Technische Universität Berlin

Forschungsberichte
der Fakultät IV – Elektrotechnik und Informatik

Formal Analysis and Verication of Self-Healing Systems: Long Version

Hartmut Ehrig[1], Claudia Ermel[1], Olga Runge[1],
Antonio Bucchiarone[2] and Patrizio Pellicione[3]

[1]Institut für Softwaretechnik und Theoretische Informatik
Technische Universität Berlin, Germany
fehrig, lieske, olgag@cs.tu-berlin.de

[2]FBK-IRST, Trento, Italy
bucchiarone@fbk.eu

[3]Dipartimento di Informatica Universita dell'Aquila, Italy
patrizio.pelliccione@di.univaq.it

Formal Analysis and Verification of Self-Healing Systems: Long Version[*] = Technical Report =

Hartmut Ehrig[1], Claudia Ermel[1], Olga Runge[1],
Antonio Bucchiarone[2] and Patrizio Pelliccione[3]

[1]Institut für Softwaretechnik und Theoretische Informatik
Technische Universität Berlin, Germany
{ehrig, lieske, olga}@cs.tu-berlin.de

[2]FBK-IRST, Trento, Italy
bucchiarone@fbk.eu

[3]Dipartimento di Informatica Università dell'Aquila, Italy
patrizio.pelliccione@di.univaq.it

March 17, 2010

Abstract

Self-healing (SH-)systems are characterized by an automatic discovery of system failures, and techniques how to recover from these situations. In this paper, we show how to model SH-systems using algebraic graph transformation. These systems are modeled as typed graph grammars enriched with graph constraints. This allows not only for formal modeling of consistency and operational properties, but also for their analysis and verification using the tool AGG. We present sufficient static conditions for self-healing properties, deadlock-freeness and liveness of SH-systems. The overall approach is applied to a traffic light system case study, where the corresponding properties are verified.

1 Introduction

The high degree of variability that characterizes modern systems requires to design them with runtime evolution in mind. Self-adaptive systems are a variant of fault-tolerant systems that autonomously decide how to adapt the system at runtime to the internal reconfiguration and optimization requirements or to environment changes and threats [7]. A classification of modeling dimensions

---

for self-adaptive systems can be found in [2], where the authors distinguish *goals* (what is the system supposed to do), *changes* (causes for adaptation), *mechanisms* (system reactions to changes) and *effects* (the impact of adaptation upon the system). The initial four self-* properties of self-adaptive systems are self-configuration, self-healing[1], self-optimization, and self-protection [22]. Self-configuration comprises components installation and configuration based on some high-level policies. Self-healing deals with automatic discovery of system failures, and with techniques to recover from them. Typically, the runtime behavior of the system is monitored to determine whether a change is needed. Self-optimization monitors the system status and adjusts parameters to increase performance when possible. Finally, self-protection aims to detect external threats and mitigate their effects [31].

Even with good reactions to both system and context changes a set of high-level goals "should be maintained regardless of the environment conditions" [11]. In other words, the joint ability to effectively react to changes without degrading the level of dependability is a key factor for delivering successful systems that continuously satisfy evolving user requirements. Consequently, a serious solution to effectively support the consistent evolution of systems at runtime can be conceived only through a well-defined formalization that provides a solid basis for verification.

In [9], Bucchiarone et al. modeled and verified dynamic software architectures and self-healing (SH-)systems (called self-repairing systems in [9]), by means of hypergraphs and graph grammars. Based on this work, we show in this paper how to formally model (SH-)systems by using algebraic graph transformations [13] and to prove consistency and operational properties. Graph transformation has been investigated as a fundamental concept for specification, concurrency, distribution, visual modeling, simulation and model transformation [13, 14].

The main idea is to model SH-systems by typed graph grammars with three different kinds of system rules, namely normal, environment, and repair rules. Normal rules define the normal and ideal behavior of the system. Environment rules model all possible predictable failures. Finally, for each failure a repair rule is defined. This formalization enables the specification, analysis and verification of consistency and operational properties of SH-systems. More precisely, we present sufficient conditions for two alternative self-healing properties, deadlock-freeness and liveness of SH-systems. The conditions can be checked statically for the given system rules in an automatic way using the AGG[2] modeling and verification tool for typed attributed graph transformation systems.

Summarizing, the contribution of this paper is twofold: (i) we propose a way to model and formalize SH-systems; (ii) we provide tool-supported static verification techniques for SH-system models. The theory is presented by use of a running example, namely an automated traffic light system controlled by means of electromagnetic spires that are buried some centimeters underneath the asphalt of car lanes.

The paper is organized as follows: Section 2 motivates the paper comparing it with related work. Section 3 presents the setting of our running example.

---

[1]following [28] we consider self-healing and self-repair as synonymous
[2]AGG (Attributed Graph Grammars): http://tfs.cs.tu-berlin.de/agg.

Section 4 introduces typed attributed graph transformation as formal basis to specify and analyze SH-systems. In Section 5 we define consistency and operational system properties. Static conditions for their verification are given in Section 6 and are used to analyze the behavior and healing properties of the traffic light system. We conclude the paper in Section 7 with a summary and an outlook on future work.

This technical report is the long version (including full proofs of the technical theorems and additional details of the traffic light example) accompanying our contribution at FASE 2010 [15].

2   Motivation and Related Work

Focusing on modeling approaches for SH-systems, the *Software Architecture* approach (SA) [25], has been introduced as a high-level view of the structural organization of systems. Since a self-healing system must be able to change at runtime, *Dynamic Software Architectures* (DSAs) have shown to be very useful to capture SA evolution [23, 17, 18]. Aiming at a formal analysis of DSAs, different approaches exist, either based on graph transformation [9, 6, 8, 5, 4, 19, 24] or on temporal logics and model checking [21, 1, 27]. In many cases, though, the state space of behavioral system models becomes too large or even infinite, and in this case model checking techniques have their limitations. Note that static analysis techniques, as applied in this paper, do not have this drawback. In addition to graph transformation techniques, also Petri nets [26] offer static analysis techniques to verify liveness and safety properties. But in contrast to Petri nets, graph transformation systems are well suited to model also reconfiguration of system architectures which is one possible way to realize system recovery from failures in self-healing (SH-)systems.

In the community of Service Oriented Computing, various approaches supporting self-healing have been defined, e.g. triggering repairing strategies as a consequence of a requirement violation [30], and optimizing QoS of service-based applications [10, 32]. Repairing strategies could be specified by means of policies to manage the dynamism of the execution environment [3, 12] or of the context of mobile service-based applications [29].

In [20], a theoretical assume-guarantee framework is presented to efficiently define under which conditions adaptation can be performed by still preserving the desired invariant. In contrast to our approach, the authors of [20] aim to deal with *unexpected* adaptations.

In contrast to the approaches mentioned above, we abstract from particular languages and notations. Instead, we aim for a coherent design approach allowing us to model important features of SH-systems at a level of abstraction suitable to apply static verification techniques.

3   Running Example: An Automated Traffic Light System

In an automated Traffic Light System (TLS), the technology is based upon electromagnetic spires that are buried some centimeters underneath the asphalt of car lanes. The spires register traffic data and send them to other system

3

components. The technology helps the infraction system by making it incontestable. In fact, the TLS is connected to cameras which record videos of the violations and automatically send them to the center of operations. In addition to the normal behavior, we may have failures caused by a loss of signals between traffic light or camera and supervisor. For each of the failures there are corresponding repair actions, which can be applied after monitoring the failures during runtime.

