

# Self-Healing Systems: Foundations and Challenges

Gabi Dreo Rodosek<sup>1</sup>, Kurt Geihs<sup>2</sup>, Hartmut Schmeck<sup>3</sup> and Burkhard Stiller<sup>4</sup>

<sup>1</sup> Universität der Bundeswehr München, Faculty of Informatics  
Werner-Heisenberg-Weg 39, 85579 Neubiberg, Germany  
[gabi.dreo@unibw.de](mailto:gabi.dreo@unibw.de)

<sup>2</sup> University of Kassel, Department of Computer Science  
Wilhelmshöher Allee 73, 34121 Kassel, Germany  
[geihs@uni-kassel.de](mailto:geihs@uni-kassel.de)

<sup>3</sup> Karlsruhe Institute of Technology, Institut für Angewandte Informatik und Formale  
Beschreibungsverfahren 76128 Karlsruhe, Germany  
[schmeck@kit.edu](mailto:schmeck@kit.edu)

<sup>4</sup> University of Zürich, Department of Informatics  
Binzmühlestrasse 14, 8050 Zürich, Switzerland  
[stiller@ifi.uzh.ch](mailto:stiller@ifi.uzh.ch)

**Abstract.** The term and characteristic of self-healing, applied to systems, is often seen from different fields of computer science, such as fault tolerance or network and service management, with diverging semantics. Since this impression was confirmed also during the first discussions of the Dagstuhl seminar on "Self-Healing and Self-Adaptive Systems", a seminar's working group on "Terminology" was formed with the objective to address the question of finding commonalities and differences in a self-healing characteristic of stand-alone and distributed systems. The outcomes of the discussion in terms of foundations, the description of the self-healing process and the identification of the main challenges of such self-healing systems are presented.

**Keywords.** Self-healing system, foundations, terminology, applications, challenges

## 1 Introduction

Inspired by biology, self-healing has evolved as a property of IT systems that enables such systems to diagnose themselves and react to faults such that a satisfactory mode of operation is restored after the occurrence of a fault or a failure. Self-healing is not confined to a certain class of IT systems but covers the broad span of computing from applications over system software to networking, from small embedded devices over Personal Computers to GRIDs. Currently, self-healing is a buzzword. Since Koopman's proposal of a "taxonomy for describing the problem space for self-healing systems" ([1]), a boost of research and development projects have started. In spite of this, the term "self-healing" is still vague, and a more precise definition is still missing.

During the Dagstuhl seminar on "Self-Healing and Self-Adaptive Systems" a working group was formed to address these issues. The objective of this short summary is to sketch the view of self-healing as discussed in the working group. This covers the circling of the term by relating it to established concepts, but also to present (i) the view of self-healing in adaptive systems, (ii) the development of a methodology in respect to self-healing, and (iii) the key challenges such systems have to take up.

## 2 A New Attempt at Terminology

Before we give our definition of self-healing, we want to further characterize and shed more light on this property by putting it into relation with other well-known terms in this problem space. Often, a clear distinction between names for concepts is a matter of viewpoint and taste. For sure, all attempts in this direction will generate fierce discussions. Nevertheless, we dare to provide the following taxonomy for self-\* systems, as visualized in Table 1. We claim that from a practical point of view it does not pay off to differentiate those terms that appear in the same cell. We argue that self-healing is (almost 100%) synonymous to self-repairing, self-regeneration and self-immunity. These terms describe a system that can make by itself all necessary recovery steps to restore its distributed behaviour to a specified mode of operations. Besides, terms such as self-optimizing related with self-tuning, self-protecting and self-managing are circulating. Already in 1973, Dijkstra ([2]) provided a definition of a self-stabilizing system in terms of that a self-stabilizing system arrives at a legitimate state in a finite number of steps regardless of its initial state.

Effective self-healing mechanisms are one of the means to increase the dependability of IT systems by reacting to faults and keeping the system operational. Of course, there is a tight relationship to the term "fault tolerance" here. So what is this relationship? Is "self-healing" reinventing the wheel?

