


# Economic Models and Algorithms for Distributed Systems

Dirk Neumann  
Mark Baker  
Jörn Altmann  
Omer F. Rana  
Editors





## AUTONOMIC SYSTEMS

---

### **Series Editors:**

Frances M.T. Brazier (VU University, Amsterdam, The Netherlands)

Omer F. Rana (Cardiff University, Cardiff, UK)

John C. Strassner (POSTECH, Pohang, South Korea)

### *Editorial Board:*

Richard Anthony (University of Greenwich, UK)

Vinny Cahill (Trinity College Dublin, Ireland)

Simon Dobson (University of St. Andrews, UK)

Joel Fleck (Hewlett-Packard, Palo Alto, USA)

José Fortes (University of Florida, USA)

Salim Hariri (University of Arizona, USA)

Jeff Kephart (IBM Thomas J. Watson Research Center, Hawthorne, USA)

Manish Parashar (Rutgers University, New Jersey, USA)

Katia Sycara (Carnegie Mellon University, Pittsburgh, USA)

Sven van der Meer (Waterford Institute of Technology, Ireland)

James Won-Ki Hong (Pohang University, South Korea)

The AUTONOMIC SYSTEMS book series provides a platform of communication between academia and industry by publishing research monographs, outstanding PhD theses, and peer-reviewed compiled contributions on the latest developments in the field of autonomic systems.

It covers a broad range of topics from the theory of autonomic systems that are researched by academia and industry. Hence, cutting-edge research, prototypical case studies, as well as industrial applications are in the focus of this book series. Fast reviewing provides a most convenient way to publish latest results in this rapid moving research area.

The topics covered by the series include (among others):

- self-\* properties in autonomic systems (e.g. self-management, self-healing)
- architectures, models, and languages for building autonomic systems
- trust, negotiation, and risk management in autonomic systems
- theoretical foundations of autonomic systems
- applications and novel computing paradigms of autonomic systems

# **Economic Models and Algorithms for Distributed Systems**

Dirk Neumann  
Mark Baker  
Jörn Altmann  
Omer F. Rana  
Editors

Birkhäuser  
Basel · Boston · Berlin

Editors:

Dirk Neumann  
Chair for Information Systems  
Kollegiengebäude II  
Platz der Alten Synagoge  
79085 Freiburg  
Germany  
e-mail: dirk.neumann@is.uni-freiburg.de

Jörn Altmann  
Technology Management, Economics  
& Policy Program  
College of Engineering  
Seoul National University  
San 56-1, Shillim-Dong,  
Gwanak-Gu, Seoul 151-742  
South Korea  
e-mail: jorn.altmann@acm.org

Mark Baker  
Research Professor of Computer Science  
ACET Centre, School of Systems Engineering  
The University of Reading  
Whiteknights, Reading  
Berkshire RG6 6AY  
UK  
e-mail: mark.baker@computer.org

Omer Rana  
School of Computer Science  
Cardiff University  
Queen's Buildings, Newport Road  
Cardiff CF24 3AA  
UK  
e-mail: o.f.rana@cs.cardiff.ac.uk

1998 ACM Computing Classification: C.2.4 [Distributed Systems]; C.2.1 [Network Architecture and Design]: Distributed networks: Network communications; C.2.3 [Network Operations]; C.4 [Performance of Systems]; H.3.4 [Systems and Software]: Distributed systems; I.2.11 [Distributed Artificial Intelligence]; K.6.4 System Management

Library of Congress Control Number: 2009931265

Bibliographic information published by Die Deutsche Bibliothek.  
Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie;  
detailed bibliographic data is available in the Internet at <http://dnb.ddb.de>

ISBN 978-3-7643-8896-6 Birkhäuser Verlag AG, Basel – Boston – Berlin

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in other ways, and storage in data banks. For any kind of use permission of the copyright owner must be obtained.

© 2010 Birkhäuser Verlag AG  
Basel · Boston · Berlin  
P.O. Box 133, CH-4010 Basel, Switzerland  
Part of Springer Science+Business Media  
Printed on acid-free paper produced from chlorine-free pulp. TCF $\infty$

ISBN 978-3-7643-8896-6

9 8 7 6 5 4 3 2 1

e-ISBN 978-3-7643-8899-7

[www.birkhauser.ch](http://www.birkhauser.ch)

# Contents

|  |   |
|--|---|
| Economic Models and Algorithms for Distributed Systems . . . . . | 1 |
|--|---|