The TLS is responsible for regulating the traffic lights in a "smart" way. In particular, it provides the following functionalities:

Switching the traffic light signal: The spires send signals created by the passing vehicles to an Electronic Control Unit (ECU) which registers the medium speed of the cars, hence the intensity of traffic. The ECU component sends the data to a component called Supervisor. The Supervisor imparts a cadence to the green and the red lights of several traffic lights in the same street. The system gives a priority if more cars are in line in the same direction.

Management of infractions: Each camera is connected to the Supervisor, which constantly controls the traffic light signal and the ECU. If a vehicle transits upon the spires while the street light is red, the Supervisor triggers the camera, which starts recording. When the light is green, it ignores the spires and does not trigger the cameras. Contextually with the infraction event, the camera sends the records to the video recorder of the *CenterOfOperations*, which stores date, time and Supervisor ID in order to avoid legal challenges.

Error management: In regular time intervals, the system is monitored for errors or failures which are sent to the *CenterOfOperations*. Errors can be the breakdown of components as well as the loss of connection from one of the components of the TLS to the rest of the system. The *CenterOfOperations* triggers a repairing procedure to recover from system faults.

Figure 1 presents the use case diagram showing the functional requirements for the TLS. *Car* is an actor of the TLS while *Data send* is the use case that represents the gathering of the data related to the traffic flow. It activates the use case traffic light and, if there is an infraction, it also activates the *infraction management* use case. The *traffic light management* administrates the switching of the traffic lights, while the *infraction management* activates the cameras when there is an infraction going on in a cross road. Both use cases activate *error management* use case whenever there is an error in the system. Use case *error management* handles the case of an error in the TLS.

In this scenario the dynamism is given by the traffic flow. Cars join and leave the system continuously and there is no way to predict it, so we cannot predict the traffic light behavior. In the following we give the TLS requirements:

• *Traffic flow:* As the cars get to the cross road, the traffic flow is stored in the *Spire* component and sent to the *ECU* component. Then, the *ECU* component forwards the information to the *Supervisor* component which will manage it;

• *Light:* For each system configuration, each crossroad has at most one green light turned on.

• *Broken camera:* the Supervisor component checks in regular time intervals if there is an error signal linked to the *Camera* component, and the *CenterOfOperations* component takes care to repair it;

• *Broken traffic light:* if there is a loss of a traffic light signal, the system handles it in order to repair it;

• *No traffic flow:* when there is no traffic flow then every traffic light must be
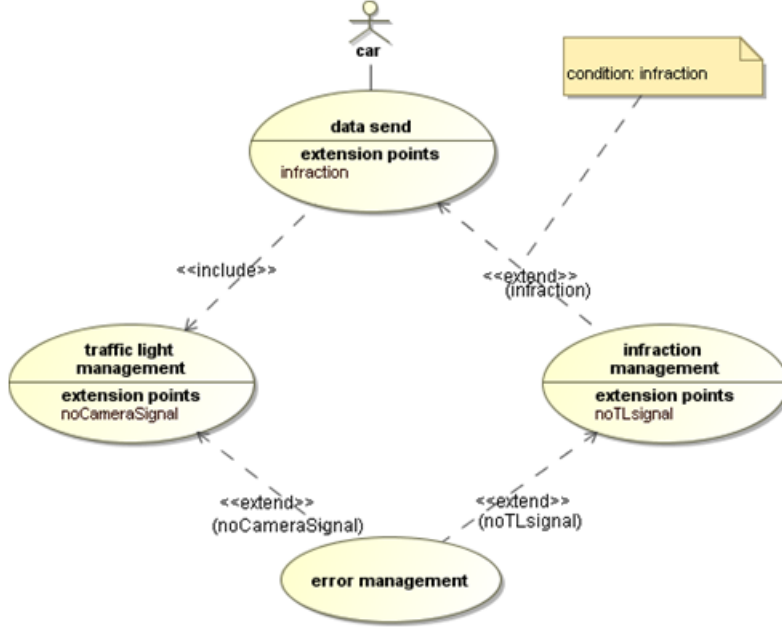
Figure 1: Use case diagram: Functional requirements for the TLS

turned to red;

• *No traffic blocking:* when there is traffic then it must be avoided that all traffic lights are red.

The aim of our TLS model is to ensure suitable self-healing properties by applying repair actions. What kind of repair actions are useful and lead to consistent system states without failures? What kind of safety and liveness properties can be guaranteed? We will tackle these questions in the next sections by providing a formal modeling and analysis technique based on algebraic graph transformation and continue our running example in Examples 1 – 6 below.

4 Formal Modeling of Self-Healing Systems by Algebraic Graph Transformation

In this section, we show how to model SH-systems in the formal framework of algebraic graph transformation [13]. The main concepts of this framework which are relevant for our approach are typed graphs, graph grammars, transformations and constraints. Configurations of an SH-System are modeled by typed graphs.

Definition 1 (Typed Graphs) *A graph $G = (N, E, s, t)$ consists of a set of nodes N, a set of edges E and functions $s, t : E \to N$ assigning to each edge $e \in E$ the source $s(e) \in N$ and target $t(e) \in N$.*

A graph morphism $f : G \to G'$ is given by a pair of functions $f = (f_N : N \to N', f_E : E \to E')$ which is compatible with source and target functions.

A type graph TG is a graph where nodes and edges are considered as node and edge types, respectively. A TG-typed, or short typed graph $\overline{G} = (G, t)$ con-

sists of a graph G and a graph morphism $t : G \to TG$, *called* typing morphism *of G. Morphisms $f : \overline{G} \to \overline{G'}$ of typed graphs are graph morphisms $f : G \to G'$ which are compatible with the typing morphisms of G and G', i.e. $t' \circ f = t$.*

For simplicity, we abbreviate $\overline{G} = (G, t)$ by G in the following. Moreover, the approach is also valid for *attributed* and *typed attributed graphs* where nodes and edges can have data type attributes [13], as used in our running example.

Example 1 (Traffic Light System) *The type graph TG of our traffic light system TLS is given in Fig. 2. The initial state is the configuration graph in Fig. 3 which is a TG-typed graph where the typing is indicated by corresponding names, and the attributes are attached to nodes and edges. The initial*


Figure 2: TLS type graph TG

state shows two traffic lights (TL), two cameras, a supervisor, and a center of operations, but no traffic up to now.

