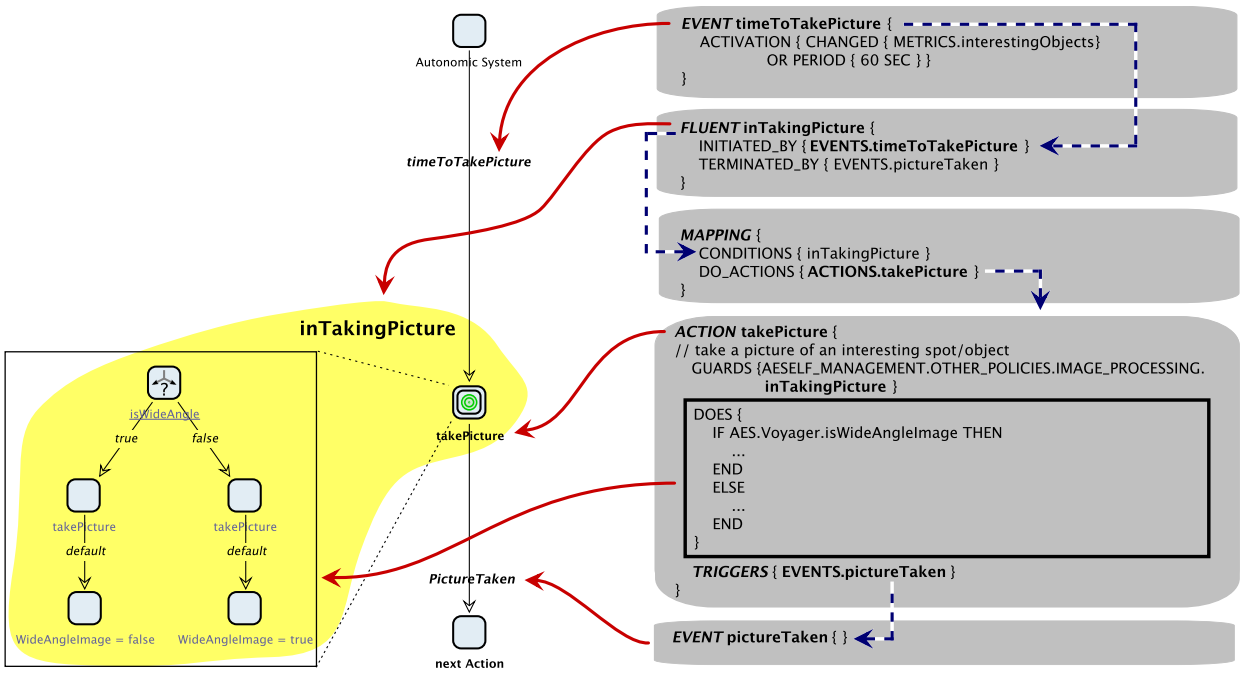

Fig. 6. *Action, Event, Fluent*, and *Mapping* in KTS behavioral model representation

ASSL-specification and of the behavioral. The InTakingPicture cloud defines the current state of the system (an atomic proposition).

AE Event. Event is the central language element in ASSL. It specifies fluents, actions, and policies globally in the AS tier and locally in the AE tier. Events could be activated by messages, other events, actions or metrics. In our behavioral model, events are mapped to homonymous Branches. In Figure 6, the behavioral model starts with the event *timeToTakePicture*, activated for interesting objects or after a time period of 60 s. It initiates the self management policy (fluent) *inTakingPicture*.

AE Self Management Policy. It defines the behavior of the autonomic system by connecting specific system states with the intended (re)action. A policy consists of two elements:

- A *fluent*, similar to a state. It is initiated (ie., that state is reached) when the system satisfies specific *conditions*. It will be terminated (left) if specific events occur. Fluent activation and termination is driven by events.
- A *mapping* of certain conditions to *actions*. The conditions test fluents: in a certain state, certain actions (in the AS or AE tier) are performed. Actions activate specified actions.

They are central to the model extraction: the information contained in a self management policy is used and useful both for model construction and for verification.

Together, fluent and mappings define the control flow, i.e. create branches with the name of the initiating event. They define all possible incoming branches of an action. The specific condition that activates the fluent is stored in the *context* of the system's model. The context represents the current global state of the system, like a global Blackboard or shared memory-mechanism. For model checking purposes, the fluent is additionally associated as atomic proposition to the corresponding node(s) of the behavioral model. This enables global model checking. The fluent can be used as preconditions of actions. They hold on all states in the region between initiation and termination.

