## Doctoral thesis (PhD) position at INRIA

# Monitoring and Diagnosis of Large Scale Systems

**Location:** SUMO team, INRIA Rennes, France
**Contact:** Eric Fabre, Eric.Fabre@irisa.fr, +33 (0) 2 99 84 73 26

**Application:** send resume, motivation letter and a list of professional references

**Context:**

The SUMO team develops formal methods for the modelling, verification and management of large scale distributed systems. Besides scalability issues, the team focuses on quantitative models and methods, abstraction and approximation techniques, and distributed algorithms.

SUMO is collaborating with leading telecommunication operators and vendors. The proposed research will support some of the theoretical developments needed for this collaboration.

**Subject:**

Large-scale systems are pervasive today. They are hidden behind numerous services we use on a daily basis, such as apps running in data centers, the data centers themselves, telecommunication networks, etc. While such systems are deployed and running, their verification and their management remain research issues. This thesis aims at providing theoretical tools to deal with such large systems obtained by assembling components. The underlying application domain is the telecom industry, with a focus on the management of software-defined networks, and further on the joint management of clouds connected by virtual networks.

The proposed research is oriented to diagnosis problems for systems modelled as automata, networks of automata, Petri nets, and the quantitative versions of these models (with time, probabilities, costs…). Diagnosis problems consist in deciding whether some specific property (e.g. presence of a fault) holds or not in a system given the observations collected on this system. It is a simple paradigm of inference problems under partial observations. Similarly, diagnosability states that as soon as the property holds, this can be detected from the following observations. Over the past decades, these problems have been well understood for different models, sometimes with quantitative features, and both in centralized and decentralized settings. However, the applicability of the existing results remains limited due to a mismatch between formal problems and the features required by some applications.

The objective of the thesis is to contribute to filling this gap, by exploring some of the following research directions:

- Multiresolution diagnosis. This consists in building hierarchical models that describe a system at different levels of resolution or granularity. This is the standard way engineers use to build complex systems. Management strategies should therefore adopt the same paradigm. One expects for example methods that perform diagnosis at coarse resolution first, and then refine their result if necessary, possibly by requesting extra observations.
- Defining and computing the diagnosability degree of a system. So far, diagnosability is a 0/1 (or logical) property. For non-diagnosable systems, one would like to understand the severity of this non-diagnosability, by means of an indicator ranging in [0,1]. This relates to information theory: how much extra information is necessary on the average to complement the partial observations collected on a system in order to make it diagnosable?
- Distributed quantitative diagnosis. In the case of distributed systems observed by distributed sensors, communication can be necessary to help one local supervisor decide on the presence

or not of a fault. One would like to quantify the amount of information exchanges that are necessary for this decision. This problem relates to security issues: how much does one component learn on the behaviour of another component, from the partial observations that are available? It also has connections to communications complexity.

- Active diagnosis. So far, diagnosis problems have mostly been stated as passive problems, where one does his best from the available observations. In large scale distributed systems, it is rather likely that all observations are not available at once, and that one has to choose which observation it should go and collect. Therefore active strategies must be designed that poll the system for new observations, possibly with extra observation costs to minimize.

The above problems are only examples of research directions that the Sumo team wants to explore. They can be adapted to the tastes and skills of the candidate. As these topics are also related to our collaborations with Telecommunication Companies, they can either be theory oriented, or application oriented, according again to the wishes of the candidate.

**Keywords:** formal methods, diagnosis, diagnosability, quantitative models, abstraction, multi-resolution system, large-scale system management, telecommunication networks

**Required skills:** Master in computer science, background in formal methods, possibly in quantitative models (e.g. stochastic systems).
A taste and skills for applications are a plus.