The dynamic behavior of SH-systems is modeled by rules and transformations of a typed graph grammar in the sense of algebraic graph transformation [13].
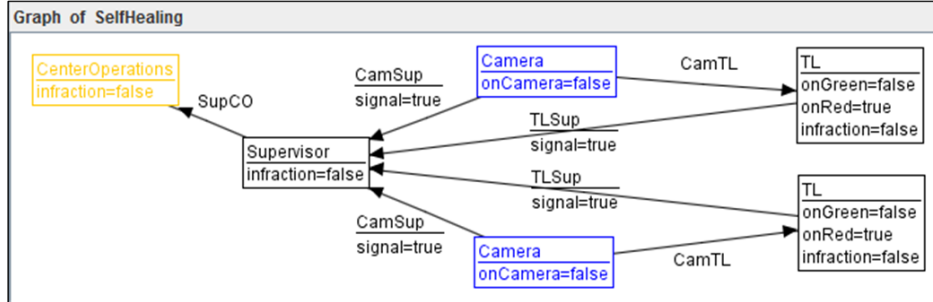

Figure 3: TLS initial state G_{init}

Definition 2 (Typed Graph Grammar) *A* typed graph grammar *$GG = (TG, G_{init}, Rules)$ consists of a type graph TG, a TG-typed graph G_{init}, called* initial graph*, and a set Rules of* graph transformation rules*. Each rule $r \in Rules$ is given by a span $(L \leftarrow I \to R)$, where L, I and R are TG-typed graphs, called* left-hand side, right-hand side *and* interface*, respectively. Moreover, $I \to L$, $I \to R$ are injective typed graph morphisms where in most cases I*

can be considered as intersection of L and R. A rule $r \in Rules$ is applied to a TG-typed graph G by a match morphism $m : L \to G$ leading to a direct transformation $G \overset{r,m}{\Longrightarrow} H$ via (r,m) in two steps: at first, we delete the match $m(L)$ without $m(I)$ from G to obtain a context graph D, and secondly, we glue together D with R along I leading to a TG-typed graph H.

More formally, the direct transformation $G \overset{r,m}{\Longrightarrow} H$ is given by two pushout diagrams (1) and (2) in the category $\mathbf{Graphs_{TG}}$ of TG-typed graphs, where diagram (1) (resp. (2)) corresponds to gluing G of L and D along I (resp. to gluing H of R and D along I).

$$
\begin{array}{ccccccc}
N & \overset{nac}{\longleftarrow} & L & \overset{l}{\longleftarrow} & I & \overset{r}{\longrightarrow} & R \\
& {\scriptstyle q}\searrow & {\scriptstyle m}\downarrow & (1) & \downarrow & (2) & \downarrow {\scriptstyle m^*} \\
& & G & \longleftarrow & D & \longrightarrow & H
\end{array}
$$

Note that pushout diagram (1) in step 1 only exists if the match m satisfies a gluing condition w.r.t. rule r which makes sure that the deletion in step 1 leads to a well-defined $TG-$typed graph D. Moreover, rules are allowed to have Negative Application Conditions (NACs) given by a typed graph morphism $nac : L \to N$. In this case, rule r can only be applied at match $m : L \to G$ if there is no injective morphism $q : N \to G$ with $q \circ nac = m$. This means intuitively that r cannot be applied to G if graph N occurs in G. A transformation $G_0 \overset{*}{\Longrightarrow} G_n$ via Rules in GG consists of $n \geq 0$ direct transformations $G_0 \Longrightarrow G_1 \Rightarrow ... \Rightarrow G_n$ via rules $r \in Rules$. For $n \geq 1$ we write $G_0 \overset{+}{\Longrightarrow} G_n$.

Example 2 (Rules of TLS) *A rule $r = (L \leftarrow I \to R)$ of TLS with NAC nac $: L \to N$ is given in Fig. 4 (interface I is not shown and consists of the nodes and edges which are present in both L and R, as indicated by equal numbers). For simplicity, we only show the part of the NAC graph N which extends L. All graph morphisms are inclusions. Rule r can be applied to graph G in Fig. 3 where the node (1:TL) in L is mapped by m to the upper node TL in G_{init}. This leads to a graph H where the attributes of TL are changed and the node Cars of R is attached to TL. Altogether, we have a direct transformation $G \overset{r,m}{\Longrightarrow} H$.*


Figure 4: TLS rule *ArrivalCarsOne*

In order to model consistency and failure constraints of an SH-system, we use graph constraints. A *TG-typed graph constraint* is given by a TG-typed graph C which is *satisfied by a TG-typed graph G*, written $G \models C$, if there is an injective graph morphism $f : C \to G$. Graph constraints can be negated or combined by logical connectors (e.g. $\neg C$). Now we are able to define SH-systems in the framework of algebraic graph transformation (AGT). An SH-system is given by a typed graph grammar where four kinds of rules are distinguished, called *system*, *normal*, *environment* and *repair* rules. Moreover, we have two kinds of TG-typed graph constraints, namely *consistency* and *failure* constraints.

Definition 3 (Self-healing System in AGT-Framework) *A Self-healing system (SH-system) is given by $SHS = (GG, C_{sys})$, where:*

- *$GG = (TG, G_{init}, R_{sys})$ is a typed graph grammar with type graph TG, a TG-typed graph G_{init}, called* initial state, *a set of TG-typed rules R_{sys} with NACs, called* system rules, *defined by $R_{sys} = R_{norm} \cup R_{env} \cup R_{rpr}$, where R_{norm} (called* normal rules*), R_{env} (called* environment rules*) and R_{rpr} (called* repair rules*) are pairwise disjoint.*

- *C_{sys} is a set of TG-typed graph constraints, called* system constraints, *with $C_{sys} = C_{consist} \cup C_{fail}$, where $C_{consist}$ are called* consistency constraints *and C_{fail}* failure constraints.

For an SH-system, we distinguish *reachable, consistent, failure* and *normal* states, where reachable states split into normal and failure states.

Definition 4 (Classification of SH-System States) *Given an SH-system SHS = (GG, C_{sys}) as defined above, we have*

1. *$Reach(SHS) = \{G \mid G_{init} \overset{*}{\Longrightarrow} G \text{ via } R_{sys}\}$, the* reachable states *consisting of all states reachable via system rules,*

2. *$Consist(SHS) = \{G \mid G \in Reach(SHS) \land \forall C \in C_{consist} : G \vDash C\}$, the* consistent states, *consisting of all reachable states satisfying the consistency constraints,*

3. *$Fail(SHS) = \{G \mid G \in Reach(SHS) \land \exists C \in C_{fail} : G \vDash C\}$, the* failure states, *consisting of all reachable states satisfying some failure constraint,*

4. *$Norm(SHS) = \{G \mid G \in Reach(SHS) \land \forall C \in C_{fail} : G \nvDash C\}$, the* normal states, *consisting of all reachable states not satisfying any failure constraint.*

Example 3 (Traffic Light System as SH-system) *We define the Traffic Light SH-system $TLS = (GG, C_{sys})$ by the type graph TG in Fig. 2, the initial state G_{init} in Fig. 3, and the following sets of rules and constraints:*

- *$R_{norm} = \{ArrivalCarsOne, ArrivalCarsTwo, RemoveCarsOne, RemoveCarsTwo, InfractionOn, InfractionOff\}$,*

- *$R_{env} = \{FailureTL, FailureCam\}$,*

- *$R_{rpr} = \{RepairTL, RepairCam\}$,*

- *$C_{consist} = \{\neg allGreen, \neg allRed, TLCamSupInfractionTrue \lor TLCamSupInfractionFalse\}$,*

- *$C_{fail} = \{TLSupFailure, CamSupFailure\}$.*

*The normal rule ArrivalCarsOne is depicted in Fig. 4 and models that one or more cars arrive at a traffic light (*1:TL*) while all of the crossing's lights are red. The NAC in Fig. 4 means that in this situation, no cars arrive at the other direction's traffic light (*3:TL*). Applying this rule, the traffic light in the direction of the arriving cars (*1:TL*) switches to green. Rule ArrivalCarsTwo (see Fig. 5) models the arrival of one or more cars at a red traffic light (*2:TL*)*

