Science of Computer Programming 110 (2015) 3–22

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico

A coordination model of pervasive service ecosystems ☆

Mirko Viroli [a,*], Danilo Pianini [a], Sara Montagna [a], Graeme Stevenson [b], Franco Zambonelli [c]

[a] *Alma Mater Studiorum – Università di Bologna, Italy*
[b] *University of St Andrews, UK*
[c] *Università di Modena e Reggio Emilia, Italy*

ARTICLE INFO

Article history:
Received 15 October 2013
Received in revised form 7 June 2015
Accepted 8 June 2015
Available online 18 June 2015

Keywords:
Coordination models and languages
Pervasive computing
Service ecosystems
Process algebras
Self-organisation

ABSTRACT

The complexity of pervasive computing systems is rapidly increasing, and this calls for service models and coordination approaches intrinsically manifesting self-organisation of component interactions. The goal of this paper is to provide a coordination model (formalised as a process algebra) capturing the relevant aspects of such systems. It should allow the behaviour of large-scale, situated, and self-organising systems to be conveniently expressed, paving the way towards their rigorous study as well as development of supporting platforms. Focusing on the recently introduced concept of pervasive ecosystems, the proposed model revolves around *(i)* the notion of a distributed and dynamic space of "live semantic annotations" (wrapping data, knowledge, and the relevant facts about activities of humans, devices, and services) upon which autonomous agents coordinate and *(ii)* a small set of chemical-resembling coordination rules that enforce mechanisms of diffusion, aggregation, decay, and bonding between such annotations.

1. Introduction

The focus of this paper is on pervasive computing systems, made of actors and components of various kinds (which we group under the term *individuals*) [1] populating our everyday environments: humans, their smartphones, software services, pervasive displays, sensors and devices spread across the environment, sources of knowledge, data and events. On the one hand, individuals all work towards achieving their own goals or executing their own tasks. On the other hand, designers of whole systems must ensure that individuals can interoperate opportunistically to achieve global coherent behaviour, often expressed in terms of provided distributed pervasive computing services. For such systems to be effective in spite of (at least partially) unpredictable working contingencies, they must overall manifest some degree of autonomous behaviour, in terms of an intrinsic level of self-adaptivity and self-organisation.

Recent works [2–4] suggest a relationship between such systems and natural ecosystems, and hence propose to address this problem by enacting some "laws" present in the infrastructure, called *eco-laws*. By carefully balancing between each individuals' local interaction abilities and the behaviour of these eco-laws, such works show that it may be possible to identify a methodology for system design in which local agent goals, global objectives and self-* properties can fruitfully

☆ This paper has been supported by the EU FP7 project "SAPERE – Self-aware Pervasive Service Ecosystems" under contract No. 256873.

* Corresponding author.
E-mail addresses: mirko.viroli@unibo.it (M. Viroli), danilo.pianini@unibo.it (D. Pianini), sara.montagna@unibo.it (S. Montagna), graeme.stevenson@st-andrews.ac.uk (G. Stevenson), franco.zambonelli@unimore.it (F. Zambonelli).

http://dx.doi.org/10.1016/j.scico.2015.06.003

coexist. This idea can be recast into a coordination problem [5,6] for multiagent systems [7]. That is, how do we design a coordination model (agent coordination primitives and coordination laws regulating agent interactions) making such a fruitful coexistence possible and leading to effective large-scale and self-* pervasive computing systems? Answering this question, which is a key topic of the roadmap of [4], is precisely the goal of this paper.

We formalise a coordination model based on the idea of continuously reflecting the presence and activities of agents in the pervasive computing system through *Live Semantic Annotations* (LSAs), which are stored in the devices forming the infrastructure, and altogether form a global network of annotations representing the virtual counterpart of the ecosystem. In many ways, LSAs resemble the tuples of tuple-space architectures [8–10] but with a number of differences in their structure (LSAs can better articulate and encode semantic information [11]), relationship with the underlying space (unlike tuples they have a unique identifier, and can therefore hold references to others by a *bonding* mechanism) and behaviour (they can be modified only by the agent who created them, and bonding is the sole means to achieve observability). The whole ecosystem behaviour is regulated by a set of rules, called *eco-laws*, which act locally on each node and its neighbourhood, structurally matching LSAs and accordingly combining and manipulating them. They enact processes of LSA diffusion, aggregation, decay, and bonding, which are necessary to achieve self-organisation [12].