## **Part I: Reputation Mechanisms and Trust**

|   |   |
|---|---|
| Ali Shaikh Ali and Omer F. Rana<br>A Belief-based Trust Model for Dynamic Service Selection . . . . . | 9 |
|---|---|

|   |    |
|---|----|
| Arun Anandasivam and Dirk Neumann<br>Reputation, Pricing and the E-Science Grid . . . . . | 25 |
|---|----|

|   |    |
|---|----|
| Georgia Kastidou and Robin Cohen<br>Trust-oriented Utility-based Community Structure in Multiagent<br>Systems . . . . . | 45 |
|---|----|

|   |    |
|---|----|
| Thomas E. Carroll and Daniel Grosu<br>Formation of Virtual Organizations in Grids: A Game-Theoretic<br>Approach . . . . . | 63 |
|---|----|

|  |    |
|--|----|
| Jürgen Mangler, Erich Schikuta, Christoph Witzany, Oliver Jorns,<br>Irfan Ul Haq and Helmut Wanek<br>Towards Dynamic Authentication in the Grid – Secure and Mobile<br>Business Workflows Using GSet . . . . . | 83 |
|--|----|

## **Part II: Service Level Agreements**

|  |     |
|--|-----|
| Mario Macías, Garry Smith, Omer Rana, Jordi Guitart and<br>Jordi Torres<br>Enforcing Service Level Agreements Using an Economically Enhanced<br>Resource Manager . . . . . | 109 |
|--|-----|

|   |     |
|---|-----|
| Tim Püschel, Nikolay Borissov, Dirk Neumann, Mario Macías,<br>Jordi Guitart and Jordi Torres<br>Extended Resource Management Using Client Classification and<br>Economic Enhancements . . . . . | 129 |
| Chris Smith and Aad van Moorsel<br>Mitigating Provider Uncertainty in Service Provision Contracts . . . . .   | 143 |
| Axel Tenschert, Ioannis Kotsiopoulos and Bastian Koller<br>Text-Content-Analysis based on the Syntactic Correlations between<br>Ontologies . . . . .  | 161 |
| <b>Part III: Business Models and Market Mechanisms</b>  |     |
| Ashraf Bany Mohammed, Jörn Altmann and Junseok Hwang<br>Cloud Computing Value Chains: Understanding Businesses and Value<br>Creation in the Cloud . . . . .                                     | 187 |
| In Lee<br>A Model for Determining the Optimal Capacity Investment for Utility<br>Computing . . . . .  | 209 |
| Melanie Moßmann, Jochen Stößer, Adam Ouorou, Eric Gourdin, Ruby<br>Krishnaswamy and Dirk Neumann<br>A Combinatorial Exchange for Complex Grid Services . . . . .                                | 221 |
| Christian Bodenstein<br>Heuristic Scheduling in Grid Environments: Reducing the Operational<br>Energy Demand . . . . .  | 239 |
| Raimund Matros, Werner Streitberger, Stefan Koenig and<br>Torsten Eymann<br>Facing Price Risks in Internet-of-Services Markets . . . . .  | 257 |

# 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, epitomise 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 punish 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 further use within peer-to-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.

# Part I: Reputation Mechanisms and Trust

## Reputation Mechanisms and Trust

Reputation mechanisms and trust as well as Service Level Agreements, addressed in the previous section are somewhat complementary. Whereas SLAs primarily encode contractual obligations between consumers and providers, reputation models enable choice of providers based on their past performance (assuming provider identity is persistent or traceable), or on their ability to deliver on these contractual obligations over time. Where “trust” is often defined between two participants, “reputation” often involves aggregating views from a number of different sources.

It is useful to note that when developing reputation mechanisms, not all aspects (i.e. capabilities offered by a provider) need to be considered as part of the reputation model – hence, depending on the context of usage, reputation may be calculated differently. This forms the basis for the reputation model from Ali and Rana in their chapter “*Belief-based Trust Model for Dynamic Service Selection*”, where reputation is calculated based on the particular context of use, or subjective belief of a participant. The authors attempt to combine various views on reputation and trust, depending on how these terms are perceived by a user. They subsequently demonstrate how trust may be used as a selection criterion between multiple service providers.

Anandasivam and Neumann continue this theme in their chapter “*Reputation, Pricing and the E-Science Grid*” by focusing on how the use of reputation can be used to incentivise a provider, essentially preventing such a provider from terminating a computational job from a client, even though the provider could make greater revenue by running an alternative computational job. Their work compares job submission with sites that do (and do not) use reputation mechanisms, and discuss how price determination can be associated with reputation – and present the associated decision model that may be used by market participants. Most importantly, they demonstrate that the correct use of price setting enables better collaborative interactions between participants.

The next two chapters focus on the formation of communities and virtual organizations in order to allow participants to maximise their reward (or “utility”). Kastidou and Cohen in their chapter “*Trust-oriented Utility-based Community Structure in Multiagent Systems*” discuss how better community structures could be established by allowing their participants to exchange reputation information. In this way, reputation may serve as either an incentive or a barrier to entry for an agent attempting to join another community. The focus of their work is

on the incentive mechanisms for communities to truthfully and accurately reveal reputation information, and the associated privacy concerns about disclosing such information to others. Their work is particularly relevant in open environments, as exemplified through file sharing Peer-2-Peer systems, where a decision about what files to share (upload/download) and from which participants, becomes significant.