*where no cars have been before, while at the same time the traffic light for the other direction (*3:TL*) shows green and there are already cars going in this direction. This rule causes a change of the traffic light colors in both directions. Rules RemoveCarsOne and RemoveCarsTwo are the inverse rules (with L and*


Figure 5: TLS rule *ArrivalCarsTwo*

R exchanged) of the arrival rules in Fig. 4 and 5, and model the reduction of traffic at a traffic light. Rule InfractionOn is shown in Fig. 6 and models the situation that a car is passing the crossroad at a red light: the signal infraction *of both the supervisor and the center of operations is set to true and the corresponding camera is starting to operate. The rule ensures that the corresponding camera is connected, using the edge attribute* signal = true *for edge* 13:CamSup*. Rule InfractionOff (not depicted) models the inverse action, i.e. the infraction attribute is set back to false, and the camera stops running.*
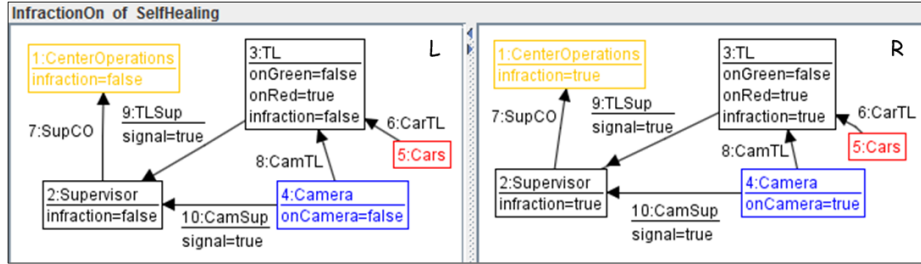

Figure 6: Normal rule *InfractionOn* of *TLS*

The environment rules are shown in Fig. 7. They model the signal disconnection of a traffic light and a camera, respectively. The repair rules (not depicted) are defined as inverse rules of the environment rules and set the signal attributes back to true*.*


Figure 7: Environment rules *FailureTL* and *FailureCam* of *TLS*

The consistency constraints model the desired properties that we always want to have crossroads with at least one direction showing red lights (¬ allGreen)

9

and avoiding all traffic lights red when there is traffic (¬ allRed). Moreover, in infraction mode we want the attributes infraction of Supervisor and CenterOfOperations and the attribute onCamera to be true all at once, and in normal mode (no infraction), we want them all to be false. The constraint $TLCamSupInfractionTrue \lor TLCamSupInfractionFalse$ models that our system is either in infraction mode or in normal mode. The corresponding constraint graphs (without negation) are shown in Fig. 8. The failure constraints TL-SupFailure *and* CamSupFailure *express that either a traffic light or a camera is disconnected (the constraint graphs correspond to the right-hand sides of the environment rules in Fig. 7).*


Figure 8: Consistency constraint graphs of *TLS*

5 Consistency and Operational Properties of SH-Systems

In this section, we define desirable consistency and operational properties of SH-Systems. We distinguish system consistency, where all reachable states are consistent, and normal state consistency, where the initial state G_{init} and all states reachable by normal rules are normal states. Environment rules, however, may lead to failure states, which should be repaired by repair rules. We start with consistency properties:

Definition 5 (Consistency Properties) *An SH-System SHS is called*

1. *system consistent, if all reachable states are consistent, i.e.*
   $Reach(SHS) = Consist(SHS);$

2. *normal state consistent, if the initial state is normal and all normal rules preserve normal states, i.e.*
   $G_{init} \in Norm(SHS)$ *and* $\forall G_0 \overset{p}{\Longrightarrow} G_1$ *via* $p \in R_{norm}$
   $[\ G_0 \in Norm(SHS) \Rightarrow G_1 \in Norm(SHS)\ ]$

10

Example 4 (Consistency Properties of TLS) *The SH-System TLS is system consistent, because for all $C \in C_{consist}$ $G_{init} \models C$ and for all $G_0 \overset{p}{\Longrightarrow} G_1$ via $p \in R_{sys}$ and $G_0 \in Consist(SHS)$ we also have $G_1 \in Consist(SHS)$. Similarly, TLS is normal state consistent, because $G_{init} \in Norm(SHS)$ and for all $G_0 \overset{p}{\Longrightarrow} G_1$ via $p \in R_{norm}$ for all $C \in C_{fail}$ [$G_0 \not\models C \Rightarrow G_1 \not\models C$]. In both cases this can be concluded by inspection of the corresponding rules, constraints and reachable states. Moreover, there are also general conditions, which ensure the preservation of graph constraints by rules, but this discussion is out of scope for this paper.*

Now we consider the operational properties: one of the main ideas of SH-Systems is that they are monitored in regular time intervals by checking, whether the current system state is a failure state. In this case one or more failures have occurred in the last time interval, which are caused by failure rules, provided that we have normal state consistency. With our self-healing property below we require that each failure state can be repaired leading again to a normal state. Moreover, strongly self-healing means that the normal state after repairing is the same as if no failure and repairing would have been occurred.

Definition 6 (Self-healing Properties) *An SH-System SHS is called*

1. *self-healing, if each failure state can be repaired, i.e.*
   $\forall G_{init} \Rightarrow^* G$ *via* $(R_{norm} \cup R_{env})$ *with* $G \in Fail(SHS)$
   $\exists \ G \Rightarrow^+ G'$ *via* R_{rpr} *with* $G' \in Norm(SHS)$

2. *strongly self-healing, if each failure state can be repaired strongly, i.e.*
   $\forall G_{init} \Rightarrow^* G$ *via* $(p_1 \ldots p_n) \in (R_{norm} \cup R_{env})^*$ *with* $G \in Fail(SHS)$
   $\exists \ G \Rightarrow^+ G'$ *via* R_{rpr} *with* $G' \in Norm(SHS)$ *and*
   $\exists \ G_{init} \Rightarrow^* G'$ *via* $(q_1 \ldots q_m) \in R_{norm}^*$,
   where $(q_1 \ldots q_m)$ *is subsequence of all normal rules in* $(p_1 \ldots p_n)$.

Remark 1 *By definition, each strongly self-healing SHS is also self-healing, but not vice versa. The additional requirement for strongly self-healing means, that the system state G' obtained after repairing is not only normal, but can also be generated by all normal rules in the given mixed sequence $(p_1 \ldots p_n)$ of normal and environment rules, as if no environment rule would have been applied. We will see that our SH-System TLS is strongly self-healing, but a modification of TLS, which counts failures, even if they are repaired later, would only be self-healing, but not strongly self-healing.*

Another important property of SH-Systems is deadlock-freeness, meaning that no reachable state is a deadlock. A stronger liveness property is strong cyclicity, meaning that each pair of reachable states can be reached from each other. Note that this is stronger than cyclicity meaning that there are cycles in the reachability graph. Strong cyclicity, however, implies that each reachable state can be reached arbitrarily often. This is true for the TLS system, but may be false for other reasonable SH-Systems, which may be only deadlock-free. Moreover, we consider "normal deadlock-freeness" and "normal strong cyclicity", where we only consider normal behavior defined by normal rules.