The proposed model is formalised as a process algebra of agents, LSAs, and eco-laws. Not only does this formalisation act as a non-ambiguous description, it is also an executable specification and can hence serve as ground for developing simulations of self-organisation mechanisms and for deriving formal proofs of behavioural properties, like self-stabilisation [13]. Few works attempted at a process algebraic formalisation of self-organisation [14,15], developing on top of well-known works on coordination and spatial computing [16,17] (see more details in Section 7). However, the one presented in this paper is—to the best of our knowledge—the first process algebra of self-organising multiagent systems capturing the clear separation of agent behaviour and self-organised interactions, hence effectively expressing the behaviour of large-scale and situated pervasive computing scenarios. Put in another way, the proposed operational semantics is the formalisation of an abstract machine that is the base of the SAPERE approach (both for execution and for simulation).

The remainder of this paper is organised as follows. Section 2 introduces the scenarios we aim to tackle by an example, showing how the basic requirements of pervasive ecosystems shape the main concepts of the proposed framework; Section 3 describes the coordination model, informally detailing agent primitives and the minimal eco-laws set adopted; Section 4 presents the formalisation of the model as a process algebra; Section 5 illustrates an example application in the context of pervasive computing systems, where self-organising spatial structures of LSAs are created in the style of spatial computing [18,19,9,20,21] and are used to help steer people in dynamic and complex environments; Section 6 discusses the expressive power of the proposed model; Section 7 presents related works, before we conclude with final remarks.

2. Concepts

Here we describe the general requirements of our coordination model, considering adaptive display ecosystems as a case [22,21,23] (Section 2.1), and the corresponding basic architectural choices (Section 2.2).

2.1. The scenario and requirements

It is a matter of fact that we are increasingly surrounded by digital displays: from those of wearable devices to large wall-mounted displays pervading urban and working environments. The main consequence is that they will soon no longer be used in a static way—e.g., cycling pre-defined commercials or news and information of general interest—but they will be made more useful and effective for both users and information/service providers by becoming general, open, and self-adaptive systems.

A first key requirement to be tackled is that of *situatedness*: pervasive ecosystems deal with spatially- and socially-situated activities of users, and should therefore be able to interact with the surrounding world and adapt their behaviour accordingly. In a display infrastructure deployed in a smartcity, one could exploit information from sensors that read profiles of people nearby a display (by interaction with their smartphones), to advertise today's events depending on the majority of profiles. Situatedness is generally achieved by infrastructures reifying data/knowledge/events in the precise point (or region) of space where they pertain, and promoting interactions based on proximity.

Another, complementary, requirement is *adaptivity*: pervasive ecosystems and their infrastructures should inherently exhibit properties of autonomous adaptation and management to survive contingencies without human intervention, global supervision, or both. In the display infrastructure, services should adapt to environment modifications in an automatic way without experiencing malfunction. For instance: the path towards a given area of interest, suggested by dynamic signs appearing on deployed displays, could be automatically computed by self-organisation so as to dynamically avoid overcrowded streets, or alternative locations may be selected if one cannot be reached quickly. Adaptivity is often achieved by designing coordination rules that by acting locally (namely, on a given network neighbourhood) make global properties emerge by self-organisation [24]—following a natural inspiration [22].

Finally, the infrastructure should intrinsically support *openness* of service production and usage. The number and classes of services provided should not be limited, rather the injection of new services should be taken advantage of by exploiting them to improve, integrate, and add further value to existing services whenever possible. In the display infrastructure, a new service to detect favourable areas based on a variety of information (traffic, pollution, anticipation of congestion)