UNIVERSITAT POLITÈCNICA DE CATALUNYA

Departament de Llenguatges i Sistemes Informàtics
Ph.D. Program: Artificial Intelligence

Ph.D. Thesis Dissertation

Structure in Artificial Societies

Josep M. Pujol

Advisors: Dr. Ramon Sangüesa and Dr. Jordi Delgado

Barcelona, June 2006

It's better to travel hopefully than to arrive.
(Robert Louis Stevenson)

Abstract

This thesis is focused on understanding the role that structures of interactions have on multi-agent systems, which are probably the prototypical instances of artificial societies. This thesis can hopefully be read as a contribution to the research area dealing with the interdependence between a system and its components. We frame our research in systems in which the components of the system are social, regardless of the fact that components might be computational entities, as it is the case of multi-agent systems. Our aim is twofold, since we hope to study the effect that these structures of interaction between agents have both at the level of individuals and at the level of the system. The leitmotif of the research presented in this thesis is the social structure, and we address the role played by this structure from different perspectives.

First, we attend to the task of drawing implicit information embedded in the relationships between individuals. To that end, we present a several algorithms that are able to extract knowledge by means of analyzing the structure of the social network. While the first algorithm relies on the analysis of the social network to infer a reputation measure for the agents, the second one is intended to identify the underlying community structure that exists in the social network. After addressing structure as a source of knowledge we turn our attention towards the effect that certain structures - patterns of interactions - have on a system's dynamics. We study which structures favor the emergence of cooperation between agents and show that certain structures, specifically complex networks, facilitate the emergence of autonomously-agreed normative behavior — a convention. Furthermore, we show that when one convention is more beneficial than alternative conventions, the same properties of the network promote the adoption of the most desirable convention. Last but not least, we also study the process of formation of complex networks, we show that agents performing a local optimization process, grounded in sociologically plausible assumptions, can arrange themselves so that they display different structures of interactions, networks being complex one of them.

Although the focus of our research is on a particular case of artificial societies, the conclusions derived from this thesis are not limited to multi-agent systems. Our research is an inter-disciplinary approach to complex social systems. We use different methodologies borrowed from Physics, Complex Systems, Sociology, Computer Science and, of course, Artificial Intelligence in order to contribute to a better understanding of social systems in general, and multi-agent systems in particular.

Resum

La tesis aquí presentada és una contribució a l'àrea de recerca que tracta d'entendre la interacció existent entre un sistema i els seus components. En particular, el domini que ens interessa és el dels sistemes multi-agent com a prototipus de societat artificial. La nostra recerca està enmarcada en sistemes eminenment socials, encara que estiguin formats per entitats computacionals com és el cas dels sistemes multi-agent. El nostre objectiu es centra en l'estudi del rol que l'estructura de les interaccions entre els agents del sistema té tant en el comportament individual dels agents com en la dinàmica del sistema. La recerca presentada en aquesta tesis té com a fil conductor l'estudi de l'efecte produït per l'estructura social a diversos nivells.

En un primer nivell ens centrem en la extracció de la informació implícita que existeix en les relacions entre agents. Per aquest fi presentem dos algoritmes capaços d'extreure coneixement tant dels agents com del sistema mitjançant l'anàlisi de l'estructura de la xarxa social que representa les interaccions entre agents. El primer algoritme infereix una mesura de la reputació dels agents a partir de la seva posició en la xarxa social. El segon algoritme troba l'estructura de comunitats encastada en la xarxa social. Posteriorment, ens centrem en l'estudi que determinades estructures tenen sobre la dinàmica del sistema. Així doncs analitzem quines són les estructures que afavoreixen l'aparició de convencions, on tots els agents accedeixen de forma autònoma a mantenir un comportament normatiu. Veurem que determinades estructures socials, en particular algunes de les propietats que són característiques de les *xarxes complexes*, faciliten l'emergència d'aquestes convencions. També comprovarem que en el cas que una de les possibles convencions sigui més beneficiosa pel sistema com a conjunt, aquestes mateixes propietats de la xarxa afavoriran que el sistema es decanti per la convenció més favorable. Per acabar, també estudiarem com les xarxes complexes, entre altres tipus de xarxa, poden ser el resultat d'un procés d'optimització local en el que el comportament dels agents està basat en assumpcions realistes des de una perspectiva sociològica.

Tot i que el domini de la tesis sigui el dels sistemes multi-agent com a paradigma de les societats artificials, les conclusions que s'en deriven no són exclusives dels sistemes multi-agent. La nostra recerca és una aproximació multi-disciplinar als sistemes socials complexos. Així doncs, diferents metodologies manllevades de la física, sistemes complexos, sociologia, la informàtica i per suposat la intel.ligència artificial són combinades per contribuir a una millor comprensió dels sistemes socials com ara els sistemes multi-agent.

Contents

Part I

Introduction and Background

Chapter 1

Introduction

This thesis aims to provide some evidence on how the structure — pattern of interactions — between individual actors living in a society affects both individual behaviour and the society as a whole (i.e. the sum of all individuals actor plus their interactions). The motto of this thesis could be described as *"structure matters"*.

The core of the research presented in this thesis is the analysis of the interdependence between social structure and social individuals. However, we do not address the problem from a single perspective. We divide the thesis into three parts which might, at first sight, seem heterogeneous, but which in fact share a common aim.

The first part is focused on extracting knowledge about a society that does not belong to any of its particular constituents but to their mutual relationships. Quoting Kevin Kelly *"an organization's intelligence is distributed to the point of being ubiquitous"* [134]. Beyond considerations on intelligence, we would add that the same applies to knowledge. Most of the knowledge of any kind of social system, community, organization or society resides in the relationships between its constituents. Therefore, methods to extract that knowledge prove useful for a better understanding of both the individual actors and the society as a whole.

The second part of our research addresses the question of how social mechanisms, such as conventions, are affected by the structure of the interactions between individuals. A convention is a situation were all actors display certain (patterns of) behaviour. In order to comply with a given convention the behaviour of individual actors will be determined by its internal state as well as by the rest of society. However, differences in the structure of a society produce different results in the actors internal state. Certain structures can facilitate the emergence and spread of conventions. Furthermore, structure also affects the normative behaviour that is established eventually.

In the last part of our thesis we deal with the problem of how social structures can be formed by autonomous self-interested actors without centralized control.

In particular, we are interested in those social structures that have been shown to improve coordination and other interesting properties such as *robustness* and *efficiency* in information diffusion [60, 13, 63, 14].

1.1 Artificial Societies

Social science defines *society* as a network of relationships between social entities. Implicit in the meaning of society is that its members share some mutual concern or interest, a common objective or some common characteristics. As such, society is often used as synonymous to the collective citizenry of a country as directed through national institutions concerned with civic welfare.

By definition *society* covers any kind of social entity. A Multi-Agent System (MAS) can be defined as systems composed of a set of agents who interact with each other to perform complex tasks [236]. This definition clearly poses a question; do Multi-agent Systems form societies? We believe that the answer is affirmative in spite of the differences between the constituents of both systems; software agents in one case and humans in the other. [80, 170]

The level of *sociality* — the quality of being social — required by the definition of society is more restrictive than the aforementioned definition of MAS. Gilbert [98] pointedly illustrates this in his *"walking together"* example. Imagine a person is walking down the road, then another person joins her and they both walk down the road. They are not walking together, they are just walking alongside. Accordingly, if two agents interact but they do not collaborate or share a set of values of goals, their interaction is not fully social, assuming that there are different levels of sociality. Therefore, the *artificial* adjective is not strictly necessary to define society, although it is used to stress the different nature of the individual actors. Other authors who prefer not to use the term society use the term Artificial Social Systems instead. [170]

Regardless of the level of sociality in agents, they do interact with each other to carry out a complex task. As a consequence of these interactions many of the problems of sociology and multi-agent systems are closely related [33]. This is the reason why social theories are called upon as a framework in which to set the research in MAS and solve the associated research questions. MAS researchers and designers face the problem of ensuring efficiency and stability at the level of the system whilst maintaining the autonomy of the agents. Sociology, on the other hand, tries to understand how a system of autonomous entities (humans) can organize itself in such a way that it forms a fairly efficient and stable system (our organizations, communities, institutions and societies). Both disciplines look at the same problem from complementary points of view [33, 219, 43, 44].

Not only MAS researchers are concerned with applying mechanisms borrowed from social science to their systems. Actually, any open system which has constituents that poses certain degree of autonomy can benefit from this approach. Thus peer-to-peer systems (P2P) [111], web-services (WS) [162], collaborative environments, recommender systems [161], business portals are following the same steps of MAS. At the conceptual level, there is a big difference between MAS, P2P, WS and the other systems, since the individual actors of the first group are software programs, while individual actors of the second are predominantly humans, with the exception of the cases in which agent act as avatars of persons in electronic communities. Despite these differences between the nature of the individual actors of each system, social mechanisms have borrowed elements from Sociology, such as reputation, trust, norms and conventions. These elements help systems to operate properly, some examples of systems that rely on social mechanisms are e-bay [1] and Slashdot. E-bay has a simple yet effective model for reputation and trust that helps users conduct (almost) safe trades. In the case of Slashdot[2] the users are readers and censors at the same time. The users censor the comments in a collaborative way, and some conventions such *"stop the flames"* and *"don't feed the troll"* are widely accepted and adhered to without any central censor that manages it all. As a consequence, Slashdot's comments usually are of decent quality. This is the result of a system where undesirable comments are filtered out without the need for any centralized censorship.

The thesis addresses computational entities -multi-agent systems- as the components that build up artificial societies. Nevertheless, most of the results, findings and insights presented are very general. Hence, they can also be applied to other systems.

1.2 Complex Social Systems and Networks

Societies, and by extension Artificial Societies, are *complex systems*. In a society global dynamics emerge from the interactions of its parts rather than from the command of a central controller or designer. *Emergence* expresses *arising* properties at the global level that are not reducible to the properties of the individuals (micro level). Formally, a *complex system* [30] refers to a system built up of many parts which are coupled in a nonlinear fashion. Because they are nonlinear, complex systems are more than the sum of their parts since they are not subjected to the principle of superposition. A non-linear connection means that the change on one side is not proportional to a change on the other side. When non-linearities

[1] http://ebay.com it is the biggest on-line market-place in the world. With a revenue of $2.17 billion in 2003

[2] http://slashdot.org

are frequent in a system, behaviours at the global (macro) level can be as unpredictable as interesting. As a consequence, the *reductionist* (top-down) approach is not sufficient to understand this kind of systems. Since a complex system is the sum of its parts and their relations cannot be broken down into sub-systems which are easier to analyze without losing part of the essence of the system. Without denying that reductionism accounts for many achievements in the understanding of complex systems, we concur with the argument that a holistic approach is necessary to complement this view.

One of the disciplines that has devoted more effort to the understanding of a particular kind of complex system is *Sociology*. Together with other Social Sciences such as Economics, Organizational Theory and Social Psychology, Sociology has made and effort over the years to find order in societies. Many theories have been proposed in order to comprehend the complexity of society, and two main streams of research within social theory can be identified. One focuses on the macro level, studying the collective behaviour. The other focuses on the micro level, studying the individual behaviour. Both approaches are antagonist; *macro-sociology* takes individual behaviour as a mere consequence of its position or role within the group. Whereas, *individualism* takes the collective behaviour as the mere superposition of autonomous individuals' behaviour. Sociology has been the playground of the discussion between these two conceptual approaches of finding order in societies. A third approach has been brought forward in Sociology. Giddens [97] proposed an intermediate theory where the relation between both levels is circular. Changes at one level lead to changes at the other level. In turn, the first level is affected by the alteration it has brought about at the second level. Both levels influence each other. A clear example of this is an ecosystem, where its inhabitants can modify the environment but at the same time the environment shapes the individuals who live in it. There exists mutual interdependency. Ecology has shown us that there is no isolation between the system and the individuals, and that the actions of the individuals affect other individuals through the environment. We support the idea that social processes are very similar to those of ecosystems, since the interactions between autonomous individuals affect the collective, which in turn bounds their individual behaviours.

The approach suggested by Giddens was not novel in science. Many disciplines, especially those involving real systems such as Ecology, Physics or Biology have come to a similar conclusion. In fact, a cross-field discipline called *complexity* [132, 59] has devoted itself to the study of complex systems. This discipline has provided insights into many different complex physical systems. However many of these insights cannot be directly translated into social sciences. This is mainly due to the difference in the nature of individuals in physical and social complex systems. Furthermore, assumptions that can be acceptable in some systems can turn out to be unacceptable

in others. For instance, many models aiming to explain certain complex phenomena in physical systems assume homogeneous populations and trivial interaction patterns, commonly represented by complete, regular or random networks. If those limiting assumptions were to be relaxed because they are too restrictive the model would easily become analytically intractable. This is the reason why Sociology, despite studying complex social systems, does not heavily rely on this approach. The assumptions required to obtain analytical solutions to social complex processes too often turn out to be implausible and unrealistic in terms of Sociology. However, this is changing rapidly thanks to a new methodology known as *agent-based modelling* [100, 99, 80, 155], which allows to build models with plausible assumptions and non-trivial interaction patterns. It also allows to formally experiment with those models, hence obtaining some insights in systems that are too complex to be approached analytically.

Nevertheless, Sociology has not downplayed the relevance of the relations between individuals. Well before complexity proposed a holistic approach, the relationships were already studied in Sociology, for example in the field of *Social Network Analysis* [163, 106, 226]. This area departed from the traditional sociological studies that focused only on the attributes of individuals. It produced an alternative view where the attributes of individuals were less important than their relationships and ties with other members of the network. Despite focusing on the relationships rather than on the attributes of the individuals, social network analysis did not intend to fill the gap between the micro and macro levels. Its aim was to add the relational data to the already existing attribute and ideational data, so that network analysis could be incorporated to variable analysis and typological analysis already used in Sociology [200].

Social network analysis shows that it is useful to assume that the relationships between components of a system can be characterized as networks, where nodes stand for the individuals and the relations are edges between nodes. In recent years network research has experienced an impressive growth thanks to the availability of very large networks covering a wide range of systems. The spread of computer resources and the rise of its power has facilitated the compilation of very large networks and their analysis. While social network analysis has been characterized by managing relatively small networks, up to some hundred nodes, nowadays we can easily analyze networks of millions. This change in scale by several orders of magnitude provided a sound basis for analyzing the statistical properties of these networks.

The analysis of these networks has revealed common patterns and regularities despite the very different nature of the systems being modelled. Networks such as the Web [12, 4], networks of sexual contacts [150], scientific collaborations [24], protein networks [125] or phone calls [9] share topological similarities such as a characteristic

connectivity distribution (e.g. a power-law [22]), or the *small world effect* [163, 232], that corresponds to an unexpected short average path length and diameter of the network while displaying a very high clustering. However, what is most interesting about these networks is that they emerge from a wide range of systems without any central control. Complex systems tend to form non-trivial structures of interactions known as *complex networks*.

So this raises at least three questions to whose answer our thesis aims to contribute. First, how such different natural and artificial systems can self-organize to display such patterns on their topology of interactions? Second, why these systems end up forming complex networks and not other structures? And third, are the properties of these non-trivial structures beneficial to the efficiency and stability of the system?

The structural similarities between many different systems lead us to think that structure plays a crucial role in the dynamics of complex systems. In our thesis we focus on how structure can benefit artificial societies composed by computational entities. If artificial societies, i.e. societies of agents, are to become a reality, they have to share the same versatility, efficiency, stability as their natural counterparts. So, the understanding of complex systems, specially of complex social systems, is required in order to build functional artificial societies.

1.3 Objectives

The main goal of this thesis is to study the relevance of social structure in multi-agent systems and other artificial societies. As we already discussed in the introduction, it is not possible to fully comprehend a system without taking into account the relationships between its individual actors. Thus, the purpose of the thesis is to contribute to a better understanding of how the structure of a network of interactions affects both systems' dynamics and individual behaviour. Furthermore, the analysis of the relationships between agents can also provide useful information about both the system and the individuals, since some information does not belong to anyone in particular but is distributed throughout the system.

With this broad aim in mind we set ourselves the task of studying the role of structure in three different areas:

1. First, we wanted to investigate which are the social structures resulting from complex dynamics in which agents change both their behaviour and their relationships. Do agents naturally arrange themselves into particular structures as it has been found in many natural complex systems? If so, this would provide some evidence that certain topologies emerge regardless of the nature of the agents.

2. The next goal was to analyze the effects of structure on the behaviour of the system as a whole. Can certain topologies facilitate the emergence and stabilization of cooperative regimes? Can topology improve coordination among agents? Those questions are crucial for multi-agent systems; cooperation and coordination are agent-level choices in autonomous agents. However, the best course of action for an agent depends on its environment. Thus, it is plausible to assume that the structure of agents' interactions have a deep effect on the emergence and maintenance of these regimes.

3. Finally, we also want to study how structure can be used to infer social measures such as reputation, authority and group membership. Those social measures are called upon in order to reduce the inherent complexity and uncertainty of open systems. Thus, we will also address the problem of how structure can be used to extract this kind of knowledge.

Although the focus of our research is on multi-agent systems, the findings, results and insights that this thesis might provide can also be applied to other complex social systems.

1.4 Structure of the Thesis

This thesis is structured in nine chapters grouped in five parts. The first part consists of this current introduction plus background chapters. The following three parts address the three main goals of our research as stated in the previous section. The last part contains the conclusions plus the annexes. Figure 1.1 sketches the structure of the thesis.

The first part summarizes and reviews concepts extensively used through the thesis. In particular, chapter 2 focuses on individual actors of a society: social agents. A brief description and some previous work on multi-agent systems and agent-based models are briefly reviewed. Moreover, social mechanisms such as reputation, trust, norms and conventions are also described. Although all these concepts are also described in the following chapters, where our actual research is presented, chapter 2 and 3 provide the basics of the problems we address in this thesis.

In chapter 3 we present how different disciplines — with their own techniques and methodologies — approach the study of social structures and networks. First, we present how social structure affect multi-agent and collaborative systems. After reviewing the Computer Science approach, the thesis turns towards a sociologically grounded discipline that has been studying the effect of structure in human societies since the 60's: Social Network Analysis. Last but not least, we also review the work of physicists in finding non-trivial interaction patterns in real complex systems,

PART I: Introduction and Background	**Chapter 1: Introduction** **Chapter 2: Social Agents** **Chapter 3: Social Structure and Networks**
PART II: Structure as a source of Knowledge	**Chapter 4: Extracting Reputation** **Chapter 5: Finding Community Structure**
PART III: System Dynamics and Structure	**Chapter 6: Emergence of Conventions** **Chapter 7: Emergence of Efficient Social Conventions**
PART IV: Dynamics of Social Structure	**Chapter 8: Complex Networks through Local Optimization**
PART V: Conclusions	**Chapter 9: Conclusions** ***Annexes***

Figure 1.1: Thesis structure

which are known as *complex networks*. The last part of the chapter is devoted solely to complex networks since concepts related to complex networks are profusely used troughout the thesis.

The second part of the thesis is about acquiring knowledge by analyzing the structure of the network. The topology of the network contains knowledge both at the individual agent level, such as reputation, and at the society level, by finding the community structure. In chapter 4 we present an algorithm that ranks agents by analyzing their position in network of interactions. We also show how this rank can be used as a measure of reputation of individual agents. In chapter 5 we present another algorithm that is able to reveal the underlying community structure of the network by analyzing the interactions between agents.

The third part of the thesis addresses the question of how network structure affects the behaviour of its constituents. Concretely we study the relation between the emergence of a social convention and network topology. As we shall see in chapter 6, the time for a convention to emerge heavily depends on the network topology. Whereas certain social structures facilitate the coordination of agents to follow a certain behaviour, other social structures make coordination difficult to achieve. Furthermore, the effect of the underlying social structure is not limited to the convergence time, i.e. the time needed for the convention to be established, but it also affects which convention is finally established. In chapter 7 we show that in the case that two conventions are not equivalent the underlying social structure plays an important role driving the dynamics towards a particular convention.

The fourth part of the thesis focuses on the dynamics of social structure. In chapter 8 we show how complex networks can emerge from a local optimization process whose rationale is grounded in the social exchange theory. Thus, we give evidence

that several social structures observed in empirical networks can be reproduced by an agent-based model with sociologically plausible assumptions.

Finally, the last part of the thesis consists of the conclusions, which can be found in chapter 9, plus the corresponding annexes.

Chapter 2

Social Agents

This chapter comprises a brief overview of agents with an special emphasis on their social dimension. First we overview multi-agent system and argue about the need for social mechanism in order to exploit the agent paradigm to its full extent. The social mechanisms that are relevant for agents: reputation trust, norms and conventions are also overviewed, placing especial emphasis on their relation to multi-agent systems. Finally, we address agent-based models as the means for a better understanding of the complex social processes that take place in multi-agent systems and other artificial societies.

2.1 Multi-Agent Systems

Many research areas exist within multi-agent system: semantics, communication languages, ontologies and knowledge representation, reasoning, planning, learning, and so on [235]. The list is very long, as it has to, since the agent paradigm is very ambitious. In fact, many of the research areas within multi-agent systems are inherited from classical Artificial Intelligence. Social aspects of agents are no exception. Long before Jennings proposed the *Social Level* [124], which was the natural extension of the Newell's *Knowledge Level* [172], Distributed Artificial Intelligence [182] was already working on some of the problems inherited later by multi-agents systems. For instance, coordination, cooperation, competition, and conflict resolution. There are different approaches to social aspects of agents: formal logic [66, 234], cognitive approaches [57, 54], emergence and self-organization [116, 33, 170].

Wooldrige and Jennings [236, 235] define multi-agent systems as systems composed by a set of agents who interact with each other to perform complex tasks. Agents are defined as computer programs capable of taking their own decisions with no external control (autonomy), based on their perceptions of the environment and the objectives they aim to satisfy. Although there is no total agreement about what

an agent really is, probably due to the ambiguous nature of the term and its use by other areas of knowledge, there is at least some common ground in defining its properties. Following [236] an agent must be:

- *Autonomous*, which is the ability of an agent to control its actions as well as its internal state.
- *Social*, an agent must be able to interact with other agents regardless of the communication channel used (message passing or interactions through the environment).
- *Reactive*, an agent must be capable to act upon changes on the perceptions of its environment.
- *Proactive*, an agent displays a goal-oriented behaviour so that its behaviour is not only dictated by its environment and by its interactions but the agent is pursuing its own goals.

Many of the divergences over agent's definitions are due to which of those properties are emphasized the most.

2.1.1 Open and Closed Multi-Agents Systems

Most of multi-agent systems research still relies on the *benevolent agent assumption* [124] that only applies when agents have common or non-conflicting goals. Benevolent agents cooperate among them to indirectly resolve a complex task by fullfilling their individual goals, which correspond to tasks that are simpler than the one addressed by the whole system. Thus, agents carry out a cooperative problem solving. Since there is no conflict between agents' and system's goals, the sum of individual efforts results in the solution of the complex task.

The design of such systems corresponds to the classical *reductionist* approach, also referred to as *top-down*, where the original problem is decomposed into smaller and simpler sub-problems until it is possible for an agent or a group of coordinated agents to solve it. The correctness, stability and efficiency of the system can be anticipated and enforced at design time.

The *top-down* approach that has provided many of the advances in science and engineering has some drawbacks when applied to multi-agent systems. *Top-down* design requires to have all possible states, transitions and interactions accounted for at design time, therefore the system is closed, no element outside the ones that were accounted for in the design can enter the system later. This same approach, so powerful in the case of cooperative problem solving, leads to closed multi-agent systems which fail to exploit the full potential of agent paradigm [115].

The *social level* in closed systems is weak, and if one considers that communication does not necessarily imply sociality, one could argue that, in fact, closed systems are inherently *asocial*. Furthermore, closed systems do not allow interactions outside the original domain, because the interaction space must be known a priori. Closed systems often take interactions as a cost to be minimized. Even autonomy is affected in closed system, while agents still have control over their internal state, they cannot interact beyond their designed interaction space. As a result agents are not really autonomous at the social level, thus agents do not have designer autonomy [219].

For agents to exploit its full potential, they must be part of an open system, where not all factors are known at design time. It is in the open system scenario where the properties of the agent paradigm makes the difference between software agents an other software paradigms [115, 124].

Unfortunately, the design of open system is far from trivial. A *bottom-up* approach is often regarded as the better methodology for open systems. Designers must place the emphasis on building agents rather than building a system. Unlike reductionism, the *bottom-up* approach does not resort to a recursive decomposition of the original problem, but an aggregation of agents with certain skills or services instead. Thus, the goals of agents are not sub-problems of the complex task intended to revolve. Agents are granted with specialized capabilities and services which are independent of system, then, the complex task can be attained by the interaction between these agents. Perhaps, the most illustrative example of this approach is using reactive agents to solve routing and load balancing in telecommunications networks [70, 131]. This is indeed a very complex problem that can be successfully addressed with a myriad of very simple agents obeying very simple rules. The optimal, or *quasi*-optimal, route is found by no agent in particular, but by the interactions among all agents in the system.

Agents design with the *bottom-up* approach is well suited for open and distributed domains; these agents build up a set of services that can be used for several purposes [115]. This approach integrates agents in a wide-open environment, where agents can interact with other agents that where not previously known. Unfortunately, this approach also has downsides, mostly originated by the domain openness. Once agents are social, that is, once they can wander and interact in an open environment, at least two very important problems arise: 1) How can an agent know about the services offered by other agents? And 2) How agents can rely on agents who are not directly intended to interact with them? To answer the first question a lot of research in service description is being carried on, for example, by the Semantic Web community [1], creating languages such as DAML-S, RDF and OWL for service description, ontology description and defining matchmaking processes. The second

[1] World Wide Web Consortium for Semantic Web, available on http://www.w3.org/2001/sw/

question has to do with agents themselves. Some assumptions such as the benevolent agent assumption [124], in which reductionism strongly relies, are no longer met in a fully open environment.

2.1.2 Social Mechanisms in MAS

One of the characteristics of open domains is that the behaviour of the entities escape the control of the designers of the systems. There is no certainty that one agent is acting as it would be desired by the designers. The *benevolence assumption* does not apply in open systems. A system that assumes that agents will always act according to the common good but has no means to enforce the proper behaviours is doomed to fail. One might expect the agents introduced by the designers to play nice and to be trustworthy. However, this cannot be ensured for alien agents upon which designers have no control whatsoever.

Open systems usually entails a social dilemma situation[2], where agents will only look for the common good as long as it serves their particular interests.

Let us illustrate it with an example. To resolve a resource allocation problem, let us say in a logistic company, one might point to a multi-agent systems where agents would model the different roles in the logistic pipeline. Agents might well collaboratively resolve the problem by negotiating for the resources, even if it was with a competitive negotiation. In any case, agents will always work as intended, since they would always be constrained by the rules imposed by the designers. However, how can a good-behaviour be enforced when any alien agent can enter into the system at any moment?

This situation should be very familiar to the reader as human society is one of the best example of an open system. We humans do act collaboratively, not out of altruism, but because of social mechanisms that bound us to do so. Human societies have created a myriad of different mechanisms to reduce the inherent complexity and the uncertainty that exist in society. Analogously to human societies, research on multi-agent systems and artificial societies is in its way to create social mechanisms to be applied to computational entities. For instance, institutions [218, 82, 194], norms [57, 66, 219], conventions [221, 203], reputation [243, 196, 244], homophyly [111, 192, 116], and trust [245]. We will further describe some of these social mechanisms in more detail in the following section.

Even if social mechanisms studied in social sciences can be called upon for building efficient and stable artificial societies, there is still a good deal to be understood about those mechanisms before being able to apply them. Ironically, multi-agent sys-

[2]A social dilemma [62] arises when cooperation is Pareto efficient but may nevertheless fail because individuals fear being "suckered" or are tempted to exploit the willingness of others to cooperate. For a more precise game theoretic definition see [113].

tems can contribute to a better understanding of the processes occurring in complex social systems. Agent-based models or agent simulations [100, 99, 80] are a powerful methodology to gain insight into these complex systems. Thus, agent models can provide results and findings that can help to better understand complex social processes that take place in society, and this knowledge can revert to the agents in terms of more stable and efficient artificial societies.

2.1.3 Electronic Communities

Despite there are arguable differences between electronic communities and multi-agent systems, we want to remark that those communities, regardless whether they resort in agents or not, also require the use of social mechanisms of the same nature of multi-agent systems.

Electronic communities evolve around a given area of activity or topic of interest. The basis for their sustainability and persistence over time is the interchange of services or information between community members. Typical interchanges include commercial transactions[3], knowledge exchange[4], service interchange and support exchange [5] among many others. The members of this communities are usually people, specially in the case of commercial electronic communities listed in the footnotes. Research is being conducted into applying multi-agent systems to those communities, in particular, in communities created for knowledge management and organizational learning purposes [58, 133, 199, 217, 167]. Multi-agent system applied to communities usually rely on the *personal* agent concept, in which the agent is a representative of a user — a person — within the community. This agent can operate while the user is not present, for instance, gathering and filtering information tailored to the needs and interests of the user.

Perhaps, the most straightforward example of electronic community is that intended to the knowledge exchange based on collaborative systems. Although collaborative systems do not always resort to agents [3, 38, 102], they do still resort to collaboration between the users of the community. This community might well be open and the benevolence of its members be second guessed. Then, following the same argumentation as in the case of multi-agent systems one can say the electronic communities require social mechanisms. Actually, many of the systems that were classified as electronic communities are now regarded as *social software*. Although the name is rather commercial, it illustrates nonetheless the need to incorporate the social dimension into the system. Thus, it is not strange to see that most of the

[3] ebay, available on http://www.ebay.com, amazon, available on amazon.com

[4] Expert-exchange, available on http://www.experts-exchange.com/, AskMe, available on http://www.askmecorp.com/

[5] Linked In, available at https://www.linkedin.com/, Orkut http://www.orkut.com

electronic communities apply rudimentary models of reputation and social networking into their systems. Unfortunately, most of the social software, especially in the case of commercial systems, only exploit their social dimension at a very superficial level, or quoting Eagle and Pentland [75]: *today's social software isn't very social.*

By placing humans into the system many of the inherent short-comings of multi-agent systems are wizened. Important open problems such as service discovery, integration and semantics that draw much attention of the multi-agent systems research are not so acute in the case of electronic communities. However, other problems, specially those related to the benevolence assumption and the openness of the system still remain. In sociological literature cohesive social structures and trust are often seen as being almost equivalent [42]. Weber in the 1920s observed that the exchange stock market form cohesive societies where trust reaches high levels. Relevant factors that favor cohesion in communities (we refer here to communities in the *conventional* sense, such as the neighbourhood, the trade-union, the city) are distance and the high cost of communication. However, these constraints are no longer valid in electronic communities, the consolidation of the Internet has reduced the cost of communications, and distance is not as determinant as it was before. Then, how can trust emerge in a world-wide community of people that might not have any interaction in the physical world? Those electronic communities require to put the social dimension into the equation if the system is ever to be fully functional. In fact, this is already happening to a certain extent: would you buy an item from e-bay to someone with a negative reputation? Or when choosing a book from amazon, would you follow the advice of a total stranger rather than someone close to you? Artificial societies, whether they are composed of computational entities alone or composed of people interacting through the Web need to dig further into the integration of the social level into the system.

2.2 Trust and Reputation

As mentioned in the previous sections, we must resort to social mechanisms such as trust and reputation in order to reduce the uncertainty and complexity of open systems where the benevolence assumption does not hold. Social mechanisms help agents in their decision making process, and are intended to help the system maintain a cooperative and coordinated regime. Needless to say that reputation and trust has been studied throughfully in social sciences [191, 144, 42]. However, the profiferation of open artificial societies extended the study of these social mechanisms to the area of computation.

Trust

Gambetta [95] defined trust as a particular level of subjective probability with which an agent will perform a particular action before the action is performed. Trust between two agents is built on the outcome of interactions, usually *positive* interactions in which both agents benefit from the interaction. Certain level of trust is required to maintain cooperative regimes under non-ideal circumstances such as those found in open systems. A straightforward example would be a *prisoners' dilemma* scenario: the ideal strategy for that game, as shown by Axelrod [16], is *tit-for-tat*, which is based on direct reciprocity. Under the tit-for-tat strategy, if an agent's counterpart played nice the agent will reciprocate, otherwise the agent will defect. Thus, after a mutually positive interaction the agents build up trust to continue cooperation until the trust is betrayed by a defection. However, in this example the concept of trust is very weak since it only depends on the last interaction. As a consequence, random errors would break down cooperation. If an agent would have had a genuine trust with its counterpart it would wait some rounds before changing to defection. For example, it would spare a certain number of defections believing that its counterpart was not actually defecting, but defections were due to an accident. The version of tit-for-tat that incorporates trust, albeit at a basic level, is known as generous tit-for-tat or *tit-for-two-tats*. This strategy forgives the first defection assuming it was an accident. Tit-for-two-tats (trust) clearly outperforms tit-for-tat (no trust, just reciprocity) in a noisy environment.

