Characterizing and checking self-healability

Marie-Odile Cordier[1] and **Yannick Pencolé**[2] and **Louise Travé-Massuyès**[2] and **Thierry Vidal**[1]

1 INTRODUCTION

Real-life complex systems are often required to offer high reliability and quality of service and must be provided with self-management abilities, even in faulty situations. They are expected to be self-aware of their current state and survive autonomously the occurrence of faults, still managing to provide the desired functionality. In other words, such systems must be self-healing [2].

Designing self-healing systems requires to be able to evaluate the joint degree of self-awareness and reactiveness. In the artificial intelligence community, these two properties are better known as diagnosability [3, 1], i.e. the capability of a system to exhibit different observables for different anticipated faulty situations, and repairability, i.e. the ability of a system and its repair actions to cope with any unexpected situation.

Checking separately diagnosability and repairability leads to a conservative assessement of self-healability. In this paper, we show that neither standard diagnosability nor repairability of every anticipated fault are necessary to achieve self-healability. Our main contribution consists of defining self-healability as a joint property bridging diagnosability and repairability, which requires a new definition of diagnosability that allows diagnosable subsets of faults to overlap, as opposed to the standard definitions which rely on a partition.

2 MAIN CONCEPTS

The presented framework, which is relevant for state based or event based systems, adopting the generic viewpoint defined in [1], is illustrated with discrete event systems [3] as our current objective is to apply it to service oriented architectures like Web Services in the framework of the WS-DIAMOND European project [4].

Observations and Faults : The set of observable events is $\mathcal{O} = \{o_1, \ldots, o_{no}\}$. Complementing $\mathcal{O}$ with the set of unobservable events $\mathcal{U} = \{u_1, \ldots, u_{nu}\}$ determines the whole set of events of the system $\mathcal{E} = \mathcal{O} \cup \mathcal{U}$. The occurrence of *basic faults* that might occur on the system are represented as specific unobservable events noted f_i. In the following, we restrict ourselves to the single fault assumption (i.e. only one fault can be present in the system at a given time). The system can then be either in a nominal mode (absence of fault) or in one of the nf fault modes. The set of all possible system modes is hence given by $\mathcal{F} = \{f_0, f_1, ..., f_{nf}\}$, where $f_0 = ok$.

$\mathcal{T}$ denotes the set of (infinite) possible trajectories (i.e. sequences of events) occurring in the system, while OBS is the set of all possible sequences of observable events. A trajectory $\tau \in \mathcal{T}$ corresponds to only one observable σ, while one σ may correspond to several disctint trajectories.

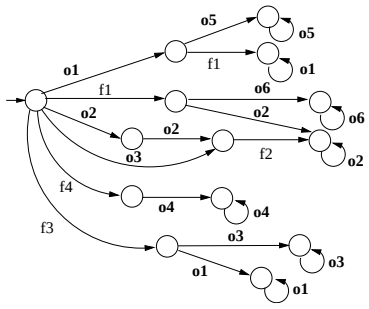

The above figure represents the global model of a discrete-event system. The set of fault modes is $\mathcal{F} = \{ok, f_1, f_2, f_3, f_4\}$. Fault events are not observable, the other events being observable. $o_1 o_5^\infty$ is both a trajectory and the observable obtained over that trajectory including an infinite sequence of o_5. $o_2 o_2 f_2 o_2^\infty$ is another trajectory yielding the observable o_2^∞. $f_1 o_2^\infty$ is yet another trajectory which, interestingly enough, yields the same observable o_2^∞, which means these two trajectories cannot be discriminated from the observations.

Macrofaults: It is not always possible to know with certainty in which mode a system is. It is often even not necessary with respect to reparability. It is why we define the concept of *macrofault* that represents the belief state referring to the system mode. A macrofault can be seen as an abstraction of system modes. For instance, if a pipe can be in the two basic fault modes leaking or blocked, it can also be said to be in an abnormal macrofault mode, where abnormal corresponds to leaking or blocked.

A *macrofault* F_j is described by a non empty set of fault modes. With our single fault assumption, an 'occurrence' of F_j means that exactly one of the faults $f_i \in F_j$ has occurred in the system. For instance, the macrofault $\{f_1, f_2\}$ represents the fact that either f_1 or f_2 has occurred . A macrofault may be a singleton ($F_j = \{f_i\}$). If all basic faults appear in a set of macrofaults $E(\mathcal{F}) \subseteq 2^{\mathcal{F}}$, then it is called a *covering set*.

