AI & Soc
DOI 10.1007/s00146-013-0531-6

ORIGINAL ARTICLE

The pursuit of computational justice in open systems

Jeremy Pitt · Dídac Busquets · Régis Riveret

Received: 15 March 2013 / Accepted: 28 November 2013

Abstract Many open networks, distributed computing systems, and infrastructure management systems face a common problem: how to distribute a collectivised set of resources amongst a set of autonomous agents of heterogenous provenance. One approach is for the agents themselves to self-organise the allocation of resources with respect to a set of agreed conventional rules; but given an allocation scheme which maps resources to those agents and a set of rules for determining that allocation scheme, some natural questions arise—Is this allocation fair? Is the allocation method effective? Is it efficient? Are the decision makers accountable? In this paper, we argue that some answers to these questions can be found in the formal characterisation of different aspects of 'justice' and that these different aspects need a principled operationalisation as policies for system management. We present a formal model and some experimental results, concluding that the different aspects are all inter-connected and that what is required is a comprehensive research programme in computational justice.

Keywords Multi-agent systems · Self-organisation · Resource allocation · Computational justice

J. Pitt (✉) · D. Busquets · R. Riveret
Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2BT, UK
e-mail: j.pitt@imperial.ac.uk

D. Busquets
e-mail: didac.busquets@imperial.ac.uk

R. Riveret
e-mail: r.riveret@imperial.ac.uk

1 Introduction

In open systems, it is a common requirement to collectivise and distribute computing resources amongst the components of the system. This requirement cuts across different scales of time and space and is a feature of distributed systems for cloud and grid computing (Ardagna et al. 2011; Birman et al. 2009; Choi et al. 2008), and sensor and vehicular networks (Ding et al. 2003; Manvi et al. 2009). The decision-making required to achieve this distribution is too fast, too frequent and too complex for (human) operator intervention, so the system components have to self-organise the distribution by, and between, themselves.

Furthermore, there is another trend towards the automation of infrastructure management systems, for example in energy grids, water irrigation and transportation systems, which critically involve active participation of (human) users (Bourazeri et al. 2012; Ferscha et al. 2011). In these applications, there is also a requirement to share physical resources amongst the infrastructure users, who can be both producers and consumers of resources (i.e. prosumers, typically found in energy grids). If that infrastructure is instrumented with inter-connected sensors and (personal) devices, producing a so-called smart infrastructure, then the (user-centric) self-determination of the resource distribution can be assisted by computational means.

In both cases, whether a 'technical' system or network composed purely of autonomous computing components, as found in grid computing or sensor networks, or a socio-technical system composed of ICT-enabled people interacting with inter-connected, instrumented (and increasingly intelligent) devices, such as Smart Grids, we see that the actors have to self-determine the resource allocation; however, they can also self-determine the rules that are used in this self-determination. This is self-organisation: a

Published online: 24 December 2013