Reputation

According to Misztal [165] reputation helps us to manage the complexity of social life by singling out trustworthy people in whose interests it is to meet promises. Actually, reputation, if deserved, reduces complexity in many ways: 1) it reduces the number of agents worth interacting with. So, reputation can be used as a filter. 2) Reputation ensures, to a certain extent, a positive interaction. 3) Reputation promotes cooperation under the menace of getting a negative review.

Although reputation and trust are seen as critical factors in order to manage complexity, there is no general and cross-domain model for reputation. The advantages described by Misztal are not free from drawbacks. For instance, the filtering induced by reputation can bring the system to a sub-optimal state where latecomer agents, despite offering better services or better reliability, have no visibility. To avoid this problem, one might suggest that reputation should only report negative interactions so that agents have a fair starting point, but then the problem of *cheap* identities arises. If changing your identity is affordable, free-riders with a hit-and-run strategy proliferate [93]. They would keep defecting until their reputation were compromised and then they would change their identity for a brand-new one ready

for a fresh new round of cheating.

Yet another drawback is that the measure of the reputation could be tampered, specially when the reputation model resorts to explicit feedback to estimate the reputation of a given agent. Reputation is subjected to possible *collusions* of a group of agents trying to undermine the reputation model of the system or the reputation of a particular agent (probably a competitor). Another type of collusion is the artificial increase of reputation of one agent who has a group of cooperator agents that report inexistent positive interactions with it. In doing so, the agent would appear as a trustworthy agent to the whole community although it only interacted with its accomplices [238, 65, 164, 244]. Yet another problem, especially acute when the reputation model is completely distributed, is *correlated evidences*. This occurs when the same interaction is reported several times from different sources [196], as a consequence, the reputation measure is altered by accounting the same interaction more than once.

Most reputation and trust models used in multi-agent systems [243, 244, 245, 196, 167] are devoted to the domain of electronic commerce, where trust and reputation are basic since business need to be run in a trustworthy environment. Another domain where reputation models are usually applied is in peer-to-peer file-sharing systems, where reputation and trust are used to minimize free-riding and to avoid profiferation of fake files [238, 128, 111]. Another example of a domain where trust and reputation has been used is in recommender and collaborative systems, where reputation usually refers to people or to agents that act on their behalf in an electronic community. Thus, reputation allows agents to filter the good from the bad, finding the reliable and knowledgeable people, agents or pieces of information in the system [130, 133, 1].

Many models have been suggested to implement the concepts of reputation and trust in on-line systems, multi-agent systems and peer-to-peer systems. Sabater and Sierra [197] did a comprehensive review of computational models of trust and reputation from an agent-based perspective.

Perhaps the most wounding criticism to trust and reputation research is the proliferation of ad-hoc models which are weakly grounded in social theory. Most reputation models address a particular scenario of a given domain focusing on the particularities of the domain rather than addressing the underlying concept of reputation. We would like to say that although we also would agree in the lack of social theory and cognitive grounding of most models, we do not believe it is possible to create a general cross-domain reputation model, as it is not possible in human societies either. Another criticism to trust and reputation models is the lack of integration between models, making comparison between different models unfeasible. Researchers in the trust and reputation area are already addressing this issue, for instance by creating the ART competition [94].

A common trait shared by most reputation models is that they rely on some form of feedback collection after an interaction, for example, in order to rate satisfaction after transactions in e-commerce environments or to rate quality and usefulness of recommendation in a recommender system. Besides the fact that feedback might be false or unfair, there are more problems related to feedback: a) it is sometimes difficult to evaluate the quality of an action or a piece of information, and b) it is too much dependent on collaboration, the lack of which is a known problem in collaborative environments [38]. All these problems, which are a lesson from the recommender and collaborative systems, are too often disregarded in reputation and trust models, with some exceptions [123].

We present in chapter 4 a mechanism that does not rely on feedback but on the position of an agent within its social network. Therefore, a reputation measure can be inferred by analyzing the topology of the network of interactions between the agents. Agents who are well known and highly regarded by the rest of agents of the community can be easily identified as highly connected nodes in the social network, in a similar fashion as relevant Web pages are identified by analyzing the structure of the Web [183, 137]. This relational information could serve as the basis for a ranking mechanism instead of having to resort to feedback or ratings provided by individual agents.

2.3 Norms and Conventions

Other social mechanisms borrowed from sociology are norms and conventions [214]. In fact, these concepts do not appear only in sociology but they are used and studied in other social sciences such as legal theory, economics, psychology, philosophy and decision theory. Different disciplines came up with different points of view about norms, Verhagen [219] provided a comprehensive description of norms from different perspectives.

Analogously to trust and reputation, multi-agent systems researchers as well as other fields of distributed artificial intelligence have investigated the feasibility of the use of norms, conventions and laws in multi-agents systems in order to simplify the agent decision-making process and to improve the effectiveness and performance of the overall system [221].

Social norms are usually meant as solutions to problems of cooperation which are rooted in the conflict between common and individual good. Multi-agent systems are no different since they need to coordinate and/or cooperate in order to work as intended, resolving the conflicts of interests among agents [57].

Even within multi-agent research there are different views on norms depending on the methodological approach followed. The formal logic approaches intend to define norms using some modal logic, such as deontic logic [66, 218]. The approaches from

cognitive science focuses on the internalization of the concept of norm into the agents rational process [57, 54, 48]. There are also approaches to social norms grounded on game-theory, in which norms are often seen as the maintenance of a cooperative regime in a social dilemma situation [214, 16, 17, 27]. Despite divergences, there is a baseline agreement: *a norm is a restriction on the set of actions available to agents.* When a social norm is restricted to one particular action it is called a convention [203]. Walker [221] defined norms and conventions as behavioural constraints. Conte and Castelfranchi [54] agreed on the interpretation of norms as behavioural constraints, although they also considered norms as ends — or goals — that agents try to fulfill.

In this thesis we restrict our work to a computational study of social norms and conventions. This approach is focused on studying the dynamics of a system based on coordination, cooperation or evolutionary games by means of agent-based simulations. These simulations show how the performance of the overall system is improved by the adoption of norms and conventions. We connect with a tradition of studies such as the control of aggression [55], the reduction of inequality among agents [195] or the facilitation of emergence of cooperation in the prisoner's dilemma [51] to cite only a few.

In most of these studies, the effects of norms have been shown to act as mechanisms that improve cooperation among agents within a system. However, there are still open issues as to how norms spread. Usually the diffusion of norms relies on the imitation of successful strategies [51]. However, in real life, norms do not only spread through imitation. As the advocates for cognitive-grounded norms would put it: diffusion is also due to recognition of norms as such by people (or agents) and an active defense of norms on the part of norm-observers [48]. Thus, norms are not only to be followed by agents but these agents should urge other agents to do so, introducing a second order choice; that is whether to comply with a norm or not and whether to urge other agents to do so. Apparently, under some settings, the more normative behaviour there is in the system the more successful the system is [48].

Another key issue in the understanding of norms is to decide on the method by which they can come to exist within a society. There are two antagonistic approaches. *Off-line* design assumes that social norms are hard-coded into the agents [55, 201]. The counterpart relies on spreading and adopting processes that may eventually lead to a normative behaviour, that is the *emergence* of a norm [221, 51, 63, 135, 89].

Obviously, the *off-line design* is not suited for open systems, in which the compliance of all the agents to the norms cannot be guaranteed by design. We are more interested in the second approach, where emergence of norms and conventions is the key. In this setting, agents should be able to choose the appropriate behaviour

despite the fact that its own decision making process is based upon imperfect and local information. By doing so, the complexity is transferred to the behaviour — strategy — update function [221]. Most research on computational study of norms has focused on conventions rather than on norms, since coordination games, where there are no *social dilemmas* [62], are simpler than cooperation games. As for norms, there is still much work to be done to explain how they do emerge, become established, internalized and eventually change [195]. Vergahen [219] did an interesting work about the internalization and spreading of norms; however emergence was not addressed. It is understandable that emergence of norms is not a main issue within the MAS community since there is still much research to be done on their representation, diffusion and specially on their compliance.

The role played by the topology of the underlying social network of a multi-agent system in the emergence of a norms and conventions has not been studied in depth either. In the nineties, some researchers such as Kittock [135], Shoham and Tennenholtz [203] and Cohen et al. [51] pointed out that topology was a key factor for the efficiency in the emergence of conventions. However, it was not until some years later that the effect of topology attracted the interest of more researchers working in norms and conventions [2, 63]. The tipping point was the finding that components of real systems exhibit non-trivial interaction patterns, which could be modelled as a network. This field known as *complex networks* [11, 232, 4] introduced a new class of networks that had properties that were not present in the idealized networks —- random, complete or regular networks — used to model the pattern of interaction between agents.

The contribution of this thesis goes into this strand of research. In chapter 6 we show how certain patterns of interactions among agents influence the time for a convention to be established. Moreover, in chapter 7, we show that the topology of the interactions also affects which convention is finally agreed upon in the case of dissimilar conventions.

2.4 Agent Based Modelling

So far we have presented how research in multi-agent system has turned to social science and borrowed concepts such as trust, reputation, norms and conventions. Those concepts are needed to design and build open systems where agents are not bounded by the benevolence assumption, thus, exploiting the multi-agent paradigm to its full potential. At the same time, researchers in social science have turned to the multi-agent systems paradigm as a powerful approach to be applied to problems involving complex social dynamics such as emergent social norms, social structure and social change.

Traditionally, sociologists have tried to understand social life as a structured

system of institutions and norms that shape individual behaviour from the top-down. In contrast, some sociologists suspect that much of social life emerges from the bottom-up [87, 18, 155, 158, 99, 56, 19, 79, 45]. To that purpose, sociologists have adopted an alternative modelling tool which was first developed by computer science and artificial intelligence: Agent Based Modeling (ABM). Software-based agents systems try to solve complex tasks by using a set of autonomous agents. Once you remove the word software there are many similarities between multi-agent systems and societies of humans. Moreover, the agent paradigm also imposes certain constraints such as imperfect and local information as well as bounded rationality [205]. Those properties are somehow seen in human nature, so it is no wonder that sociologists have taken this approach to be an interesting field in order to gain further understanding of complex social dynamics.

Although the main goals of both research communities differ, there is common ground for a cross-fertilization of both disciplines. Multi-agent researchers are applying social mechanisms in order to improve the stability and performance of their systems. Sociologists are applying the agent paradigm as a framework to study and to explain certain social phenomena.

Epstein [80] proposes agent based simulation for explaining macroscopic social patterns by generating — or as he says *"growing"* — them in agent models. According to him, the main contribution of agent based systems into social science is to facilitate explanation. Agent based modeling is not the only framework to tackle complex social problems but it is probably the latest and it is gaining momentum thanks to its explanatory power [79, 19]. Game theory, evolutionary game theory, mathematical economics, rational choice theory, political science, theories of learning in strategic interaction, socio-physics and statistical models have been continuously improving and being applied to wider domains.

However, those methods have in their analytical foundation its strongest and weakest point. The analytical approach offers sound formal proof but the price to pay is that the domain is rather limited due to the imposed constraints necessary to achieve the analytical solution. Thus, analytical models yet beautiful, are difficult to be applied in realistic social systems. For instance, a completely analytical game theory approach requires to impose unrealistic patterns of interactions so that the formal proof, i.e. to find the Nash equilibrium (or a distinguished Nash equilibrium) of some game, can be derived. Those analytical demonstrations are only possible in confined domains and imposing strong limiting assumptions which do not correspond to the complex social subject to study. Therefore, their explanatory power, which is unarguable on paper, is diminished when confronted with reality, where most of the assumptions are not meet. Moreover, experimental results often do not match the results predicted by the analytical solution revealing that the part of the problem left out by the assumption was indeed relevant [42].

Epstein, in [78], goes even further proposing agent-based models as a new paradigm: *"...it is no sufficient to demonstrate that if a society of rational (homo economicus) agents were places in the pattern, no individual would unilaterally depart, which is the Nash equilibrium condition. Rather, to explain a pattern, one must show how a population of cognitively plausible agents, interacting under plausible rules, could actually arrive at the pattern on time scales of interest".*

Agent based models using computational simulations are flourishing due both to limitations of formal and analytical approaches and to the rise of computing power. Furthermore, agent based modeling is a formal method that educes in a systematic and rigorous way macroscopic implications (at the level of a society) from assumptions about the microscopic level (the behaviour of the agent).

Some researchers, specially those fond of mathematical *"hard" science*, are skeptical about agent based models. Some of the main criticisms are enumerated and argued in [78]. Epistemological concerns of agent-based modelling are also addressed in [26, 171, 91]. Most of the criticisms, however, are not aimed at the agent-based modeling framework itself. As every discipline in its early stages, agent-based modelling is affected by certain degree of unrigorousness by some of its proponents. To deal with those problems, which are not due to the framework itself but to some abuses of it, prominent researches on the field have proposed a set of methodological principles [87] as well as some books of best practice [100].

2.4.1 Rigorousness of Agent Based Computational Models

Flache and Macy [87] proposed a set of methodological principles. Some of those principles are summarized below.

- The model must be grounded on existing sociological theory. Simulation provide a rigorous methodology for studying effect of different microfoundations on macro dynamics. The model must not be based on observed phenomena but from well grounded theory. Then, two important related pitfalls — often used against agent based models — can be avoided: 1) the desired empirical patterns can be generated by a large set of different models, therefore, it is possible to create a model that generates any empirical pattern. 2) if a model generates the desired empirical pattern but the mechanism responsible remains unknown (without theoretical backup), therefore, the model does not have any explanatory power.

- The model should be as simple as possible. Pressure to make models more realistic (and agents more cognitively sophisticated) can become as hard to interpret as the natural phenomena they try to explain.

- Experimentation should be theoretically motivated. The space of parameters must be explored carefully and systematically. Furthermore, the manipulation of parameters must be theoretically motivated.

- The results should be replicated independently. In order to guarantee the validity of the model its implementation should be replicated, and systematically checked. This is an important step since the models whose replication have been shown to produce different results are often used as an example against the rigourousness of agent based models. As a defense against these criticism it would suffice to say that the author of the experiment whose results are not replicable is the one to be blamed, not the experimental tool.

- The results must be robust. The statistical robustness and significance of the results must be ensured.

- Compare and align models. When possible, the model should be compared to similar models in the literature. This would help to validate the replicability of models already published. Furthermore, alignment would help to identify those assumptions that cause differences in the behaviour of the models. Ideally, this would lead to the theoretical integration of competing models and, eventually, to a more general model.

The research presented in this thesis is partially based in agent-based models and simulations (chapters 6 to 8). A close inspection of the models reveals that they meet the methodological principles proposed by Flache and Macy. We specially stressed the independent replication of the results yielded by our models by making our models available to other researchers as well as by including the pseudo-code of the models (see appendix A).

Chapter 3

Social Structures and Networks

In the previous chapter the discussion has focused on the agent level, the individuals of artificial societies. In this chapter the focus is on the social structure where the agents are embedded. We will briefly discuss the approach to social structure from three disciplines: computer science, sociology and physics. Later on we will present some concepts on networks which will be used extensively in the forthcoming chapters.

3.1 Social Network Analysis

A social network is a map of all relevant ties between actors in a group [226, 225, 200, 233]. Nodes are the individual actors and ties - edges - are the social relationships between the actors. These relations can express a wide range of interactions: friendship, acquaintanceship, trading relations, support exchange, collaboration and many others.

Modern social networks analysis (SNA) dates from the 60's and it is the result of the convergence of earlier traditions, including sociometric analysis, group dynamics [169] and social anthropology [25]. These trends took a social approach to analysis of their particular domains: they transformed interactions and, in particular, the structure of those interactions into explanatory mechanisms that accounted for the processes taking place in the society or the group under examination. By doing so they brought together mathematics - especially graph theory - and social theory. After some years, a well-developed methodology of social network analysis was consolidated. Scott in chapter 2 of [200] and Wasserman and Faust in chapter 1 of [225] review how that development took place.

Social network analysis studies the influence of the social structure on social order. Its practitioners departed from traditional sociological studies to which the only thing that matters is the attributes of individual actors. Social network analysis

produces an alternate view where the focus is placed on the relationships and ties between individuals rather than on the attributes of individuals themselves. Social network theory has come a long way in explaining social processes that classical sociology failed to explain.

For instance, influence within organizations depends more on the centrality of the individual rather than its actual job title [107]. Centrality[1]; it measures how influential individuals are in, for example, the transmission of information, and it depends on being connected to individuals who are influential as well.

Weak ties, on the other hand, are those links that connect people socially distant, that is, individuals who belong to different social circles and therefore manage different types of information. Granovetter found that weak ties [106] gave access to many useful resources for actors involved in processes such as job hunting: information on new jobs is better obtained from sources who do not belong to the same social group and therefore do not manage the same redundant information. The mechanism brought by weak ties is alternative, yet complementary, to that of betweenness.

Related to this, success for firms or in job performance also depends heavily on the relationships between actors that belong to different departments in the organizational structure. If an individual has relationships in two different groups he has access to resources which are not accessible for an actor with relationships in a single group. Therefore, individuals with more diverse relationships have a competitive advantage when compared to those embedded in more homogeneous groups [41].

Social capital [151] focuses precisely on this phenomenon. Social capital comprises the pool of resources available to an individual through its social network. The concept is defined in opposition to human capital, which is the amount of resources acquired by an individual in the course of training and experience. Social capital is not embedded in individuals but in their connections. And, being a form of capital, it provides returns that would be unattainable in its absence. This is what folk-wisdom calls *networking*. Social capital affects both individuals and groups. Putman [189], for example, related social capital to the health of civic movements and democracy. Even if this work has arisen some methodological concerns, it provides a good argument about the role networks play in facilitating certain outcomes that are crucial for the workings of society [103].

Social network theory has also been applied to innovation diffusion [215]. The structure of the social network provides insights into how new knowledge or trends will spread through a society and how fast. If members of a society form a close-

[1]There are many definitions of centrality, although the most common is the betweenness centrality [92]. Betweenness centrality of node i is formally defined as the number of shortest paths between any two nodes that contains node i.

knit group, with a few members who belong to more than one group, then the society is closed and innovation spreads very slowly. Conversely, if society contains many individuals who act as bridges between groups, innovation and knowledge will spread much faster. Thus, social network analysis tells us that open societies favor innovation. However, it also warn us that too many relations outside the group might eventually destroy the concept of social group itself, and as a consequence, endanger cooperation between members due to lack of trust [42, 41].

There are two main streams of research within SNA. One of them is based on empirical research, where information is mostly gathered through questionnaires and interviews which are analyzed by using social network analysis methods. This field-study oriented approach is very common due to the proximity of SNA to classical sociological studies. However, networks obtained by this approach tend to be small, since collecting data on bigger networks would require a lot of resources (time and money). As a consequence, the results of this approach do not confront macro properties of the networks and how they could affect the actor's behaviour. Instead, they focus on egocentric networks, where society is perceived as a composition of relations of many individual actors. This approach provides empirical results that either back up or question social theory, especially when longitudinal data is collected to better understand the evolution of a particular social group or system.

The other main stream within SNA is focused on methodology and analytical models. Unfortunately, most standard statistical methods cannot be directly applied to SNA because of the existence of mutual dependency between actors and the ties in a social network. The creation of statistically sound methods forms an important part of the research conducted by social network researchers [206, 207, 68].

Because of the statistical complexity involved in the analysis of the data and the difficulty to get large datasets, most social network studies end up as comparisons between different static networks and can only be applied in descriptive analysis. Stokman and Doreian stated in [211, 69] that important questions such as evolution and dynamics were mainly neglected in SNA. Nevertheless, this is no longer true since recent work in SNA is just addressing the evolution and dynamics within social networks [215, 193, 208, 210]. Addressing these issues is a necessary step towards a better understanding of the emergence of social order and of the interdependency of the individual actor and society as a whole. However, some of the work on dynamics and evolution of social networks rely on analytical approaches, mainly statistics, which suffers from the same drawbacks as game-theoretical approaches (see section 1.2): In order to obtain an analytical solution strong assumptions or limited domains have to be added into the model. Thus, unlike economics or decision theory based on game theory, most of the research on complex systems dynamics cannot be directly applied to SNA.

As we hinted in section 2.4, agent based models might offer a solution to over-

come these problems. Physics is also contributing to social network research by applying their own methodology for complex systems to very large social networks. Unfortunately, the core of social network analysis is not very fond[2] neither of the physics approach to social networks - sometimes refereed to as *socio-physics* - nor of the agent based models.

3.2 Patterns in Networks

The use of networks as a representation of the interactions between the entities of a system is nothing new. For instance, Biology [168, 187] has extensively used food-webs to better understand ecosystems. Networks of firms are commonly studied in Economics and Management [20, 222, 140]. Organizational theory [166, 8, 46] has used networks - mostly hierarchies - to understand the processes occurring within a company. Sociology, together with other social sciences, developed social network analysis as a framework to explain social processes by the relationships between people. There are many examples in different disciplines where networks are used as a means to represent the system subjected to study.

The networks studied, however, were usually very small networks, ranging from a dozens to a few hundreds nodes maximum. As a consequence, statistical analysis of networks was limited to small scale networks. Furthermore, there was little intuition of the fact that networks modelling different systems might exhibit non-trivial statistical regularities. Therefore cross-domain studies on networks did not exist, instead each discipline focused on their particular set of interests.

Nevertheless, during the late nineties the availability of huge networks, such as the Web or the Internet, together with the rise of computing power lead to a collective effort to gather and analyze networks from very different domains. The research on networks underwent a change and expanded to statistical analysis of large-scale networks. Using methodologies borrowed from Physics such as statistical mechanics, which had been proven to be very useful.

Many distinct systems were modelled as networks. Their analysis unveiled striking regularities in network properties despite of their different nature. Tables 3.1 and 3.2 summarize the characteristics of several empirical networks from very different systems.

Most analyzed networks, regardless of the system they modelled, exhibited surprising invariant characteristics. For instance, the average path length was surprisingly low, networks were formed from cliques of highly interconnected nodes, the distribution of connectivity had the so called long tail effect. This long tail is a

[2]To further illustrate this heated debate we recommend to visit the blog entry in Crooked Timber regarding this issue. http://crookedtimber.org/2005/05/19/isolated-social-networkers/

situation in which a few nodes had a high connectivity while most nodes had a relatively low connectivity. The characteristics found in empirical networks endow the networks with very interesting properties such as efficiency in propagation, community structure, high redundancy, resilience against random errors, and many others which will be discussed in section 3.2.1.

The question of why such different systems display very similar interaction patterns attracted the interest of many researchers. What do systems composed by web pages, proteins in living organisms, scientific communities, routers on the Internet, movie actors, phone calls, trophic relations in an ecosystem, firms in the stock exchange market or neurons in the brain, have in common? The systems are very different in nature by many parameters: social and non-social systems, artificial (human made) and natural systems, cognitive and non-cognitives entities (here we refer to the individuals the system is composed of).

However, in spite of the differences in nature of these systems, the networks that modeled their interactions exhibited some strikingly similar patterns. They all displayed non-trivial wiring schemes that have been given the name of *complex networks*.

It is even more surprising the fact that these patterns have a direct effect, mostly positive, on very important macroscopic processes taking place in these systems, such as diffusion and robustness. Systems have somehow organized themselves to a setting that contributes to several interesting properties, from the efficiency of information diffusion to the resilience against random errors. Furthermore, this phenomenon is observed in many different systems that have nothing in common besides the fact that they are build up of autonomous entities without any central planner or designer.

Aside from the hundreds of research papers and books, there are several books targeting a more general audience [230, 40, 21, 119]. Such books are always an indication of ground-breaking findings and stipulate the relevance of the field.

3.2.1 Summary of Complex Networks

This section is a brief overview of the main characteristics, properties and models of complex networks which are discussed in the rest of this thesis. For a more comprehensive review we strongly suggest any of these excellent reviews and books about complex networks [11, 175, 71, 186].

Characteristics

These are the main characteristics taken into account when studying complex networks,

- *Degree distribution.* The degree of a node i is the number of connections (edges) of that node (k_i). Degree is often called connectivity. In the case of directed networks we must distinguish between in-degree and out-degree. The first being the number of connections pointing towards the node and the later being the number of connections that have their origin in the node. Not all the nodes have the same degree. The spread of degrees is characterized by a distribution function $P(k)$. For complex networks the connectivity distribution can be classified into two groups: 1) the degree distribution is a power-law: $P(k) \sim k^{-\gamma}$, (such networks are called scale free [22]). Or 2) the degree distribution is an exponential: $P(k) \sim Ae^{-k}$. In both cases the frequency per degree k decreases rapidly. In the later case the probability of having nodes with an extremely large degree is not negligible. Furthermore, power-law distributions can display an exponential cut-off due to finite size effects [181, 186]. See figure 3.1 for a graphical representation of different connectivity distributions.

- *Diameter* (d) is the maximal distance between any pair of nodes. Distance is the number of edges of the path between two nodes, roughly speaking, the number of hops to get from node i to node j. In the case that the network is a disconnected one (a network is made up of several isolated clusters, also referred to as connected components) the diameter is often defined as the maximum diameter of its clusters. However, in strict terms the diameter of a disconnected network is infinite.

- *Average path length* (l), also known as *characteristic path length*, is the average distance between any pair of nodes. For random networks the average path length as well as the diameter (since it is the length of the longest path) scales logarithmically with the size. Formally, $l_{rand} \sim \frac{ln(N)}{ln(\langle k \rangle)}$. This implies that even for very large networks we expect to find a path between any two pairs of nodes of reasonable length. Networks where the average path scales more slowly than the actual size for the network are required for an efficient communication. Complex networks also scale logarithmically with the size of the network.

- *Clustering coefficient* (c) measures the cliquishness of a network. The network clustering coefficient is the average clustering coefficient of all nodes of the networks, $c = \sum_1^N c_i$ and $c_i = \frac{2E_i}{k_i(k_i-1)}$, where i is a node having a connectivity k_i, and E_i is the number of edges that exist between the nodes that have an edge with i. Roughly speaking, this coefficient is the ratio between the number of edges contained in the set composed of the neighbours of a node and all the possible edges that that set could possible contain. The clustering coefficient

ranges from zero to one and can be intuitively be seen as the probability that edge jz exists provided that edges ij and iz exist. Thus, the clustering coefficient is a measure of transitivity (or triangular closure) among the nodes of the network. When the clustering coefficient is high, which is the case for complex networks, it means that a community structure exists. There are groups of nodes that are densely connected with members of the group forming a clique.

- *Degree correlations* measures how nodes are connected with each other in terms of their degree. Correlations can be either positive or negative. If correlation value is close to 0 there exist no correlation. A negative correlation means that *high-degree* nodes tend to connect to *low-degree nodes* and vice versa, hence, the mixing pattern of nodes is dissortative. On the other hand, when correlation is positive *high-degree* nodes tend to connect to *high-degree* nodes, and *low-degree* nodes tend to connect to *low-degree* nodes. Therefore, positive correlation implies that nodes tend to pair up to nodes of the same kind, forming an assortative mixing pattern. Pastor-Satorras et al [184] proposed to measure degree correlation based upon the average degree of the neighbours of a node as a function of the degree k of the given node. Newman [174] proposed an even more compact measurement, called *assortativeness*, based on calculating the Pearson correlation coefficient of the degrees at both ends of the edges [174]. Degree correlation has been found in many empirical complex networks. The striking fact about correlation is that only those complex networks that are also social networks display a positive correlation. This leads to think that formation of social networks differs from formation of other complex networks since it is governed by an assortative mixing process.

Tables 3.1 and 3.2 summarize the characteristics of several empirical complex networks. These tables have been extracted from Albert and Barabási [11]. Table 3.1 is focused on the average path length and the clustering coefficient of empirical complex networks compared to those obtained from random networks.

Table 3.2 focuses on the connectivity distribution of the empirical complex networks, which follows a power-law of different exponents depending on the network. Table 3.2 also summarizes the average path length found in the empirical network compared to the expected average path length of a random network or an analytical power-law network.

Properties of Complex Networks

Complex networks exhibit some properties that are remarkable, specially if one takes into account that these networks are the results of the self-organization of a myriad

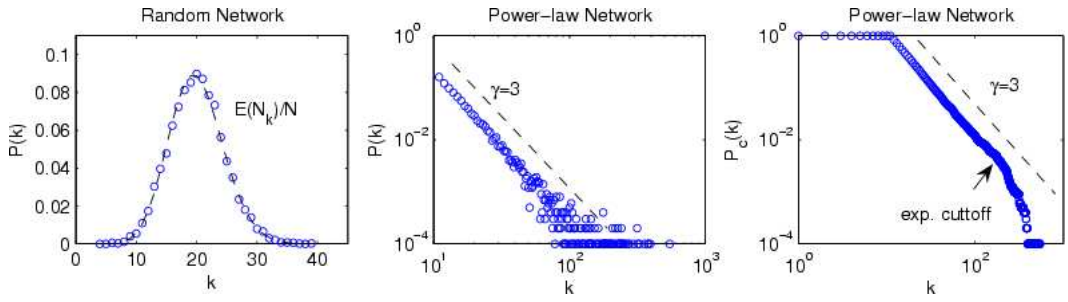

Figure 3.1: Comparison between connectivity distribution and analytical predictions. The left sub-figure corresponds to the connectivity distribution of a random network of 10000 nodes and connection probability $p = 0.002$ using the Erdös-Rény model [81]. The average connectivity degree is $\langle k \rangle = 20$. The plot compares the normalized frequency of nodes with connectivity $\frac{N_k}{N}$ with the expected value of a Poisson distribution $\frac{E(N_k)}{N}$ (dashed line). The central figure corresponds to a power-law network generated with the Barabási-Albert model [22]. The network has the same number of nodes and edges as in the case of the previous random network. We can observer the differences of connectivity distribution between random and power-law networks. The dashed line correspond to the analytical prediction $P(k) \sim k^{-3}$. The right sub-figure is the cumulative connectivity distribution $P_c(k)$ of the central sub-figure. We can observe how the potential distribution has an exponential cut-off due to finite size effect

Table 3.1: Characteristics of several empirical complex networks. Size is the number of nodes of the graph, $\langle k \rangle$ is the average node degree, l is the average path length, and C is the clustering coefficient. L_{rand} and C_{rand} are the average path length and the clustering coefficient of a random graph of the same size and the same average degree. After Albert and Barabási [11].

Network	*Size*	$\langle k \rangle$	l	l_{rand}	C	C_{rand}
WWW, site level, undir.	153127	35.21	3.1	3.35	0.1078	0.00023
Internet, domain level	3015-6209	3.52-4.11	3.7-3.76	6.36-6.18	0.18-0.3	0.001
Movie actors	225226	61	3.65	2.99	0.79	0.00027
LANL co-authorship	52909	9.7	5.9	4.79	0.43	1.8×10^{-4}
MEDLINE co-authorship	1520251	18.1	4.6	4.91	0.066	1.1×10^{-5}
SPIRES co-authorship	56627	173	4.0	2.12	0.726	0.003
NCSTRL co-authorship	11994	3.59	9.7	7.34	0.496	3×10^{-4}
Math. co-authorship	70975	3.9	9.5	8.2	0.59	5.4×10^{-5}
Neurosci. co-authorship.	209293	11.5	6	5.01	0.76	5.5×10^{-5}
E. coli, substrate graph	282	7.35	2.9	3.04	0.32	0.026
E. coli, reaction graph	315	28.3	2.62	1.98	0.59	0.09
Ythan estuary food web	134	8.7	2.43	2.26	0.22	0.06
Silwood Park food web	154	4.75	3.40	3.23	0.15	0.03
Words, co-occurence	460902	70.13	2.67	3.03	0.437	0.0001
Words, synonyms	22311	13.48	4.5	3.84	0.7	0.0006
Power grid	4941	2.67	18.7	12.4	0.08	0.005
C. Elegans	282	14	2.65	2.25	0.28	0.05

Table 3.2: Scaling exponents characterizing the degree distribution of several scale-free networks. Size is the number of nodes of the graph, $\langle k \rangle$ is the average degree, κ is the cut-off for the power-law scaling. For directed networks, the indegree γ_{in} and the outdegree γ_{out} has been listed separately. The columns l_{real}, l_{rand}, and l_{pow} compare the average path length of real networks with power-law degree distribution and the predictions of random-graph theory. After Albert and Barabási [11].