Repairs : A repair plan is defined in a simplified way as, for our purpose, only the existence of such repair plans and their matching to (basic) faults is relevant.

The set of available repair plans is denoted $\mathcal{R} = \{r_1, ..., r_{nr}\}$. The predicate $Repair$ relates repair plans to (macro)faults: $Repair(r_k, F_i)$ means that applying the repair plan r_k brings back the system into a nominal state, under the condition that the system is in one of the modes described by the macrofault F_i [4] For instance,

[1] IRISA/INRIA/Université de Rennes1; Campus de Beaulieu, F-35042 Rennes cedex, France, email: marie-odile.cordier, thierry.vidal@irisa.fr
[2] LAAS-CNRS; Université de Toulouse; 7, Avenue du Colonel Roche, F-31077 Toulouse cedex, France, email: yannick.pencole, louise@laas.fr
[3] We assume the *liveness* of the observations [3].
[4] r_{ok}, the (void) repair plan such that $Repair(r_{ok}, ok)$, is assumed to exist.

the repair plan r_1 such that $Repair(r_1, \{f_1, f_2\})$ can be executed only if either f_1 or f_2 has occured.

Having a repair plan for a macrofault is equivalent to having a repair plan for all the basic faults belonging to the macrofault, hence the following property : $Repair(r_k, F_j) \equiv \forall f_i \in F_j, Repair(r_k, \{f_i\})$.

3 SELF-HEALABILITY

Self-healability is intuitively defined by *"A system is self-healing if, and only if, after the occurrence of any basic fault, a diagnosis is issued that automatically raises a repair plan fitted to the fault."* Behind this intuitive definition, two properties of the system are hidden: diagnosability and repairability.

Diagnosability : Diagnosability relies on the notion of *fault signatures* [1]. Intuitively, a fault signature is the association between a fault and a set of possible observables.

We use the following notations :

- The predicate $yields(f_i, \sigma)$ means that there exists at least one trajectory in which $f_i \in \mathcal{F}$ is present and that yields the observable $\sigma \in OBS$. The predicate $yields$ can be generalized to macro-faults: $yields(F_j, \sigma)$ means that it $\exists f_i \in F_j$ such that $yields(f_i, \sigma)$. σ is then called an elementary signature, or *e-signature* of the fault F_i,
- $MF(\sigma)$ is the (unique) macrofault containing all faults that may yield σ, i.e. $MF(\sigma) = \{f_i$ such that $yields(f_i, \sigma)\}$. MF can be generalized to sets of e-signatures: $MF(\Sigma) = \bigcup_{\sigma \in \Sigma} MF(\sigma)$.

In this work, we are not interested in checking that any basic fault can be diagnosed, but we are interested in finding the level of diagnosability of a system. This is why the partition of faults classically used is replaced by a set of macrofaults possibly sharing common faults. Still, each macrofault must be associated to distinct observables and the corresponding sets of observables need to form a partition. Hence the following new definition for diagnosability that extends the classical definition and is suitable for self-healability.

Definition 1 (Diagnosability of a set of macrofaults) *The covering set $E(\mathcal{F})$ is* diagnosable*, noted $Diagnosable(E(\mathcal{F}))$, iff there exists a partition $\pi = \{\Sigma_1, \ldots, \Sigma_m\}$ of the observables OBS such that: $E(\mathcal{F}) = \{MF(\Sigma_j), \Sigma_j \in \pi\}$.*

Example: A first straightforward set of macrofaults is $E_\top(\mathcal{F}) = \{\mathcal{F}\} = \{\{ok, f_1, f_2, f_3, f_4\}\}$ in which faults are indistinguishable: obviously it is diagnosable, the partition being $\pi_\top = \{\{o1o5^\infty, o_1^\infty, o_2^\infty, o_3o_2^\infty, o_3^\infty, o_4^\infty, o_6^\infty\}\} = \{OBS\}$.
The set of macrofaults $E_1(\mathcal{F}) = \{\{ok\}, \{f_1, f_2\}, \{f_1, f_3\}, \{f_4\}\}$ is diagnosable with $\pi_1 = \{\{o1o5^\infty\}, \{o_2^\infty, o_3o_2^\infty\}, \{o_1^\infty, o_3^\infty, o_6^\infty\}, \{o_4^\infty\}\}$. Note that $E_1(\mathcal{F})$ also corresponds to another partition $\pi_2 = \{\{o1o5^\infty\}, \{o_2^\infty, o_3o_2^\infty, o_6^\infty\}, \{o_1^\infty, o_3^\infty\}, \{o_4^\infty\}\}$. $E_2(\mathcal{F}) = \{\{ok\}, \{f_1\}, \{f_2\}, \{f_3\}, \{f_4\}\}$ is not diagnosable because there are some cases in which f_1 and f_2 cannot be discriminated, there is no partition of observables associated to it (the same for f_1 and f_3).