The fluent in our example is activated by the *timeToTakePic-*

ture-event and the overall status of the autonomous system is changed to *intakingPicture*. This change activates an action: *takePicture* which is specified in the Mapping section of the self management object.

The self management policy which connects the event to actions is additionally used to annotate the nodes in the behavioral model with atomic propositions (AP). The name of the AP is equal to the name of the fluent. They can later be used for model checking.

AE Action. Actions are routines performed by AE or AS (global and local). In our behavioral model, they are the second essential element: the nodes of our behavioral model, named as the action. The different elements of an action are used to describe the nodes and for verification purposes. Action *parameters* become parameters of a node, the *does* part represents the body of a node. It can be a single action (then the node is an atomic node), but for complex *does* it is an entire behavioral model. We then model them as a SLG hierarchy, as in Figure 6: the node *takePicture* has a corresponding sub-model, presented on the left. The *guards*, *returns* and *outcomes* are used for verification. We offer two possibilities for verification:

- The Localchecker uses the Guard to verify if an Action could be executed within the current system state (defined by the fluents and stored in a global context).
- We can use a model checker to verify relations of nodes and actions expressed as temporal logic constraints. GEAR uses internally the modal mu-calculus [14] enriched with forward and backward modalities, so it is best equipped e.g to express dataflow properties, or other behavioral constraints like e.g. CTL formulas.

The specified action in Figure 6 contains a guard which must conform to the AP annotated at the node.

How to define the outgoing branches of a node depends on the information found in the action's specification: Actions can use events; triggers are communication functions to communicate with the autonomic system and its elements. We thus have several possibilities to detect outgoing branches.

- a trigger statement in the specification of an action will create an event which introduces the next fluent and/or action, and is comparable to an outgoing branch,
- event statements in the Does part are added as possible outgoing branches,
- if communication functions are used, we follow the chain from the function to the communication channel to the events which will be activated by a specific message in the channel. It is not unusual that more than one event will be created from one message.

The *takePicture* action of Figure 6) is closed by a new event *pictureTaken*. This is specified in the triggers section and represents the outgoing branch of this node. The new event will again initiate a fluent, and it terminates the *inTakingPicture* fluent. Therefore, the next action has another AP.

AE Outcomes, AE Behavioral Models, and AE Recovery Protocol. These elements are not yet treated in depth. They will become relevant when applying the SHADOWS methodology. In short, **AE Outcomes** are post-conditions of actions or behavioral models - they are useful for verification purposes. An **AE Behavioral Model** is comparable, from the model generation point of view, to a further mapping in the self management policy. It consists of conditions, a do element where an action is activated, and outcomes. We can model the behavioral model similarly to an action (atomic or hierarchical). Condition and outcomes become the pre- and post-condition and the action is the implemented behavior.
An **AE Recovery Protocol** should guarantee fault-tolerant operation of the autonomous system (e.g. create snapshots, log messages, consistency checking). A recovery protocol specification is rather complex, and it is specified in a separate submodel.

IV. Verifying the Voyager's Behavioral Model

Figure 7 contains the behavioral model of the Voyager II spacecraft. Note that the error handling graph at the right was not part of the original ASSL specification.

A simple verification issue that immediately emerges is whether the system takes care of an error-handling process whenever picture pixels are transmitted. This can be easily expressed in CTL [7] as

$$\mathsf{AG}(\mathsf{inProcessingPicturePixels} \Rightarrow \mathsf{EF}(\mathsf{errorHandling}))$$

This formula can be interpreted as follows:

> *Wherever the system evolves to (the* AG*-part), whenever picture pixels are about to be processed (the atomic proposition* inProcessingPicturePixels*) it follows that the system has an option to evolve into an error-handling process (the* EF(errorHandling)*-part).*

Since the original model of Figure 7 does not support any kind of self-healing capabilities, this property does not hold.

Therefore, in a first attempt to reconcile model and property, we added an error-handling routine directly in the model. We slightly changed the design manually, by refining the sendImgPixelMsg action, originally atomic, to an entire routine. Now, if problems during the transmission process occur, the system tries to resend those picture pixels that were not transmitted correctly. If the problem still exists afterwards, the system is halted and needs manual interaction from ground control.