Definition 7 (Deadlock-Freeness and Strong Cyclicity Properties) *An SH-System SHS is called*

1. *deadlock-free, if no reachable state is a deadlock, i.e.*
   $\forall G_0 \in Reach(SHS) \ \exists\, G_0 \overset{p}{\Longrightarrow} G_1 \ via \ p \in R_{sys}$

2. *normal deadlock-free, if no state reachable via normal rules is a (normal) deadlock, i.e.* $\forall G_{init} \Rightarrow^* G_0 \ via \ R_{norm} \ \exists\, G_0 \overset{p}{\Longrightarrow} G_1 \ via \ p \in R_{norm}$

3. *strongly cyclic, if each pair of reachable states can be reached from each other, i.e.* $\forall G_0, G_1 \in Reach(SHS) \ \exists\, G_0 \Rightarrow^* G_1 \ via \ R_{sys}$

4. *normally cyclic, if each pair of states reachable by normal rules can be reached from each other by normal rules, i.e.*
   $\forall G_{init} \Rightarrow^* G_0 \ via \ R_{norm} \ and \ G_{init} \Rightarrow^* G_1 \ via \ R_{norm} \ we \ have \ \exists\, G_0 \Rightarrow^* G_1$
   via R_{norm}

Remark 2 *If we have at least two different reachable states (rsp. reachable by normal rules), then "strongly cyclic" (rsp. "normally cyclic") implies "deadlock-free" (rsp. "normal deadlock-free"). In general properties 1 and 2 as well as 3 and 4 are independent from each other. But in Thm. 3 we will give sufficient conditions s.t. "normal deadlock-free" implies "deadlock-free" (rsp. "normally cyclic" implies "strongly cyclic" in Thm. 4).*

6 Analysis and Verification of Operational Properties

In this section, we analyze the operational properties introduced in section 5 and give static sufficient conditions for their verification.

First, we define direct and normal healing properties, which imply the strong self-healing property under suitable conditions in Thm. 1. In a second step we give static conditions for the direct and normal healing properties in Thm. 2, which by Thm. 1 are also sufficient conditions for our self-healing properties. Of course, we have to require that for each environment rule, which may cause a failure there are one or more repair rules leading again to a state without this failure, if they are applied immediately after its occurrence. But in general, we cannot apply the repair rules directly after the failure, because other normal and environment rules may have been applied already, before the failure is monitored. For this reason we require in Thm. 1 that each pair (p, q) of environment rules p and normal rules q is sequentially independent. By the Local Church-Rosser theorem for algebraic graph transformation [13](Thm 5.12) sequential independence of (p, q) allows one to switch the corresponding direct derivations in order to prove Thm. 1. For the case with nested application conditions including NACs we refer to [16]. Moreover, the AGG tool can calculate all pairs of sequential independent rules with NACs before runtime.

Definition 8 (Direct and Normal Healing Properties) *An SH-System SHS has the*

1. *direct healing property, if the effect of each environment rule can be repaired directly, i.e.* $\forall G_0 \overset{p}{\Longrightarrow} G_1 \ via \ p \in R_{env} \ \exists\, G_1 \overset{p'}{\Longrightarrow} G_0 \ via \ p' \in R_{rpr}$

2. *normal healing property, if the effect of each environment rule can be re-paired up to normal transformations, i.e.* $\forall G_0 \overset{p}{\Longrightarrow} G_1$ *via* $p \in R_{env}$ $\exists$ $G_1 \Rightarrow^+ G_2$ *via* R_{rpr} *s.t.* $\exists G_0 \Rightarrow^* G_2$ *via* R_{norm}

Remark 3 *The direct healing property allows one to repair each failure caused by an environment rule directly by reestablishing the old state G_0. This is not required for the normal healing property, but it is required only that the repaired state G_2 is related to the old state G_0 by a normal transformation. Of course, the direct healing property implies the normal one using $G_2 = G_0$.*

Theorem 1 (Analysis of Self-healing Properties) *An SH-System SHS is*

   I. *strongly self-healing, if we have properties 1, 2, and 3 below*

  II. *self-healing, if we have properties 1, 2 and 4 below*

   1. *SHS is normal state consistent*

   2. *each pair* $(p, q) \in R_{env} \times R_{norm}$ *is sequentially independent*

   3. *SHS has the direct healing property*

   4. *SHS has the normal healing property*

Proof I. Given $G_{init} \Rightarrow^* G$ via $(p_1, \ldots p_n) \in (R_{norm} \cup R_{env})^*$ with $G \in Fail(SHS)$ we have $n \geq 1$, because $G_{init} \in Norm(SHS)$ by 1. By sequential independence in 2. we can switch the order of $(p_1, \ldots p_n)$, s.t. first all normal rules $p_i \in R_{norm}$ and then all environment rules $p_i \in R_{env}$ are applied. As example let us consider $G_{init} \Rightarrow^+ G$ via $(q_1, p_1, q_2, p_2, q_3)$ with $q_i \in R_{norm}$ and $p_i \in R_{env}$. Then sequential independence leads by the Local Church-Rosser theorem to equivalent sequences in subdiagram (1), (2), (3) respectively.

By the direct healing property of 3, we have $p_1', p_2' \in R_{rpr}$ with $G_4' \overset{p_1'}{\Longrightarrow} G_3'$ and $G \overset{p_2'}{\Longrightarrow} G_4'$. With $G' = G_3'$ we have $G \Rightarrow^+ G'$ via $(p_2', p_1') \in R_{rpr}^*$ and $G_{init} \Rightarrow^* G'$ via $(q_1, q_2, q_3) \in R_{norm}^*$ where (q_1, q_2, q_3) is the subsequence of $(q_1, p_1, q_2, p_2, q_3)$ of all normal rules and normal state consistency of 1. implies $G_{init}, G' \in Norm(SHS)$. Note that in the repair sequence (p_2', p_1') the (possible) failures caused by $p_1, p_2 \in R_{env}$ are repaired in opposite order. In general the sequence $(p_1, \ldots p_n)$ $(n \geq 1)$ contains at least one rule in R_{env}, because otherwise $G \notin Fail(SHS)$ by 1. This implies that we have a repair sequence $G \Rightarrow^+ G'$ via R_{rpr} of length $n \geq 1$. Hence SHS is strongly self-healing.