Avizienis defined fault tolerance as "mechanisms and techniques that enable to deliver its specified service despite the presence of faults" ([4]). A popular fault-tolerance mechanism is redundancy of components and voting on results, as used, for example, in Triple-Modular-Redundancy (TMR) designs. Clearly, a TMR device would not necessarily be called self-healing, because it lacks reactivity that aims at healing (i.e. repairing) the system. Hence, not all fault-tolerant systems are self-healing systems. Furthermore, fault-tolerance aims at keeping the system running at 100% of its designed functionality, while self-healing can mean that after the healing the system operates at less than 100%. If a system is able to self-heal completely the occurrence of faults such that the operation of the system is not impaired at all, it could be called fault-tolerant. Thus, we are inclined to say that the answer to the question raised by Koopman at the above mentioned ICSE 2003 WASD "Was all self-healing also fault tolerance?" depends on the degree of degradation after a self-healing activity: If the self-healing mechanism in a system is such that it always restores the system to 100% functionality, then the system is fault-tolerant.

**Table 1.** An Attempt to a Taxonomy of Self-\* Systems

<b>self-managing</b> <i>self-organizing</i> <i>self-adapting</i>	A system continuously perceives its own state and the state of its environment and reacts to certain events in order to maintain a high degree of usefulness without a human in the loop	
	<b>self-healing</b> <i>self-repairing</i> <i>self-regenerating</i> <i>self-immunity</i>	A system can make by itself all necessary recovery steps to restore its disturbed behavior to a specified mode of operation.
	<b>self-optimizing</b> <i>self-tuning</i>	A system optimises its use of resources; it may decide to initiate a change in an attempt to improve its performance or quality of service ([3]). This optimization action may seem proactive from the perspective of a self-healing system design, but nevertheless it is reactive because it monitors the performance and decides to act when some specified condition is reached.
	<b>self-protection</b>	A system protects itself from malicious attacks as well as potential erroneous states but also from end users who inadvertently make software changes, for example, by deleting an important file ([3]).
	<b>self-stabilizing</b>	I call a system self-stabilizing when, regardless its initial state, it is guaranteed to arrive at a legitimate state in a finite number of steps. ([2])

### 3 Self-Healing Process

In order to explain the term "self-healing", the process of a self-healing action needs a careful investigation. Fig. 1 visualizes this process, in which the deviation from the desired system behavior is plotted against time, as it progresses for a given system in operation.