A game-based approach as presented in Bakera et al. [15] would do much more than just allowing the identification of the missing recovery mechanism in the original specification. Enabling this investigation for self-healing and self-healing enactment is our aim. A domain-dependent guidance also enables to pinpoint that part of the model which is best-suited for integration of recovery mechanisms. Due to space limitation, we cannot discuss this process in detail in this contribution.

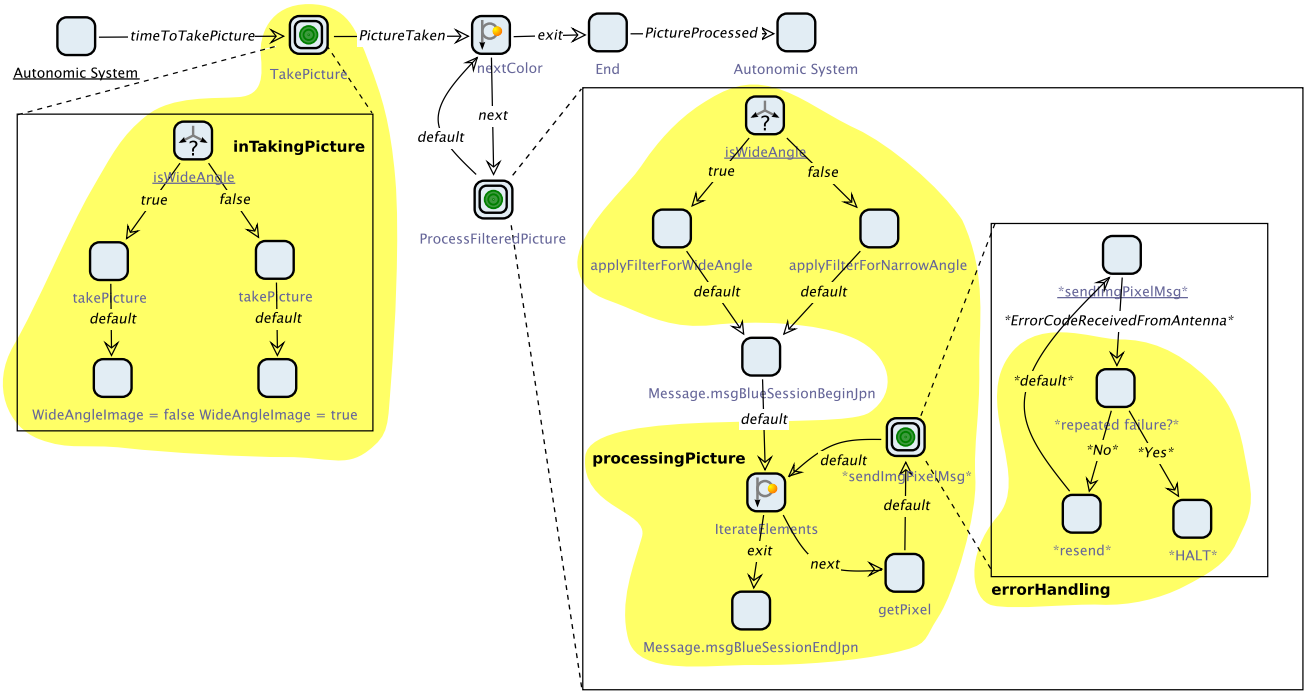

Fig. 7. Behavioral model of the picture transmission process. Bottom right: a new error handling recovery mechanism.

A. *Enabling Model based Self-healing*

Within SHADOWS, we adopt a model-based approach, where models of desired software behavior direct the self-healing process. This allows for life cycle support of self-healing applicable to industrial systems. We contribute to SHADOWS a number of enabling technologies for model-driven self-healing residing in the functional part of the architecture. Our technologies deal with self-healing issues at design-time for ensuring functional correctness, i.e. correctness with respect to the system's behavior over time. For this, we apply among others a game-based model-checking approach as a powerful technique for the verification, diagnosis and adaptation according to desirable temporal properties that the system's behavior must exhibit.