II. We can proceed as above up to the point that first all normal and then all environment rules are applied. As shown in our example the normal healing

property 4 leads first to a repair transformation $G \Rightarrow^+ G_5$ via R_{rpr} with $G'_4 \Rightarrow^*$ G_5 via R_{norm}, then we can switch rules in (4) according to 2. Finally the normal healing property leads to $G_5 \Rightarrow^+ G_6$ via R_{rpr} with $G'_5 \Rightarrow^* G_6$ via R_{norm}.

$$G_{init} \xoverset{q_1}{\Rightarrow} G_1 \xoverset{p_1}{\Rightarrow} G_2 \xoverset{q_2}{\Rightarrow} G_3 \xoverset{p_2}{\Rightarrow} G_4 \xoverset{q_3}{\Rightarrow} G \xoverset{R_{rpr}}{\Longrightarrow^+} G_5 \xoverset{R_{rpr}}{\Longrightarrow^+} G_6 = G'$$

Altogether, we obtain for $G' = G_6$ a repair transformation $G \Rightarrow^+ G_5 \Rightarrow^+$ $G_6 = G'$ and a normal rule transformation $G_{init} \Rightarrow^* G'_3 \Rightarrow^* G'_5 \Rightarrow^* G_6 = G'$, which implies $G' \in Norm(SHS)$ by 1. In general $G \in Fail(SHS)$ implies that we have at least one environment rule in the given sequence and hence a repair transformation $G \Rightarrow^+ G'$ of length $n \geq 1$. Hence SHS is self-healing. $\square$

In the following Thm. 2 we give static conditions for direct and normal healing properties. In part 1 of Thm. 2 we require that for each environment rule p the inverse rule p^{-1} is isomorphic to a repair rule p'. Two rules are isomorphic if they are componentwise isomorphic. For $p = (L \leftarrow I \rightarrow R)$ with negative application condition $nac : L \rightarrow N$ it is possible (see [13] Remark 7.21) to construct $p^{-1} = (R \leftarrow I \rightarrow L)$ with equivalent $nac' : R \rightarrow N'$. In part 2 of Thm. 2 we require as weaker condition that each environment rule p has a corresponding repair rule p', which is not necessarily inverse to p. It is sufficient to require that we can construct a concurrent rule $p *_R p'$ which is isomorphic to a normal rule p''. For the construction and corresponding properties of inverse and concurrent rules, which are needed in the proof of Thm. 2 we refer to [13].

Theorem 2 (Static Conditions for Direct/Normal Healing Properties)   *1.*
    An SH-System SHS has the direct healing property, if for each environment rule there is an inverse repair rule, i.e. $\forall p \in R_{env} \exists p' \in R_{rpr}$ with $p' \cong p^{-1}$

    2. An SH-System SHS has the normal healing property if for each environment rule there is a corresponding repair rule in the following sense: $\forall p = (L \leftarrow K \rightarrow R) \in R_{env}$ we have

        a) repair rule $p' = (L' \leftarrow^{l'} K' \rightarrow^{r'} R')$ with l' bijective on nodes, and

        b) an edge-injective morphism $e : L' \rightarrow R$ leading to concurrent rule $p *_R p'$, and

        c) normal rule $p'' \in R_{norm}$ with $p *_R p' \cong p''$

Proof 1. Given $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ with NACs $nac_{i,L} : L \rightarrow N_{i,L}(i \in I)$ the inverse rule is given by $p^{-1} = (R \xleftarrow{r} K \xrightarrow{l} L)$ with corresponding NACs $nac_{i,R} : R \rightarrow N_{i,R}$. Now given $p \in R_{env}$ with $G_0 \xoverset{p}{\Longrightarrow} G_1$, we have by assumption $p' \in R_{rpr}$ with $p' \cong p^{-1}$ and by construction of p^{-1} also $G_1 \xoverset{p^{-1}}{\Longrightarrow} G_0$ and hence also $G_1 \xoverset{p'}{\Longrightarrow} G_0$.

2. The concurrent rule $p^* = p *_R p'$ is given by

$$
\begin{array}{ccccccccccc}
L & \xleftarrow{l} & K & \xrightarrow{r} & R & & & L' & \xleftarrow{l'} & K' & \xrightarrow{r'} & R' \\
{\scriptstyle id}\downarrow & (5) & \downarrow{\scriptstyle id} & (6) & \searrow{\scriptstyle id} & & {\scriptstyle e}\swarrow & & (7) \downarrow & & (8) \downarrow \\
L^* = L & \leftarrow & K & \longrightarrow & R & \longleftarrow & & K'' & \rightarrow & R^* \\
& & & \searrow & & (9) & & \swarrow & & \\
& & & & & K^* & & & &
\end{array}
$$

where (5), (6) are trivial pushouts, pushout-complement K'' exists in (7) because e is edge-injective and l' bijective on nodes (see prop. a)) s.t. the gluing condition is satisfied. Moreover, (8) is constructed as PO and (9) as pullback (PB).

Given $G_0 \overset{p,m}{\Longrightarrow} G_1$ with m injective (see Remark 4) and $p \in R_{env}$ we have

$$
\begin{array}{ccccccccccc}
L & \xleftarrow{l} & K & \xrightarrow{r} & R & \xleftarrow{e} & L' & \xleftarrow{l'} & K' & \xrightarrow{r'} & R' \\
{\scriptstyle m}\downarrow & (1) & \downarrow & (2) & \searrow{\scriptstyle n} & {\scriptstyle m'}\swarrow & (3) & \downarrow & (4) & \downarrow \\
G_0 & \leftarrow & D_0 & \longrightarrow & G_1 & \longleftarrow & & D_1 & \rightarrow & G_2
\end{array}
$$

where POs (1), (2) are given, leading to injective comatch n, m' defined by $m' = n \circ e$ is edge-injective by property b), PO-complement D_1 in (3) exists, because the gluing condition is satisfied, and (4) is constructed as PO. Hence, we have an R-related transformation sequence $G_0 \overset{p,m}{\Longrightarrow} G_1 \overset{p',m'}{\Longrightarrow} G_2$. Using the Concurrency Theorem (see [13] Thm 5.23) we have a corresponding transformation $G_0 \overset{p^*,m}{\Longrightarrow} G_2$ via the concurrent rule p^*. According to property c) we have $p'' \in R_{norm}$ with $p^* \cong p''$ and hence also $G_0 \overset{p''}{\Longrightarrow} G_2$ via R_{norm}. Finally, $G_1 \overset{p'}{\Longrightarrow} G_2$ via $p' \in R_{rpr}$ is the required transformation $G_1 \Rightarrow^+ G_2$ via R_{rpr}. The sequence $G_0 \overset{p,m}{\Longrightarrow} G_1 \overset{p',m'}{\Longrightarrow} G_2$ is R-related because we have the following diagram where PO (3) splits into

$$
\begin{array}{ccccccccccc}
L & \xleftarrow{l} & K & \xrightarrow{r} & R & & L' & \xleftarrow{l'} & K' & \xrightarrow{r'} & R' \\
{\scriptstyle id}\downarrow & (5) & \downarrow{\scriptstyle id} & (6) & \searrow{\scriptstyle id} & {\scriptstyle e}\swarrow & (7) \downarrow & & (8) \downarrow \\
L^* = L & \leftarrow & K & \longrightarrow & R & \longleftarrow & K'' & \longrightarrow & R^* \\
{\scriptstyle m}\downarrow & (1) & \downarrow & (2) & \downarrow{\scriptstyle n} & (9) & \downarrow & (10) & \downarrow \\
G_0 & \longleftarrow & D_0 & \longrightarrow & G_1 & \longleftarrow & D_1 & \dashrightarrow & G_2
\end{array}
$$

POs (7) and (9) using the PO-PB decomposition property (see [13]) with n injective. Finally (4) splits into POs (8) and (10). $\qquad\square$

Remark 4 *In part 2 of the theorem we require that $p \in R_{env}$ is applied with injective match m. If all rules in R_{sys} are required to be applied with injective matches, then we have to require that e in* Condition 2 b) *is injective. If p and p' have NACs we have to require that $m \models NAC(p)$ implies $m' \models NAC(p')$, where $m' = n \circ e$ and n is the comatch in $G_0 \overset{p,m}{\Longrightarrow} G_1$ as shown in the proof.*