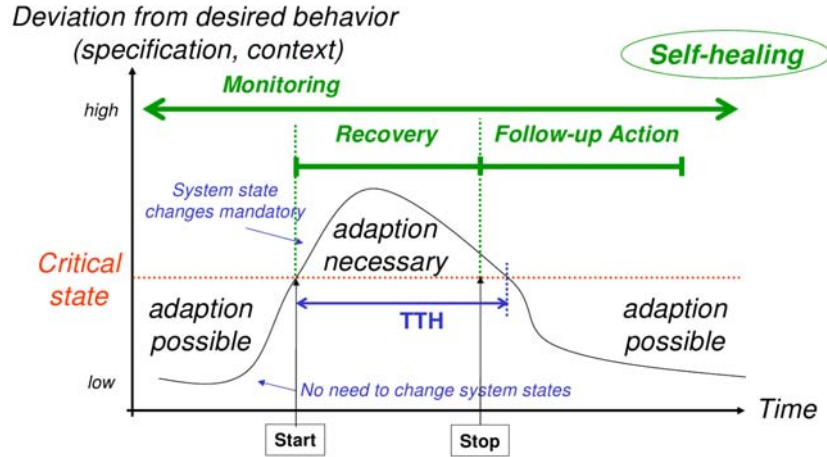


Fig. 1. Self-Healing Process

Over time, the criticality of a system state can change. If the state of a self-healing system exceeds some threshold (critical state), then it is necessary to start recovery actions to bring the system back to "normal" state, resp. a state that is "below critical state". The criticality of a state is a measure of the deviation of a system from the desired behavior, which is a function of the specification and the context. It is important to note that if the deviation from the desired behaviour is below the critical state, no adaption is needed, but nevertheless possible. However, as soon as the deviation from the desired behavior of the system is above the threshold, an adaptation is necessary. At this point, the recovery actions need to start. An output of the discussion in the working group was also that recovery actions can stop although the system state is not yet back below the threshold. We define the time between the start of the recovery actions and the time when the system is no longer above the critical state as the *time to heal (TTH)*. We distinguish between recovery and follow-up actions. The recovery actions restore the behavior of the system below or close to the critical state, whereas the follow-up actions may refer to additional changes in the environment etc. in terms of proactive actions.

Let us summarize the discussions so far in this paper by the following definition of a self-healing system:

*A system is showing the self-healing characteristic if it is able to monitor and heal itself from the inside, which requires the ability of this system to decide about and perform recovery actions to return itself to a behavior conforming to its initial specification, especially without external interference.*

## 4 Challenges of Self-Healing

Based on the discussion so far and the described self-healing process, the following challenges were identified:

- Monitoring: Monitoring a system refers to the detection of relevant events by means of polling the system state or by reports from the system (e.g. traps). This addresses the ability to:
  - receive any type of events from inside and outside a self-healing system and log it
  - correlate "inside" and "outside" events according to temporal, spatial and functional constraints; this is especially necessary when the cause of a fault is the result of an intentional external malicious attack
  - instrument applications and develop new types of monitoring systems
- Specification of the desired behavior (critical state, context): Is it possible that a self-healing system reacts differently in another context (i.e. exhibits situation awareness)?
- Fault analysis and decision taking: When is it necessary to take what recovery actions?
- Recovery action: What kind of actions is needed / available to recover a certain kind of system? What degree of recovery is sufficient to "heal" the system? What does recovery mean in service-oriented computing based on service-level agreements?
- Learning: Self-learning involves the definition of the "initial settings" as well as the process of how to "correctly" learn from detected, diagnosed and repaired faults. A self-healing system cannot be independent from the outside. The learning capabilities are in a similar way influenced as the analysis and decision steps.
- Pro-activeness: In fact, we have two perspectives on a self-healing system: (i) the global view with the inter-dependencies and the (ii) local view with the intra-dependencies. Should a system only react on indicators (sensors reporting a fault) or should something like a "near failure state" be recognized and reacted on before a fault occurs?
- Follow-up actions: What follow-up actions are appropriate to improve the stability of a system?

## 5 Summary and Preliminary Conclusions

This short paper did discuss the reasons for addressing the notion of self-healing characteristic of systems, while assuming that such a meaning may be different

for "traditional" systems compared to networked and distributed systems. In particular, a more detailed investigation of fault tolerance and related areas is necessary to ensure that the self-healing characteristic can be utilized in the future without any major discrepancies, if used for a stand-alone or a networked system.

Thus, the term itself was introduced and described by related terms and characteristics, which are in use in many areas of computer science. The key to defining a self-healing characteristic was found by determining the process for applying self-healing mechanisms. This led to a definition of a self-healing system.

Additionally, such a view on a networked system raises a number of further aspects, which were determined as challenges for self-healing systems, and which have not seen in all cases an existing solution as of today. Therefore, those aspects constitute a list of open research issues..

## Acknowledgments

We are grateful to the participants of the international Dagstuhl seminar on "Self-healing and self-adaptive systems" for many inspiring presentations and discussions.

## References

1. Koopman, P.: Elements of the self-healing system problem space. 1st Workshop on Self-Healing Systems (WOSS'02), ACM Press (2002)
2. Dijkstra, E.: Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, I (1974)
3. Huebscher, M.C., McCann, J.A.: A survey of autonomic computing—degrees, models, and applications. *ACM Computing Surveys*, Vol. 40, No. 3, ACM Press (2008)
4. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, Vol. 1., No. 1, (2004)