In particular, we show how to model the several abstraction levels of the system's behavior in a uniform and formal but intuitive way. This happens in term of processes in the jABC framework [22], a mature, model-driven, service-oriented process definition platform. Subsequently, we leverage the formality of these models to prove properties by model checking. In particular we exploit the interactive character of game-based model checking to show how to discover an error, then localize, diagnose, and correct it. Design-time healing technologies that naturally emerge when dealing with self-adaptive systems, as in the context of the SHADOWS project, demand for a deeper insight of design-time faults to effectively identify and overcome them.

The use of models rather than code is already a significant step towards the understandability of the actual behavior's descriptions to non programmers, like the engineers, in charge of designing a space module. This enables e.g. early discovery of misbehaviors, hazards, and ambiguities via design-time analysis. We strive to improve the diagnostic features making them as detailed as necessary yet as intuitive as possible.

For this purpose we use GEAR [15], a model checker capable of the full modal μ-calculus temporal logic with a rich user interface that allows for pinpointing problems in system design. This is achieved by interactively exploring the problem space in a game-based way. The game-based nature of GEAR's verification algorithm supports the system designer at design-time to interactively explore the problem space upon property mismatches.

In case of the Voyager mission case study such properties can be used to check for complete picture transmission to the four antennas in case of transmission interrupts. Further the verification process is able to assure the application of all four color filters before picture transmission. In addition it is essential for the picture transmission to send closing notification signals of transmission endings to the antennas. This as well can be assured by the aforementioned Model Checking techniques.

If problems occur in the verification task one immediate result of the game-based algorithm of the Model Checker is an interactive counter-example. This counter-example both pinpoints the problem of the property mismatch and provides a strategy encoded into the counter-example to adapt and self-heal the system. GEAR [15] elaborates on the application of this technique on an ESA mission example.

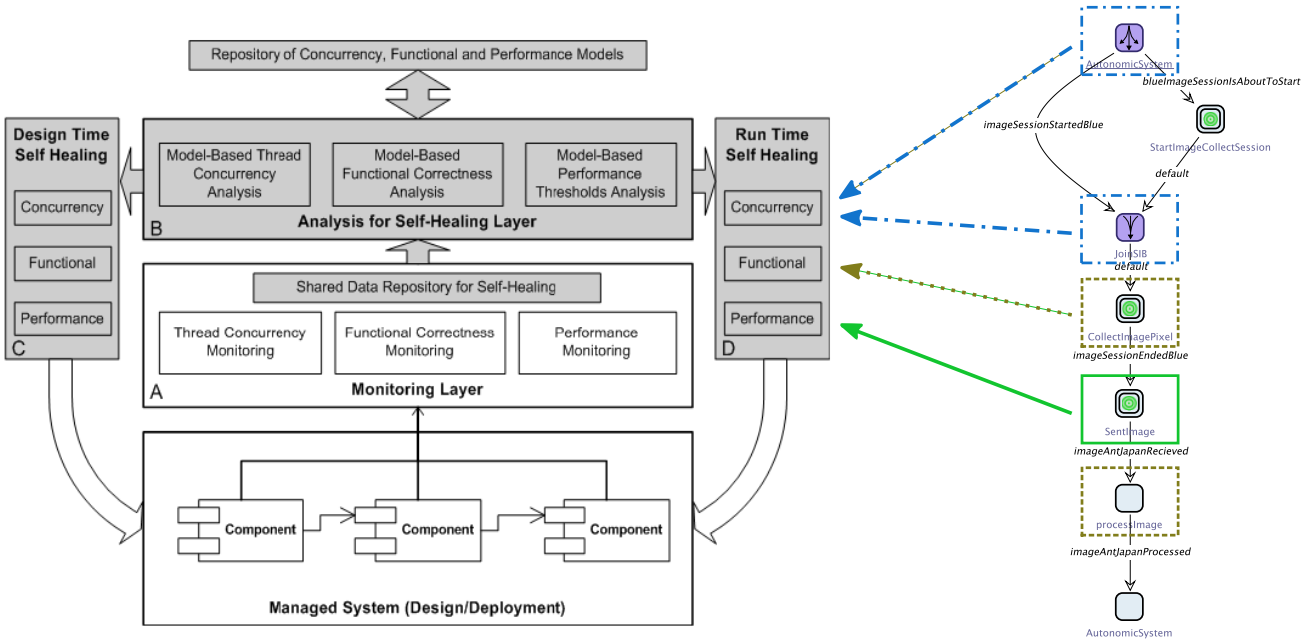

Fig. 8. Shadows architecture [21] (left) and self-healing annotated Voyager model (right)