Network	*Size*	$\langle k \rangle$	κ	γ_{out}	γ_{in}	l_{real}	l_{rand}	l_{pow}
WWW	325729	4.51	900	2.45	2.1	11.2	9.32	4.77
WWW	4×10^7	7		2.38	2.1			
WWW	2×10^8	7.5	4000	2.72	2.1	16	8.85	7.61
WWW site	260000			1.94				
Internet, domain	3015-4389	3.42-3.76	30-40	2.1-2.2	2.1-2.2	4	6.3	5.2
Internet, router	3888	2.57	30	2.48	2.48	12.15	8.75	7.67
Internet, router	150000	2.66	60	2.4	2.4	11	12.8	7.47
Movie actors	212250	28.78	900	2.3	2.3	4.54	3.65	4.01
Co-authors, SPIRES	56627	173	1100	1.2	1.2	4	2.12	1.95
Co-authors, neuro.	209293	11.54	400	2.1	2.1	6	5.01	3.86
Co-authors, math.	70975	3.9	120	2.5	2.5	9.5	8.2	6.53
Sexual contacts	2810			3.4	3.4			
Metabolic, *C.Elegans*	778	7.4	110	2.2	2.2	3.2	3.32	2.89
Protein, *S.cerev.*	1870	2.39		2.4	2.4			
Ythan estuary	134	8.7	35	1.05	1.05	2.42	2.26	1.71
Silwood Park	154	4.75	27	1.13	1.13	3.4	3.23	2
Citation	783339	8.57			3			
Phone call	53×10^6	3.16		2.1	2.1			
Words, co-occurence	460902	70.13		2.7	2.7			
Words, synonyms	22311	13.48		2.8	2.8			

of autonomous nodes interacting with each other without planning, instead of being the result of any organized design process.

One of the more interesting properties of complex graphs it is the so-called *small world* effect, coined by Milgram's experiment in the 60's [163]. Milgram's work empirically showed how people were much closely connected than expected. The popular expression of *degrees of separation* — everybody can contact everybody else in the world only by using acquaintances — is derived from his work. This surprising effect has been found in empirical analysis of complex networks, and it is indeed a characteristic feature of them. For a network to be small world its average path length and diameter must be low, even for very large networks. Thus, it is possible for a node to reach any other node in the network within few hops.

Random networks have a short diameter and average path length, in fact, it increases logarithmically with the size of the network. A direct consequence of having a low average path length is efficiency in processes of propagation and diffusion, since the path between any two nodes is short. Complex networks behave as random

networks, therefore, they are very efficient in propagation: from innovation [60] through human diseases [185] to computer viruses [178]

However, average path length is not the only factor to be taken into account. Unlike average path length, the clustering coefficient of complex networks and random networks is completely different. Random networks have no internal structure, which results on a low clustering coefficient (tends to zero as network size grows). On the other hand, complex networks, specially social networks, often exhibit a strong community structure. The pattern of interaction is not random, it is the result of some undergoing process. This process might correspond to the transitivity studied in social networks analysis and sociological theory. In short, *small world networks — those networks that display the small world effect — are as efficient as random networks in terms of information propagation while displaying a strong community structure.*

There is no exact measure for *small worldliness*. A network is considered to meet the small world condition if when compared to a random network of the same number of nodes and edges its average path length is similar to that of a random network and its clustering coefficient is much larger [232, 228]. Formally, $C >> C_{rand}$ and $l \simeq l_{rand}$. Walsh [223] proposed a related measure that gives a discrete value: $(C/C_{rand})/(L/L_{rand})$. According to Walsh, a network is small world if the ratio is greater than 1.

Yet another striking property of complex networks is their robustness against errors. It is advantageous for a system — networks are systems — to be fault-tolerant in order to maintain its stability. It is undesirable that when a power line fails a total black out follows. Nor is it desirable that the brain, or our body, stops working with deadly consequences because some neurons, cells or proteins fail to operate properly. Neither is it desirable for a village to disintegrate and fall apart because some of its inhabitants pass away. Internet or peer-to-peer systems do not crash because routers break or people sharing files shutdown their computers or delete their files. It is evident that those systems are robust and can cope with random failures and keep working. Many complex systems, and complex networks, display a surprisingly high degree of tolerance to errors.

Redundancy has been often assumed as the cause of robustness [220, 49]. An informal definition of redundancy is the existence of many alternative paths that can preserve communication even if some nodes fail. However, recent studies on complex network have shown that the degree distribution of the nodes has a strong effect on the robustness of a network. Cohen et al. [53] and Albert et al. [13] showed how the Internet is very robust against random errors, Jeong et al. [126, 125] showed that simple organisms persist after a drastic intervention in their metabolic and genetic network. Watts in [229], and Holme and Kim [117] studied the effect of power-law networks but from a dynamic perspective, where failure of one node affects other

nodes beyond the topological sense. For instance, if a router of the Internet fails, the traffic will be redirected to other route, this might produce an overload of traffic on some other routers, ending up in a cascade failure. Again scale-free networks were more robust than random networks.

Complex networks that have a scale-free connectivity distribution are very robust against random errors. As we have seen in the previous section, many of the real networks that grow without supervision are in fact scale-free networks. However, scale-free networks also have an *Achilles heel* [13], because even if they are very robust against random errors, they are vulnerable against directed attacks. The power-law distribution implies that there are very few highly connected nodes (hubs), so when these are targeted, the whole system will collapse since the path between most pairs of nodes in the network pass through these hubs. We must remember, though, that attacks are artificial, they do not occur in nature. So, it is understandable that the system does not adapt itself to cope with those attacks as it does against random errors, which they do occur in nature.

Another interesting property of complex networks is that they display a characteristic *mixing pattern* between nodes. In other words, the probability of nodes to pair up depends on the type of these nodes. This is a phenomena well studied in sociology under the name of assortative mixing, or *homophyly*. For instance, people tend to interact with similar people, where similarity can be defined by many criteria such as social status, expertise, race or education.

Complex networks also exhibit a mixing pattern between nodes. Pastor-Satorras et al. [184] proposed a measure called *degree correlation* to capture the type of mixing that occurs in complex networks. They showed that complex networks often display a negative correlation by degree of connectivity, meaning that *high-degree* nodes were often connected to *low-degree* nodes. Thus, nodes of different types, which were defined by their degree, tended to pair up more frequently than nodes of the same type. The correlation degree of most models prior Pastor-Satorras work was zero, therefore, without a mixing pattern. Analysis of empirical complex networks revealed a negative correlation, so nodes were arranged with dissortative mixing.

Unlike other complex networks, social networks have a completely different mixing pattern. The nodes of a social network exhibit a positive degree correlation. Therefore, the network displays assortative mixing. *High-degree* nodes tend to pair up with *high-degree* nodes, and the same goes for the *low-degree* nodes. The first benefit that was derived from studying a network's mixing pattern — assortative or disassortative — is that it establishes a clear boundary between social and non-social complex networks [174]. This fact is very important since it implies that the distributed process that leads to the formation of social networks differs from the process underlying the growth of other non-social networks such as technological, biological or information networks. As we show in the next section, most of the

models that have been proposed to explain complex network formation cannot reproduce assortative mixing. In some sense this should come as no surprise, since the underlying assumptions were not plausible from the standpoint of social theory.

In spite of this, there are models that achieve assortative mixing and positive degree correlations [180, 31]. For instance, Newman in [180] presented a growth model that achieved a great part of the *assortativeness* [174] found in empirical social networks, although not all. Newman himself suggested that the assortativeness that was unaccounted for might well be caused by the underlying social processes of the system. Arenas et al. [15] also suggested that formation of social networks could be driven by some sort of optimization process.

Models

Research in complex networks have proposed dozens of models aiming to explain the formation and evolution of this interesting class of networks. Most models can be classified into groups depending on the rationale behind their connection scheme. However, among models of the same group there are many differences. Most of the differences are due to incremental improvements by reproducing empirical data more accurately or by being more general or by capturing yet another characteristic of complex networks that remained unnoticed until the date. Let us briefly summarize the models that we consider more relevant. For a more comprehensive review of complex networks models see [11, 175].

Watts and Strogatz in his pioneering article [232] presented a model that reproduced the small world property found in many empirical networks. Their model was based on a stochastic rewiring process. It worked by starting from a regular graph, a ring lattice network, which has a very high clustering coefficient but it also has a high average path length and diameter. For regular graphs the average path length does not grow logarithmically as it does for random graphs. Then, each edge of the graph was randomly rewired with a probability p. Watts-Strogatz model was a one-parameter model, so they tried the experiment for different values of p. When $p = 0$ the final graph was the same ring lattice, for $p = 1$ the final network was a random network, therefore, it had short average path length but also low clustering. As it can be seen in figure 3.2, as the rewiring probability p increases both the clustering and the average path length decrease, however clustering decreases slowlier and for p in the range [0.01..0.1] the average path length is so low that it is comparable with the average path length of a random network. The clustering coefficient, however, is still big enough, almost as the clustering of the original regular network. Hence, the small world property can be obtained from a regular graph provided a little fraction of the edges (between 1 and 10 percent) are randomly rewired.

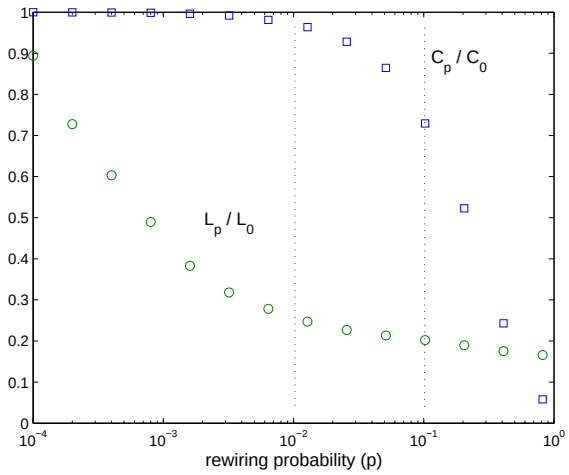

Figure 3.2: Average path length L_p and clustering coefficient C_p for the Watts-Strogatz model [232]. The data is normalized for the values of a regular graph C_0 and L_0. The values L_1 and C_1 correspond to those of a random network when $p = 1$. We can observe how in the range [0.01..0.1] the average path length drops drastically, being similar to the average path length of a random network. Clustering coefficient in that range is still comparable to the clustering coefficient of a regular network.

Preferential Attachment

After the Watts and Strogatz [232] model, other models followed aiming not only to reproduce other properties of complex networks, but to help understand how networks came to have those properties. Perhaps the most well-know model is the Barabási and Albert model of *preferential attachment* [22], which is able to reproduce the power-law (scale-free) degree distribution. This model is based on network growth, so it assumes that the network is continuously growing in size. At every time step a new node is added, and it gets connected to a set of nodes that already exist in the network. These nodes are chosen with a probability that depends on its own degree, technically, $\Pi(k_i) = \frac{k_i}{\sum_j k_j}$. This model generates a network where the degree distribution follows a power law with an exponent $\gamma = 3$, $P(k) \sim k^{-3}$. The original model of Barabási and Albert attracted an exceptional amount of attention in the literature. Albert and Barabási themselves proposed in [10] an extension were edges could be rewired. By doing so, the model is able to generate a power-law degree distributions with an exponent ranging in the [2..∞] interval. The extended Albert-Barabási model is described as follows:

The algorithm starts with m_0 isolated nodes, and perform at every step one of these three actions:

1 With probability p add m ($\leq m_0$) new links. We pick two nodes randomly.

The starting point of the link is chosen uniformly and the end point of the new link will be chosen according to the following probability distribution:

$$\Pi_i = \frac{k_i + 1}{\sum_j (k_j + 1)}$$

where Π_i is the probability of selecting the i-th node, and k_i is the number of edges of node i. This process is repeated m times.

2 With probability q, m edges are rewired. That is, we repeat m times: Choose (uniformly) at random one node i and a link l_{ij}. Delete this link. Choose another (different) node k with probability $\{\Pi_l\}_{l=1...N}$ and add the new link l_{ik}.

3 With probability $1 - p - q$ add a new node with m links. These new links will connect the new node to m other nodes chosen according to $\{\Pi_l\}_{l=1...N}$.

Once the desired number N of nodes is met the algorithm stops. The networks generated with this algorithm are scale-free *random* networks, random is because there is no correlations among edges [184].

It can be shown [10] that in the limit of large N, when $p = q$, this algorithm ends up with a graph with connectivity distribution

$$P(k) \propto (k+1)^{-\left(\frac{2m(1-p)+1-2p}{m}+1\right)}$$

that can be approximated, when $k >> 1$, by $P(k) \propto k^{-\gamma}$ where $\gamma = \frac{2m(1-p)+1-2p}{m}+1$.

Other researchers achieved the same results using other analytical approaches [73]. As a consequence they could fit the obtained distribution to the degree distribution found in the actor collaboration network. Dorogovtsev and Mendes [72] proposed a model where new edges were added between old nodes and existing edges can be removed. Amaral et al. [181] introduced constraints such as age, cost and capacity into the preferential attachment. Bianconi and Barabási [28] introduced competition among the nodes, thus a fitness value is assigned to each node. The probability of receiving connections is, following Bianconi and Barabási, a function of the degree (the preferential attachment) plus the fitness (attractiveness) of the nodes(attractive). Another interesting model based on the preferential attachment was proposed by Krapivsky et al. [143]. Most of the work on preferential attachment considered the connection probability, Π, to be linear, Krapivsky et al. generalized that to non-linear, covering the supra-linear, sub-linear as well as the original linear preferential attachment. All those models based on the original idea of preferential attachment aimed to a better fit with the empirical networks, and henceforth, explain the underlying process of network formation. However, these models, yet elegant and beautiful, have some assumptions that are difficult to justify such as the

assumption of global knowledge that is required to calculate the connection probability. The main criticisms to the preferential attachment models is with regard to their unrealistic assumptions. For instance, they assume to have perfect and global information about the system, which for sure does not hold in open systems such as human societies. Thus, it is obvious that the motives that drive the autonomous actors of a system are not a consequence of preferential attachment, although this phenomenon appears as a by-product.

Edge Redirection

Another group of models for the understanding of complex networks formation are those based on *edge redirection*, *edge copying*, or *transitive closure*. The different names come from different authors who proposed the same underlying idea independently [138, 142, 209, 76, 145]. This group of models do not rely on preferential attachment to explain the process that generates complex networks. For instance Krapivsky and Redner [142] proposed a model where nodes copy part of the connections of other nodes. When a node enters the network, it randomly chooses another node to connect to, then with a given probability the node connects to the neighbours of that node or it goes to the *ancestor* of the node and repeats the process. To be precise, the authors do not use the term copy but edge redirection, but as we said a similar idea was proposed by different authors, for instance, Kleinberg et al. [138] and Kumar et al. [145]. Its rationale for their model was inspired on how links are placed on the Web. They assumed that a page copied partially the links from other existing pages of the same topic. Notice that this approach is much more plausible since it only requires local information.

Furthermore, these models are able to reproduce the preferential attachment phenomenon without relying on the global connection probability of the preferential attachment model. Nodes with high degree are more likely to be found than nodes with low connectivity, since high degree nodes are more visible to those nodes that enter into the system. Thus, the dynamics of the model itself reinforces the phenomena of preferential attachment. The more connections a node has the more connections is likely to have in the future. The saying *"the rich gets richer"*, which is folk wisdom, entirely applies here as well as in the case of the preferential attachment model. Nevertheless, edge redirection models do not require implausible assumptions to operate.

Optimization

Yet another group of models to explain the emergence of complex network are those based on optimization. These models are not as common in the literature as the previous ones. Their rationale is that complex networks are the result of a opti-

mization process. We find the idea very interesting since it is directly related to the main properties of complex networks: robustness and efficiency in the diffusion of information. Ferrer and Solé [83] show how a system built up from nodes and edges that optimize both average path length and density (number of edges) ends up having a complex network topology. This approach although interesting also suffers from non-plausible assumptions since it requires global knowledge. The optimization process is led by a fitness function that requires system level measurements which in principle are unknown at the micro level. Ferrer and Solé run a hill-climbing optimization process that minimizes the average path length as well as the density (edges are costly) and the result is a complex network with the small world property. Another method based on optimization is the *highly optimized tolerance* (HOT) proposed by Carslon in [47]. This model is related to self-organized criticality. The optimization is based on finding the best topology to be fault-tolerant by minimizing the redundancy. For instance, this model would find the optimal distribution of food suppliers in a city so that not everybody depends on the same food-supplies, since its lost would cause a famine, but the number of food suppliers is kept to a minimum for the sake of market efficiency. We believe that this model's assumptions ignore the decision making process at the micro level (the individual node) therefore is it difficult to explain the emergence of complex networks where deploying and removing connections are the result of a decision process.

The model we present in chapter 8 falls into the group of models based on optimization. Nevertheless, the nodes of our model -agents- perform a local optimization process grounded in social exchange theory under local and imperfect information. With the local optimization model we see that several kinds of complex networks can be obtained by modifying an exogenous parameter of the system called *harshness*, which makes the optimization process easy or difficult depending on its value. Thus, agents trying to optimize their partners can self-organize in different kinds of complex networks depending on the harshness of the systems.

Part II

Structure as a Source of Knowledge

Chapter 4

Extracting Reputation

This part of the thesis is devoted to methods for extracting knowledge embedded in the relationships between the components of a system. To that end, two different algorithms that resort to the analysis of the structure are presented. The first one, in the current chapter, assesses the relevance of a node by its position within the network. This ranking of relevance is indeed a good approximation to the agents' reputation. The second one, presented in the following chapter, uses the topology of the network to retrieve its underlying community structure.

Both algorithms are intended to extract knowledge — or information — about the system by means of analyzing the structure of the interactions alone. This information, which is inherently social, can be called upon for building efficient and stable systems. We focus specifically on artificial societies populated by agents, although the same algorithms can be applied to other domains. In particular, to other systems composed of autonomous individuals living in an open environment (see sections 1.1, 2.1 and 2.2 for a more developed argument about complex artificial systems and their need of social measurements and mechanisms).

4.1 Description

The World Wide Web is beyond a doubt the biggest complex system subject to study. This vast eco-system of information is composed of billions of pages and even a greater number of links. The Web is a completely distributed and open system where anybody can create pages and links according to their particular interests. However, the spectacular growth of the Web in the last decade has ended in an overflow of information that cannot be managed without pre-processing and filtering.

In order to introduce some order in this awkward amount of information search engines proposed to rank web pages by relevance. The first approach to rank pages was focused in their content. However, this approach proved soon to be inadequate

to find the relevant information on the Web. Therefore, an alternative approach based on the analysis of the links rather than the analysis on the pages was rapidly introduced.

This step was a paradigm shift; from using the information contained in web pages to rely on relationships among web pages. If one takes a web page as an individual, a link between web pages might be considered as some form of relation between individuals, so that some kind of social awareness comes up. Hence, ranking algorithms for the Web might also be applied to other systems represented as a network, such as Social Networks. Those algorithms can provide information about relevant individuals by identifying their position in the society.

Several algorithms are now in use that consider the World Wide Web as a graph. We are interested in assigning a score to each node. A possible way to proceed is to consider that a node i in this graph (i.e. a page or a document linked to it) is more important (that is, gets a better position in the ranking) than a different node j if it is linked to nodes that rank higher than the nodes connected to j. A node *authority* results from the number and quality of the nodes that point to it. Essentially, the main idea of this method lies in the transfer of confidence on the quality of a node through a link. This intuitive and elegant idea underlies some well-known ranking algorithms, such as *HITS* [137] and *Pagerank* [183].

Following Adamic and Adar [5] the Web is also a web of people, therefore the relationships of trust that evolve among people might be subjected to the same analysis of the relationships between web pages. Social Network analysis [226] (see section 3.1 for further discussion) suggests that relationships between individuals can be regarded as a measure of trust. The location of a given member of a community within a social network may be used to infer some properties on his/her degree of reputation in the community. For instance, experts who are well-known and highly regarded by most of the members in a community are easily identified as highly connected nodes in the social network. This relational information could be used as the basis of a ranking mechanism instead of using ratings created on isolated information about an individual. Consequently, the very same methods that have proven successful for the Web could be applied to social networks to retrieve reputation of individuals.

To that end, we devised a ranking algorithm called *NodeRanking*. This algorithm was designed to meet the requirements of multi-agent systems. First, it operates using only local information. The process is distributed through all agents — or other computation entities — of the system so that global knowledge, such as the adjacency matrix of the network, is not required. Second, the algorithm adapts dynamically to different network structures. The characteristics of social networks differ to those of the Web as discussed in section 3.2.1. Those differences are manifested in differences on the rankings issued by different ranking algorithms, as we will show in the ex-

periments section. Our algorithm unlike other ranking algorithms such as *Pagerank* and *HITS* take some particularities of the structure into account.

4.2 Ranking Algorithms

Before presenting our algorithm we review two of the most well-know representatives of ranking algorithms in graphs: *Pagerank* [183] and *HITS* [137]. These algorithms will be used in the experiments to test the performance of our *NodeRanking* algorithm.

4.2.1 Overview of *Pagerank*

The main idea behind *Pagerank* [183] is that *good* nodes point to or are pointed by *good* nodes and its strategy is that of a modified random walker. Let L the adjacency matrix of the graph and P such that:

$$P_{ij} = \frac{l_{ij}}{\sum_k l_{ik}} \tag{4.2.1}$$

that is, P is a stochastic matrix ($\sum_k P_{ik} = 1$) derived from L. Now, to avoid the *rank sink* problem [183], some probability of jumping to other nodes in the graph is introduced, in such a way that cycles on the graph (that would generate an erroneous increment of authority) may be broken. Thus, the full stochastic matrix is:

$$M = e\left(\frac{1}{N}\vec{1}\vec{1}^T\right) + (1-e)P \tag{4.2.2}$$

where $\vec{1} = (1 \ldots 1)^T$, N is the number of nodes and e is the jumping probability ($0 \leq e \leq 1$). *Pagerank* depends on e and its value must be fixed *a priori* ($e = 0.15$ is the recommended value [183]). The equilibrium distribution, that is, the stationary state of the Markov chain defined by the transition matrix M, can be obtained by calculating the principal eigenvector of the matrix M^T.

4.2.2 Overview of *HITS*

In *HITS* algorithm [137], each page has both a hub score y_i and an authority score x_i. The rationale for *HITS* is that a node with high authority is pointed to by many nodes with a high hub score and a node with high hub points to many nodes with high authorities. The final scores of every node can be obtained through an iterative process:

$$\vec{x}^{(t+1)} = L^T L\vec{x}^{(t)} \vec{y}^{(t+1)} = LL^T\vec{y}^{(t)} \tag{4.2.3}$$

from a given $x^{(0)}$ and $y^{(0)}$. Finally, the solution for authority, that is, the stationary state of the equation 4.2.3, is the principal eigenvector of the $L^T L$ matrix. Also, the principal eigenvector of LL^T matrix contains the hub scores.

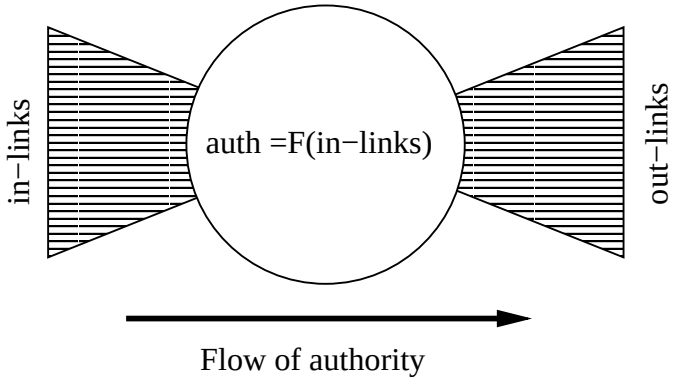

Figure 4.1: Flow of authority in *NodeRanking*.

```
do
    n = getNode(g)
    do
         passAuthority(n)
         nnew = getNextNode(n,g)
         n = nnew
    while (nnew!=null)
while (!converge())
```

Figure 4.2: *NodeRanking* algorithm

4.2.3 *NodeRanking* algorithm

The underlying idea of *NodeRanking* is the same one used by *Pagerank* and *HITS*. Roughly speaking, each node of a graph has an associated *degree of authority* (which is always positive), which depends on the *degree of authority* of its neighbours. Initially, all nodes are set to have the same authority, as the algorithm proceeds, the authority of a node is recalculated as a function of the authority of the nodes that point to it (*in-nodes*). The underlying idea is that authority of a node is propagated through its *out-nodes*. The outcome of the algorithm is that the authority of a node, depends on the authority of those nodes who point to it. This idea, sketched in figure 4.1, is the very same idea used by the previously reviewed algorithms.

This approach, however, is sensitive to cycles in the graph, which can result in a deadlock in the authority propagation. To solve this problem, known as *rank-sink*, we followed a similar solution to that proposed by Page et al. in [183] (i.e. with a very small probability it is possible to go from a given node to any other node of the network). This effect is achieved by introducing the *jumping* probability that avoids

the stochastic processes to get trapped in a cycle of the graph.

NodeRanking follows a *random walker strategy* to explore the graph. It starts from a randomly selected node, and proceeds by selecting one of the nodes that can be reached through *out-edges*. The algorithm proceeds as described in figure 4.2. Some functions of the algorithm require further explanation:

- **getNode (Graph g):** returns a randomly (according to a uniform distribution) chosen node.

- **getNextNode(Node n, Graph g):** returns one of the *out-neighbors* nodes of node n. Each node has a set (can be empty) of *out-edges* that point to other nodes. Thus, *getNextNode(n,g)* returns the next node m to be visited from n. This node is selected with a probability that is calculated as a function of the weight of the edge between n and m:

$$Pr_{choose}(n \to m) = \frac{w_{n \to m}}{\sum_{\forall l \in outNodes(n)} w_{n \to l}} \tag{4.2.4}$$

  where $w_{n \to m}$ is the weight of the link connecting n and m. This probability would be P_{nm} following the notation introduced when *Pagerank* was described. When *getNextNode()* returns a node set to *null* the path is *broken*. There are two cases where the path may break: the first one is when the algorithm arrives at a node that has been visited in the previous k steps (for an *a priori* fixed value of the parameter k); the second case depends on the jumping probability. The path is broken with a probability $Pr_{jump}(n)$:

$$Pr_{jump}(n) = \frac{1}{\#outEdges(n) + 1} \tag{4.2.5}$$

  Nodes with fewer *out-edges* have a greater probability of breaking the path. This could be seen as a walker that gets bored because of the reduced range of choices. By doing this, the deadlock problem previously mentioned due to cycles in the graph is solved.

- **passAuthority(Node x):** this function assigns part of the authority of node x to all these nodes that x points to. That is,

$$\Delta auth(y) = \left(\frac{Pr_{choose}(x \to y) \times auth(y)}{F_{rw}^{y}} \right) \tag{4.2.6}$$

  where $auth(y)$ is the authority of the node y and F_{rw}^{y} is a normalization factor to maintain the authority within a limited range of values. Without this factor, values calculated with equation 4.2.6 would tend to infinity since the authority of a node gets higher and higher as the algorithm proceeds. It is associated

to the particular walk of the random walker (hence the rw subscript). The factors of more frequently visited nodes grow faster than the values of the less visited ones. The growth of these factors is monotonously increasing. Thus, we can insure convergence towards a finite value. Factor F^y_{rw} is initialized for every node as the sum of the authority of all nodes in the graph. The initial authority of a node has to be positive, and factor F^y_{rw} must be ≥ 1.

- **converge():** this function is a stationarity test on all the nodes in the graph. Each node n stores its last increment of authority $\Delta auth(n)$. The increment of authority tends to 0 since F^n_{rw} factor grows monotonously with time. The function *converge()* tests the state of each node and, if $\Delta auth(n) < \epsilon$ the node will be considered as stationary. When all the nodes of the graph are stationary the algorithm stops. Actually, the convergence function does not test all the nodes because it would not be very efficient. The event of becoming stationary is notified by the nodes themselves. The parameters of the algorithm that have been used along all the experiments are: $k = 4$ and $\epsilon = 10^{-6}$.

4.2.4 Comparisons

Most ranking algorithms based on the analysis of the links share the same underlying idea of a random walker process. The algorithm presented in this chapter is not different, although its implementation and behaviour differ from those of *Pagerank* and *HITS* algorithms. Let us then, overview the differences between our algorithm and other ranking algorithms, in particular *Pagerank* and *HITS*.

Distribution Using Only Local Information

HITS and *Pagerank* are based on finding out the stationary state of a linear dynamical system characterized by the variance-covariance matrix in the case of *HITS*, and the transition probability matrix in the case of *Pagerank*. The adjacency matrix must be available for these algorithms to operate. Moreover, the principal eigenvector, which is calculated iteratively, need to be frequently normalized. Consequently, these algorithms are also dependent on synchronization. Even though these algorithms can be easily parallelized to reduce both computational and spatial cost, access to the adjacency matrix is still required. For systems composed by autonomous entities, such as multi-agent systems, the assumption of global information is highly inadvisable, since it goes against the limitations of the individual agents of the system (see section 2.1).

Unlike *Pagerank* and *HITS* our algorithm *NodeRanking* uses only local information. Each node i only needs to know the nodes that it points to, and these ones have to be aware of i's convergence towards the stationary state. In order to re-

trieve the results of the ranking process, a centralization point is required, but even in this process the communication proceeds unidirectionally (from the nodes to the controller) so it is not necessary at all to have information about the whole graph. *Pagerank* and *NodeRanking* are almost identical in the underlying idea. Both of them follow the same random walker strategy. Actually, the transition probability matrix of *NodeRanking* can be defined as follows:

$$M = J\frac{1}{N}\vec{1}\vec{1}^T + \left(\vec{1}\vec{1}^T - J\right)P \tag{4.2.7}$$

where $\vec{(1)}$ is a vector of 1, N is the number of nodes, P is the adjacency matrix normalized by rows and J is the jumping probability matrix defined as a 0 matrix where the diagonal contains the jumping probability of a node, i.e. J_{ii} contains node i jumping probability (this was defined in equation 4.2.5). The results obtained by *Pagerank* when finding out the stationary state of the Markov chain defined by the transition probability matrix M, i.e finding out the principal eigenvector of the matrix M^T, are equivalent to the results obtained by *NodeRanking* algorithm. The advantage of *NodeRanking* is that it is not necessary to know the adjacency matrix of the graph to run the algorithm.

Dynamic Jumping Probability

Unlike *Pagerank* the jumping probability of *NodeRanking* is dynamic and only depends on the connectivity of nodes (see equation 4.2.5). This is not the case of *Pagerank* where the jumping probability is fixed regardless of the connectivity of the node or the network topology. The jumping probability of *Pagerank* is set to 0.15 (values in the $[0.1 \ldots 0.2]$ range) , those values were found experimentally by testing the algorithm on the Web. However, as we will discuss in the following section, the performance of *Pagerank* is affected when applied to other networks rather than the Web due to their different topologies.

4.3 Experiments about ranking, reputation and relevance

We performed two series of experiments with the *NodeRanking* algorithm. The first series aimed towards testing how the reputation of community members might be derived from the topology of the social network. The second series of experiments was carried out to test the quality of the ranking yielded by *NodeRanking* in a Web-like network.

