Economic Models and Algorithms for Distributed Systems
, 1–3
Book series: Autonomic Systems
© 2009 Birkhäuser Verlag Basel/Switzerland

Economic Models and Algorithms for Distributed Systems

Modern computing paradigms have frequently adopted concepts from distributed systems. The quest for scalability, reliability and cost reduction has led to the development of massively distributed systems, which extend organisational boundaries. Voluntary computing environments (such as BOINC), Grids (such as EGEE and Globus), and more recently Cloud Computing (both open source and commercial) have established themselves as a range of distributed systems.

Associated with this advance towards cooperative computing, the paradigm of software agents generally assumes that cooperation is achieved through the use of obedient agents that are under centralised control. In modern distributed systems, this main assumption is no longer valid. On the contrary, cooperation of all agents or computing components is often necessary to maintain the operation of any kind in a distributed system. Computer scientists have often considered the idea that the components of the distributed system are pursuing other selfish objectives, other than those that the system designer had initially in mind, when implementing the system. The peer-to-peer file sharing systems, such as BitTorrent and Gnutella, epitomises this conflict of interest, because as low as 20% of the participants contribute more than 80% of the files. Interestingly, various distributed systems experience different usage patterns. While voluntary computing environments prospered through the donation of idle computing power, cooperative systems such as Grids suffer due to limited contribution from their participants. Apparently, the incentive structure used to contribute to these systems; can be perceived differently by the participants.

Economists have also demonstrated research interest in distributed systems, exploring incentive mechanisms and systems, pioneered by Nobel-prize winners von Hayek and Hurwicz in the area of incentives and market-based systems. As distributed systems obviously raise many incentive problems, economics help complement computer science approaches. More specifically, economics explores situations where there is a gap between individual utility maximising behaviour and socially desirable deeds. An incorrect balance between such (often conflicting) objects could lead to malfunctioning of an entire system. Especially, cooperative computing environments rely on the contribution of their participants. Research test beds such as EGEE and PlanetLab impose regulations on the participants that contribute, but the enforcement of these institutions is informal by the loss of reputation.

While such a system is dependent on the reputation of the participants that work in academia, a commercial uptake has been limited. In the past, it became evident that cooperative computing environments need incentive mechanisms that reward contribution and punishes free-riding behaviour. Interestingly, research on incentive mechanisms in distributed systems started out in economics and computer science as separate research streams. Early pioneers in computer science used very simple incentive mechanisms in order to align individual behaviour with the socially desirable deeds. The emphasis was on the implementation of these mechanisms in running computing environments. While these studies demonstrate that it is possible to combine the principles of economics in sophisticated (Grid) middleware, it has also become evident that the mechanisms were too simple to overcome the effects of selfish individual behaviour. Interestingly, research in economics pursued a diametrically opposing approach. Abstracting from the technical details of the computing environments, were sophisticated mechanisms were developed that demonstrated desirable economic properties. However, due to the abstract nature of these mechanisms a direct implementation is not always possible.

It is, nevertheless, interesting to see that these initially different research streams have been growing together in a truly inter-disciplinary manner. While economists have improved their understanding of overall system design, many computer scientists have transformed into game theory experts. This amalgamation of research streams has produced workable solutions for addressing the incentive problems in distributed systems.

This edited book contains a compilation of the most recent developments of economic models and algorithms in distributed systems research. The papers were selected from two different workshops related to economic aspects in distributed systems, which were co-located with the IEEE Grid 2007 conference in Austin and with the ACM MardiGras 2008 conference in Baton Rouge. The extended papers from these events have been added to by projects being funded by the European Union, which in particular, address economic issues in Grid systems. As Grid computing has evolved towards the use of Cloud infrastructure, the developed economic algorithms and models can similarly be utilised in this new context – in addition to also use within Peer-2-Peer systems.

This book inevitably emphasises computing services, which look at the economic issues associated with contracting out and the delivery of computing services. At the outset of each service delivery the question arises, which service request will be accommodated at what price, or is it even provided free of charge. As these issues are spawned around business models and in particular around markets as a special kind of business model, the first chapter is devoted to the exploration of these questions. Once it has been determined, in order to resolve which service request should be accepted, a formal contract needs to be defined and mutually signed between service requester and provider. The second chapter of the book deals with aspects of service-level agreements (SLAs). One particular emphasis is on how infrastructure providers (e.g. Cloud vendors) maximise their profit, such that the Quality of Service (QoS) assertions specified in the SLA are always maintained. In the last phase of the transaction chain stands the enforcement of the SLAs. In case of detected SLA infringements (which may be by the client or the provider, but with a focus generally on the provider), penalty payments will be need to be paid by the violating provider. If the services are small-scale, it is in many cases too costly to enforce penalty payments by law. Thus, there is a need to enforce the SLAs without formal legal action; otherwise the contracts would prove to be worthless. A current practice is to establish trust among the service providers by means of reputation systems. Reputation systems embody an informal enforcement, where the SLA violators are not punished by the requester, whose SLA was breached, but by the community, which may subsequently limit use of the service offering from the respective provider. The design of reputation mechanisms is often quite difficult to undertake in practice, as it should reflect the actual potency of a provider and not be politically motivated.

http://www.springer.com/978-3-7643-8896-6

Economic Models and Algorithms for Distributed Systems
Neumann, D.; Baker, M.; Altmann, J.; Rana, O.F. (Eds.)
2010, VI, 270 p., Softcover
ISBN: 978-3-7643-8896-6
A product of Birkhäuser Basel