Technology	Tool	Covered Issues	Scope
Model Checking	GEAR	Functional	Design-Time
Model Checking	Java Pathfinder	Functional	Design-Time
Regression Testing	BCT	Functional	Runtime
Testing	ConTest	Concurrency	Runtime
Monitoring	Panacea	Performance	Runtime
Monitoring	TPTP	Performance	Runtime

TABLE I
SHADOWS TECHNOLOGIES AND TOOLS TAILORED TO HANDLE ISSUES THAT COMMONLY OCCUR IN AUTONOMOUS SYSTEMS.

V. CONNECTING THE SHADOWS DESIGN-TIME AND RUN-TIME TECHNIQUES

In SHADOWS we have developed a rich set of self-healing techniques that span concurrency, performance, and functional issues. Table I summarizes the main technologies and related tools that were developed and used in SHADOWS to cope with issues that appear in designing autonomous systems like those from aerospace contexts. These technologies are suitable to handle several kinds of issues on different kinds of abstraction levels. A detailed description of the organization of the SHADOWS self-healing platform and of the underlying methodology is provided in [12].

As shown in Table I, these techniques, however, strike at runtime. They are currently driven or triggered by annotations in the applications, which are for the moment produced mostly manually and inserted in the application's source code. In order to develop the full benefits of the SHADOWS platform, it would be necessary to link the runtime power of the healers to annotations placed automatically in the design time artifacts, typically models.

In fact, the behavioural model extraction we just described is in our opinion the right bridge between

- the high-level specification provided in a domain specific language like ASSL,
- behavioural models that capture the functionality and that spot the possible critical locations in the behaviour, and
- the annotations that are necessary in order to include the issue monitoring mechanism.

As shown in Figure 8(left), the SHADOWS architecture is organized in this fashion. It foresees the use of models to link the design-time and the run-time aspects of the healing platform, and it provides (annotation-driven) issue monitoring as the mechanism that implements this link. As shown on the right on a fragment of the Voyager model derived from the ASSL specifications, the SLGs are adequate models for this abstraction because they present a suitable granularity: they are coarse enough to be still abstract, and fine enough to support spotting where to place which kind of annotation. We see that concurrency monitoring is well placed at fork/join locations, where multiple threading happens, functional monitoring (e.g for runtime extension and adaptation) is suitably introduced where complex functions are foreseen, and performance monitoring is useful at locations that may be subject to timeouts, as here the conclusion of receiving an image.

Once the annotations are properly inserted, the model-to-code generation provided by Genesys can transfer/link the annotations to the SHADOWS runtime. How these runtime techniques are organized and how they jointly work in practice is summarized in [12].

In Table II we show the mapping of the specialized tools used in SHADOWS with the ASSL tiers. The AS and AE tiers

Tool	ASSL (Sub-) Tier
GEAR	AS/AE Service Level Objectives AE Behavioral Models (high-level)
Java Pathfinder	AE Service Level Objectives AE Behavioral Models (code-level)
BCT	AE Behavioral Models (runtime in test and field) AE Metrics AE Actions, AE:Events
ConTest	AE Friends AE Behavioral Models (runtime in test)
Panacea	AE Self Management AE Metrics
TPTP	AS/AE Metrics AE Outcomes

TABLE II
TOOLS FOR AUTONOMOUS SELF-HEALING AND THEIR RELATED ASSL TIERS.

are linked in many ways and under several specific aspects to the SHADOWS platform.

VI. RELATED WORK

Nowadays, there is a growing consensus that model checking is most effective as an intelligent and early error-finding technique rather than a technique for guaranteeing correctness. This is partly due to the fact that specifications that can be checked automatically through model checking are necessarily partial in that they specify only certain aspects of the system behavior. Therefore, successful model checking runs, while reassuring, cannot guarantee full correctness. Rather, model checkers are increasingly conceived as elaborate debugging tools that complement traditional testing techniques.

Various model checkers are used to verify aerospace systems. Java Pathfinder [8] developed at NASA Ames is a prominent representative for verifying smaller systems. It assists developers at the Java code level, and therefore addresses a later phase than to our approach. We aim at assertions on interactions between components or of the system as a whole, with a focus on demanding properties.