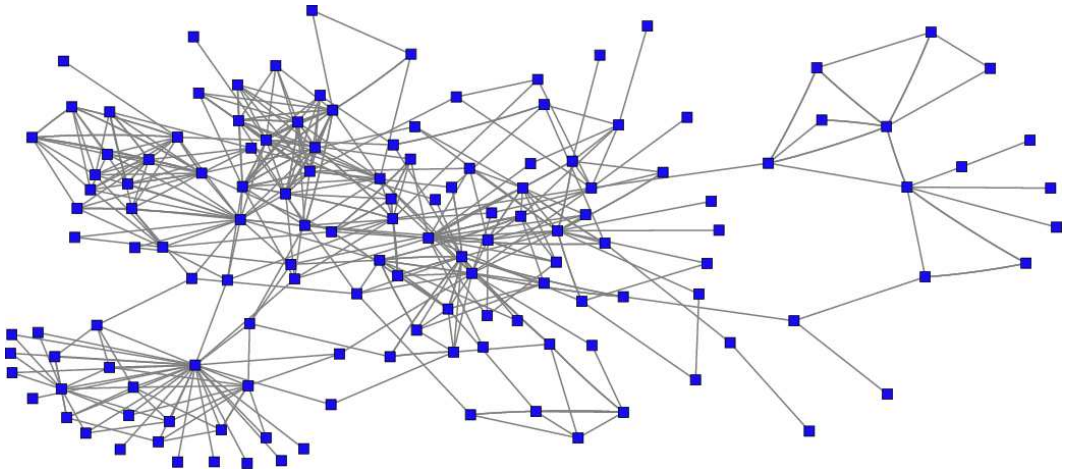

Figure 4.3: Social network of the Software Department at UPC

4.3.1 Extracting Reputation from Social Networks

NodeRanking was applied to the social network of an experimental community formed by the members of the Software Department at the Universitat Politécnica de Catalunya, herein UPC. Figure 4.3 is a snapshoot of the social network used in our set of experiments.

Social Network of the Software Department at UPC

This social network was generated automatically by our *NetExpert* system [198]. This system was part of the *Collaboratory* [217], which was a multi-agent recommender system aimed to facilitate the knowledge sharing and diffusion in a community of researchers. Basically, when new content was available in the system it was distributed to those people susceptible to be interested in it. This classical collaborative filtering was extended with the addition of social networks provided by *NetExpert*. The social networks were built by analyzing the personal web pages and deploying links between researchers that either were co-authors of a paper or whose names appeared in other researchers' personal web pages.

According to the measure of *small worldliness* proposed by Walsh [223] the social network of UPC is a small world, the Walsh's ratio is $(C/C_{rand})/(L/L_{rand}) = 10.4$. Furthermore, the connectivity distrution found in the UPC's social network decays exponentially, as it can be observed in figure 4.4 (because of the low number of nodes in the graph, we used the accumulated frequency instead of the simple frequency in order to reduce fluctuations). These measurements clearly point to the fact that the UPC Social Network is indeed a small world network.

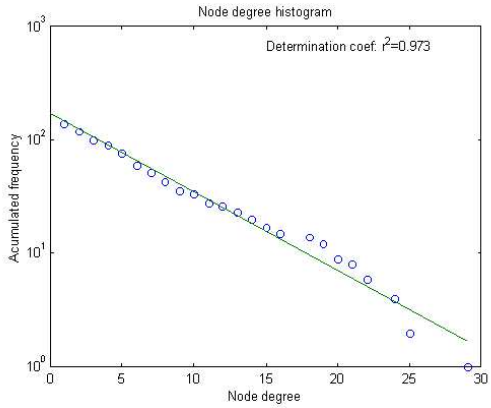

Figure 4.4: This figure shows the node degree histogram in a logarithmic scale on the frequency. Because of the low number of nodes in the UPC Social Network, the accumulated frequency has been used instead of the simple frequency.

As we reviewed in section 3.2.1 some of the characteristics of social networks, such as the small world property, can be used to improve graph algorithms that did not take into account the special properties of the complex networks. In particular, one must take into account the special characteristics and properties of these networks when designing algorithms of search and propagation, some examples are [223, 5, 6, 231, 7, 237].

Inferring Reputation from the Social Network

For the experiments on social networks a fragment of the community (34 members) was selected randomly to become the test set. The ranking obtained by means of *NodeRanking* was called $Rank_{NodeRanking}$. It was compared with the results obtained by applying *Pagerank* with $e = 0.15$ and also with the results of the *HITS* algorithm.

Pagerank and *HITS* work with unweighted graphs since they were devised to work on the Web (a generalization to weighted graphs of these algorithms would be trivial). Weighted edges may be useful, if available, because the normalized weight of an edge gives more information than the fact that the edge exists or not. The UPC Social Network has weighted edges, so two rankings instead of one were built for each algorithm: $Rank_{PageRank}$ and $Rank_{HitsAuth}$ if weights are not taken into account, or $Rank_{PageRank(w)}$ and $Rank_{HitsAuth(w)}$ otherwise.

In order to validate the resulting rankings they had to be compared against a real and accepted measure of importance for this type of network. The community behind the social network is a group of researchers. It is important to test both algorithms on a social network because the alternative (to test them on the web)

Correlation coefficient	$Rank_{cite}$	$Rank_{cite-self}$
$Rank_{cite}$	1.0	0.983
$Rank_{cite-self}$	0.983	1.0
$Rank_{NodeRanking}$	0.687, $s = 8.6 \times 10^{-4}$	0.621, $s = 0.011$
$Rank_{PageRank(w)}$	0.535	0.486
$Rank_{PageRank}$	0.521	0.495
$Rank_{HitsAuth(w)}$	0.412	0.383
$Rank_{HitsAuth}$	0.342	0.323

Table 4.1: Correlation between rankings

poses several problems. On the one hand, it is difficult to get a representative sample of resources from the web. On the other hand, it is difficult to analyze the quality of a resource on the net. Finally, the underlying idea of authority flux that is common to both algorithms is better represented in the links of a social network (in the sense of being a more significative measure) than in the links existing on the Web. A simple and recognized measure of *authority* in scientific communities is the one obtained through the rankings of an independent scientific publication ranking agency.

Therefore, the citation indexes for each of the 34 randomly selected members of the Software Department were compared against the ratings that *NodeRaking* yielded ($Rank_{NodeRanking}$), against the ratings that *Pagerank* yielded ($Rank_{PageRank}$, and $Rank_{PageRank(w)}$) and also against the ratings that *HITS* yielded ($Rank_{HitsAuth}$ and $Rank_{HitsAuth(w)}$).

CiteSeer[1] was used as a source for citation index values. The research papers of each member of the community were retrieved from *CiteSeer* with the corresponding number of citations and self-citations for each of them. Two rankings for the members of our test community were calculated. $Rank_{cite}$ sorted researchers by number of citations and $Rank_{cite-self}$ sorted researchers by number of citations without counting self-citations. These rankings can be considered as reference rankings, the closest ones to real reputation measures in the scientific community. To compare the quality of the ratings obtained by *NodeRanking*, *Pagerank*, *HITS* and the reference rankings built using *Citeseer*, the correlation coefficient between rankings was used as a similarity measure.

Table 4.1 and figure 4.5 summarize the correlation values between reference, the desired, rankings: $Rank_{cite}$ and $Rank_{cite-self}$ and the rest. The ranking $Rank_{NodeRanking}$ is the average of twenty executions of the algorithm. Means and standard deviations are also given. Variability is due to the asynchronicity of the authority transfer

[1]Citeseer, available at http://citeseer.ist.psu.edu/

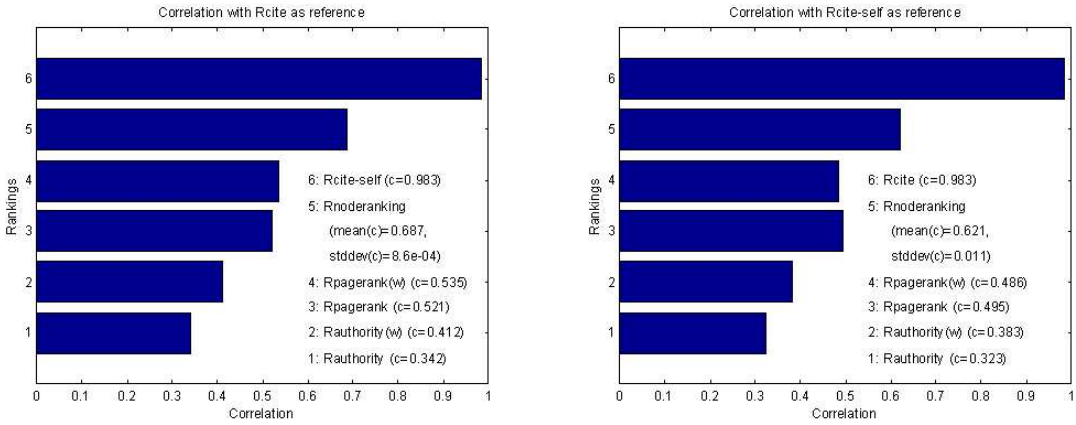

Figure 4.5: Correlation between ranking. Left sub-figure shows correlations with $Rank_{cite}$, right sub-figure shows correlations with $Rank_{cite-self}$

process within *NodeRanking*:

Regardless the reference ranking there is always the same order among the generated rankings results. $Rank_{NodeRanking}$ values are always better than $Rank_{PageRank}$ and $Rank_{PageRank(w)}$ values, and these ones are better than the $Rank_{HitsAuth(w)}$ and $Rank_{HitsAuth}$. Notice that the rankings obtained by the *HITS* algorithm are better if weights of the edges are taken into account. However, in any case the correlation is high, the correlation between two random ranking is close to 0. Therefore, correlations of 0.621 or 0.687, yet not perfect, are indeed interesting. One must take into account very different nature of the ranking yield by the ranking algorithm and the ranking build from the number of citations.

The rankings obtained by *NodeRanking* are better than the ones obtained by *Pagerank*. Both algorithms follow the same random walker strategy with the exception of the authority updating mechanism (it is asynchronous in *NodeRanking* and synchronous in *Pagerank*). One might expect the results to be more similar. The reason for a better performance of *NodeRanking* may be in its ability to adapt itself to the graph topology.

In figures 4.6 and 4.4 the distribution of (*out-edges*) degrees can be observed. It follows an exponential distribution which may be used to compute the average jumping probability. Let us call this degree distribution $d_{out-connectivity}$:

$$\overline{Pr_{jump}} = \sum_i d_{out-connectivity}(i) Pr_{jump}(i) \tag{4.3.8}$$

The average jumping probability of our social network is 0.5314. So *NodeRanking*, with its dynamical adjustment of the jumping probability, is able to adapt to different graph topologies. *Pagerank* is not able to calculate such good rankings for

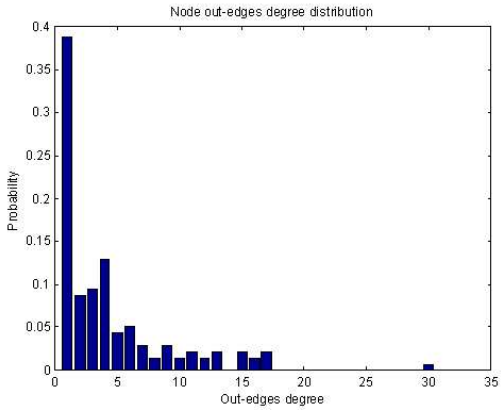

Figure 4.6: Social Network *out-edges* degree distribution

social networks. One must take into account the differences between the Web and social networks. For instance, the clustering coefficient is typically higher in social networks, which display a stronger community structure. The effect of close-knit communities was described by Lempel and Moran [148] as a factor that worsens the calculation carried out by *Pagerank*. Moreover, the connectivity distribution of the social network used in the experiments is different from that found in the Web. The exponential distribution of the connectivity implies having a higher number of nodes with lower connectivity. We modified *Pagerank* to work with a with a jumping probability of $e = 0.5414$ and results obtained by *Pagerank* were very similar to those obtained by *NodeRanking*. Thus, our algorithm is better suited to work with other networks besides Web-like networks. Furthermore, the fact that *NodeRanking* only uses local information is another advantage over *Pagerank*.

4.3.2 Extracting Relevance from the Web

As it was mentioned before, *NodeRanking* is able to adapt itself to the graph topology. In order to check this ability, another experiment was performed to compare *NodeRanking* against *Pagerank*. Now the graph used for testing purposes is a graph with some of the topological properties of the World Wide Web.

Many efforts have been devoted to generating graphs that reproduce the topological properties of the Web as we saw in section 3.2.1. For this experiment we chose the model proposed by Klemm and Eguíluz [139]. This model generates graphs with power-law connectivity distributions with exponent around $\gamma = 2.25$, which is the exponent observed experimentally in the World Wide Web [185]. The parameters of the graph are $m = 8$, $a = 2$ and $N = 25000$ (see [139] for details).

The correlation between rankings obtained by *Pagerank* and *NodeRanking* in a

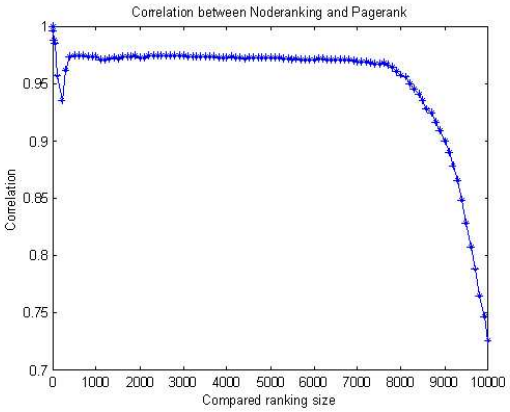

Figure 4.7: This figure shows the correlation between the ranking obtained by PageRank and the ranking obtained by *NodeRanking* in Klemm and Eguíluz's model of scale-free graphs.

web-based graph is very high as can be observed in figure 4.7. In the x axis there is the size of the compared ranking. The correlation of the first ten nodes is 0.9964. The correlation of the first one hundred is 0.9564. Finally, the correlation of the first nine thousand nodes is 0.8892. It can be concluded that the rankings obtained in huge graphs that have a scale-free connectivity distribution are almost identical for both algorithms.

4.4 Discussion

In this chapter we presented an algorithm that finds the relevance of nodes by means of the network structure. The rankings by relevance provided by our algorithm were compared to other rankings yielded by *Pagerank* and *HITS* in two networks with different structural properties; the first one was a social network from a real community, and the other was a Web-like network generated with the Klemm-Eguíluz model of scale-free networks.

The results of the experiment performed on a real community — an academic community — were compared to an established measure of reputation in Academia, i.e the publication impact by the number of citations. The results seem to indicate that the ranking by relevance yielded by *NodeRanking* is a good approximation to the reputation of a researcher within its community. Other ranking algorithms, such as *Pagerank* or *HITS*, do not obtain such an acceptable approximation. A second experiment has been performed in a graph that follows the Web connectivity distribution. The results are highly correlated with the ranking obtained by *Pagerank*.

NodeRanking is able to rank the nodes of a graph regardless its topology and by only using local information. The algorithm does not require to have the adjacency matrix to operate. Furthermore, the algorithm is not biased to a particular topology since it is able to adapt itself to different network structure by dynamically changing its jumping probability.

Besides the algorithm, an important contribution of this chapter is to show that analysis of the structure can reveal social information such as reputation, as we have seen in this chapter, or community structure, that will be addressed in the following chapter.

Reputation is an important topic in Artificial Societies, and in particular in multi-agent systems and electronic communities (see section 2.2 for a discussion about trust and reputation). The presented algorithm illustrates that structure can be used to infer knowledge about individuals without resorting to analysis of their internal behaviour. Thus, by analyzing the social position of an agent within the system a good approximation to its reputation can be drawn using our algorithm. This method, yet not perfect, has an important advantage with respect to other ways of measuring reputation. Namely, it does not require users to be continuously and explicitly issuing ratings [245]; a method that is seen as a burden on users and eventually a reason for poor performance of collaborative systems.

Other multi agent systems that use social networks either do not use them for reputation measurement, as is the case of *ReferralWeb* [133] or still rely exclusively on rating feedback from users as [243] does. This one has only been tested on a simulated community as opposed to the test we carried out on a real one (another example of these tests on simulated communities is [245]). Work is under way to state under which patterns of interactions among agents a social network graph structure arises [67], and to use structure to improve agent design and operation.

Chapter 5

Finding Community Structure

In the previous chapter we presented an algorithm to find out the relevant vertices of a network by analyzing its structure, as a consequence a measure of agents' reputation could be drawn depending on their position in the social network.

To delve further into the knowledge embedded in the structure of networks we devised a clustering algorithm to extract the community structure within a network. *Community structure* refers both to the cohesive groups within a network and to the relationships between those groups. Finding the underlying communities in a network allow us to identify those groups of nodes — or agents — whose interactions occur mostly with members of the same group, thus, forming a close-knit community. The relationships between these communities provide further understanding of the structure of the network, and by extension, of the systems that the network models.

5.1 Description

Clustering plays a key role in the analysis and exploration of data. In short, clustering is the method by which meaningful groups within collections of data are revealed. These clusters are intended to group individuals — or samples — who are similar to each other so that the hidden structure within the collection of data is extracted, resulting in a valuable acquisition of knowledge about the original data. Data-mining and machine learning are disciplines that extensively work with clustering, specially, with datasets composed by individuals and attributes. Their goal is to identify groups of individuals which are similar according to their attributes. Thanks to the recent collective effort on analyzing and compiling very large networks, there is a growing interest in methods based on the structure — topology — of the networks rather than on the individuals' attributes.

This turn toward structure has been possible because of the characterization of many systems as networks (see section 3.2). One of the regularities found in

complex networks [11, 175] is the high cliquishness of the network [232], which leads to the fact that there are groups of vertices that are very interconnected among them with few interactions outside each group. Therefore, there is an implicit community structure within complex networks.

Girvan and Newman [101] proposed an algorithm to extract the community structure from complex networks that has become one of the most used among the researchers in this area. From that important work a branch of research on complex networks has turned into clustering algorithms to retrieve the community structure in those networks. To evaluate the accuracy — or quality — of a community structure yielded by a clustering algorithm, Newman and Girvan devised a quantitative measure called *modularity* Q. Although there are other quantitative measures [129], modularity is widely accepted in the physics community. Q is defined in [179] as

$$Q = \sum_i \left(e_{ii} - a_i^2\right) \tag{5.1.1}$$

Modularity is the addition of the modularity of all the groups, $Q = \sum_i q_i$. Thus, for each group i that contains k vertices the modularity is calculated as the fraction of edges that have both ends pointing at vertices in group i, e_{ii}. The fraction of *intra-group* edges is confronted with the fraction of edges of that group, a_i, which are edges whose end points belong to at least one of vertices in i. This successful measure has been adopted not only to benchmark the accuracy of the clustering but also as the *fitness* value for clustering algorithms based on optimization. Finding the partition of groups that maximizes Q is believed to be a *NP-hard* problem, which makes a brute force exploration impossible for networks bigger than dozens of vertices. However, several search heuristics can be applied to explore the huge space of states in order to find a good partition. Following this approach, many algorithms have investigated different exploration heuristics to find the community structure while maximizing Q. Newman proposed in [176] a hill-climbing heuristic to create the hierarchy following an agglomerative strategy. The baseline is that every single node is a cluster, then the pair of clusters whose union produces the biggest increment in Q are merged into a single one. The process is repeated until only one cluster remains. By following the merging operations, the hierarchy that reveals the community structure is built. However, a *hill-climbing* heuristic cannot escape a sub-optimal maximum. Therefore, other search heuristics were devised. For instance, Guimerà and Amaral [108] proposed a simulated annealing approach. Duch and Arenas [74] proposed an algorithm based on Extremal Optimization. Both algorithms were able to extract the community structure more accurately in terms of modularity, although they were not as efficient as Newman's Fast algorithm [176]. Newman has very recently proposed another clustering algorithm [177] which outperforms the previous algorithms in both modularity and efficiency, although it is

not as efficient as his previous fast algorithm [176]. Danon et al. [61] have also very recently presented a modification of Newman's Fast algorithm that while maitaining its computation efficiency yields more accurate partitions in terms of modularity.

Modularity optimization methods are neither the first nor the only ones to work on clustering in complex networks. In [101] Girvan and Newman reviewed classical hierarchical clustering algorithms on networks, showing that some classical distance measures were not well suited to work with complex networks. While the review done by Girvan and Newman in [101] was essentially correct, it overlooked two relevant areas that were already working in clustering of complex networks. Sociology was addressing clustering in Social Network Analysis [225]. On the other hand, Computer Science was also working in clustering of a particular instance of complex networks: the Web. Gibson et al. [96] and Kumar et al. [146] addressed clustering based on the analysis of the links between Web pages. A common tool used to address clustering on the Web is spectral analysis. However, this technique is applicable to any kinds of network, for instance, newsgroups [34] and protein networks [39]. Spectral analysis has also been used in many other areas besides clustering. For instance, in work-load distribution between processors [204] and to find the relevant vertices of a network [183, 137]. Obviously, not all clustering is limited to spectral analysis. Flake [88] proposed an alternative approach based on minimum cut-trees over expanding networks that worked over the Web and could be applied to other kind of networks. However, we find particularly interesting clustering methods based on random walks [112, 37], which can be seen as a particular case of spectral analysis. The underlying idea behind clustering using random walks is very intuitive: if a random walker starts from a given node, it will tend to visit more often vertices that belong to the same community of this initial node. So, provided that a community structure exists, a random walker will spend most of the time stuck within the community it started from.

Our algorithm is a combination of spectral analysis and modularity optimization. We adopted this combination in order to achieve a good compromise between efficiency and accuracy of the clustering. Spectral analysis is used to reduce the number of initial vertices of the network: by means of a set of random walkers we create an initial partition of the network formed by a number of groups much smaller than the initial number of vertices. Consequently, the number of merge operations required to build up the hierarchy is reduced. Asymptotically our algorithm has a complexity $O(n^2)$, which is the same complexity of Newman's fast algorithm [176]. However, in terms of computational cost it is more efficient since the complexity can be decomposed as $O(ns) + O(s^2)$, where n is the number of vertices and s is the number of groups in the initial partition produced by the random walkers. Despite s being smaller than n it is not upper-bounded by a sub-linear function of n, so that the complexity remains $O(n^2)$. Yet, it is clearly more efficient and allow us

to analyze very large networks in reasonable time while maintaining high quality clustering.

5.2 Algorithm

The proposed algorithm, henceforth PBD (after the initials of the authors), consists in an agglomerative hierarchical clustering where the initial groups are those produced by an initial partition of the network. The first step of the algorithm consists of a process of s random walkers traversing the network. The transition probability matrix M is defined as

$$M = (A + I)D^{-1} \tag{5.2.2}$$

Where I is the identity matrix and D is a diagonal matrix of the form $D_{ii} = 1 + \sum_j A_{ij}$. Thus, M_{ij} is the probability to go to node j from node i. The process carried out by the random walkers is defined by

$$G^{t+1} = M'G^t, \tag{5.2.3}$$

G^t is the matrix that contains the probability distribution of each random walker, G^t_{ij} is the probability that the random walker j is at node i at time t. Usually, the process is repeated iteratively until the stationary state is reached. However, we are interested in the transient state for all random walkers, consequently the process is repeated until we obtain G^T, where T is set to 3. Therefore, each random walker has done three jumps, which is the minimum number of hops to complete the shortest path to the origin point. Once the stochastic process is finished, each node i is classified into the group j which corresponds to the largest column at row i in G. Through this process, the initial n vertices are classified in approximately s groups and all the vertices of the same groups share that they were visited the most by the same random walker. Consequently, this means that they have a high degree of neighbours in common which implies a community. While this method is far from perfect, it allows us to drastically reduce the initial number of groups. The final number of groups might not correspond exactly to s since random walkers could preclude others. A random walker i is precluded when all vertices by a random walker i are also visited more often by other random walkers, consequently the visited nodes are classified into others groups rather than the group started by i. Further, since the markov process is only iterated T times there is no guarantee that all vertices will be visited at least once, in this case an extra group with a single node is created.

The partition of the network heavily depends on which vertices are seeds — origins — of the random walkers. This problem is very related to classical clustering

algorithms [121, 122] such as *k-means* [90]. How many seeds are required and where to place them is an open question [36]. We propose a straight forward heuristic that selects which vertices will be the seeds for the random walkers, i.e. to define G^0. Let R be the fraction of the most connected vertices chosen as seeds. If $k_i \geq z$ a random walker will start at node i, where the connectivity threshold z is defined as the maximum connectivity that makes the partition composed of the most connected nodes larger or equal than R, $\sum_{j=z}^{j \leq max(k)} p(k_j) \leq R$, where $p(k_j)$ is the fraction of vertices with connectivity k_j. The parameter R allows us to decide approximately the number of seeds, although the initial number depends on the structure of the network. For $R = 1$ there would be too many seeds for the algorithm to be efficient. On the contrary, $R \sim 0$ would be very efficient but the partition would be ill-constructed. In our experiment we set R to $\frac{1}{5}$ obtain good results for a wide variety of networks, as shown in Table 5.1 and 5.2. Future work will look into different heuristics to choose the seeds, the quality-efficiency tradeoff of our algorithm is really dependant in this process, and other heuristics more elaborate might provide better results than our current straight-forward selection rule.

The complexity of finding the seeds for the random walkers is $O(n)$. The connectivity distribution and the connectivy threshold z can be computed in linear time respect to the number of vertices. In fact, the first approach we tried was to get the $n \cdot R$ most connected nodes, which would entail a sort operation with cost $O(nlogn)$. This option was discarded in favor of the connectivity threshold due to the extra cost which our algorithm intends to minimize.

The stochastic process defined in equation 5.2.3 has the iterative multiplication of two matrices, M and G, of dimension $n \times n$ and $n \times s$ respectively. However, thanks to the sparseness of both networks the cost can be reduced from $O(n^2 s)$ to $O(ms)$, where m is the number of edges. For each random walker j its probability distribution can be calculated in the worst case scenario with cost $O(m)$. Thus, the final cost can be considered $O(ns)$ because the number of edges scales with n in the limit of large n.

Once the initial partition is created, the algorithm builds an agglomerative hierarchical clustering. This method consists in creating a series of partitions of the data: $C_s, C_{s-1}, ..., C_1$, where first C_s consist of s single clusters (groups), and the last C_1, consists of a single group containing all the individuals. The method iteratively joins the two individuals or clusters (groups of individuals) which are most similar. Thus, after $s - 1$ join operations the clustering is complete, and the result is a binary tree known as *dendrogram* that reveals the underlying structure of the data.

Let us say that the initial partition yielded s groups, despite the fact that it is a upper bound since some groups might be empty because their random walkers were precluded by others. For each group j, the contribution to the total modularity, that

is $q_j = e_{jj} - aj^2$, can be calculated in linear time $O(s)$. The group that contributes the least to the total modularity Q — let us say j such that $j = argmin_k(q_k)$ — is selected to be joined to the group that maximizes the increment of modularity as defined in the following equation,

$$\triangle Q = (2e_{ij} + e_{ii} + e_{jj}) - (a_i + a_j)^2 - (e_{ii} - a_i^2) - (e_{jj} - a_j^2) \tag{5.2.4}$$

The increment in total modularity is the modularity of the merged group $(2e_{ij} + e_{ii} + e_{jj}) - (a_i + a_j)^2$ minus the contribution to the modularity of both groups; q_i and q_j. Equation 5.2.4 can be reduced trivially to Eq. 2 of Newman's Fast Algorithm [176],

$$\triangle Q = 2e_{ij} - 2a_i a_j = 2(e_{ij} - a_i a_j) \tag{5.2.5}$$

In the event that two candidates, i_1 and i_2, have the same effect over the total modularity, the candidate group chosen will be the one that has the least modularity, $min(q_{i_1}, q_{i_2})$. Thus, groups with low modularity are preferred in the merge operation.

The merge operation can be performed in the worst case scenario in linear time with respect to the current number of groups, thus $O(s)$. Furthermore, the operation needs to be done $s-1$ times. Therefore, the complexity of building up the hierarchy is $O(s^2)$. The search heuristic proposed is extremely *greedy* since it only takes into consideration pairs of groups, provided that one groups is fixed. Conversely, Newman's fast algorithm calculates the gain of modularity for each possible pair of groups. Besides that, other algorithms based on modularity optimization usually have even more expensive search heuristics that allow a better exploration at the expense of efficiency. Our proposal was designed to focus on efficiency, as it can be seen in the heuristic decisions made by the algorithm. However, as we will show in the experiments section, this focus on efficiency does not necessarily imply a loss of quality of the clustering.

5.2.1 Parallelization

To reduce even further the execution time of the algorithm its parallelization could be easily implemented. The stochastic process as defined in equation 5.2.3 can be carried out in parallel by different computers or processors. To calculate the probability distribution of a given set of random walkers, only the transition probability matrix is required. Thus, the matrix G of dimension $n \times s$ could be split column by column into a set of smaller matrices of dimension $n \times \varsigma$, where $\varsigma \ll s$. This would drastically reduce the cost of the stochastic process. Unfortunately, the modularity optimization step cannot be parallelized so easily. Thus, trivial parallelization would only affect the spectral analysis part of the algorithm. Although this part is

the most expensive part in the algorithm $O(ns)$, the modularity optimization $O(s^2)$ would still be executed sequentially. Therefore, the asymptotic cost of the PBD algorithm would still be $O(n^2)$.

5.2.2 A Working Example with Zachary's Network

To illustrate our algorithm we include an execution on the Zachary network [246], which is a well-known dataset in the literature of community extraction. In figure 5.1 we can find the network at the different stages of the execution of the algorithm. Figure 5.1.a shows which vertices, labelled 2, are seeds of the random walkers, 16 in total. Thus, s is 16 compared to the original 34 vertices. In table 5.2 the relation between network size n and the number of random walkers s for a wide range of networks is shown. Figure 5.1.b shows the initial partition of the networks produced by the random walker process, described in equation 5.2.3. The initial 34 vertices are grouped into 13 groups. This partition has a modularity Q of 0.1547. From that point on the algorithm starts the modularity optimization stage governed by equation 5.2.4. At each step, two of the remaining groups are joined according to equation 5.2.4. Figures 5.1.c and 5.1.d show the network divided into 4 and 2 groups respectively.

The community structure of Zachary's is better seen in figure 5.2. The maximum modularity is obtained by the partition in 4 communities, achieving Q=0.3937. However, the original empirical work on the Zachary Karate Club [246] found two communities: those aligned with the instructor and those aligned with the administrator. The division into two groups produced by PBD algorithm produces a high modularity Q=0.3718. Also, the two original communities found empirically correspond to the communities found by the algorithm with the exception of node 10, which is misclassified.

It is evident that figure 5.2 is not a dendrogram over all the vertices of the network, but a dendrogram over the initial communities. As a consequence of the random walker process, the structure between vertices belonging to the same initial community remains unknown. However, this loss is negligible since the relevant high-level structure is not affected as can be seen in figure 5.2.

5.3 Experiments

In order to further analyze our algorithm we chose a set of ten different networks of different sizes, ranging from 34 to 498925 vertices. The networks modeled a wide spectrum of systems. There are social networks, such as the Zachary Karate Club [246] and the social network of the Software Department (LSI) at the Technical University of Catalonia (see section 4.3). Scientific collaboration networks such as

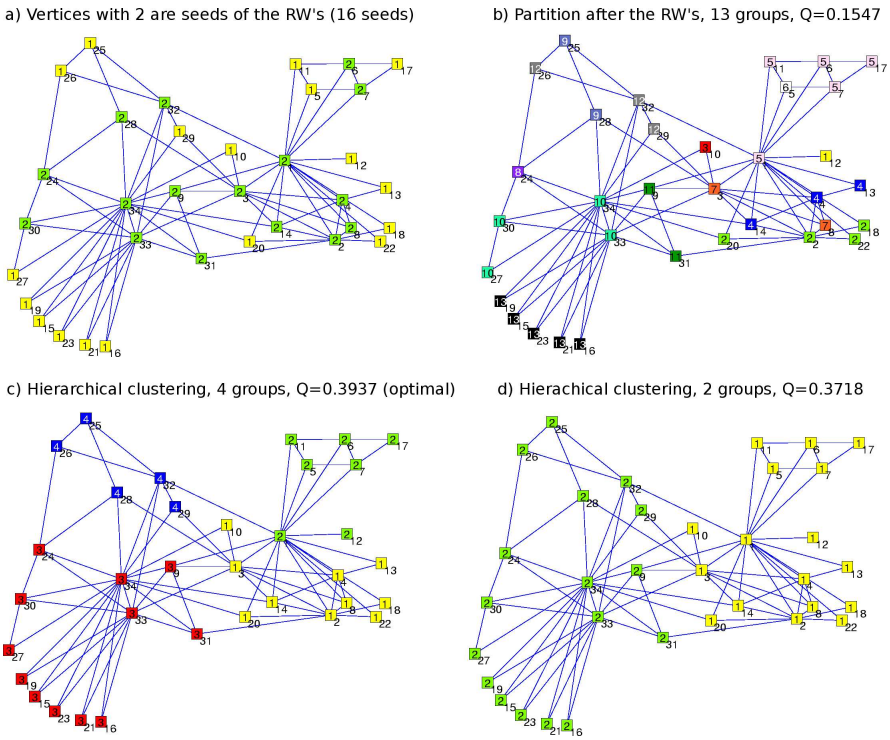

Figure 5.1: Working of the PDB algorithm in the Zachary network. Sub-figure *a* shows the initial vertices (labelled with 2) which are seeds for the random walkers vertices labeled. Sub-figure *b* shows the initial partition of the network into communities produced by the random walker stage. Sub-figures *c* and *d* show different partitions created by the modularity optimization process. The optimal partition, whose modularity is maximal, is shown in sub-figure *c*.