The chapter from Carroll and Grosu entitled “*Formation of Virtual Organizations in Grids: A Game-Theoretic Approach*”, has a similar focus. They consider the formation of Virtual Organizations (VOs) which involves the aggregation of capacity from various service providers –which has a similar scope, although a different focus (on application/job execution, rather than community structure) to the notion of communities in the chapter by Kastidou and Cohen. They discuss incentive mechanisms that would enable self interested Grid Service Providers (GSPs) to come together to form such VOs using a coalitional game-theoretic framework. They demonstrate how given a deadline and a budget, VOs can form to execute particular jobs, and then dissolve. They use Myerson’s cooperation structure to achieve this, and rely on the assumption that GSPs exhibit welfare maximising behaviours when participating in a VO.

A last chapter in this section looks more at the payment issue emphasizing the business perspective of cooperative computing infrastructures. The paper “*Towards Dynamic Authentication in the Grid -Secure and Mobile Business Workflows using GSet*” by Mangler, Schikuta, Witzany, Jorns, Ul Haq and Wanek introduce the use of gSET (Gridified Secure Electronic Transaction) as a basic technology for trust management and secure accounting in cooperative computing environments.

# A Belief-based Trust Model for Dynamic Service Selection

Ali Shaikh Ali and Omer F. Rana

**Abstract.** Provision of services across institutional boundaries has become an active research area. Many such services encode access to computational and data resources (comprising single machines to computational clusters). Such services can also be informational, and integrate different resources within an institution. Consequently, we envision a service rich environment in the future, where service consumers can intelligently decide between which services to select. If interaction between service providers/users is automated, it is necessary for these service clients to be able to automatically chose between a set of equivalent (or similar) services. In such a scenario trust serves as a benchmark to differentiate between service providers. One might therefore prioritize potential cooperative partners based on the established trust. Although many approaches exist in literature about trust between online communities, the exact nature of trust for multi-institutional service sharing remains undefined. Therefore, the concept of trust suffers from an imperfect understanding, a plethora of definitions, and informal use in the literature. We present a formalism for describing trust within multi-institutional service sharing, and provide an implementation of this; enabling the agent to make trust-based decision. We evaluate our formalism through simulation.

## 1. Introduction

The existence of online services facilitates a novel form of communication between individuals and institutions, supporting flexible work patterns and making an institutional's boundaries more permeable. Upcoming standards for the description and advertisement of, as well as the interaction with and the collaboration between on-line services promise a seamless integration of business processes, applications, and online services over the Internet. As a consequence of the rapid growth of on-line services, the issue of trust becomes significant. There are no accepted

techniques or tools for specification and reasoning about trust. There is a need for a high-level, abstract way of specifying and managing trust, which can be easily integrated into applications and used on any platform. The need for a trust-based decision becomes apparent when service consumers are faced with the inevitability of selecting the *right* service in a particular context. This assumes that there is likely to be a service-rich environment (i.e. a large number of service providers) offering similar types of services. The distributed nature of these services across multiple domains and organizations, not all of which may be trusted to the same extent, makes the decision of selecting the *right* service a demanding concern, especially if the selection proves to be automated and performed by an intelligent agent.

We present a formalized approach to manage trust in online services. Our work contributes the following to the research in this field: (1) a detailed analysis of the meaning of trust and its components; (2) a trust model based on a socio-cognitive approach; (3) a trust adaptation approach; (4) an approach for service selection based on trust (using different criteria). The remainder of this article is structured as follows. First, we provide an overview of related work (Section 3.). We then present a brief overview of methodology we apply for deriving the formalism, in Section 4.. In Section 5. a discussion of the trust system and its components is presented. In Section 7. we present our approach, and the evaluate it in Section 8..

## 2. Motivations

In order to exemplify our trust formalism we will apply it to a particular scenario, based on the Faehim (Federated Analysis Environment for Heterogeneous Intelligent Mining) toolkit [8]. The aim of the Faehim project is to develop machine learning Web Services and combine them using the Triana workflow engine for Web Services composition. The scenario involves a user confronted with the inevitability of selecting a machine learning Web Service within the workflow. The potential number of suitable services is large, and services are deployed with different qualities, i.e. speed, reliability, etc. The scenario makes use of multiple such services (such as a regression technique, a clustering technique, etc). In such a scenario, the user should make a trust-based selection that enables service prioritization based on their beliefs about service quality. It is intended that the user should select a service that most matches his trust preferences or policy.

## 3. Related work

The general notion of trust is excessively complex and appears to have many different meanings depending on how it is used. There is also no consensus in the computer and information sciences literature on what trust is, although its