For model checkers to be useful as debugging tools it is important that failing model checking attempts are accompanied by appropriate *error diagnosis* information that explains *why* the model check has failed. Model checkers may in fact also fail *spuriously*, i.e., although the property does not hold for the investigated abstraction it may still be valid for the real system. In order for model checking to be useful, it should therefore be easy for the user to rule out spurious failures and to locate the errors in the system based on the provided error diagnosis information. Therefore, it is important that error diagnosis information is easily accessible by the user.

Currently, ASSL provides a consistency checking mechanism to validate specifications of autonomic systems against correctness properties. Although proven to be efficient with handling consistency errors, this mechanism cannot handle logical errors. Another submission to this workshop [26] proposes a different model checking approach for ASSL, based on Labelled Transition systems and LTL properties.

For linear-time logics, error diagnosis information is conceptually of a simple type: It is given by a (possibly cyclic) execution path of the system that violates the given property. Thus, in case model-checking fails, linear-time model checkers like SPIN [9] compute an output in form of an *error trace* that represents a violating run, and is therefore valuable for the subsequent diagnosis and repair. The situation is more complex for properties that embody recovery issues. These claim for more demanding properties expressible in branching-time logics like CTL or the modal μ-calculus. Such logics do not just specify properties of single program executions but properties of the entire execution tree, comprising the local of decision points. Hence, meaningful error diagnosis information for branching-time logic model checking cannot be represented by linear executions in general. This is where games help.

ESA's FORMID Project (FOrmal Robotic Mission Inspection and Debugging) aimed at creating a development environment for the verification and analysis of robotic missions [4]. Unfortunately the system is solely concerned with predefined property patterns for safety, liveness, and conflict-freedom of the system. Therefore it is unable to handle more demanding properties as they typically arise during system modelling of complex systems that deal with self-healing and recovery issues.

VII. SUMMARY AND CONCLUSION

In this paper we have shown how to translate parts of an ASSL specification for autonomic systems into a behavioral model. This task implied to map the ASSL specific *self-management policy, action*, and *event* parts that made up the system to corresponding counterparts in a behavioral system model that is based on a Kripke-Transition-Structure.

We applied this translation step to the NASA Voyager II mission case study. This case study constitutes a picture transmission process that sends picture pixels taken by the Voyager spacecraft to four antennas on earth. These antennas in turn forward transmitted picture pixels to Voyager's mission base where the complete image will be reconstructed.

The translation step opened up several options for verifying issues related to e.g. recovery issues. After having detected the absence of a recovery mechanism upon transmission error within the system specification we may leverage the game-based verification Model Checker GEAR to fix this problem. A game-based exploration of the problem space as already suggested a tool supported enhancement of the model-driven verification process [15] can help in identifying those parts of the model that need adaptation to overcome this specific problem. However, we did not elaborate on this exploration here since the translation of the specification is still incomplete.

We have previous experience of automatic generation of control flow graphs from a language's Structured Operational Semantics [17] (SOS). In [5] we showed how to do it for a process algebra, later extended for object oriented languages. Accordingly, we plan to examine the SOS for ASSL provided

in [23] and possibly take it as a starting point for an SOS-driven generation of the SLGs. This way, the palette of model analyses developed in the jABC and the self-healing specific techniques developed in SHADOWS would become immediately applicable to all ASSL descriptions.

Acknowledgement

This work was carried out partly within the SHADOWS Project funded by the EU under contract No. 035157.

REFERENCES

[1] M. Bakera, T. Margaria, C. Renner, and B. Steffen. Property-driven functional healing: Playing against undesired behavior. In *10th CONQUEST*, 2007.

[2] M. Bakera, T. Margaria, C. Renner, and B. Steffen. Game-Based Model Checking for Reliable Autonomy in Space. *Journal of the American Institute of Aeronautics and Astronautics (AIAA)*, to appear.

[3] M. Bakera, C. Wagner, T. Margaria, E. Vassev, M. Hinchey, and B. Steffen. Component-oriented behavior extraction for autonomic system design.

[4] G. Bormann, L. Joudrier, and K. Kapellos. FORMID: A formal specification and verification Environment for DREAMS. In *Proc. 8th ESA ASTRA Workshop*, 2004.