Repairability : A macrofault F_j is repairable if and only if there exists a repair plan that repairs it. $Repairable(F_j) \equiv \exists r_k$ such that $Repair(r_k, F_j)$. The repairability of a set of macrofaults is then defined as the repairability of all the macrofaults in the set.

Definition 2 (Repairability) *A set of macrofaults $E(\mathcal{F})$ is repairable, noted $Repairable(E(\mathcal{F}))$, iff $\forall F_j \in E(\mathcal{F})$ $Repairable(F_j)$.*

Example : If the only repair plan is r, with $Repair(r, \{f_1, f_3\})$, we indeed get Repairable($\{f_1, f_3\}$), and also Repairable($\{f_1\}$) and Repairable($\{f_3\}$). However, the system is not repairable since the faults f_2 and f_4 are not repairable.

Self-healability : Our definition for self-healibility directly derives from the definitions of diagnosability and repairability.

Definition 3 (Self-healing set of macrofaults) *A set $E(\mathcal{F})$ is* self-healing *iff it is diagnosable and repairable, i.e. $SelfHealing(E(\mathcal{F})) \equiv Diagnosable(E(\mathcal{F}))$ and $Repairable(E(\mathcal{F}))$.*

Definition 4 (Self-healing system) *A system is* self-healing *iff there exists a self-healing covering set $E(\mathcal{F})$.*

Example : If Repairable(ok), Repairable($\{f_1,f_3\}$), Repairable($\{f_1,f_2\}$) and Repairable(f_4), then the set $E_1(\mathcal{F}) = \{\{ok\}, \{f_1, f_2\}, \{f_1, f_3\}, \{f_4\}\}$ is diagnosable and repairable. The system is self-healing. If Repairable(ok), Repairable($\{f_1,f_3\}$), Repairable(f_2) and Repairable(f_4) then the system is not self-healing as there does not exist a repair plan for $\{f_1, f_2\}$.

Due to lack of space, the algorithm to check whether a system is self-healing is not given.

4 CONCLUSION AND PERSPECTIVES

The main contributions of this paper are first a new and original definition of diagnosability which allows to diagnose possibly overlapping sets of non-discriminated faults, and then using that definition, to propose a thorough and integrated definition of the *self-healability* of a dynamic system. Interestingly enough, diagnosability of each basic fault is not required but what is needed is a diagnosability level that can be matched to the existing repairs. As far as we know, it is the first time that such a definition is issued.

We are currently applying our work to web services in the framework of the WS-DIAMOND European project [4], in which we investigate a number of extensions to address more sophisticated and realistic cases, mostly in terms of the characterization of repair plans, their properties and conditions of applicability. One of the problems is how to deal with multiple faults that may appear sequentially. Another interesting issue refers to temporal conditions that may restrict the applicability of repairs and be in conflict with the time needed to diagnose a fault.

REFERENCES

[1] M.-O. Cordier, L. Travé-Massuyès, and X. Pucel, 'Comparing diagnosability in continous and discrete-event systems', in *17th International Workshop on Principles of Diagnosis*, eds., C.A. González, T. Escobet, and B. Pulido, pp. 55–60, (June 2006).

[2] D. Ghosh, R. Sharman, H. R. Rao, and S. Upadhyaya, 'Self-healing systems survey and synthesis', *Decision Support Systems*, **42**(4), 2164–2185, (2007).

[3] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, 'Diagnosability of discrete event system', *IEEE Transactions on Automatic Control*, **40**(9), 1555–1575, (1995).

[4] The Ws-DIAMOND team, 'Ws-DIAMOND: Web services DIagnosability, MOnitoring and DIagnosis', in *18th International Workshop on Principles of Diagnosis, DX'07*, pp. 243–250, Nashville (TN, USA), (May 2007).