



# Self-Adaptation in Highly Distributed Dynamic Systems

Ilias Gerostathopoulos, Sebastian Götz, Filip Krikava, Adnan Shahzada,  
Romina Spalazzese

## ► To cite this version:

Ilias Gerostathopoulos, Sebastian Götz, Filip Krikava, Adnan Shahzada, Romina Spalazzese. Self-Adaptation in Highly Distributed Dynamic Systems: part of Report from the GI Dagstuhl Seminar 14433: Software Engineering for Self-Adaptive Systems. Report from the GI Dagstuhl Seminar 14433: Software Engineering for Self-Adaptive Systems, Oct 2014, Dagstuhl, Germany. 2014, <www.dagstuhl.de/14433>. <hal-01108475>

**HAL Id: hal-01108475**

**<https://hal.inria.fr/hal-01108475>**

Submitted on 19 Feb 2015

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Self-Adaptation in Highly Distributed Dynamic Systems

*part of* Report from the GI Dagstuhl Seminar 14433 - Software Engineering for Self-Adaptive Systems

Ilias Gerostathopoulos<sup>1</sup>, Sebastian Götz<sup>2</sup>, Filip Křikava<sup>3</sup>, Adnan Shahzada<sup>4</sup>,  
and Romina Spalazzese<sup>5</sup>

<sup>1</sup> Charles University in Prague, Czech Republic [iliasg@d3s.mff.cuni.cz](mailto:iliasg@d3s.mff.cuni.cz)

<sup>2</sup> Technische Universität Dresden, Germany [sebastian.goetz@acm.org](mailto:sebastian.goetz@acm.org)

<sup>3</sup> INRIA Lille / University of Lille 1, France [filip.krikava@inria.fr](mailto:filip.krikava@inria.fr)

<sup>4</sup> Malmö University, Sweden [romina.spalazzese@mah.se](mailto:romina.spalazzese@mah.se)

<sup>5</sup> Politecnico di Milano, Italy [adnan.shahzada@polimi.it](mailto:adnan.shahzada@polimi.it)

This breakout group focused on identifying the challenges of performing self-adaptation in highly distributed dynamic systems. This is a pressing issue in self-adaptive systems research, as proposed “smart” systems are increasingly built out of disparate entities (sensors and actuators) that feature a close connection to the physical world – so-called cyber-physical systems (CPSs). Examples are numerous: intelligent vehicle navigation, fleets of autonomous robots, emergency coordination systems, to mention just a few. CPSs are typically distributed at the physical space and feature no firm boundaries – they are open-ended. They are composed of loosely connected entities, that are often mobile. Grafting such systems with self-adaptive capabilities is a distinct challenge, which projects itself in all phases of the autonomic loop.

After debating on the different problems that can arise, the group focused on the issue of efficient information sharing between the entities of a self-adaptive CPS and how this affects the overall utility of the system. Our main assumption was that there is no central coordination point in our target CPS, but rather each entity follows its own decisions based on its partial view over the rest of the system. We also assumed that entities form ad-hoc collaboration groups to achieve common system-level objectives. Each partial view is then updated according to the information shared between the collaborating groups. In this frame, a separate autonomic loop is instantiated for each entity and consists of (i) monitoring the changes in the environment/nearby entities, (ii) creating a partial runtime model capturing the state of the world as viewed by the entity in question, (iii) deciding based on the partial runtime model and the adaptation logic, (iv) enacting the necessary changes to the runtime behavior of the entity in question. In our discussions, we considered the case of a “smart crossroad”, where each car decides whether to cross or wait at a crossroad according to its view over the positions and intentions of the other cars in the crossroad. In this case, the ad-hoc collaboration group between the cars that approach the crossroad has the objective to avoid collisions and increase the safety of the whole system.

Moving into the more concrete questions, the following challenges have been discussed in more detail:

1. **How to disseminate the information necessary to build the individual partial runtime models in an efficient way?** The problem is that communicating the information of each entity with every other entity in the CPS overloads the network and threatens the privacy of individual entities. Hence, there is a need for an efficient dissemination strategy for partial model building. Apart from the obvious yet naive strategy where every entity shares all its information with each and every one (mentioned above), we identified two more elaborate strategies: (i) each entity in a collaboration group shares all its information only with the group members; (ii) each entity in a collaboration group shares only the information that are relevant to its objectives and the objectives of the group only with the group members, i.e. shares only a fraction of its information.
2. **How to design runtime abstractions in order to support explicit information exchange between the entities of a collaboration group?** We have looked into the technical aspects for partial model sharing. Essentially, this is a problem of model querying with role-based access control and model merging. The models of the entity that represent its own view should allow to express what information are accessible to whom and what information are to be exchanged in each collaboration group. Once the information has been gathered, the next step is to find a way to merge it into the entity's existing partial view (at the receiving end). This becomes challenging when we consider that the same information can be provided by multiple entities and with different quality and precision. The merging has to therefore involve some trust schemes, temporal aspects as well as to support versioning and transactions.

### **Acknowledgment**

This abstract is the result of a working group of the GI Dagstuhl Seminar 14433.<sup>6</sup>

---

<sup>6</sup> <https://www.hpi.uni-potsdam.de/giese/public/selfadapt/2014/12/17/report-from-the-gi-dagstuhl-seminar-14433/>