[5] V. Braun, J. Knoop, and D. Koschützki. *cool: A control-flow generator for system analysis*. Technical Report MIP-Bericht Nr. 9801, Faculty of Mathematics and Informatics, University of Passau, Germany, 1998.

[6] M. W. Browne. Technical magic converts a puny signal into pictures. *NY Times*, 1989.

[7] E. M. Clarke and E. A. Emerson. Design and Synthesis of Synchronization Skeletons Using Branching Time Temporal Logic. In *Logics of Programs — Proc. 1981 (LNCS Volume 131)*, pages 52–71. Springer-Verlag: Heidelberg, Germany, 1981.

[8] K. Havelund and T. Pressburger. Model Checking JAVA programs using JAVA PathFinder. *STTT*, 2(4):366–381, 2000.

[9] G. J. Holzmann. *The SPIN Model Checker: Primer and Reference Manual*. Addison-Wesley, Boston, Massachusetts, USA, 2003.

[10] ITU. General recommendations on telephone switching and signaling - intelligent network: Introduction to intelligent network capability set 1, Recommendation Q.1211. Technical report, Standardization Sector of ITU, Geneva, March 1993.

[11] ITU-T. Recommendation Q.1203. "Intelligent Network - Global Functional Plane Architecture". Technical report, Standardization Sector of ITU, October 1992.

[12] G. Jung, T. Margaria, C. Wagner, and M. Bakera. Formalizing A Methodology for Design- and Runtime Self-healing. In *7th IEEE Int. Conf. and Worksh. on Engineering of Autonomic and Autonomous Systems (EASe10)*, page this volume. IEEE CS, March 2010.

[13] K. Kapellos. MUROCO-II: FOrmal Robotic Mission Inspection and Debugging. Technical report, European Space Agency, 2005.

[14] D. Kozen. Results on the propositional μ-calculus. In *ICALP*, volume 140 of *LNCS*, pages 348–359, Aarhus, Denmark, 12–16 July 1982. Springer-Verlag.

[15] M. Bakera, T. Margaria, C. Renner, B. Steffen. Tool supported enhancement of the model-driven verification process. *ISSE, Journal on Innovations in Systems and Software Engineering - a NASA Journal*, 3(5):211–228, Sept. 2009.

[16] M. Müller-Olm, D. A. Schmidt, and B. Steffen. Model-Checking: A Tutorial Introduction. In *SAS*, pages 330–354, 1999.

[17] G. D. Plotkin. A Structural Approach to Operational Semantics. Technical Report DAIMI FN-19, University of Aarhus, 1981.

[18] Ralf Nagel et al. jABC. http://www.jabc.de.

[19] SHADOWS. A self-healing approach to designing complex software systems. https://sysrun.haifa.ibm.com/shadows/.

[20] O. Shehory. SHADOWS: Self-healing complex software systems. In *23rd IEEE/ACM Int. Conf. on Automated Software Engineering - Workshops (ASE)*, pages 71–76. IEEE, 2008.

[21] O. Shehory, S. Ur, and T. Margaria. Self-healing technologies in SHADOWS: Targeting performance, concurrency and functional aspects. In *10th (CONQUEST)*, 2007.

[22] B. Steffen, T. Margaria, R. Nagel, S. Jörges, and C. Kubczak. Model-Driven development with the jABC. In *Hardware and Software, Verification and Testing*, pages 92–108, 2007.

[23] E. Vassev. *Towards a Framework for Specification and Code Generation of Autonomic Systems*. PhD thesis, Department of Computer Science and Software Engineering, Concordia University, Montreal, Canada, 2008.

[24] E. Vassev and M. Hinchey. Modeling the image-processing behavior of the nasa voyager mission with assl. In *Proceedings of the 3rd IEEE International Conference on Space Mission Challenges for Information Technology (SMC-IT'09)*.

[25] E. Vassev and M. Hinchey. ASSL Specification Model for the Image-processing Behavior in the NASA Voyager Mission. Technical report, Lero - The Irish Software Engineering Research Center, 2009.

[26] E. Vassev, M. Hinchey, and A. Quigley. Model checking for autonomic systems specified with ASSL. In *Proc. of the First NASA Formal Methods Symposium*.