Cond-mat [173] and the Erdös collaboration network [1]. Citation networks such as Scientometrics [2]. Affiliation network among Spanish top director boards [3]. Network of relations between words such as WordNet [4]. Metabolic networks such as the *C. Elegans* [126]. A portion of the Web from the Notre Dame University dataset [12]. Finally, the last type of network was the movie collaboration [5] network, again

[1] Erdös Number Project. The network contains scientists with Erdös number less than or equal to 2 up to year 2002. http://www.oakland.edu/enp/thedata.html

[2] Network from Garfields's collection of citation networks. Available at http://www.garfield.library.upenn.edu/histcomp/index.html

[3] Data provided by Prof Fabrizio Ferraro, from the project *Small Worlds Of Corporate Networks* at IESE Business School, University of Navarra

[4] Network obtained from the Pajek Network dataset. Available at http://vlado.fmf.uni-lj.si/pub/networks/data/

[5] This network is a bipartite network (two-mode). Consequently, vertices have different roles and thus the structure extracted by the algorithms might be biased. We did not transform it to

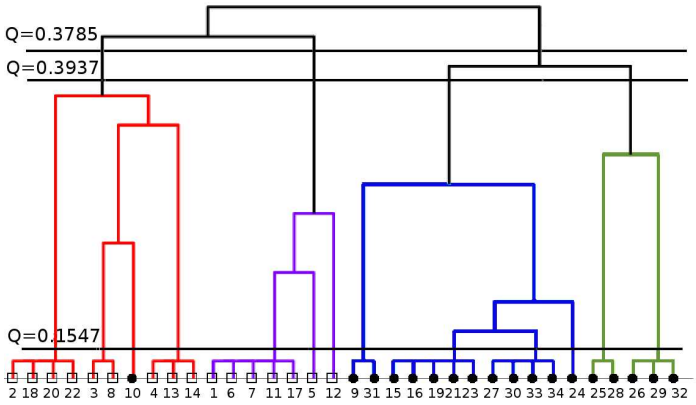

Figure 5.2: Community structure of Zachary network produced by PDB algorithm. Circles and squares over the individuals denote to who they align to after the karate club broke up. Those who aligned with the instructor are represented by circles, and those who aligned with the administrator are represented by squares.

obtained from the Notre Dame University dataset [22]. In all the networks we only worked with the biggest connex component, removing all multiple relations and self-reference edges.

Comparing Modularity

Table 5.1 summarizes the highest modularity achieved by Newman's Fast algorithm (Q_N) [176] and our algorithm (Q_{PBD}). For eight out of the ten tested networks the PBD algorithm produces a higher modularity, and the maximum difference in favor of the PBD algorithm is in the cond-mat network. Thus, in general the PBD algorithm yields a slightly better modularity than the Newman's Fast algorithm. However, as mentioned in the introduction, there exist in the literature other algorithms based on modularity optimization that also outperform the Newman Fast algorithm [176]. In table 5.1 we also included the results obtained by the extremal optimization algorithm (EO) by Duch-Arenas [74]. In this case the maximum modularity obtained by EO outperforms in two of three cases the modularity obtained using PDB, and in all the available cases it outperforms the modularity obtained using Newman's Fast Algorithm. However, the complexity of the algorithms that use elaborated search heuristics is superior to the complexity of both Newman's Fast algorithm and PDB algorithm. For instance, EO's complexity is $O(n^2 log^2(n))$. Thus, we can conclude that PDB has a good balance between efficiency and quality.

a one-mode network to maintain the maximum of vertices. Thus, the results on this network are only relevant when speaking about the algorithm's efficiency

Network	*Size (n)*	Q_N	g_N	Q_{EO}	g_{EO}	Q_{PBD}	g_{PBD}
Zachary	34	0.3807	3	0.4188	4	0.3937	4
LSI	139	0.6428	6	–	–	0.6604	6
C. Elegans	453	0.40	10	0.4342	12	0.4164	7
Directors Board	598	0.8046	21	–	–	0.8273	16
Scientometrics	2678	0.5555	24	–	–	0.5629	10
Erdös (2002)	6927	0.6723	57	–	–	0.6817	20
Cond-Mat	27519	0.6653	324	0.6790	647	0.7251	44
Word-Net	75606	0.7963	453	–	–	0.7885	47
WWW ND	325729	0.9273	2192	–	–	0.9272	83
Actors ND	498925	0.7243	2113	–	–	0.7297	14

Table 5.1: Comparison between maximum modularity Q and number of communities g obtained by Newman's fast algorithm (N), Duch and Arenas Extremal Optimization algorithm (EO) and PBD algorithm. Results from EO algorithm are limited to those published in [74]

Comparing the Number of Communities

Clustering, though, does not only depend on obtaining the partition that maximizes modularity. The number of communities — or groups — contained in the optimal partition is also very important. In table 5.1 the number of communities is denoted by g_N, g_{EO}, g_{PBD}. We can clearly observe an order $g_{EO} > g_N > g_{PBD}$. Provided that maximum modularity does not greatly differ between partitions, the huge differences in the number of communities obtained by the different algorithms are indeed striking and requires further study.

Provided we assume that modularity Q is a good measure for community structure, we must take for granted that two partitions with similar modularity are equally accurate. Thinking otherwise would lead us to the conclusion that modularity is not a good measure. Finding a *bogus* partition that yielded higher modularity than a *good* partition would mean that modularity is not representative of the structure. Consequently, it could not be used as the fitness variable for the optimization. However we do not believe that this is the case.

So, if we assume that modularity is a good measure, what happens when two partitions having very similar modularity have a very different number of groups? A first approach would be to think that the partition with smaller number of groups is more general than the partition with the larger number of groups. Thus, partitions with smaller number of groups would be, in principle, more interesting for several reasons. (i) They would provide a more general perspective on the underly-

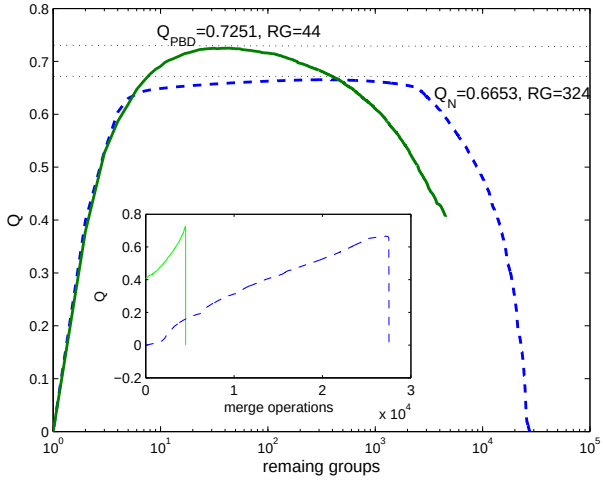

Figure 5.3: Modularity Q over the execution of the algorithms for the cond-mat network. PBD algorithm is shown as the solid line and Newman's Fast algorithm (N) as the dashed line. Notice that the number of merge operations required for both algorithms is different. PBD algorithm requires approximately the 22% of the merge operations performed by the N algorithm. Both the figure and the inset contain the same information, although the main figure is plotted in a log scale to magnify the last steps of the algorithm when maximum modularity is found.

ing structure of the network, since they would be able to find a meaningful partition at a higher level of the structure. (ii) A small number of groups would simplify the analysis of the obtained results. And (iii), general or *high-order* partitions could always be re-clustered to further analyze the structure of a particular group if a more detailed — fine-grained — analysis were required. The opposite could not be done, otherwise the algorithm would have detected the more general partition with higher modularity.

Table 5.1 shows that PBD algorithm yields more general partitions while having similar or better modularity. This effect is specially acute in large networks. One might conclude from the results of the experiments that both EO and N algorithms undergo an unnecessary over-specialization. For instance, let us we take the optimal partition of the cond-mat network. PBD provides the highest modularity and the smallest number of partitions, which is 647, 324 and 44 using EO, N and PBD algorithms respectively.

Figure 5.3 shows the evolution of Q in the cond-mat network using N and PDB algorithm. Notice that the modularity obtained by the N algorithm is very close to its maximum modularity when the number of remaining groups is in the range

between 40 to 1000. Once the partition contains more than 1000 groups it does not increase or decrease substantially until the number of groups in the partition reaches 40, at that point, modularity starts to decrease abruptly. This range of stable modularity contains a partition with 44 groups which is the maximum modularity found by PBD algorithm. This behaviour is not observed in the evolution of modularity for the PBD algorithm, which keeps increasing modularity until it gets to the maximum and then it starts a quick descent. Being in a *plateau* where modularity is neither increasing or decreasing significantly might suggest that the search heuristic of N algorithm could be stuck in a local sub-optimal state.

A more detailed study of the internal behaviour of N and PBD algorithm is described in figure 5.4, where the evolution of the normalized size of groups to be merged is depicted. Each merge operation joins group i and j in a new group z. The size of the two groups chosen by the algorithm to be merged is expressed as $r_{mo} = \frac{min(s_i,sj)}{s_i+s_j}$. The ratio r_{mo} has values in the range $(0..\frac{1}{2}]$, when r_{mo} is close to zero means that there is one group that is clearly bigger than its counterpart. When $r_{mo} = \frac{1}{2}$ both groups have the same size. It is easy to see how the N algorithm tends to join groups of very different size, whereas the PBD algorithm tends to do the opposite, i.e. it prefers groups of similar size. Consequently, the N algorithm produces at the very beginning groups of extremely large size by joining a single vertice to a very large group, when this group cannot accept more vertices a new group is started. This behaviour is clearly observed at operation 6000, 10000 and 15000 approximately, creating three massive groups of 13000 vertices and leaving the rest of groups composed mostly by individual vertices. This might be the reason why the search heuristic fails to increase modularity in the *plateau* of figure 5.3. Approximately from operation 16000 to 26000 this behaviour is not so evident, where groups of similar same size are often merged. However, the behaviour that merges very dissimilar groups appears again between operations 26000 and 27000, overlapping the *plateau* of modularity. At this point, big groups cannot be merged obtaining a gain of modularity, therefore, the only possible option left to be explored by the search heuristic is to merge *left-over* groups of very few vertices into the existing big groups. As a consequence, the search heuristic gets stuck in a situation where changes of modularity are minimum and there is no escape from the sub-optimal partition. The bias of the N algorithm towards creating very large groups has also been very recently reported by Danon et al. in [61]. We can see how this behaviour is not present in the PBD algorithm. The greedy search heuristic used by PBD algorithm forces the worst group to be the one going to be merged, and this results in groups having similar sizes. By doing so, the search heuristic has many possible combinations to explore, avoiding getting trapped too early in a local sub-optimum as it happened in the N algorithm. That is the reason why the PBD algorithm is able to find in the cond-mat network a higher modularity partition

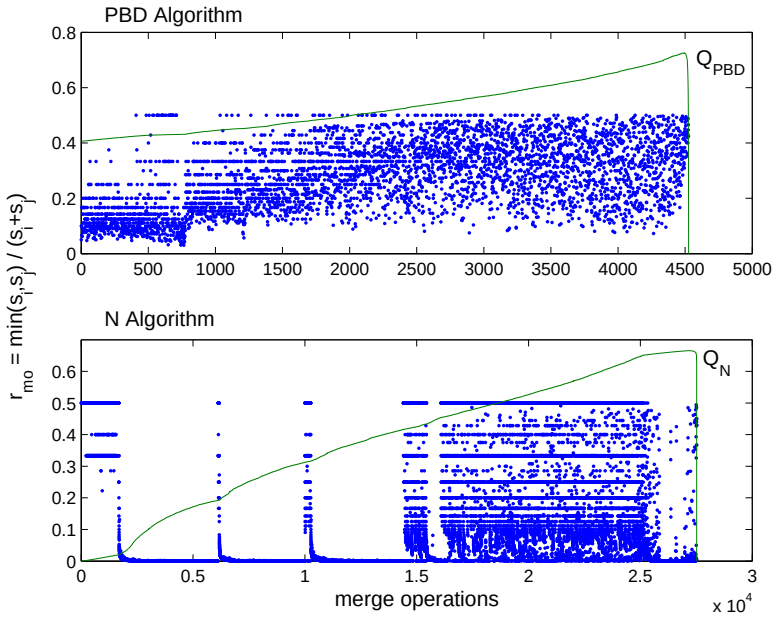

Figure 5.4: Behaviour of the merge operation of the PBD (above) and the N (below) algorithms. For each merge operation the normalized ratio of the size of groups to be merged is calculated as $r_{mo} = \frac{min(s_i,sj)}{s_i+s_j}$, where s_i and s_j are the size of group i and j respectively. The merge operation creates a new group $z = i \cup j$ of size $s_z = s_i + s_j$. The line corresponds to the evolution of modularity.

which also has less groups. As for the EO algorithm, we could not carry out the same analysis but it is plausible that the resulting large number of groups can be attributed to its divisive clustering strategy.

In order to analyze with more detail the partitions yielded by the N and the PDB algorithms the distance between the optimal partitions for the cond-mat network are calculated. Gustafsson et al [110] reviewed different distance measures to compare partitions of the same network. Since we want to look into the hypothesis that the partition yielded by PDB (P_{PBD}) is more general than the partition yielded by N algorithm (P_N) we chose the m_{div} measure, which is the minimum number of divisions to be applied to partitions A and B to obtain the partition C defined in [109] as,

$$C = \bigcup_{i=1}^{|A|} \bigcup_{j=1}^{|B|} (a_i \cap b_j) \tag{5.3.6}$$

Partition C is the union of all possible intersections between the groups in A and B. We rename C as P_{N-PBD}. The distance between partition P_{PDB} and P_{N-PBD}

Network	*Size (n)*	t_N	t_{PBD}	s_{PBD}
Zachary	34	0.002	0.014	16
LSI	139	0.003	0.015	42
C. Elegans	453	0.026	0.064	118
Directors Board	598	0.038	0.031	125
Scientometrics	2678	1.6	0.320	619
Erdös (2002)	6927	3.14	2.6	2155
Cond-Mat	27519	125.8	11.2	6224
Word-Net	75606	490.6	204.1	38701
WWW ND	325729	10932.1	1775.6	86908
Actors ND	498925	34208.3	3326.3	118897

Table 5.2: Comparison between CPU-time t (in seconds) between Newman's Fast algorithm and the PBD algorithm. It also includes the number of random walkers required to create the initial partition s_{PBD}

is 894, which is the number of divisions to be applied to P_{PBD} in order to obtain P_{N-PBD}. The distance between P_N and P_{N-PBD} is 614. Thus, the total distance is 1508, which is the sum of both distances. In order to have a baseline comparison for the distance between P_{PBD} and P_N we created a random partition P_{rand} with the same cardinality as P_{PBD}. The random partition was replicated 30 times and so was the measurement, the average distance between P_N and P_{rand} was 7166.7 with a standard deviation of 33.65.

Comparing Efficiency

After comparing the quality of the clustering produced by the PBD algorithm we must turn our attention towards its performance. As we already mentioned, all design decisions were biased towards improving the efficiency so that the algorithm could cope with medium and large networks, which other algorithms cannot handle in reasonable time. Table 5.2 summarizes the run-time (cpu-time) of our algorithm compared to Newman's Fast algorithm, which is the reference algorithm due to its efficiency. We are perfectly aware of the problems related to comparison of algorithms based on run-time instead of only considering their complexity. In order allow a fair comparison between both algorithms we implemented them from scratch optimizing them to the best of our abilities. Needless to say the runs were executed on the same desktop computer (Pentium 4, 3Ghz) exclusively devoted to the experiment. Table 5.2 shows that PBD is much faster than Newman's Fast algorithm for networks bigger than a thousand vertices. Conversely, the PBD algorithm is slower than the N algorithm for small networks. This is due to the two sequential processes — spec-

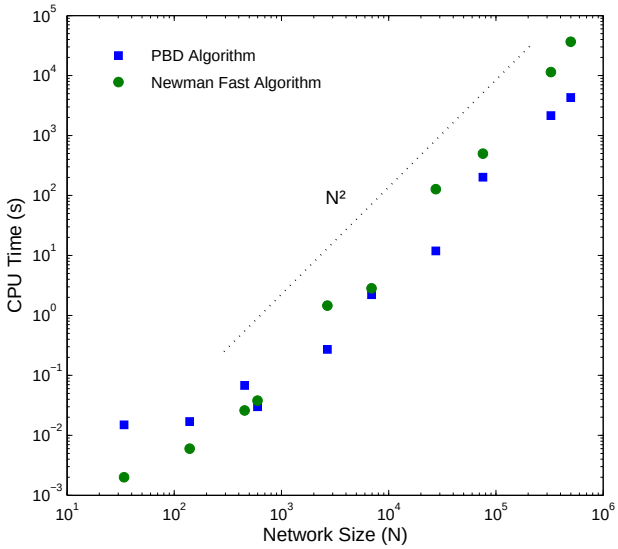

Figure 5.5: Comparison between the efficiency in CPU-time in seconds for 10 networks of different sizes. The solid circles show the results for the Newman's Fast algorithm. The results of the PBD algorithm are represented by solid squares.

tral analysis and modularity optimization — that take place in the PBD algorithm. The difference in execution time heavily depends on the number of random walkers (s_{PDB}) required for the first stage of PDB algorithm. The smaller the ratio between network size and the number of random walkers, the faster the PDB algorithm is.

Figure 5.5 shows graphically the relation between network size and the running time of both algorithms already seen in table 5.2. The PDB algorithm clearly outperforms Newman's Fast algorithm, often by one order of magnitude. However, the asymptotic quadratic behaviour of both algorithms is evident, which lead us to think that our algorithm cannot scale to very large networks of millions of nodes. Resorting to parallelization would allow us to analyze networks of a few million nodes in an acceptable time, however, PDB will undoubtedly become too slow for very large networks. Nonetheless, it allows to shift the network size threshold far enough to be useful for medium and large networks. Reduction of approximately one order of magnitude allows to shift from minutes to seconds, or from hours to minutes. For instance, the clustering of the largest network we had access to was reduced from approximately 9 hours to just one.

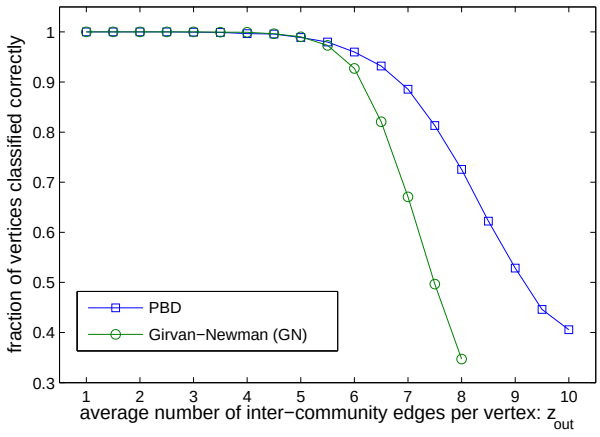

Figure 5.6: Fraction of vertices correctly classified using computer-generated networks. The circles show the result of the Girvan-Newman algorithm (GN). The squares show the results of the PBD algorithm. Each point is an average over 50 networks.

5.3.1 Computer-generated Networks

To conclude, we include experiments carried out with the computer-generated networks first proposed by Girvan and Newman [101], which have become a common testbed in the field. Those networks are constructed with 128 nodes divided into four groups of the same size. For each node 8 edges are deployed. With probability P_{in} the edge is connected to a node that belongs to the same group chosen at random. Otherwise, the edge is connected to a node that does not belong to the same group. Thus, the average degree of a node is 16. Accordingly to the nomenclature of Girvan-Newman, we will use z_{out}, which is the number of *inter-community* edges per node. It is important to notice that for $z_{out} = 12$ the network is totally random, that is, without community structure. Another important value for z_{out} is 8, since it marks the boundary between having more *inter-community* than *intra-community* edges. The quality measurement is the fraction of vertices that are correctly classified, explained in more detail in [176].

In figure 5.6 we can see the results obtained by our algorithm compared to the results yield by the Girvan-Newman (GN) algorithm based on edge betweenness [101]. We decided to use the GN algorithm instead of the Newman's Fast algorithm since it obtains slightly better results and it is the reference algorithm for this particular experiment [101]. Although the results of GN outperform those obtained by the N algorithm it is not a suitable option for medium and large networks due to its complexity, that is $O(n^3)$.

The GN algorithm correctly detects the communities until values of $z_{out} = 6$ are

reached. From this point on, the quality of the communities decreases very quickly. On the other hand, the PBD algorithm detects the communities very well up to values of $z_{out} = 7$, from that point on the performance starts to decay, although the pace is more steady than in the GN case. As already mentioned, other algorithms based on modularity optimization perform much better in the computer-generated networks example. For instance, Duch and Arenas extremal optimization algorithm [74] starts to decline at $z_{out} = 8$. However, its increase on clustering quality is done at the expense of efficiency. So, it is unadvisable in the case of large networks.

5.4 Discussion

In this chapter we have presented an algorithm to extract community structure from networks by using a combination of different existing methods. First, the algorithm uses spectral analysis, via the multiple random walker process, to reduce the dimensionality of the network by creating the initial partition of the network into communities. Then, a modularity optimization process with an extremely greedy search heuristic is applied to extract the underlying structure of the network.

Experiments show that our algorithm outperforms Newman's Fast algorithm both in clustering quality and efficiency. Newman's Fast algorithm is the reference algorithm in terms of efficiency, and while asymptotically both algorithms are $O(n^2)$, PDB algorithm is always faster in computation time for medium and large networks, as it has been shown in the experiments.

The reason behind this is the reduction of dimensionality provided by the random walker process, such that the cost of PBD can be expressed as $O(ns)$ where $s \ll n$. Furthermore, experiments also show that PDB retrieves more *general* community structure than other algorithms. The number of existing communities in the partition with maximum modularity is notably smaller in the case of PBD. This fact leads us to think that other algorithms tend to unnecessarily over-specialize their clustering.

In summary, the presented algorithm is an interesting choice when analyzing medium and large networks. The structure of large networks can be found in reasonable time, from seconds for a network of 27k vertices to less than one hour for a network of 500k vertices. The gain in efficiency does not come with a loss in the quality of the clustering as the maximum modularity obtained by the algorithm is comparable to the reference algorithms in the literature. Throughout the background chapters of this thesis we saw that complex social systems, such as multiagent systems, can be characterized by networks. The analysis of these networks can provide useful knowledge both at the scale of individuals and at the scale of the system. As a matter of fact, one must take into account that the existing knowledge of a system is not only contained in its constituents but rather on the relationships

between constituents of the system [134]. Consequently, the need to take structure into consideration is out of question for a comprehensive understanding of the system. With that in mind, we devised two different algorithms that extract knowledge only resorting to the structure.

In the previous chapter we showed how the network structure can be used to infer the relevance and reputation of an individual from its position in the network. In the current chapter we used the network topology to identify the underlying community structure; groups as well as relationships between those groups can be revealed by analyzing the structure through clustering.

The knowledge retrieved by the methods here presented, as well as by other related methods in the literature, can give a better understanding of the system under study. However, our aim is not limited to a descriptive analysis, our goal is to use structure to retrieve information that is inherently social so that it can be used by multi-agent systems designers and the agents themselves to build an open, stable and full-fledged artificial societies. Needless to say that the two algorithms presented do not suffice to achieve such an ambitious goal. However, finding social aspects such as reputation and groups membership can contribute to leverage the uncertainty and complexity of open systems like artificial societies.

Part III

System Dynamics and Structure

Chapter 6

Emergence of Conventions

In the previous part of the thesis, structure was used to extract relevant knowledge of the system which is embedded in relationships rather than in individuals. Therefore, one might argue that some *social* aspects of systems must be tackled by looking at the pattern of interactions, that is to say, the topology of the network. Despite the trivial look of this statement one must consider it in contrast to the common approach to the study of systems, where the focus is placed on the reductionist study of the components taking — wrongly — the system as the simple addition of all its components, and so, neglecting the interactions between those components.

Lazebnik [147] clearly illustrated this problem by wondering whether a biologist could ever fix a broken radio; *"what matters for the radio — as well as for the cell or organism — is not only what is there [components] but, perhaps more importantly, how they are connected"*. Lazebnik remark is not limited to Biology though, many other areas of knowledge suffer from the same methodological bias.

Analysis of the structure allows us not only to retrieve hidden information about a given system, but also to understand its internal dynamics. In this part of the thesis we want to emphasize the role played by the structure in the system's dynamics. The pattern of interactions to be followed by the individual actors is determinant in what respects the long-time behaviour of a dynamical system, and multi-agent systems are no exception.

6.1 Description

The study of social conventions in human societies is more than forty years old [149, 214], see section 2.3 for a comprehensive review. Social conventions provide society with rules of behavior in which all agents agree, and so they make easy to deal with situations where several individuals may have conflicting goals. In this sense, multi-agent systems (MAS) are not different. They also need some (either ex-

plicit or implicit) rules to avoid conflicts and to facilitate interactions among agents. One means of implementing such rules in multi-agent systems are *conventions*. Essentially, a convention gives a common choice of action to all agents in conflicting situations.

Two ways of introducing conventions in MAS have been explored: off-line design, where every agent has the conventions "hard-wired" from the beginning, and the on-line design, also known as emergent design, where the collective of agents decides, through interaction, which are the most suitable conventions given the current state of the system. The former design strategy is clearly unsuitable in dynamical and changing environments, where one cannot know *a priori* which will be the conditions under which the system will operate (this has been argued in [203, 221]). In this case, the dynamical nature of on-line conventions appears to be most appropriate. It is not difficult to think of other situations where these non-fixed conventions may be an advantage, for example in the case of agents with changing goals.

In the simplest multi-agent system one can think of, every agent may interact with every other agent. This means that the underlying topology is a network with an all-to-all connectivity pattern, that is, a complete network. However, this is not very realistic. It is far more accurate to assume some restrictions in the pattern of interactions that an agent may develop. We can think of different possibilities: Regular graphs, lattices, etc. This has already been (partially) analyzed, since emergence of conventions in MAS with topological restrictions has been studied in regular graphs and lattices [135, 136, 221]. This work is quite interesting, since it shows that the underlying MAS topology is important in the efficiency of the emergence of conventions; however, regular topologies are not very realistic either. If we pay attention to the topology of *real* networks, we will find out that most of them have a very particular topology: they are *complex networks* with non-trivial wiring schemes. This particular topologies are found in many empirical systems as mentioned in section 3.2.

One of the most suitable environments for a MAS, the Internet, is among the most prominent complex networks found in the real world. Complex networks are well characterized by some special characteristics (see subsection 3.2.1) such as the connectivity distribution (either exponential or power-law) and the *small world* property [181, 232].

In this chapter we study the efficiency of the emergence of social conventions in MAS with a complex underlying topology. We followed the conceptual framework introduced by Shoham & Tennenholtz [201, 202, 203] and our measure of efficiency will be one of those introduced in the work of Kittock [135]: the time it takes to reach a 90% of the agents in the system to use the same convention.

6.2 Network Models

There are several models of graphs we are going to use to represent the underlying pattern of interaction between agents in our MAS.

As we pointed out in section 3.2, recent discoveries on empirical networks lead us to think that regular or complete graphs are not the most realistic environment for MAS. Complex networks [12, 13, 22, 53, 181, 227] are found to be most suitable for MAS, since real systems composed by autonomous entities without a global planner form complex networks. So, its natural to choose those networks to model the underlying pattern of interactions between agents. The graphs we will use as the pattern of interaction between agents are:

- *Complete* graph K_N, it the complete N-nodes graph, where every node is adjacent to the all the nodes but itself. Thus, connectivity k is $N-1$. The number of edges of the graph is $\frac{N(N-1)}{2}$, which is a clearly unrealistic when compared to empirical networks.

- *Regular* graphs or lattices $C_{N,k}$, it is the graph on N nodes such that node i is adjacent to nodes $(i+j)modN$ and $(i-j)modN$ for $1 \leq j \leq k$, where k is the connectivity.

- *Small world* graphs W_N, these are highly clustered graphs (like regular lattices) with small characteristic path lengths (like random graphs) [232, 227]. This is the small world property. We will choose the Watts-Strogatz model as model of small world graphs.

- *Scale-free* graphs S_N^γ, these are graphs with a connectivity distribution $P(k)$ (the probability that a node has k adjacent nodes) of the form $P(k) \propto k^{-\gamma}$. We will choose the Albert-Barabási extended model as model of scale-free graph.

As it was mentioned in subsection 3.2.1, there are many models of complex networks. We have chosen Watts-Strogatz [232] model for small world graphs and the Albert-Barabási [10] *extended* model of scale-free graphs. We decided for the later since it gives us some control over the exponent γ of the graph. The underlying idea is that of growth with preferential connectivity, where the most "popular" nodes get most of the links. This model was built on a simpler one [22, 23], able to generate graphs with exponent $\gamma = 2.9 \pm 0.1$ (by setting $p = q = 0$ in the algorithm detailed in section 3.2.1 we recover this previous model).

Albert-Barabási extended model of scale-free graphs have not the small-world property, although the graphs are random. We will see below that it is of interest to us to study the consequences of these properties (small world and randomness)

separately. So we have chosen another graph model to work with: The Watts-Strogatz model. This model starts with a $C_{N,K}$ graph and then rewires each link at random with probability P. In fact, for $P = 0$ we have $W_N = C_{N,K}$ and for $P = 1$ we have a completely random graph (but not scale-free). For intermediate values of P there is the "small world" region, where the graph is highly clustered (which means it is *not* random) but with a small characteristic path length (a property shared with random graphs). The Watts-Strogatz model does not generate scale-free graphs, since the distribution $P(k)$ associated to these graphs is exponential [23].

6.3 Social Conventions in MAS

First we will describe some general properties of the systems we are going to study. Details will require separate subsections. The MAS we deal with are extremely simple, but also *necessarily* simple if we want to get to any conclusion about its dynamics. The use of these simple settings in MAS theory has been largely discussed in [203, 221], to which we refer for more information.

Our MAS will consist of N agents on a graph, where every agent will be located on a node of the graph. Its adjacent agents will be called its *neighbors*. Every agent will be in one out of two states (or actions), called A and B. The system will evolve in time, and at each time step one agent will be selected at random, for state updating [1]. Different rules to update agent's state will define different systems. In this note we will study two different rules: the *generalized simple majority* rule and the *highest cumulative reward* rule (both detailed in the following subsections).

Quoting Shoham and Tennenholtz [203] *"A social law is a restriction on the set of actions available to agents. A social law that restricts the agents' behavior to one particular action is called a social convention."* In our case a social convention will be reached if all the N agents are either in state A or in state B. From [135] we will get the performance measure we use to evaluate how fast conventions arise in our systems, it is the *convergence time* T_c: the convergence time for a given level of convergence c is the earliest time at which $C_t \geq c$, where C_t is the convergence of a system at time t, that is, the fraction of agents using the majority action (either A or B). In this note we will focus on the study of the average time to a fixed convergence (we set c to 90%, following [135]).

[1] The dynamics we use is *asynchronous*, following previous work [135, 136, 203]. We will depart from Walker and Wooldridge formalization [221] because the dynamics they use (their function r, used to define a *run*) imposes a *synchronous* dynamics, where all agents interact at once. This is, at least, problematic. It is well known that some "emergent" properties of synchronous systems are not due to the system itself, but to global correlations introduced by this synchronous update [120, 153].