Remark 5 *By combining Thm. 1 and Thm. 2 we obtain static conditions ensuring that an SH-System SHS is strongly self-healing and self-healing, respectively.*

Example 5 (Direct Healing Property of TLS) *TLS has direct healing property because "RepairTL" rsp. "RepairCam" are inverse to "FailureTL" resp. "FailureCam" and each pair $(p, q) \in R_{env} \times R_{norm}$ is sequentially independent according to the dependency matrix of TLS in Fig. 9.*

|	R_{norm}						R_{env}		R_{rpr}	
first \ second	1: Ar...	2: Ar...	3: Re...	4: Re...	5: Inf...	6: Inf...	7: Fai...	8: Fai...	9: Re...	10: R...
1: ArrivalCarsOne	0	1	1	0	0	0	0	0	0	0
2: ArrivalCarsTwo	0	0	0	1	1	0	0	0	0	0
3: RemoveCarsOne	2	0	0	0	0	0	0	0	0	0
4: RemoveCarsTwo	0	1	1	0	0	0	0	0	0	0
5: InfractionOn	0	0	0	0	0	2	0	0	0	0
6: InfractionOff	0	0	0	1	2	0	0	0	0	0
7: FailureTL	0	0	0	0	0	0	0	0	1	0
8: FailureCam	0	0	0	0	0	0	0	0	0	1
9: RepairTL	2	2	2	2	1	1	1	0	0	0
10: RepairCam	0	0	0	0	1	1	0	1	0	0

Figure 9: Dependency Matrix of TLS in AGG

In the following Thm. 3 and Thm. 4 we give sufficient conditions for deadlock-freeness and strong cyclicity which are important liveness properties. Here we mainly use a stepwise approach. We assume to have both properties for normal rules and give additional static conditions to conclude the property for all system rules. The additional conditions are sequentially and parallel independence and a direct correspondence between environment and repair rules, which should be inverse to each other. Similar to sequential independence, also parallel independence of rules (p, q) can be calculated by the AGG tool before runtime.

Theorem 3 (Deadlock-Freeness) *An SH-System SHS is deadlock-free, if*

1. *SHS is normally deadlock-free, and*

2. *Each pair $(p, q) \in (R_{env} \cup R_{rpr}) \times R_{norm}$ is sequentially and parallel independent.*

Proof Given $G_{init} \Rightarrow^* G_0$ via R_{sys} we have to show the existence of $G_0 \overset{p}{\Longrightarrow} G_1$ via some $p \in R_{sys}$. Similar to proof of Thm. 1, sequential independence allows one to construct an equivalent transformation sequence $G_{init} \Rightarrow^* G_0' \Rightarrow^* G_1$ via R_{norm} in the first and via $(R_{env} \cup R_{rpr})$ in the second part. Now normal deadlock-freeness

16

implies the existence of $G'_0 \overset{p}{\Longrightarrow} G'_1$ via some $p \in R_{norm}$. Now parallel independence allows one to construct by the local Church-Rosser theorem $G_0 \overset{p}{\Longrightarrow} G_1$ and $G'_1 \Rightarrow^* G_0$ via $(R_{env} \cup R_{rpr})$ for some G_1. $G_0 \overset{p}{\Longrightarrow} G_1$ with $p \in R_{norm} \le R_{sys}$ is the required transformation. □

Theorem 4 (Strong Cyclicity) *An SH-System SHS is strongly cyclic, given*
   I. properties 1 and 2, or
   II. properties 1, 3 and 4 below.

1. *For each environment rule there is an inverse repair rule and vice versa.*

2. *For each normal rule there is an inverse normal rule.*

3. *SHS is normally cyclic.*

4. *Each pair $(p, q) \in (R_{env} \cup R_{rpr}) \times R_{norm}$ is sequentially independent.*

Proof I. Given conditions 1 and 2 it is sufficient to show that for each $G_{init} \Rightarrow^*$ G_n via R_{sys} there is also a reverse sequence $G_n \Rightarrow^* G_{init}$ via R_{sys}. In fact, for $G_{init} \Rightarrow^* G_n$ via $(p_1, \ldots, p_n) \in R^*_{sys}$ we have $G_n \Rightarrow^* G_{init}$ via $(p_n^{-1}, \ldots, p_1^{-1}) \in R^*_{sys}$.

II. Given conditions 1, 3 and 4, and $G_{init} \Rightarrow^* G_1$ via R_{sys} and $G_{init} \Rightarrow^* G_2$ via R_{sys}. By sequential independence (cond. 4) we can split the sequences as shown in (1) and (2). For each direct transformation $G_3 \overset{p}{\Longrightarrow} G_4$ in $G'_1 \Rightarrow^* G_1$ via $(R_{env} \cup R_{rpr})$ we have $p' \in (R_{rpr} \cup R_{env})$ with $p \cong p'^{-1}$ leading to a direct transformation $G_4 \overset{p'}{\Longrightarrow} G_3$ with $p' \in (R_{rpr} \cup R_{env})$. Hence we obtain a transformation sequence $G_1 \Rightarrow^* G'_1$ via $(R_{rpr} \cup R_{env})$. Since SHS is normally cyclic (cond. 3), we obtain $G'_1 \Rightarrow^* G'_2$ via R_{norm} and altogether $(G_1 \Rightarrow^* G_2) = (G_1 \Rightarrow^* G'_1 \Rightarrow^* G'_2 \Rightarrow^* G_2)$ via R_{sys}.

This shows that SHS is strongly cyclic. □

Remark 6 *In part I of Thm. 4, we avoid the stepwise approach and any kind of sequential and parallel independence by the assumption that also all normal rules have inverses, which is satisfied for our TLS.*

Example 6 (Strong Cyclicity and Deadlock-Freeness of TLS) *We use part I of Thm. 4 to show strong cyclicity. Property 1 is satisfied because "FailureTL"*

*and "RepairTL" as well as "FailureCam" and "RepairCam" are inverse to each
other. Property 2 is satisfied because "ArrivalCarsOne(Two)" and "Remove-
CarsOne(Two)" as well as "InfractionOn" and "InfractionOff" are inverse to
each other. Moreover, deadlock-freeness of TLS follows from strong cyclicity by
remark 2. Note that we cannot use part II of Thm. 4 for our example TLS,
because e.g. ("RepairTL", "ArrivalCarsOne") is not sequentially independent.*

7 Conclusion

In this paper, we have modeled and analyzed self-healing systems using algebraic
graph transformation and graph constraints. We have distinguished between
consistency properties, including system consistency and normal state consis-
tency, and operational properties, including self-healing, strongly self-healing,
deadlock-freeness, and strong cyclicity. The main results concerning operational
properties are summarized in Fig. 10, where most of the static conditions in
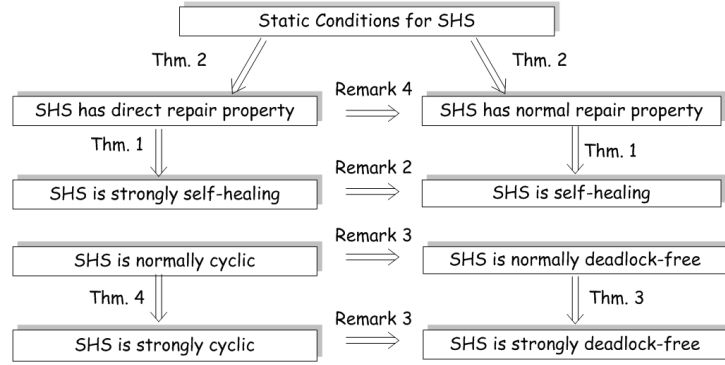Thms. 1- 4 can be automatically checked by the AGG tool.