6.3.1 Generalized Simple Majority (GSM)

We generalize herein the simple majority rule, as was defined in [221]. We have N agents on a graph, so we have a well defined neighborhood for every agent. The initial state of the system is a random state (either A or B) for every agent. Now, at every time step one agent, say the j-th, is chosen randomly. Let us suppose that agent j has k neighbors and that k_A neighbors are in state A (so there are $k - k_A$ neighbors in state B). If agent j is in state S, let $\bar{S}$ be the complementary state. Thus, agent j will change to state $\bar{S}$ with probability

$$f_\beta(k_{\bar{S}}) = \frac{1}{1 + e^{2\beta(2k_{\bar{S}}/k-1)}} \tag{6.3.1}$$

This rule generalizes simple majority since for $\beta \to \infty$ we recover the change of state only when more than $k/2$ neighbors are in state $\bar{S}$. (see figure 6.1). There is no

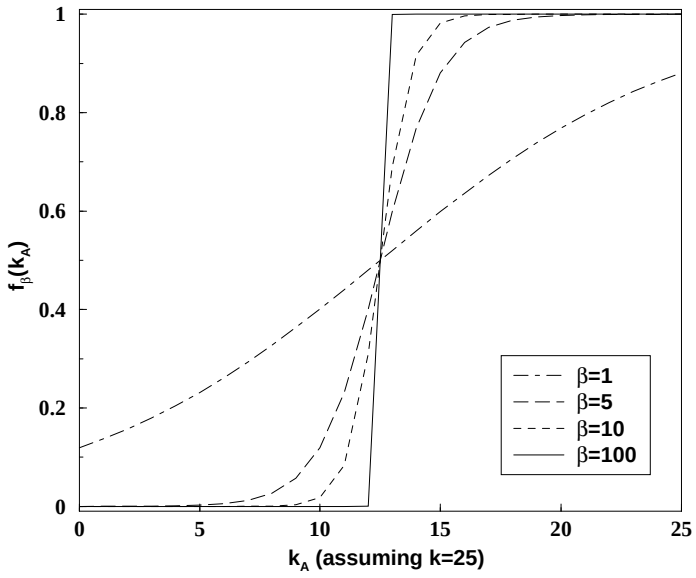

Figure 6.1: Generalized simple majority rule: If an agent chosen for updating is in state B and has 25 neighbors, this figure shows how the probability of changing state varies as a function of the number of neighbors in state A, for different values of β.

theorem assuring convergence in the emergence of conventions in this system, but we can provide some analytical evidence that this is the case. We use what in physics is called a *mean-field* argument [185]. Let $N_A(t)$ be the number of agents in state A at time t and $\rho(t) = N_A(t)/N$ be the density of agents in state A. We will assume the following *homogeneity* condition: for every agent with k neighbors, the number

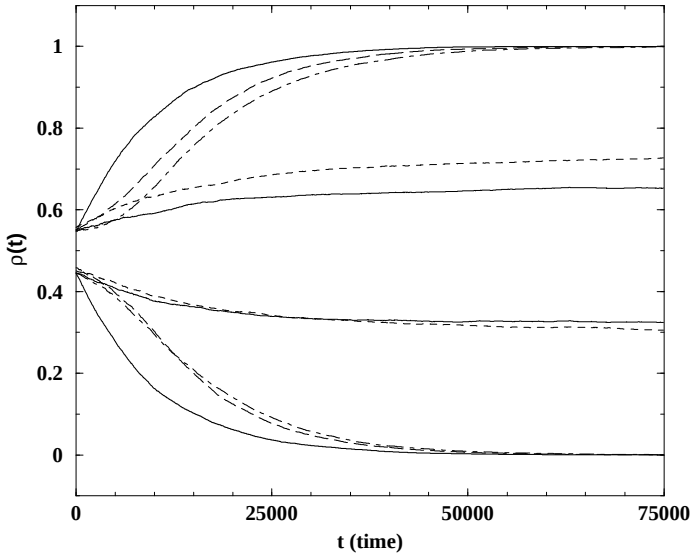

Figure 6.2: Evolution in time of the density of agents in state A for the system defined in subsection 6.3.1. Several graphs have been used ($N = 10^4$ and $\beta = 10$): K_N (thick solid line), $C_{N,12}$ (solid line), $S_N^{2.5}$ (dot-dashed line, $m_0 = 4$, $m = 2$, $p = q = 0.4$), S_N^3 (long-dashed line, $m_0 = 7$, $m = 6$) and W_N (dashed line, $P = 0.05$ and $K = 12$, inside the small world region). We observe a fast convergence in the graphs that fulfill the homogeneity condition, that is, the scale-free graphs and the K_N graphs.

of neighbors in state A is $k_A(t) \simeq k\rho(t)$. This condition is completely fulfilled for K_N graphs (obviously), and approximately fulfilled for S_N^γ graphs and W_N graphs with $P \to 1$, since these are random graphs. Thus, the mean-field equation for $\rho(t)$ can be written as

$$\frac{\partial \rho(t)}{\partial t} = (1 - \rho(t)) f_\beta(\rho(t)) - \rho(t) f_\beta(1 - \rho(t)) \tag{6.3.2}$$

that is, the variation of $\rho(t)$ is the fraction of agents in state B that change to state A (first term, right side of the equality) minus that fraction of agents that switch form state A to state B (second term, right side of the equality). After some arrangements, this equation reads

$$\frac{\partial \rho}{\partial t} = -\rho + \frac{1 + e^{2\beta(2\rho-1)}}{2 + e^{-2\beta(2\rho-1)} + e^{2\beta(2\rho-1)}} \tag{6.3.3}$$

finally, with the change $x(t) = 2\beta(2\rho(t) - 1)$ we get to

$$\frac{\partial x}{\partial t} = -x + 2\beta \frac{e^x - e^{-x}}{2 + e^x + e^{-x}} \tag{6.3.4}$$

We want to study the stable fixed-points of $x(t)$, since these will give us information on the final state of the system. Thus, we must find the solutions of $\partial_t x(t) = 0$. It can be shown that, for $\beta \gg 1$ the only stable fixed-points of this equation are $x_1 = -2\beta$ and $x_2 = 2\beta$, that is, $\rho_1 = 0$ (state B is the reached convention) and $\rho_2 = 1$ (state A is the final state of all the agents). Let us point out that this result implies the convergence to a social convention in systems using the simple majority rule, as defined in [221]. Initial conditions will break the symmetry of the solutions, that is, an initial fraction of agents in state A slightly larger than the initial fraction of agents in state B will get the system to a $\rho = 1$ stationary state, and vice-versa (see figure 6.2). In this work we will not study the effect of β, setting $\beta = 10$.

6.3.2 Highest Cumulative Reward (HCR)

The framework in which we will work here was introduced by Shoham and Tennenholtz [201, 202, 203] some time ago, though it is in frequent use nowadays (see [35, 50, 154, 63] for example). In this work we will adapt from [203] the definitions and theorems we need, not dwelling on justifications of this formal framework (it was eloquently done in [203]). We will focus on *coordination* games [149, 202].

A payoff matrix G 2×2 defines a 2-person 2-choice symmetric coordination game if G has the form

		Agent j	
		A	B
Agent i	A	x	u
	B	v	y

Figure 6.3: Payoff Matrix of the Game G

where $x > v$ and $y > u$

Essentially the idea is that every player has two available actions, say A and B. If both players play A, both players receive a payoff of x. If they play B they receive a payoff of y. When the players do not agree, for example, player 1 plays A and player 2 plays B, the former receives a payoff of u and the latter a payoff of v; the remaining situation is symmetric. The condition on the entries of G makes clear that to play the same action is the best choice. Specifically, we will use the *pure coordination game* [149] G, where $x = y = +1$ and $u = v = -1$.

Once defined the game we need to define the players. Our MAS will be composed of N agents (every agent is a player) that will interact with other agents, playing the game G once per interaction. What we are interested in is whether the dynamics of this system makes all the agents reach a social convention. In our particular setting, this means that we want to know whether all the agents will end up playing one of

the two possible actions of the game G, say A and B.

Following Kittock [135], every agent, say the k-th, will be characterized by a *memory* M_k of size m (same size for all the agents) and an action a_k (to play the next time agent k is selected, so the value of a_k is either A or B). The memory M_k will record some information on the m last plays of the agent k: The value of the position i of the memory M_k will be a tuple $< a_k^i, p_k^i, t^i >$ where t^i is the time the i-th play took place, a_k^i is the action played by agent k and p_k^i is the payoff received ($1 \leq i \leq m$). However, in this work we will not study the effect of memory, setting $M = 1$.

We must also define the dynamics of the system (a variant of $n-k-g$ stochastic social games [203] where we will take into account the underlying topology). At every time step t, a pair of agents will be selected to play the game G, where one of them will be randomly chosen and the other will be one of its neighbors, according to the underlying graph. They will receive a payoff (either $+1$ or -1) depending on their actions. Let us assume that at time t, agents k (with memory M_k and action a_k) and l (with memory M_l and action a_l) are chosen to play. Every agent will receive a certain payoff, say p_k and p_l. Now, agent k must decide which action it is going to play next time it is chosen, as a function of its memory M_k, the action a_k played and the payoff received p_k. It uses the *Highest Cumulative Reward* rule. Agent k will compute the payoff received for using action A in the last m plays in which it has been involved: $P_A^k = \sum_{i:a_k^i=A} p_k^i$, where P_B^k is defined in the same way. Agent k will add p_k to either P_A^k or P_B^k, depending on a_k. Now, agent k can decide: Next time it is chosen to play, the action chosen by the agent k will be either A if $P_A^k > P_B^k$, B if $P_B^k > P_A^k$ or a_k otherwise. Finally agent k updates its memory, deleting the oldest entry and adding the tuple $< a_k, p_k, t >$ (agent l will do the same thing, the rest of the system will do nothing). For the system as defined above Shoham and Tennenholtz [203] provided a general theorem[2] that guarantees the convergence of our system to a stable social convention.

Kittock [135] studied numerically the efficiency of the emergence of conventions in regular graphs $C_{N,K}$ and K_N. His main result was that the underlying topology has a profound effect on the efficiency with which conventions emerge, and he conjectured that this efficiency depends essentially on the *diameter* of the graph.

[2]Theorem 12 in [203]. Given a 2-person 2-choice symmetric coordination game with dynamics as defined and using the HCR action selection rule: i) $\forall \epsilon > 0$ there exists a bounded number Γ, such that if the system runs for Γ iterations then the probability that a social convention will be reached is greater than $1 - \epsilon$. ii) Once the convention is reached, it will never be left.

6.4 Convergence Time of Social Conventions

Once we know that social conventions will emerge in the systems we are interested in, we would like to know how fast these conventions will be reached. From our numerical work (see figures 6.4 and 6.5, these figures are representative of results obtained with different sets of parameters) we may conjecture that $T_{90\%} = O(N^3)$ for $C_{N,K}$ graphs (which was already observed by Kittock [135]) and $T_{90\%} = O(NlogN)$ for complex graphs and K_N graphs (this is the lower bound predicted analytically in [203, 212]) for the HCR rule. Results for the GSM rule are $T_{90\%} = O(N^3)$ for $C_{N,K}$ graphs and $T_{90\%} = O(N)$ for complex graphs and K_N graphs. Besides, we observe that an underlying small world graph makes the system less efficient than an underlying scale-free graph, despite they have the same behavior.

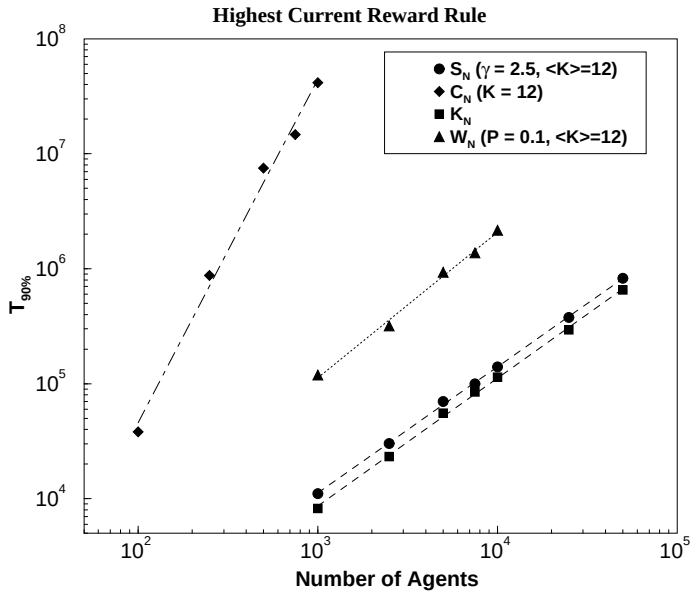

Figure 6.4: HCR rule: $T_{90\%}$ vs. N, averaged over 25 samples for each N. Several graphs have been used: K_N, $C_{N,12}$, $S_N^{2.5}$ ($m_0 = 4$, $m = 2$, $p = q = 0.4$) and W_N ($P = 0.1$ and $K = 12$, inside the small world region). All the graphs have the same average connectivity per node (except K_N, for obvious reasons). In all cases $m = 1$.

Kittock's conjecture provides us with a partial explanation of the observed behavior. According to [135], the efficiency of the emergence of social conventions depends on the diameter of the graph. The diameter of $C_{N,K}$ grows linearly with N [227] but the diameter of complex graphs grows *logarithmically* with N [181], hence the difference between the growth of $T_{90\%}$ in regular graphs and complex graphs. However, the precise relation between the linear growth of the diameter in regular

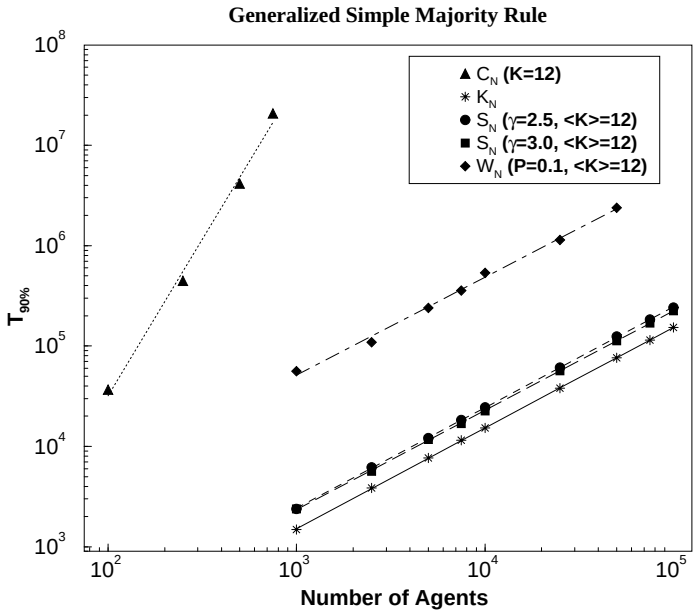

Figure 6.5: GSM rule: $T_{90\%}$ vs. N, averaged over 25 samples for each N. Several graphs have been used: K_N, $C_{N,12}$, $S_N^{2.5}$ ($m_0 = 4$, $m = 2$, $p = q = 0.4$), S_N^3 ($m_0 = 7$, $m = 6$, $p = q = 0$) and W_N ($P = 0.1$ and $K = 12$, inside the small world region). All the graphs have the same average connectivity per node (except K_N, for obvious reasons). In all cases $\beta = 10$.

graphs and the $O(N^3)$ behavior of $T_{90\%}$ for both rules, and between the logarithmic growth of the diameter in complex graphs and the $O(NlogN)$ behavior of $T_{90\%}$ for the HCR rule, $O(N)$ behaviour for the GSM rule, remains to be fully justified by means of analytical arguments.

Now, the efficiency of the emergence of social conventions in systems with underlying scale-free graphs is almost as good as with K_N graphs, despite having a constant (with respect to N) average connectivity. Notice that K_N are optimal in respect to $T_{90\%}$, provided Kittock's conjecture is correct, since this quantity depends on graph diameter and this equals 1 for K_N graphs. Besides, underlying small world graphs are less efficient than scale-free graphs, despite the same behavior with respect to N. This is so because what is important here is the randomness (in the sense mentioned above) of the scale-free graphs we have used, since randomness reduces the graph diameter.

However, the small world property seems to have no effect on $T_{90\%}$. We may perform some experiments to test this hypothesis. The Watts-Strogatz model allows us, by means of the parameter P, to go from small world graphs (small P) to random

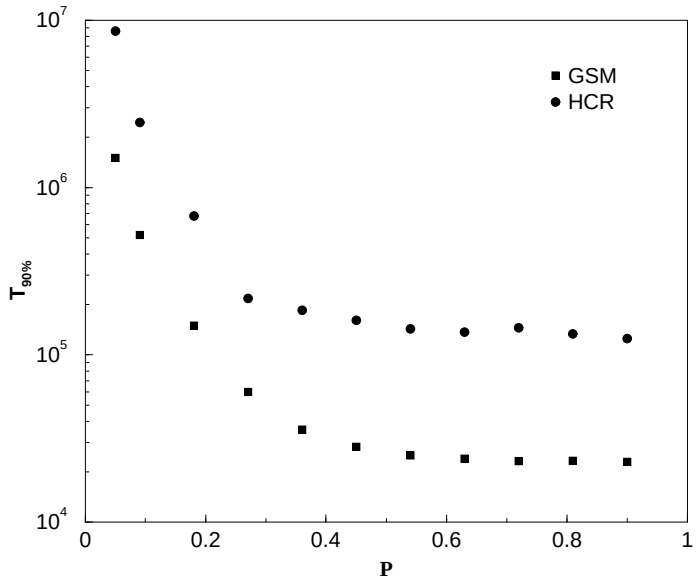

Figure 6.6: $T_{90\%}$ vs. P for systems using the two different action update rules studied in this note, with underlying graphs generated with the Watts-Strogatz model. Parameters are $N = 10^4$ and $K = 12$ ($\beta = 10$ for the GSM rule, $m = 1$ for the HCR rule).

graphs ($P \to 1$) with an exponential $P(k)$. Thus, measuring $T_{90\%}$ on the Watts-Strogatz model with varying P will make clear the importance of randomness.

We see in figure 6.6 that $T_{90\%}$ decreases with P, without noticing the small world zone for small P: The graph becomes more and more random and the system becomes more and more efficient. This fact makes clear, again, that the diameter seems to be the important factor in the efficiency with which conventions are reached. Thus, our results are fully consistent with Kittock's results.

6.5 Discussion

In this chapter we defined a simple MAS with which to study the efficiency of the emergence of social conventions in complex networks. On the one hand we have defined MAS with the action update rule called the generalized simple majority rule, providing analytical evidence of convergence to a social convention, and, on the other hand, we have studied the well-known MAS with the highest cumulative reward rule as action update rule. On both systems we have performed a numerical study of $T_{90\%}$ as a function of N and, in graphs defined according to the Watts-Strogatz model, of P. Our results on both systems are consistent with the hypothesis

that the diameter of the graph underlying the MAS is of essential importance in the efficiency with which conventions are reached [135]. We have found a topology that makes the system as efficient as the K_N graph but at a lower cost, where the cost is the average number of links per node.

Some questions are still open: It remains to be analytically justified the precise relation we have found among the growth of the diameter for different classes of graphs and the behavior of $T_{90\%}$. The almost identical behavior of $T_{90\%}$ for the systems studied should also be explained, since the substantial differences between both systems are obvious from their definitions.

Despite the open questions just mentioned, the effect that particular structures has on the dynamics of the system is evident. The time for a convention to emerge — agents coordinating themselves in a certain action — deeply depends on the diameter of the network modeling the interactions between agents. Complex networks are very adequate models of graphs for that matter; they display a short diameter that grows logarithmically with network's size, therefore, favoring the emergence of coordination in reasonable time even in the case of systems composed by tens of thousands of agents.

There are many ways to extend the work introduced in this chapter. This study may be repeated with cooperative games, since these games were also considered by Shoham and Tennenholtz [201, 202, 203] and Kittock [135]. Furthermore, preliminary results indicate that the behavior of cooperative games in complex networks is far from trivial [2, 227]. Also, it would be interesting to study MAS playing stag-hunt games [216, 190] in complex networks. Another possible extension, which is actually addressed in the next chapter, is to consider coordination when coordination in a given convention *Pareto-dominates* the other. Thus, agents would be better off choosing the efficient action in terms of payoff. However, this might not be always possible resulting in coordination in a sub-optimal convention. Next chapter studies how the underlying topology also influences the convention which is eventually established by the agents.

Chapter 7

Emergence of Efficient Social Conventions

7.1 Description

In the previous chapter we studied the emergence of conventions from local coordination processes without a central authority. We showed that the elapsed time for the convention to be established depended on the underlying structure of interactions, in particular, on the diameter and characteristic path length of the network. In this chapter we further develop and extend previous work to address how and under what conditions emerging conventions are also socially efficient, i.e. better for all agents than potential alternative conventions.

The existence of norms, and by extension, social conventions is perhaps one of the most fundamental problems that social sciences have ever tackled [27]. Intuitively, a social convention might be regarded as any rule of behavior, that is, a behavioral constraint [221]. As such, they simplify people's decision making problem by dictating how to act under certain situations. Therefore, social conventions help to reduce the complexity and uncertainty, specifically, when the environment is open and dynamic. Which is exactly the situation were Distributed Artificial Intelligence and Artificial Societies dwell.

Shoham and Tennenholtz [201] introduced the notion of *emergence of conventions*, as opposed to those conventions designed a priori by a *central* authority. A group of agents that through the game's payoff were able to coordinate themselves in a certain action, which with time, turned normative, since deviation from that action was sanctioned by game's payoff. They showed that self-coordination can appear without a *central* authority, thus, that social convention can emerge from disorder. Since then, many researchers have studied and improved his seminal work, proposing other strategies besides HCR [221], studying the time needed before the convention

is established [135] and adding complex interaction patterns among agents, as seen in the previous chapter and in [63, 64].

As anticipated in the previous chapter's discussion, an important question remains open: When two potential different conventions exist which one will be established at the end? This question is not applicable to the original work of Shoham and Tennenholtz [201], since the game was a *pure* coordination game, where no action Pareto-dominates the other. Nevertheless, the question applies to all coordination games where one action is more efficient[1] than the other. The discussion in social sciences gives no clear answer to why and under what conditions efficient conventions may prevail. As Bendor and Swistak [27] pointed out, there are at least two conflicting positions in sociology, the *strong functionalism* thesis and the *rational choice* view. *Strong functionalism* claims that norms and conventions exist because they are functional for the group, that is, they yield optimal collective outcomes. According to this perspective, the system should always end up in an efficient convention. However, the *functionalist* approach has been criticized for its lack of a microfoundation. Adherents of the *rational choice* view on norms argue that individuals adopt norms only when it serves their self-interest to do so. This includes the possibility that mechanisms such as social control may stabilize conventions that are an individually efficient response to the given constraints, but are socially inefficient.

7.2 Formal Model

We will follow the conceptual framework introduced by [201], already used in the previous chapter.

7.2.1 The Coordination Game

A set of N agents must choose to play one of two possible actions: either A or B. Accordingly to its current action, or state, an agent interact with its neighbors receiving an outcome defined in payoff matrix G.

		Agent j	
		A	B
Agent i	A	x	u
	B	v	y

Figure 7.1: Payoff Matrix of the Game G

[1] The action that Pareto-dominates the others is socially efficient, since it is better for all agents than potential alternative action

The payoff matrix G defines a 2-person 2-choice symmetric coordination game provided that $x > v$ and $y > u$.

The condition on the entries of G makes clear that to play the same action is the best choice. It is trivial to demonstrate that the game G have two Nash-equilibrium, both agents playing either A or B. Most work has focused in the study of *pure coordination games*: where $x = y = +1$,$v = u = -1$ [221, 135, 203, 63]. Our approach differs from related work since we will not restrict ourselves to *pure* coordination games.

Coordination in action A will be at least as profitable as coordination in action B: $x \geq y$. Thus, game G is defined as follows: $v = u = -1$, $y = +1$, and $x = \alpha$ provided that $\alpha \geq 1$. When $\alpha > 1$, coordination in B is a sub-optimal solution since there exists a Pareto-efficient solution that Pareto-dominates B, which implies, coordination in action A. Therefore, coordination in (A, A) or (B, B) are still two Nash-equilibrium solutions of the coordination problem, but, depending on α, coordination on A will be more efficient than coordination on B.

7.2.2 Action Selection Rule and Dynamics

Our MAS is composed of N agents that interact only with its neighbors, playing the game G once per interaction. Every agent, say the kth, has memory M_k that records the M last interactions of agent k. The value of the position i of the memory M_k is the tuple $\langle a_k^i, p_k^i, t^i \rangle$, where a_k^i stands for the action played by k, p_k^i stands for the payoff received after paying action a_k^i, and t^i denotes the time the interaction took place. The initial action of the agents is set randomly with a probability r_B, which is the density of agents playing action B in the beginning.

Following [201] we will use the *Highest Cumulative Reward* (HCR) action selection rule. Intuitively HCR says: if the accumulated payoff obtained from playing A is bigger than that from playing B then keep on playing A, otherwise change to action B. The HCR rule is very appropriate since it provides: 1) Locality: the selection function only depends on the agent's personal history. No global knowledge of the system is required, not even the payoff matrix of the game. 2) Adaptability: the agent learns from its experience without assuming further cognitive capabilities. These characteristics are very important in complex and open systems, such as MAS.

The dynamics of the system are as follows. At each time step t, an agent k is randomly chosen. Once the agent is chosen, or activated, it plays the game G with an agent randomly chosen from k's neighborhood, say agent l. The result of the interaction is stored into agent k's memory M_k, removing the oldest entry if necessary. Finally, agent k must decide whether to change its action or not. To do so it uses the *Highest Cumulative Reward* rule (see section 6.3.2).

The system ends once all agents are playing either action A or B, which means, that a convention on either A or B have been established. Notice that, unlike what was done in the previous chapter, the system does not stop until the 99% of agents are coordinated playing the same action. In the previous model (chapter 6) the system stopped once the 90% of agents were playing the same action. Following the terminology proposed in [135], the convergence time is $T_{99\%}$, instead of $T_{90\%}$ as it was in the previous chapter.

7.2.3 Underlying Topology

In a similar fashion as it was done in the previous chapter, several models of graphs are used as the underlying topology for our MAS. These graphs determine each agents' neighbourhood. These models correspond to those used in chapter 6 with a couple of exceptions: 1) the extended Albert-Barabási [10] model of scale-free graphs has been replaced by the Barabási-Albert model [22]. And, 2) the complete graph K_N has been replaced by a random graph R_N [81].

- *Random graphs:* $R_N^{\langle k \rangle}$, where N is the number of nodes, and $\langle k \rangle$ is the average connectivity, that is, the average size of node's neighborhood. Random graphs have a clustering coefficient that tends to zero, and the average path length grows logarithmically in function of N, the number of nodes. We chose the classical model of random graphs of Erdös-Rényi [81]

- *Regular graphs:* C_N^k, regular graphs display an extremely high clustering coefficient, while its average path length and diameter grows linearly. Which means, that for big graphs the average path length is very long, which does not match with networks empirically studied. However, regular graphs display the close-knit property due to its high clustering coefficient, which does match with empirical studies.

- *Small world graphs:* $W_N^{\langle k \rangle, p}$, these are *highly-clustered* graphs (like regular graphs) with small average path length (like random graphs). This is the small world property. We chose the Watts-Strogratz [232] model as model of small world graphs. p is the rewiring probability in the Watts-Strogatz model.

- *Scale-free graphs:* $S_N^{\langle k \rangle, -\gamma}$, these are graphs with a connectivity distribution $P(k)$ of the form $P(k) \propto k^{-\gamma}$. The connectivity degree, the number of neighbors of a node, decays as a potential law. This favors the so-called *fat-tail* phenomena, that is few nodes with an extreme high connectivity. The clustering coefficient of the scale-free network is much smaller than the small world graphs as shown in table 7.1. Unlike previous chapter we chose Barabási-Albert model [22] as a model of scale-free graph.

	d	cc	apl
C_{1000}^{10}	100	0.666	50.450
$W_{1000}^{<10>,0.1}$	8	0.492	4.480
$S_{1000}^{<10>,-3}$	5	0.0433	2.963
$R_{1000}^{<10>}$	5	0.0081	3.269

Table 7.1: Graph characteristics: d is the graph diameter, cc is the clustering coefficient and apl is the average path length. The order of the graph is 1000 nodes, and the average connectivity $\langle k \rangle$ is 10.

As we have repeatedly mentioned throughout the thesis (specially in section 3.2), recent studies on empirical networks show us that neither regular nor random graphs appear in nature. Noticing this, Delgado et al [63, 64] studied the effect of *complex* networks (small world and scale-free networks), based on previous work of [201, 135]. The most relevant conclusion was that complex networks where as efficient as the complete graph in terms of time to reach a convention. Delgado [63] showed empirically that the emergence of a convention is almost linear for *complex* graphs compared to the cubic ($O(N^3)$) for regular graphs, already shown by [135]. We follow Delgado [63] approach and use realistic network topologies as the underlying agent's interaction pattern.

At this point, it is important to remark that many empirical social networks such as the collaboration network among actors, co-authorship networks in Mathematics and Neuroscience and many others are classified as scale-free networks. This is because its connectivity distribution decays as a potential law, although an exponential cut-off is also present. However, unlike many scale-free models, these empirical social networks also exhibit a very high clustering coefficient: 0.79, 0.59, 0.76 respectively.

Table 7.1 summarizes the clustering coefficients and their characteristic path lengths of the graph models used in the experiments. Notice that empirical social networks are much more clustered than scale-free networks yield by the Barabási-Albert [22] model. But we still use this model since scale-free models are focused on reproducing the connectivity distribution of a network rather than its clustering coefficient. This fact must be kept in mind while analyzing the results and conclusions. See section 3.2.1 for a review on complex networks.

7.3 HCR-Model Experimental Results

First of all we will introduce a classification of the graphs in terms of their clustering coefficient. Regular and small world graphs will be called *highly-clustered* graphs, and scale-free and random graphs will be called *low-clustered* graphs. For the sake of

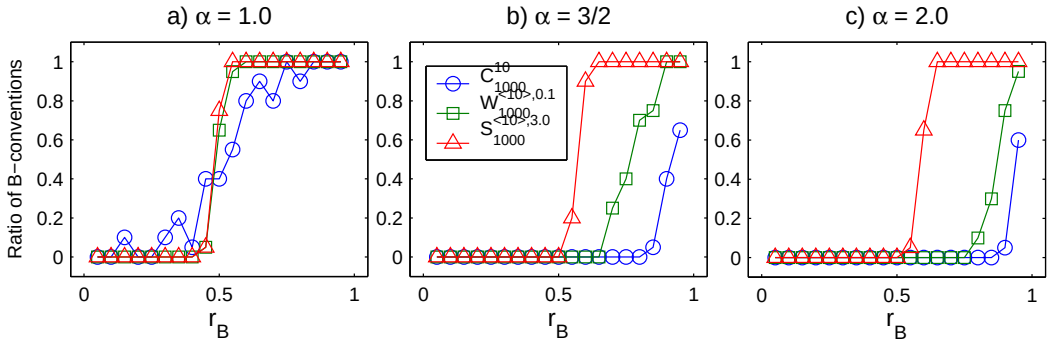

Figure 7.2: Ratio of conventions converging to action B. The x-axis shows the initial density of agents playing action B, r_B, ranging from .05 to .95. The y-axis shows the ratio of B-conventions, that is, the number of replicas that ended up in all agents playing B over the total number of replicas (20). In sub-figure a) there is no efficient convention, $\alpha = 1$. In sub-figures b) and c) all agents playing action A which is the efficient convention ($\alpha = \frac{3}{2}$ and $\alpha = 2$ respectively).

clarity we did not include results on random graphs in the figures, since they behave as scale-free graphs.

Also before going into the details we must mention that one can see two clear phases: 1) all the replicas converging to A, and 2) all the replicas converging to B. There is as well a space between these two phases where the system converges, with some replicas going to A and some to B, which we call *transitional* space.

This space is wider or narrower depending on the underlying topology and α. There is a critical point r_B^* that sets the boundaries between the two phases. The *transitional* space can be defined as ε, such that the result is: $r_B^* \pm \varepsilon$. It is straightforward to see that the *transitional* space is narrow for almost all the cases. This fact led us to suspect a clear phase transition scenario. However, when $\alpha = 1$ and the underlying topology is a regular graph the transitional space is so wide that it might not be considered as a phase transition. we provide some explanation about this behaviour in section 7.5.

Now, let us comment on the results of the experiments with the HRC model. On one hand, we have the results when there is a *pure* coordination game. That is, when coordination on action A is as efficient as coordination on action B since both yield the same payoff. These results are displayed in sub-figure 7.2.a. When the initial number of agents playing B is less than half the population, that is, $r_B < \frac{1}{2}$, the system ends up establishing the convention on B. And, when $r_B > \frac{1}{2}$ establishing the convention on A. For instance, when $r_B = .65$ the system tends to converge to convention B regardless of the underlying topology. The ratio of B-conventions is

1.0 for random, scale-free and small world graphs. And it is 0.9 for regular graphs, which means that 18 of 20 replicas ended up on B (still in the transitional space).

On the other hand, we obtain the results when coordination on action A is more efficient than B, which is the case shown in sub-figures 7.2.b and 7.2.c. We will use an example to make clear what happens with the results in the case of $\alpha > 0$. Let $\frac{3}{4}$ of the population follow the B action, and $\frac{1}{4}$ follow the A action. In this case, although B is the initially chosen action for most of the agents, coordination in A is more efficient since yields a better payoff. So, which will be the final convention agreed by the whole population? The answer depends on 1) how much better off, i.e. more efficient, is action A over B, denoted by α. And 2) the underlying topology. If $\alpha = \frac{3}{2}$ (sub-figure 7.2.b) the final convention will be B when the underlying topology corresponds to a random or scale-free graph. It will be A with a regular graph, and it could be both with a small world graph. By increasing the efficiency of coordination on A to $\alpha = 2$ (sub-figure 7.2.c) both small world and regular graphs converge to convention A, whereas random and scale-free graphs still converge to convention B. The explanation behind this result is a bit surprising, *low-clustered* graphs seem very sensitive to the initial population density, whereas *highly-clustered* graphs behave in the opposite way. That is, they are more sensitive to the efficiency of a particular action.

From these results one can come up to the following conclusion: *low-clustered* agent communities where a convention already exists will not be infected by a set of agents who play a new action, even though the new convention is more efficient. Conversely, *highly-clustered* agent communities can be infected by a new action if the new action is more efficient, replacing the current convention by the convention on the efficient action.

Therefore, *highly-clustered* agent communities are more innovative, or adaptive, since a new action can spread and finally be established as a convention. The drawback would be that this community would be unstable due to its receptiveness to new actions. Full coordination would not exist during the transient period prior the new convention is finally established. In contrast, *low-clustered* agent communities are very stable since the infection with a new convention is unlikely to happen, but, on the other hand, they are reluctant to adopt new actions even though these new action are more efficient. In this way, they become very conservative and static communities. We must remark that studies on empirical social networks have shown that these type of networks are very clustered [11].

7.4 Convergence Time

We will not provide a proof of the system's convergence. Nevertheless, throughout all the simulation runs, with their corresponding replicas, the system has always

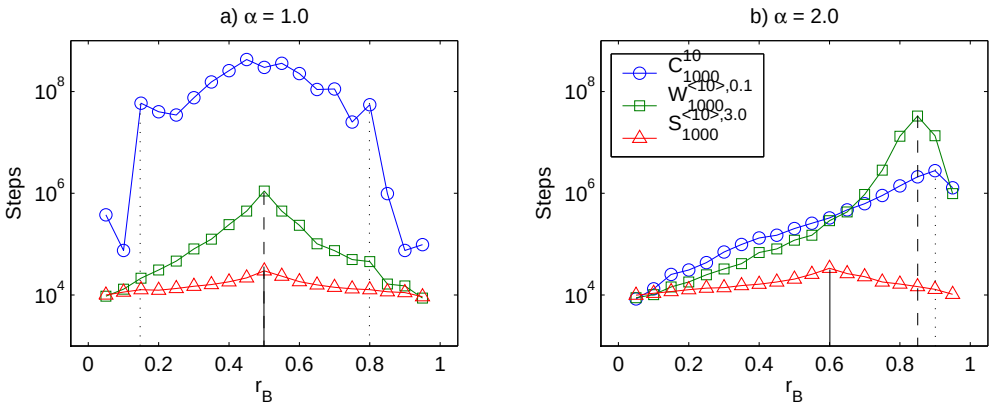

Figure 7.3: Convergence time ($T_{99\%}$). The value is the mean over 20 replicas. The x-axis shows the initial density of agents playing action B, r_B, ranging from .05 to .95. The $y-$axis shows the number of interactions (steps) before reaching a convention. Note the logarithmic scale. In the left sub-figure both conventions, A or B. In the right sub-figure A is the optimal convention since $\alpha = 2$.

converged to a convention with an upper bound of $O(N^3)$.

On the left part of figure 7.3, α is set to 1. Therefore both conventions are equally efficient. It can we observed that the regular graphs takes much longer to converge than the rest of graphs. This fact is completely consistent with the findings of [135] and [63]. Notice the existence of peaks in $r_B = \frac{1}{2}$ in all the graphs except the regular ones. These peaks are close to the critical point $r_B^* = \frac{1}{2}$ where a phase transition takes place. This is a typical behavior when the system is close to a transition boundary, it takes longer to converge [240]. The regular graph, however, is a particular case since it shows a *plateau* rather than a peak. This is due to its wide *transitional* state, in which different replicas of the same setting can converge to different conventions.

On the right part of figure 7.3 the convention A is more efficient that convention B. Playing action A yields a payoff $\alpha = 2$, when playing action B it yields a payoff 1. In this case the *plateau* of the regular graph do not appear because its *transitional* space narrows, as it is shown in figure 7.2.c. Thus, we see only the peaks in r_B where the transition phase takes place ($r_B^* = \{.6, .85, .90\}$) of scale-free, small world, and regular graphs respectively, which correspond to the transition phases observed in figure 7.2.c.

Again, we observe that the convergence is more inefficient in time when the initial r_B is close to r_B^*. It is important to remark that the regular graph has improved its efficiency in reaching the convention, that is, it is more efficient than the small world graph when $\alpha = 2$. As was shown by Peyton Young [241], convergence to a

risk-dominant equilibrium, which in our case corresponds to the Pareto-efficient one, is surprisingly rapid provided a close-knit (clustered) graph. The scale-free graph is always the most efficient graph in terms of convergence no matter what α. However, as we have discussed in section 7.4, scale-free graphs do not allow convergence to the efficient convention unless the initial number of agents playing A is very big, 40%, compared to the 15% needed with a small world graph, with α set to 2.

7.5 Analytical Model

In order to provide some analytical evidence about the critical points we will try to reduce our initial model to a more suitable one for analysis.

First of all, we replace the role of past interactions stored in agent i's memory (M_i) by the current state of agent i's neighbors. Notice that in our model we set the memory size to the average connectivity, $M = \langle k \rangle$. Instead of applying the HCR-rule over M_i we will transform it as follows. Let us take k_S as the number of neighbors in the same state as agent i, and $k_{\overline{S}}$ the number of neighbors in the opposite state. Therefore, the update is performed if $k_{\overline{S}} > \phi k_S$, where ϕ is the payoff yield by the payoff matrix G when playing (S, S).

Therefore, the switch from playing A to playing B is done when $k_B > \alpha k_A$, since $G(A, A) = \alpha$. Similarly, the switch from playing B to play A is done when $k_A > k_B$, since $G(B, B) = 1$.

The probability of updating the state is defined by equation 7.5.1. We decided to introduce some stochasticity in order to compensate for the fact that, now, the update is calculated by the current state of the neighbors and not by the interactions, as the HCR-rule did.

$$f_{\beta,\phi}(k_{\overline{S}}) = \frac{1}{1 + e^{\beta(\phi - (\phi+1)\frac{k_{\overline{S}}}{k_S + k_{\overline{S}}})}} \tag{7.5.1}$$

We will use what in physics is called a *mean-field* argument [185]. Let $N_B(t)$ be the number of agents playing action B at time t, and $\rho(t) = \frac{N_B(t)}{N}$ be the density of agents playing B. A first approach is to assume the following *homogeneity* condition: for every agent with k neighbors, the number of neighbors in state B is $k_B(t) \simeq k\rho(t)$. This condition is completely fulfilled by random graphs, and approximately fulfilled by scale-free and small world graphs (when $p \to 1$).

Nevertheless, this *homogeneity* condition is not fulfilled by regular and by small world graphs for low values of p. What breaks the *homogeneity* is the clustering coefficient. For *low-clustered* graphs this condition holds since the global density of agents in state S corresponds to the proportion of neighbors in state S. Intuitively, the clustering coefficient can be defined as the probability that a node i and a node

j have a link provided that node l has a link to both i and j. Thus, when clustering tends to 0 the node's neighborhood is a good sample of the graph. However, when the clustering coefficient is high the node's neighborhood is not a sample of the graph, since its neighbors form a *clique*, a *close-knit* group.

Therefore, we propose a new *homogeneity* condition that takes clustering (*cc*) into account. Let us define *cc* as the clustering coefficient, provided that *cc* is the probability of agent l's neighbors being also neighbors, $(1 - cc)k$ is the number of neighbors which are not in the l's *clique*, and to whom the previous *homogeneity* condition holds. Therefore, for an agent playing A with k neighbors, the number of neighbors in the opposite state (B) is $k_A \simeq (1 - cc)k\rho(t)$, which is the *cc-biased-homogeneity* condition.

Now we can write an equation for the evolution of $\rho(t)$. First, notice that the variation of $\rho(t)$ after a small time interval Δt is proportional to Δt, that is

$$\rho(t + \Delta t) = \rho(t) + \frac{\partial \rho(t)}{\partial t}\Delta t + O(\Delta t^2) \tag{7.5.2}$$

Then, we can neglect the $O(\Delta t^2)$ term (since we perform a continuum approximation $\Delta t \to 0$) and compute the variation of $\rho(t)$ as the balance between the agents switching from action A to B and the agents switching from action B to A. On one hand, the fraction of agents in A (that is, $1-\rho(t)$) that change to state B in a time interval Δt is the product $(1 - \rho(t))f(\rho(t))\Delta t$, provided Δt is small enough; on the other hand, the fraction of agents that switch from action B to A in Δt is $\rho(t)f(1 - \rho(t))\Delta t$, also for small Δt. Thus after $\Delta t \to 0$, the mean-field equation for $\rho(t)$ can be written as

$$\frac{\partial \rho(t)}{\partial t} = (1 - \rho(t))f(\rho(t)) - \rho(t)f(1 - \rho(t)) \tag{7.5.3}$$

After substitution of $f_{\beta,\phi}$ to which the *cc-biased-homogeneity* condition has been applied. Setting $\beta = \langle k \rangle$ and $\phi = (S, S)$ the equation reads

$$\frac{\partial \rho}{\partial t} = \frac{1 - \rho}{1 + e^{\langle k \rangle(\alpha - (\alpha+1)(1-cc)\rho)}} - \frac{\rho}{1 + e^{\langle k \rangle(1-2(1-cc)(1-\rho))}} \tag{7.5.4}$$

We want to study the stable fixed points of 7.5.4 since these will give us information about the final state of the system. Thus we must find the solutions of $\frac{\partial \rho}{\partial t} = 0$. As we can see in figure 7.4, stable fixed points are $\rho^* \sim 0$, and $\rho^* \sim 1$ (these have been computed numerically) and the unstable fixed-point lies in the interval $(0, 1)$. In table 7.2 we display the unstable fixed points of the density equation, which are possible critical points of the HCR model provided that our assumptions are correct (see above). Furthermore, in table 7.2 we find a comparison between the analytical unstable fixed points and those critical points coming out from the simulation of the HCR model. Notice that the plausibility of the simplifying assumptions behind our

$\rho^*(r_B^*)$	$\alpha = 1$	$\alpha = \frac{3}{2}$	$\alpha = 2$
cc=0.666	0.5 (0.5)	0.831 (0.90)	$\nexists$ (0.95)
cc=0.492	0.5 (0.5)	0.68 (0.7)	0.827 (0.85)
cc=0.0433	0.5 (0.5)	0.566 (0.55)	0.614 (0.60)
cc=0.0081	0.5 (0.5)	0.560 (0.55)	0.607 (0.60)

Table 7.2: Fixed points ρ^* of the analytical model. The table shows the critical point for which $\frac{\partial \rho}{\partial t} = 0$. In brackets the critical points observed experimentally with the HCR model (r_B^*) (figure 7.2). The parameters used to model the graph are the clustering coefficient cc, and the average connectivity $\langle k \rangle$ set to 10.

analytical model is supported by the agreement between analytical and experimental results (remind that r_B was sampled with a resolution of .05).

Figure 7.4 shows the variation in ρ, $\frac{\partial \rho}{\partial t}$, for different α and cc. We can see the effect of α enlarging the basin of attraction of convention A. On the other hand cc has the effect of reducing the amount of variation due to the effects of the *cc-biased-homogeneity* condition. Consequently, the time elapsed to reach a convention will be longer, and fluctuations in initial conditions will have a bigger impact. As a matter of fact, the experimental results on the HCR model show that these two consequences apply for highly-clustered graphs such as regular graphs. The convergence time is much higher compared to non-clustered graphs. And the transitional space, where the system can converge either to B or to A for the same initial parameters, is wider.

7.6 The Role of Imitation

In this section we will modify the model based of the HCR-rule introducing an imitation propensity i_S, which is the probability that, after a dyadic interaction where at least one agent that was playing action S, both agents end up playing action B regardless of the HCR action updating rule. By introducing an imitation propensity we model the effect of having an *attractive* action which is more likely to be chosen.

Imitation is considered as the key factor of the adoption of norms, and by extension, conventions. How does the existence of an *attractive* action affect the final convention reached by the agents? And what if the *attractive* action happens to be sub-optimal? In our model coordination on action A is the most efficient solution (if $\alpha > 1$). However, action B (if $i_B > 0$) might become a better replicator since it can be adopted by imitation as well as adopted by the learning process (HCR-rule).

The effect of i_B heavily depends on the clustering coefficient of the underlying topology. For instance, when $\alpha = 2$ and $i_B = 0.4$ the clustered communities adopt

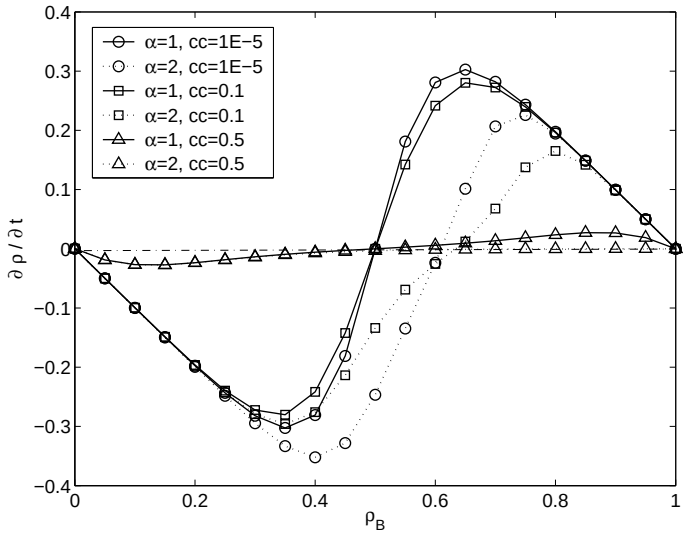

Figure 7.4: Study of the fixed points of the density equation 7.5.4. ρ is the density of agents playing B. The clustering coefficient is set to $cc = \{1E-05, 0.1, 0.5\}$. α is set to $\{1, 2\}$

the *attractive* convention (B) over the efficient one (A) regardless of the initial density (r_B), while *non-clustered* communities will still adopt A provided $r_B < 0.25$. Therefore, it might be derived that *non-clustered* communities are more resilient against *attractive* conventions in favor of efficient conventions. This result would seem to contradict the previous claim, that clustered communities are better off in converging to the efficient convention. However, when i_B is small enough, for example when $i_B = 0.1$ the opposite effect is observed, clustered communities keep on converging to the *efficient* convention regardless of the initial density of agents. This two-fold behavior is perfectly clear in the case of regular graphs (left column of figure 7.5). There is a threshold, i_B^*, under which the system ends up in the efficient convention, and over which the system ends up in the *attractive* convention. For example $i_B^* = 0.2$ when $\alpha = 2$ and the underlying topology is the regular graph. Notice, that this threshold is not found in the case of the small world graph whose clustering coefficient is high, although not so high as the case of the regular graph. However, even without the threshold we observe a similar behaviour as in the regular graph although it appears more progressively.

If we compare it against the behavior of *low-clustered* graphs, we find again that the system is more resilient to the *attractive* action invasion for low values of i_B. For high values the contrary happens: the system is very receptive to an invasion of agents playing the *attractive* action. We must state the fact that the model is interesting for small values of i_B. For high values of i_B, the dilemma of the agent that

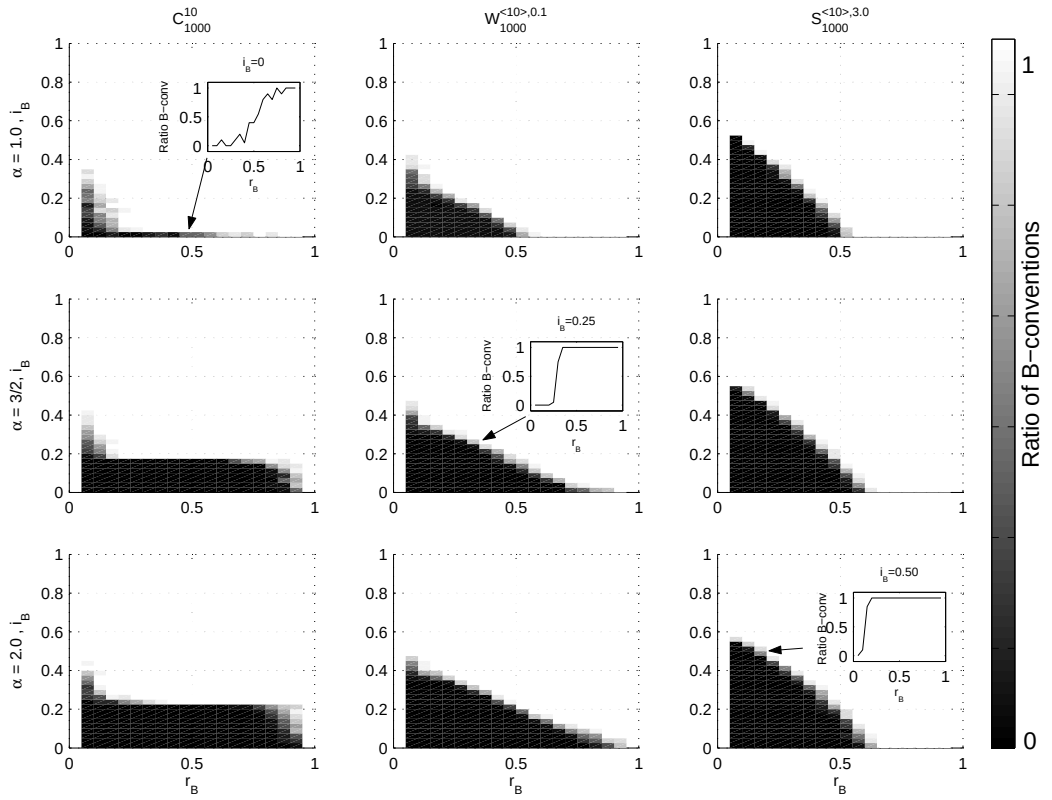

Figure 7.5: Effect of imitation propensity (i_B) on the ratio of conventions on action B. i_B ranges from [0.0 0.975] with a .0275 resolution. r_B ranges from [0.05 0.95] with a 0.05 resolution. The sub-figures within the figures are for the sake of clarity: setting the parameter i_B we observe the evolution of B-conventions depending on r_B, which correspond to the way the previous results were shown (when $i_B = 0$).

has to choose between the efficient or the *attractive* action dissapear and becomes an epidemic spread model [185].

7.7 Discussion

The research presented in this chapter aimed to identify under which conditions an efficient convention is established in a networked community. We have shown that the key factor is the clustering of the underlying social network.

When communities are *highly-clustered* the system converges to the Pareto-efficient action even though the initial population choosing that action was clearly a minority. This suggests that the efficient convention is a stable convention because it cannot be invaded by a set of agents playing another sub-optimal action. However,

a sub-optimal convention can be replaced by a set of agent playing an action that yields a better payoff. Accordingly, when the clustering coefficient is high the system always converges to the most efficient convention and this convention is stable. On the other hand when clustering tends to zero the adopted convention depends solely on the density of agents following an action. If the majority of agents play the sub-optimal action the inefficient convention will be established, and it will be stable. To back up our findings, we provided an analytical approximation that reproduces the results observed in our model based on the HCR-rule. To do so, we had to introduce a new *homogeneity* condition which let us work in clustered graphs, where the classical mean-field *homogeneity* condition is not met.

In accordance with the *strong functionalism* thesis from classical sociology, we have found that in certain graphs the agent system was capable to find and maintain the optimum in the stable state. However, this only applies to *highly-clustered* communities, which resemble many empirical social networks. At the same time, our model also corresponds in two respects to the view that *rational choice* theorists in sociology take on social norms. First, we have shown that global efficiency arises from individual goal oriented actions. Second, we have found that under certain conditions optimizing individual actions fail to generate socially efficient outcomes, a problem that is central to the contemporary discussion about the emergence of conventions and norms.

To conclude, our results seem to correspond more with a rational choice on norms than with the strong functionalism thesis. We have shown that socially optimal conventions can arise from individual optimization, but there is no guarantee that this happens. In this sense, our model matches well the ample evidence of examples of suboptimal conventions, for example in market processes. We believe that a part of the explanation for this may lie in the competition between optimization and imitation that we have addressed in our model. We have shown that imitation processes make it possible for a sub-optimal yet *attractive* action to overthrow the efficient action, and become stable, provided that its *attractiveness* is high enough to be worthy of imitation.

Part IV

Dynamics of Structure

Chapter 8

Complex Networks through Local Optimization

Throughout this thesis, structure has been presented both as a source of information about the system and as a major force driving the dynamics of the system. We feel we have give enough evidence to support our main claim, i.e., that *structure matters*. Besides that, we have also shown that certain patterns of interactions facilitate the rapid emergence of efficient social conventions without relying on centralization or planning of any sort. Furthermore, these convenient patterns of interactions arise in empirical networks that model real complex systems (see section 3.2). Therefore, in order to further explore the role of structure in artificial societies an important question remains to be addressed: which are the processes leading to the formation of such particular network structures?

In this part of the thesis we address the dynamics of structure by presenting a model based on local optimization of a social exchange process. The model presented in this chapter, as opposed to many other models of complex networks reviewed in section 3.2.1, is completely grounded in plausible assumptions, such as the use of local and imperfect information by the agents. Hence, autonomous agents, pursuing the selfish goal of finding the best set of partners to interact with, can organize themselves and exhibit a wide range of structures, including complex networks.

8.1 Macrostructures and Microprocesses

As we argued in section 3.2.1 and will continue in this section, previous modeling work has mainly relied on rather mechanistic and sociologically implausible assumptions about the processes that may generate complex networks.

Accordingly, it remains an open question from the point of view of social sciences whether and if so how complex networks can arise from the behavior of real

social agents. This question, however, is not limited to the scope of social sciences. Throughout the first part of this thesis we have exposed why we believe that multi-agent systems and other artificial societies should emphazise their social dimension. Thus, a better understanding of the social processes occurring in nature can directly benefit the design and operation of artificial societies.

If complex networks can plausibly be the product of social interactions, then their impressive efficiency and robustness may help to better answer some crucial questions in the social sciences, and by extension in what we could naively call *artificial social sciences*.

For example, how is trust possible between strangers? [239, 157]. Some researchers have argued that the efficient diffusion of reputational information through networks may be a decisive mechanism for the emergence and stabilization of trust [191, 42]. Or why can some groups quickly and effectively organize collective actions, such as strikes or demonstrations, while others fail to do so? This has likewise been related to network structures that are efficient for information transfer and recruitment processes between group members [160, 104, 105].

But why and under what conditions can societies or groups develop in the first place the networks they need to sustain trust and facilitate collective action? We believe that this question can not be answered in a satisfactory manner as long as our models of network formation rely on behavioral assumptions that are implausible as models of human social agency.

The invariants (small world, power law, etc.) found in many empirically studied social networks suggest that humans tend to organize themselves in a very particular way at the macroscopic level. Physicists have provided many models that can reproduce with high accuracy these empirical data. Nevertheless those models are based on assumptions either unrealistic from the sociological point of view (global knowledge), or they consider the agent as a mindless actor in the game. Is the emergence of structure in human societies sufficiently explained by a model of dummies interacting with other dummies? Or can this emergence be better explained as the consequence of a plausible sociological microprocesses? We believe that sociologically plausible microassumptions will not only allow to reproduce the results of previous models, but they will also help to identify new sociologically meaningful hypothesis about the conditions and patterns of complex social networks.

Physicists' models of the emergence of complex networks have mainly been formulated in the tradition of network theories informed by graph theoretic concepts. This work reflects the predominant structuralist perspective in social network research in mainly focusing on the effects of structure on individuals and largely leaving implicit the individual actions that bring structures about. In our research, we take a radically different view and emphasize individual agency as the driving force of network formation. Accordingly, we use in the following the term *agent* to refer

to the individual actors in social network, in contrast to graph theoretic approaches who often refer to individuals as *nodes* in a graph. Our actors are *agents* in the sense of agent-based modeling. In this view, an agent has four defining characteristics: autonomy, reactivity, proactivity and social ability [236, 100]. Autonomous agents have control over their own goals, behaviors, and internal states. Reactive agents perceive and react to their environment. Proactivity refers to an agent's ability to take initiative to change aspects of their environment in order to attain individual goals. Finally, social ability refers to the capacity to influence other agents in response to the influences received.

8.1.1 Previous Models

The two most prominent models of the individual level mechanisms that have been proposed to underlie the emergence of complex networks are the mechanisms of *stochastic rewiring*, described in the seminal work by Watts and Strogatz [232], and the mechanism of *preferential attachment* and *uniform growth*, proposed by Barabási and Albert [22]. This work has inspired a range of subsequent models, particularly of the emergence of power-law networks. Let us suggest to the reader to revise section 3.2.1 for a more comprehensive review of models of complex network formation.

Stochastic Rewiring

The model of stochastic rewiring has been designed to show how relatively few random changes in a network can transform an initially fully locally clustered network with large (or even infinite, cf. [228]) average path-length, into a *small world* structure that is still locally clustered, yet characterized by short average distances between the nodes. Local clustering refers here to a large degree of overlap between the ego-networks of adjacent nodes, as in small villages where residents are mainly connected to other residents and have few external contacts. Watts and Strogatz [232] showed that a small probability of *rewiring* a randomly selected tie to some randomly selected new target in the network is sufficient to bring about a radical change in the global features of the network, from a pattern with low speed of information propagation to a structure that is highly efficient for information diffusion through network ties. The main reason is that a single rewired tie does not change fundamentally the local structures of ego-networks, but it is sufficient to bridge the distance between two otherwise distant, or even disconnected, local units. Later elaborations proposed social processes such as geographical migration within or between countries as examples of this mechanism [228]. While the work by Watts and others was extremely important to demonstrate how efficient network structures may emerge from simple local processes, the authors have paid little attention to the question why and under what conditions the mechanism of stochastic

random rewiring may be a plausible result of social behavior, i.e. of the individual motives and cognitions that lead social actors to make or break relationships. It is, for example, questionable whether geographical migration may produce stochastic random rewiring. Geographical migration is often characterized by locally correlated patterns of behavior, where the experience that one pioneer migrant makes at his new residence is communicated back to old acquaintances, potentially triggering a new wave of migration to the same location. By contrast, random migration would spread new ties across the entire population, presumably with very different consequences for the network structures that may evolve. Accordingly, we believe that the conditions under which small world structures may plausibly arise in social networks can be better identified when network changes are explicitly derived from sociologically plausible assumptions about the individual goals and behavioral heuristics that underlie network changes at the microlevel.

Preferential Attachment

The model of preferential attachment also leaves the individual level decision process widely implicit. Barabási and Albert [22] developed their model to explain how in a growing network, like the network of page references in the world wide web, a power-law degree distribution emerges, where, broadly, the number of nodes that have a particular degree decreases with the degree to the power of some constant. The mechanism of preferential attachment assumes that new actors who enter the network prefer to relate to those existing nodes that already have a high degree. A plausible reason for preferential attachment, which generalizes beyond the example of the Internet, may be that actors strive for social status. New actors entering a group may expect that they maximize their own social status by relating to the most popular peers, generating preferential attachment. While the mechanism of preferential attachment may seem better motivated by social structures than random rewiring, it relies on two assumptions that seem implausible for a large range of social networks. First, the assumption that actors have public and full access to reliable knowledge about the structure of the whole network, in particular the in-degree of any other member of the network. Second, the assumption that actors are cognitively capable to process this information accurately. These assumptions are required to obtain a parametric and analytical solution that can reproduce some empirically observed networks. However, we argue that due to the behavioral implausibility of the micro-mechanism, the model of preferential attachment *reproduces* rather than *explains* complex network structures. This criticism extends to recent more sophisticated models of preferential attachment, like the one proposed by Walsh [224]. In this model, a new node that is added to the network is attached to a node in the existent network following a probability distribution that depends on

the degree of the nodes present in the network. While this model is very successful in generating graphs with the power-law property, it also assumes that agents know the complete distribution of degrees of the graph and process this information accurately when choosing to whom they want to relate.

Behavioural Models

Several researchers have addressed the issue of behavioral models underlying complex network dynamics. *Transitive linking* was in particular proposed by Krapivsky and Redner [142] and by Ebel et al. [76] as a mechanism where agents use solely information about their local ties. In a nutshell, transitive linking assumes that new ties are most likely formed between actors who have a common acquaintance in the network. This work could show that *preferential attachment* behavior may arise from the transitive linking process, even when individual actors have no global knowledge about the network. However, like the original preferential attachment model, transitive linking does not make the underlying social mechanism explicit. While transitive linking may be a plausible description for realms such as co-authorship networks, it seems, for example, less straightforward for the dynamics of help or advice exchange. Why and under what conditions should two actors who have exchanged professional advice with the same third party, become more likely to help or advice each other? While their common acquaintance may have needed advice of both of them, the advise givers may not necessarily need each others' help. Clearly, the decision to enter an exchange relationship depends not only on having common acquaintances, but also on the proper match of demand and supply in exchange relations.

Recently, social network researchers increasingly call for sociologically more plausible models of complex network emergence. For example, Robins et al. [193] made important progress by showing how computational methods allow to simulate distributions of Markov Random Graphs including small word structures, based on assumptions about simple and strictly local processes. Their approach specifies local dynamics in terms of statistical interdependencies between ties in the network neighborhood of a node. While this work mainly aims at testing whether a given model of a local process can explain observed complex network structures, it does not explicitly address individual motivations or cognitions underlying local dynamics.

Optimization Processes

To derive the emergence of complex networks from more fundamental principles of individual behavior, a model of some *optimization process* is presented in which actors make or break ties to attain specified individual goals, while agents are at the same time restricted by their current position in the network. With the approach of *highly optimized tolerance* (HOT), Carlson and Doyle [47] proposed system op-

timization as explanation of the emergence of complex system structures including complex networks. This work shows that in general power-law distributions of certain individual features may emerge when a system is optimized by a selection process resembling *survival of the fittest*. HOT models can be extended to network design, for example with the goal to make the network both efficient for information distribution and robust against randomly distributed risks of the destruction of single nodes. However, we argue that the HOT approach is problematic as basis for a model of the emergence of unplanned social networks, because the locus of decision making in this model is not the individual agent, but a central network optimizer. Ferrer and Solé [83] move a step further and explain complex network emergence from individual optimization. In their approach individuals aim to optimize global structural measures of the network, such as density and average path length. Their model may be adequate to describe a situation in which individuals' interests perfectly coincide with collective goals of, e.g., a social group or organization to which actors belong. But Ferrer and Solé neglect the partial conflict between common goals and individuals' interests that characterizes network formation in a range of areas, such as inter-firm technological cooperation [188] or social support exchange [84]. In these realms, network formation is at least partly driven by competition between agents for relationships with attractive exchange partners, while at the same time agents are restricted in the time and effort they can invest in network contacts. For example, Bonacich [32] shows that in a research team, a *communication dilemma* may arise, when status gains or bonus payments reward exceptional individual performance, while at the same time the team as a whole may compete against other teams. The team as a whole may benefit when all members aim to maximize the global efficiency of information exchange networks between them. At the same time, individual team members face incentives to withhold information and collude instead in dyadic exchanges or within small cliques to outperform their colleagues. Clearly, a model which assumes that agents optimize on global network properties is a misleading template for this type of situations where network change faces agents with a social dilemma [62].