Figure 10: Operational properties of self-healing systems

All properties are verified for our traffic light system. Note that in this
paper, the consistency properties are checked by inspection of corresponding
rules, while the operational properties are verified using our main results. Work
is in progress to evaluate the usability of our approach by applying it to larger
case studies. As future work, we will provide analysis and verification of con-
sistency properties using the theory of graph constraints and nested application
conditions in [16]. Moreover, we will investigate how far the techniques in this
paper for SH-systems can be used and extended for more general self-adaptive
systems.

References

[1] N. Aguirre and T. S. E. Maibaum. Hierarchical temporal specifications of
dynamically reconfigurable component based systems. *ENTCS*, 108:69–81,
2004.

[2] J. Andersson, R. Lemos, S. Malek, and D. Weyns. Modeling dimensions of self-adaptive software systems. In *Software Engineering for Self-Adaptive Systems*, pages 27–47. Springer, 2009.

[3] L. Baresi, S. Guinea, and L. Pasquale. Self-healing BPEL processes with Dynamo and the JBoss rule engine. In *ESSPE'07*, pages 11–20. ACM, 2007.

[4] L. Baresi, R. Heckel, S. Thone, and D. Varro. Style-based refinement of dynamic software architectures. In *WICSA'04*. IEEE Computer Society, 2004.

[5] B. Becker, D. Beyer, H. Giese, F. Klein, and D. Schilling. Symbolic invariant verification for systems with dynamic structural adaptation. In *Int. Conf. on Software Engineering (ICSE)*. ACM Press, 2006.

[6] B. Becker and H. Giese. Modeling of correct self-adaptive systems: A graph transformation system based approach. In *Soft Computing as Transdisciplinary Science and Technology (CSTST'08)*, pages 508 – 516. ACM Press, 2008.

[7] Y. Brun, G. Marzo Serugendo, C. Gacek, H. Giese, H. Kienle, M. Litoiu, H. Müller, M. Pezzè, and M. Shaw. Engineering self-adaptive systems through feedback loops. *Software Engineering for Self-Adaptive Systems*, pages 48–70, 2009.

[8] A. Bucchiarone. *Dynamic software architectures for global computing systems*. PhD thesis, IMT Institute for Advanced Studies, Lucca, Italy, 2008.

[9] A. Bucchiarone, P. Pelliccione, C. Vattani, and O. Runge. Self-repairing systems modeling and verification using AGG. In *WICSA'09*, 2009.

[10] G. Canfora, M. D. Penta, R. Esposito, and M. L. Villani. An approach for qos-aware service composition based on genetic algorithms. In *GECCO*, pages 1069–1075, 2005.

[11] B. H. C. Cheng, H. Giese, P. Inverardi, J. Magee, and R. de Lemos et al. Software Engineering for Self-Adaptive Systems: A Research Road Map. In *Software Engineering for Self-Adaptive Systems*, Dagstuhl Seminar Proceedings, 2008.

[12] M. Colombo, E. D. Nitto, and M. Mauri. Scene: A service composition execution environment supporting dynamic changes disciplined through rules. In *ICSOC*, pages 191–202, 2006.

[13] H. Ehrig, K. Ehrig, U. Prange, and G. Taentzer. *Fundamentals of Algebraic Graph Transformation*. EATCS Monographs in Theor. Comp. Science. Springer, 2006.

[14] H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors. *Handbook of Graph Grammars and Computing by Graph Transformation, Volume 2: Applications, Languages and Tools*. World Scientific, 1999.

[15] H. Ehrig, C. Ermel, O. Runge, A. Bucchiarone, and P. Pelliccione. Formal analysis and verification of self-healing systems. In D. Rosenblum and G. Taentzer, editors, *Proc. Intern. Conf. on Fundamental Aspects of Software Engineering (FASE'10)*, LNCS, pages 139–155. Springer, 2010.

[16] H. Ehrig, A. Habel, and L. Lambers. Parallelism and Concurrency Theorems for Rules with Nested Application Conditions. *EC-EASST*, 2010. To appear.

[17] J. Floch, S. Hallsteinsen, E. Stav, F. Eliassen, K. Lund, and E. Gjorven. Using architecture models for runtime adaptability. *IEEE Software*, 23(2):62–70, 2006.

[18] D. Garlan and B. Schmerl. Model-based adaptation for self-healing systems. In *WOSS '02*, pages 27–32. ACM, 2002.

[19] D. Hirsch, P. Inverardi, and U. Montanari. Modeling software architectures and styles with graph grammars and constraint solving. In *WICSA*, pages 127–144, 1999.

[20] P. Inverardi, P. Pelliccione, and M. Tivoli. Towards an assume-guarantee theory for adaptable systems. In *SEAMS*, pages 106–115. IEEE Computer Society, 2009.

[21] H. Kastenberg and A. Rensink. Model checking dynamic states in groove. In *Proc. Model Checking Software*, volume 3925 of *LNCS*, pages 299–305. Springer, 2006.

[22] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.

[23] J. Kramer and J. Magee. Self-managed systems: an architectural challenge. In *FOSE*, pages 259–268, 2007.

[24] D. L. Métayer. Describing software architecture styles using graph grammars. *IEEE Trans. Software Eng.*, 24(7):521–533, 1998.

[25] D. Perry and A. Wolf. Foundations for the Study of Software Architecture. *SIGSOFT Softw. Eng. Notes*, 17(4):40–52, 1992.

[26] W. Reisig. *Petri Nets: An Introduction*, volume 4 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1985.

[27] A. Rensink, A. Schmidt, and D. Varr'o. Model checking graph transformations: A comparison of two approaches. In *Proc. Int. Conf. on Graph Transformation (ICGT'04)*, volume 3256 of *LNCS*, pages 226–241. Springer, 2004.

[28] G. D. Rodosek, K. Geihs, H. Schmeck, and S. Burkhard. Self-healing systems: Foundations and challenges. In *Self-Healing and Self-Adaptive Systems*, number 09201 in Dagstuhl Seminar Proceedings. Germany, 2009.

[29] E. Rukzio, S. Siorpaes, O. Falke, and H. Hussmann. Policy based adaptive services for mobile commerce. In *WMCS'05*. IEEE Computer Society, 2005.

[30] G. Spanoudakis, A. Zisman, and A. Kozlenkov. A service discovery framework for service centric systems. In *IEEE SCC*, pages 251–259, 2005.

[31] S. R. White, J. E. Hanson, I. Whalley, D. M. Chess, A. Segal, and J. O. Kephart. Autonomic computing: Architectural approach and prototype. *Integr. Comput.-Aided Eng.*, 13(2):173–188, 2006.

[32] L. Zeng, B. Benatallah, M. Dumas, J. Kalagnanam, and Q. Z. Sheng. Quality driven web services composition. In *WWW*, pages 411–421, 2003.