To sum up, we argue that a sociologically plausible model of the emergence of complex exchange networks should have the following features. First, the locus of decision making about network changes is at the level of individual actors. Second, individual actors make or break ties to optimize their individual goals based on bounded rational individual decision heuristics. Third, individual agents use for this optimization only local and imperfect knowledge of network characteristics. Finally, common goals and individual interests are with respect to desirable network characteristics at least partly in conflict with each other, i.e. network members face a social dilemma situation.

In the following, we propose a model that has the features listed above. We will

show that our approach is not incompatible with the classical *preferential attachment* model, but can generate preferential attachment behavior under certain conditions. However, relational change is in our model just a consequence of optimization at the agent level. An example for this could be the emergence of scientific collaboration networks, where scientists do not choose their co-authors based on their connectedness in the network, but in terms of whom they perceive as the best partner to conduct research with. In this example, preferential attachment is not the underlying and driving mechanism, but it may be a consequence of the underlying local optimization process. The empirical study presented in section 4.3 that compared the network of website references in a scientific community to scientists' publication scores suggests this possibility. The experiments show a substantive correlation between journal publication scores and scientists' centrality in the network of website references.

8.2 Local Optimization in Social Exchange Networks

Our model is based on social exchange theory and assumptions of bounded rationality [29, 118]. As baseline, we take a model of the dynamics of social exchange networks that was originally proposed by Flache and Hegselmann [114, 84, 85, 86]. Their model assumes that agents seek to find and keep attractive exchange partners in a population where agents are to some degree free to exit from ongoing exchange relations, differ in attractiveness as exchange partners, for example due to variation in the amount of material resources at their disposal, and have only limited access to and information about potential new partners. These assumptions reflect empirical results on social exchange, which show for example that in relations of gift exchanges, actors select exchange partners partly on basis of the amount of resources they may expect to attain in the exchange [141]. Similarly, in certain industries, technology cooperations between firms constitute network patterns in which firms seek to establish relationships with partners both based on technological attractiveness and status in the industry of potential partners, as Podolny and Page [188] suggest. Finally, in networks of advice exchange, actors often differ in their degree of expertise and thus attractiveness as advice givers. For consistence of our terminology, we will use in the remainder *advice* or *support* to refer to the exchange commodity, but we wish to point out that the scope of our analysis generalizes to most forms of social or professional help exchange and collaboration networks.

To model individual goals, we assume that agents gain from being supported but incur some loss if they provide support themselves. More precisely, the larger the amount of help an agent receives, the larger his gain is. Conversely, the more help an agent provides, the larger his loss. This implies in particular that agents try to avoid exchange with partners who are either not capable or not willing to give good advice.

We assume that agents aim to optimize their exchanges in terms of these goals under imperfect, local information without initial knowledge about others' characteristics or knowledge about the global network structure. Moreover, agents in our model are strongly adaptive, i.e. they acquire knowledge only in the course of interaction and apply simple search heuristics to select potential new partners. As we will show, these assumptions do not preclude the emergence of complex networks.

8.2.1 Memory and Interaction

The network consists of a population of N agents. Each agent i is characterized by his individual attractiveness for other agents, α_i, representing for example his expertise, material wealth or beauty ($0 < \alpha_i < 1$). However, before interaction, every agent is ignorant not only of others' attractiveness but also of its own. This expresses that people learn their *market value* only through interaction with exchange partners. For simplicity, we assume that attractiveness is initially distributed uniformly and remains constant throughout the simulation. However agents' perception of their own and others' attractiveness changes through interaction.

The second attribute in which individual agents differ from each other is their memory, representing the knowledge they acquired in previous interactions. More specifically, at any point in time the memory of agent i contains for a number of other agents j the recollection i has of the payoff o_{ij} that he attained in previous interactions with j. To express bounded rationality, we assume that agents remember only a small subset of all system members. Formally, memory is represented as

$M_i^t = \{(o_{ij}, t_{ij})\}_{j \in J_i^t}$,

where t_{ij} represents the time point at which the memorized interaction took place and J_i^t represents the subset of all agents that i *memorizes*, $J_i^t \subseteq \{1..N\}$. One time step corresponds to an interaction round in which all the agents are activated once in random order to make decisions (asynchronous random activation). The subset of network members which an agent can remember is limited in size by a maximum memory size M_C, i.e. $\#J_i^t \leq M_c \; \forall t \forall i$. We model the maximum memory size M_C as a parameter that is equal for all agents, because we are mainly interested in effects of structural conditions that shape individuals' learning capacities through their ability to store information. Such a condition is for example the capacity of the information storage technology that is available in a society [159]. Obviously, agents have in our model only access to a partial view of the network, and their knowledge remains local, i.e. it depends exclusively on the individual experiences of agents.

To model the process through which agents gradually get to know other network members, we assume that the agents represented in the memory of an individual i fall apart into two distinct subsets, the *known* agents K_i^t, and the *unknown* agents

U_i^t. An agent j is unknown to i when i does not know the characteristics of j but is at least aware of j's existence. Agents of whose existence i is not aware are not represented at all in i's memory. The total set of all network members of whose existence i is aware is denoted J_i^t, where $J_i^t = U_i^t \cup K_i^t$ and $U_i^t \cap K_i^t = \emptyset$. To have a well-defined baseline situation, we assume that at the beginning of the simulation no one is *known* to anybody else, but each agent has a number of M_o randomly chosen network members he is aware of, but who are classified as *unknown*. Only in the course of the simulation, i will learn about the existence of more agents and they may even become *known* for i through interaction. Accordingly, the recollection (o_{ij}, t_{ij}) of the payoff that i received after interaction with j at time point t is only defined in i's memory when j is known, i.e. $j \in K_i^t$. Both the memory as a whole and its partition into subsets of known and unknown agents may change over time, but it remains always bounded by the maximum memory capacity M_c.

In the interaction process, the agents are activated in random sequence. Once activated, an agent tries to establish Q interactions with agents contained in his memory. Following previous models of complex network dynamics, for example [22, 228, 224], we assume that the initiation capacity Q is a system-wide parameter that represents universal technological or social constraints (i.e. communication technology or politeness norms) on the number of interactions an agent can initiate simultaneously. This parameter implements one aspect of the harshness of the exchange situation actors face. The smaller the initiation-capacity relative to the total number of exchanges in which an agent can be involved, the harder it is to find a sufficient number of attractive partners.

The agent selects the targets of his interaction initiations to optimize the outcomes he expects to attain from the interaction. To reflect bounded rationality, we assume that decision making about action initiation is entirely adaptive or backward-looking, as done by Macy and Flache [156]. That is, agents derive their expectations solely from the experience represented in their memory. Details of the optimization procedure will be explained below, when we turn to individual decision heuristics.

Interactions are dyadic and both agents need to agree before an interaction actually happens. Correspondingly, agents also face a limitation of the number of interactions they can be simultaneously involved in. We call this limitation their *interaction-capacity* (C). More specifically, the interaction-capacity C restricts the maximum number of interactions per agent per time unit. Agent's interaction-capacity must be equal or smaller than their memory-capacity, $C \leq M_c$. We introduce the interaction-capacity parameter because in many social network dynamics actors have limited resources to build up and maintain social ties, particularly when an interaction implies a certain *work load* or *time investment*. Notice that most previous models in the literature on complex networks take it for granted that actors have a potentially unlimited capacity to entertain simultaneous relationships.

This applies in particular to the model of preferential attachment. This assumption may be applicable for situations where an interaction does not imply a work load or time investment for both parties, such as deployment of a link from one web page to another. However, we consider unlimited capacity an exception rather than the rule in social network dynamics, where interactions usually require the allocation of some limited resource.

8.2.2 Actions, Outcomes and Payoffs from Social Interaction

Following Flache and Hegselman [85, 86], we model the actual exchange between agents as a support game. For simplicity, both actors have only two decision options in the constituent (one-shot) game, to provide advice (*C*ooperate), and to not provide advice (*D*efect). To further simplify, we assume that the interacting agents make these decisions simultaneously and independently. Accordingly, the outcomes of one interaction between two agents i and j can be DD, mutual defection, unilateral support of i by j or vice versa, DC or CD, or mutual cooperation, CC. The effects of outcomes on agents' goal attainment depend on the attractiveness levels of both participants. The larger the attractiveness of the focal agent, α_i, and the smaller the attractiveness of his interaction partner, α_j, the more advice i will give (and j receive), if i decides to actually cooperate with j (C).This expresses that in advice exchange a recipient may maximize his profit by getting advice from a very knowledgeable person, say a *guru*, while the guru may receive little new knowledge from an average advice seeker.

Technically, we assume that i's cooperation benefits the recipient with $\alpha_i(1 - \alpha_j)$ *units* of advice. However, when the focal actor fails to support his interaction partner, he provides zero units of advice. Conversely, i receives (and j gives) $(1 - \alpha_i)\alpha_j$ units if j decides to support i (C). At the same time, i receives no advice, when j declines to give support (D).

The payoffs that accrue to both actors from the exchange, p_{ij} and p_{ji}, follow from the balance of their respective costs and benefits. To model payoffs, mutual defection, DD, is used as the baseline outcome that yields a payoff of zero to both participants. In the outcome DD, actors neither receive advice nor provide it. Technically, we model i's gain from receiving advice from j, G_{ij} and i's loss as a result of giving advice to j, L_{ij}, as follows:

$$G_{ij} = (1 - \alpha_i)\alpha_j B \tag{8.2.1}$$

$$L_{ij} = \alpha_i(1 - \alpha_j)E \tag{8.2.2}$$

The parameters B and E are positive constants that weigh the benefit, B, of receiving one unit of advice against the effort costs, E, of providing the unit. It is a

central assumption in our analysis that self-interested individuals may in principle benefit from advice exchange. To ensure this, we assume that the benefits of receiving a unit of advice exceed the costs of *producing* it. The cost to benefit ratio, $\frac{E}{B}$, parameterizes a further aspect of the *harshness* of the exchange situation actors face. The higher this ratio, the more severe the loss agents incur if they are unilaterally exploited compared to the potential benefits of mutual cooperation. In technology cooperations, harshness may vary, for example, between sectors with different firm sizes, where in sectors with small firms, agents are more vulnerable to financial losses caused by others' defection than in sectors with large firms. Similarly, in academic environments with high emphasis on competition and individual publication scores, advice giving or collaboration are potentially more risky and costly than in more collaborative environments.

Figure 8.1 illustrates the incentive structure that ensues for the constituent support game. The first entry in a cell refers to the payoff of the row player i, the second entry indicates the payoff of the column player j. Notice that the game is not necessarily symmetrical. Players obtain different outcomes in symmetrical choice combinations, unless they are equally attractive.

		Agent j	
		C	D
Agent i	C	$(G_{ij} - L_{ij}\ ,\ G_{ji} - L_{ji})$	$(-L_{ij},\ G_{ji})$
	D	$(G_{ij},\ -L_{ji})$	$(0,\ 0)$

Figure 8.1: Payoff Matrix of the Support Exchange Game

Figure 8.1 shows that the support game may be a true social dilemma where cooperation conflicts with self-interest. There is nothing that guarantees reciprocation within one iteration. Particularly when players expect a short-term relationship and conditions are harsh, they may be tempted to withhold support. Exploiting a partner who provides advice is in the short run the most profitable outcome for a selfish agent and to be exploited by a partner who fails to give advice is least attractive, regardless how attractive the players are. Clearly, these incentives may face self-interested actors with a particularly difficult social dilemma, the Prisoner's Dilemma (PD). At the same time, the game is not necessarily a PD. In a PD, both players prefer *mutual* cooperation (CC) to *mutual* defection (DD) despite incentives to defect *unilaterally*. In our support game, it is possible that only the less attractive player may be interested in mutual support. More precisely, when the focal player is not attractive enough in comparison with his counterpart, then it may be impossible for the focal player to receive enough advice to be compensated for the investment in supporting his partner ($G_{ij} < L_{ij}$). To be precise, the support game is a Prisoner's Dilemma if and only if for both players it holds that $(1 - \alpha_i)\alpha_j B > \alpha_i(1 - \alpha_j)E$.

Notice that our model also implies that the capability to provide advice and the need for advice are inversely related. The more advice an agent needs in a certain period of time, the less he can give to others in the same period. We feel that this assumption is plausible for many social exchange relations. For example, consider the effects of variation in expertise on the individual neediness for advice in an academic community. Expert members may both have more knowledge to share with others and less need for advice themselves, as compared to less knowledgeable members. For another example, in a rural village only some farmers may be wealthy enough to afford expensive farming machinery. These farmers do not need to borrow others' machines, but they might lend their machinery to less wealthy farmers.

To model local and imperfect knowledge of agents, we assume that B and E, as well as the attractiveness levels are not explicitly known. That is: agents pursue their goals facing a high degree of uncertainty. We explain in the next subsection which knowledge agents actually do have once they decided to enter into an interaction.

8.2.3 Agent's Decision Heuristics

When an agent i is activated, he first selects another agent k to propose an interaction in the current simulation round t.

Following recent models of *socially rational* decision heuristics, such as the proposed by Lindenberg in [152], we assume that agents both seek to optimize their expected returns and tend to search for improvement through new, untried courses of action. Technically, this is modeled with an exploration probability e_i^t with which an agent within i's current memory J_i^t is chosen randomly. When the focal agent optimizes rather than explores, the interaction partner with the best expected payoff is chosen from the memory of known others, K_i^t. If there is more than one known agent with the same maximum expected payoff, one of the optimal agents is chosen at random.

Initially, agents select interaction targets at random, i.e. the exploration probability is 1. We assume that the tendency to explore declines as agents get to know more other network members. This implies that over time the exploration probability differs between agents. In particular, popular agents quickly decrease their exploration probability because they receive many interaction proposals, and, accordingly, learn faster about their environment than agents who receive less interaction proposals. Technically, we model the exploration probability at time point t as the ratio of the number of unknown agents in i's memory to the square of the number of all agents i is aware of at that point in time.

$$e_i^t = \left(\frac{\#U_i^t}{\#J_i^t} \right)^2 \tag{8.2.3}$$

After an interaction has been initiated, the targeted agent j needs to decide whether to accept the proposal of the initiator i. We assume that j will only accept the interaction if j has not reached its interaction-capacity C of concurrent interactions and j is aware of i's existence. If, furthermore, the initiator is not known to j, then the proposal will be accepted. If the initiator i is known to j, then j will accept the proposal if and only if he expects the outcome to be not negative.

If the interaction proposal is accepted, agents i and j play one round of the support game. For simplicity, we exclude in the present version of our model strategic free riding and assume instead that agents are *benevolent*. We plan to take into account strategically opportunistic behavior in future work. For the time being, our assumption implies that an agent is cooperative, but not altruistic. To be precise, i cooperates with j if and only if the expected *net benefit* from i's point of view is positive, $(G_{ij} - L_{ij}) \geq 0$. In particular, when an agent has no prior experience with his counterpart, he can reliably assess the expected net-benefit from mutual cooperation. This assumption reflects that usually in social exchange agents will learn about each others' characteristics *after* a mutual decision to interact with each other and *before* they enter into the actual exchange. Accordingly, agents may actually defect in the exchange despite their benevolence. This happens, when they find out that their interaction partner seems to be less attractive than they may have thought before they entered the interaction. For example, when two scientists decide to talk some time together on a conference because they are interested in each others' work, one of them may find out in the course of their discussion that the other is less knowledgeable than expected. As a consequence, he may invest little effort in actually giving useful advice and rather utter some general and vague comments about the colleagues' work.

After an interaction is completed, agents update their respective memories. If at that point in time an entry for the interaction partner does not yet exist in an agent's memory, then a free memory slot will be allocated for this purpose. If there is no free slot available because the memory is full, then the agent will *forget* some other previous interaction partner and replace the corresponding old memory entry with the information about the most recent interaction. We assume that agents are most likely to forget previous partners that were experienced as unattractive. More technically, the agent k who will be forgotten by i is selected according to the following rule:

$$k = argmin_{n \in J_i^t}(\mid o_{in} \mid \mid (t_{ik} < t)) \tag{8.2.4}$$

where t represents the current point in time.

In case that i has prior experience with his most recent interaction partner, the corresponding memory entry will be updated to average value of the most recent

payoff and the payoff stored in the memory provided that i cooperated in its interaction with j. Notice that this implies that the weight of old memories declines after every new interaction with the same partner. Conversely, when i defects in its interaction with j the expected outcome will be set to $-p_{ij}$. This reflects our assumption that agents are not strategic. While a unilateral defection is actually a positive outcome for the agent, the negative memory entry after a unilateral defection implies that agents do not try to exploit that partner again, but avoid instead future interactions because they failed to establish a mutually profitable interaction with j. Technically,

$$(o_{ij}, t_{ij}) = \begin{cases} (\frac{o_{ij}+p_{ij}}{2}, time) & \text{if, } t_{ij} > 0 \\ (p_{ij}, time) & \text{if, } t_{ij} = 0 \\ (-p_{ij}, time) & \text{if, } G_{ij} - L_{ij} < 0 \end{cases} \tag{8.2.5}$$

where t_{ij} refers to the current simulation time and p_{ij} denotes the payoff attained in the most recent interaction. When an interaction proposal is rejected, the memory is also updated according to 8.2.5, where the payoff p_{ij}, is set to zero.

How can agents ever get to know new network members, when their exploration only is restricted to the people they are already aware of? In previous work on the diffusion of reputations in social networks two different mechanisms are used or sometimes combined, social learning through observation [242] or gossip [191, 42]. Gossip models assume that reputation information spreads as a byproduct of interactions. To model gossip, we assume in our analysis that agents exchange some knowledge from their respective memories as an act of deference or courtesy after they have experienced a mutually profitable support interaction, that is: $p_{ij} > 0$ and $p_{ji} > 0$. More precisely, in this *explicit* memory exchange, the interaction partners i and j each tell their counterpart which other agent in their current memory has the best expected outcome. More technically, the referred agent k is found as $k = argmax_n(\mid o_{in} \mid \mid (k <> j \wedge t_{ik} > 0))$. To distinguish effects of explicit knowledge exchange from those of social learning through observation, we use as an alternative assumption *implicit* knowledge exchange. We assume that implicit knowledge exchange provides less reliable information than explicit gossip. While i can observe with whom j is currently interacting, i does in implicit knowledge exchange not know the payoffs that j derives from this interaction. Hence, in this mode i picks a randomly chosen agent from j's memory with whom j is currently interacting.

To model bounded rationality and uncertainty also in the knowledge exchange process, we assume that the knowledge an agent acquires always reflects the subjective perceptions of the sender of information, but not necessarily the interests of the recipient. Concretely, agent i gives to agent j the expected outcome from an interaction with a third agent k from i's perspective, o_{ik}. Neither the recipient nor

the sender have the capability to assess which agent is most suitable from j's point of view. Finally, we assume that in future interactions the new knowledge acquired is also subject to an updating process similar to the process described above, with the one difference that the time-point t_{ik} to which the information refers in the recipient's memory is set to 0, to express that the knowledge of i about k stems from referral or observation rather than direct experience. Technically,

$$(o_{ik}, t_{ik}) = \begin{cases} (\frac{o_{ik}+o_{jk}}{2}, 0) & \text{if, } t_{ik} = 0 \text{ and } \exists o_{ik} \\ (o_{jk}, 0) & \text{if, } t_{ik} = 0 \end{cases} \tag{8.2.6}$$

Knowledge stemming from referees or observation only applies when *first-hand* knowledge is not available (denoted by $t_{ik} = 0$). Once direct experience knowledge about agent k exists ($t_{ik} > 0$), agent i stops updating the expected outcome o_{ik} based on *third-parties* experience.

8.3 Results

We are mainly interested in the effects of two key sets of conditions in our model on the topological structure of emergent social exchange networks. The first set of conditions refers to the harshness of the exchange game, in particular the cost to benefit ratio in exchanges and the difficulty to access exchange partners. The other conditions pertain to the learning capability of the agents, notably their memory size, exploration probability and knowledge exchange mechanism. Due to limitations of space, we can not fully explore all corresponding parameters in the present paper. We present in the following results from experiments that manipulated the ratio of costs to benefit in exchanges $\frac{E}{B}$ to vary the harshness of game. Furthermore, we varied the type of memory exchange (henceforth ME) to study effects of learning capability. To explore robustness, we also varied population size, N, the capacity of agent's memory, M_c, the interaction-capacity C and the initiation-capacity Q.

To have a well defined baseline scenario, we fix the remaining parameters of the model to values that make the emergence of complex network possible, but not trivial. The number of agents initially present in the memory, M_o, is 150. Finally, the attractiveness of agents is uniformly distributed in the range between 0.05 and 0.95 with a precision of 10^{-3}.

Our main interest is whether our local interaction model (henceforth LO-model) can generate similar complex network structures than previously proposed mechanisms that assume global knowledge or perfect maximization or both. In order to compare our results with other models we use the original model of Barabási and Albert [22]. This model is based on *preferential attachment* and *uniform growth* and it is representative for the family of preferential attachment models which assume

Network	N=1000		N=5000		N=10000	
	l	C	l	C	l	C
Random	3.27	0.0085	3.97	0.0016	4.30	0.00093
BA	2.98	0.037	3.52	0.010	3.76	0.0044
$LO_{(E/B=0,I)}$	2.47	0.16	3.22	0.05	3.56	0.034
$LO_{(E/B=0,E)}$	2.51	0.20	3.16	0.085	3.42	0.065
$LO_{(E/B=3/16,I)}$	3.66	0.052	4.09	0.038	4.29	0.032
$LO_{(E/B=3/16,E)}$	3.71	0.034	4.14	0.030	4.37	0.029
$LO_{(E/B=8/16,I)}$	3.98	0.012	4.45	0.015	4.68	0.017
$LO_{(E/B=8/16,E)}$	3.91	0.027	4.57	0.040	4.84	0.038

Table 8.1: General characteristics of emergent networks. LO refers to our local optimization model. BA denotes the Barabási-Albert model. *Random* refers to a random network using the Erdos-Rényi model. Initiation-capacity Q, corresponding to half the average degree, $\frac{\langle k \rangle}{2}$, is set to 5. The interaction-capacity C is set to 150. The memory capacity M_c is set to 200. The population size, N, varies across $\{1000, 5000, 10000\}$, the cost to benefit ratio ($\frac{E}{B}$) varies across $\{0, \frac{3}{16}, \frac{8}{16}\}$. Finally, we vary memory exchange between explicit and implicit, indicated as $\{E, I\}$, respectively. For each condition the table charts average path length l and the clustering coefficient C of the emergent network, together with the average path length and clustering coefficient of the random and the Barabási-Albert network.

that agents have complete and perfect knowledge about the structure of the network.

To compare structural features of networks, we use in the following always the network structures that evolved after 100 units of time. After this number of interaction rounds, the dynamics of agents' exploration probability specified in equation 8.2.3 assure that there is practically no more exploration so that the presence of an undirected edge (i, j) indicates a stable mutual support relationship between i and j. Exploration rates tend to zero over time, because the ratio between unknown agents and total number of known agents in individuals' memory decreases monotonically due to learning through direct experience or by referral.

Table 8.1 gives an overview of the effects of a range of simultaneous manipulations. We varied the population size, N, across three levels $\{1000, 5000, 10000\}$. The cost to benefit ratio ($\frac{E}{B}$) changes from a game that is no social dilemma at all with zero costs of effort, to a mild social dilemma where the ratio is $\frac{3}{16}$ to a fairly harsh game in which the costs are half the benefit, i.e. $\frac{8}{16}$. Moreover, the table illustrates effects of variation between explicit and implicit memory exchange, indicated as $\{E, I\}$, respectively. Table 8.1 shows for every condition two main characteristics of the simulated networks, their average path length l and their clustering coefficient

C. These are compared in table 8.1 to the corresponding characteristics of random networks of equal size, l_{rand} and C_{rand} and equal average degree. The average degree $\langle k \rangle$ corresponds to twice the initiation-capacity Q which is set to 5 in all conditions.

As Watts and Strogatz [232] have argued, this comparison allows to identify whether our networks have complex network features, particularly the small world property. To explain, the average path length l indicates the average number of network ties that form the shortest path between two nodes in the network. The clustering coefficient C measures the extent of overlap between the ego-networks of related nodes. Complex networks are characterized by an average path length that is similar to graphs with a random structure, while at the same time the degree of clustering is much higher. Moreover, in complex networks the average path length increases logarithmically in network size, just as in random networks. A comparison of the columns for the average path lengths l_{LO} and the clustering coefficient C_{LO} generated by the local interaction model with the corresponding properties of random graphs, l_{rand} and C_{rand} shows that our model indeed generates networks with *small world* properties. More specifically, we find that across all conditions the LO-networks have similar average path lengths compared to random graphs of the same size and average degree, $l_{LO} \simeq l_{rand}$, and the clustering of LO-networks is much higher compared to these random graphs, $C_{LO} \gg C_{rand}$.

Table 8.1 furthermore shows that the LO-networks also satisfy the small world criterion proposed by [223]. According to this criterion, $\frac{l_{rand}C}{lC_{rand}} \gg 1$ in a small world network. Table 8.1 also highlights an important difference between the Barabási-Albert mode (BA) [22] and the local optimization model (LO). The clustering coefficient of networks generated by the local optimization model is much higher compared to the corresponding BA-networks, $C_{LO} \gg C_{BA}$. As discussed in section 3.2.1, these high clustering coefficients can be seen as an indication that our model corresponds well to the structures that many social networks exhibit. Remarkably, like the Barabási-Albert model, most other models in the literature fail to reproduce clustering coefficients that are comparably high.

To shed light on the specific types of complex networks generated by the local optimization model, we studied the emergent degree distributions. Figure 8.2 shows the degree distribution of networks produced with the LO model under varying cost to benefit ratios and knowledge exchange mechanisms. The horizontal axes in the graphs represent the degree, scaled logarithmically, and the vertical axes indicate the frequency and the degree frequency with which the corresponding degree occurs in the network, also scaled logarithmically. The typical pattern of a power-law complex network in such a graph is a straight declining line, indicating that the number of nodes with a particular degree decreases with the degree to the power of a constant. The figure 8.2 shows that depending on the exact conditions of the exchange, different kinds of distributions are observed. We can identify three clearly

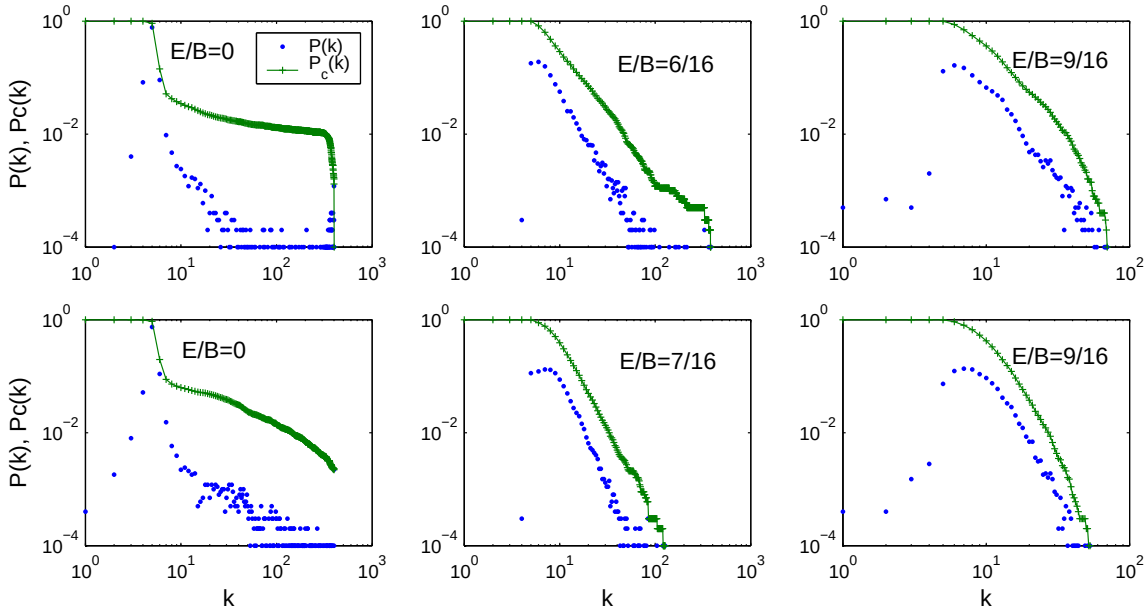

Figure 8.2: Degree distribution obtained through Local Optimization. The dots are the degree distribution $P(k)$ (frequency of nodes with degree k). The line with the plus sign is the cumulative degree distribution $P_c(k)$. Parameters: Q, or $\frac{\langle k \rangle}{2} = 5$, $N = 10000$, $M = 500$, $C = 400$. First row: explicit memory exchange. Second row: implicit. Cost to benefit ratio: $\frac{E}{B} = \{0, \frac{6}{16}, \frac{9}{16}\}$, for explicit memory exchange, and $\frac{E}{B} = \{0, \frac{7}{16}, \frac{9}{16}\}$ for implicit

distinct types:

1. star-like distribution (left sub-figures within figure 8.2, $\frac{E}{B} = 0$): These networks can be classified as *central-periphery networks*, where a small subset of nodes forms a highly clustered core network, and the majority of remaining nodes connects solely to the core. In other words, most nodes have either a small or a high number of relationships, but nodes with intermediate degree are rare.

2. potential distribution (central sub-figures within figure, 8.2, $\frac{E}{B} = \frac{6}{16}$ and $\frac{E}{B} = \frac{7}{16}$): In these networks, low degrees are most frequent and high degrees are least frequent. More specifically, the characteristic property of a potential distribution is its power-law structure, i.e. $P(k) \sim k^{-\gamma}$, where γ is the exponent that indicates the rate of decline of degree frequency. These networks, also called *power-law networks*, or *scale-free networks* are by many authors seen as the paradigmatic case of complex networks, due to their amazing robustness and efficiency (see section 3.2.1). Most of the complex network literature heavily focuses on this type of networks. *Power-law* networks generally have

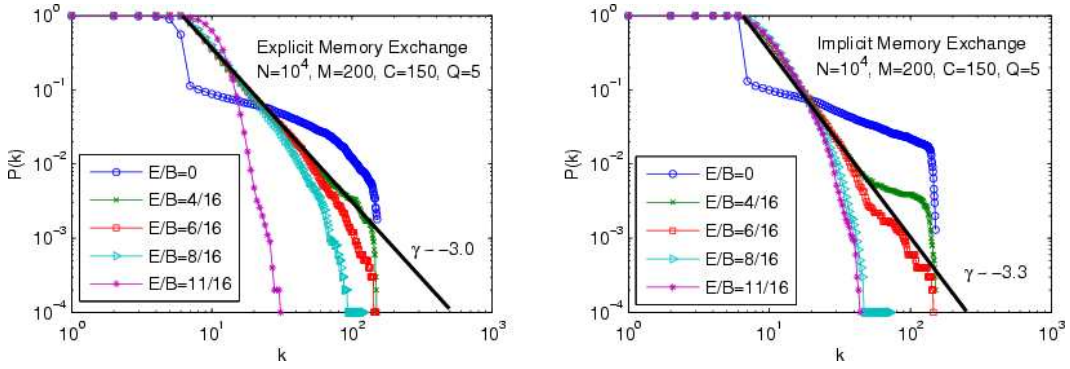

Figure 8.3: Cumulative degree distribution obtained through Local Optimization. Parameters: Q, or $\frac{\langle k \rangle}{2} = 5$, $N = 10000$, $M_c = 200$, $C = 150$. Left figure: explicit memory exchange. Right figure: implicit. Cost to benefit ratio: $\frac{E}{B} = \{0, \frac{4}{16}, \frac{6}{16}, \frac{8}{16}, \frac{11}{16}\}$. The solid line is the analytical power-law degree distribution with $\gamma = -3$ in the left figure, and $\gamma = -3.3$ in the right figure

small world properties, but not all networks with small world properties are necessarily *power-law*. Potential distributions allow nodes to have a very high degree, therefore, the degree distribution displays the *heavy-tail* effect. That is not a unique feature of power-law networks, it is also observable in star-like networks, although their degree distribution do not follow a straight line in the log-log scale.

3. exponential distribution (right sub-figures in figure 8.2, $\frac{E}{B} = \frac{9}{16}$): For exponential distributions the probability density is $P(k) \sim \theta^{-k}$. In practice, exponential degree distributions are similar to power-law distributions, with the main difference that the frequency of degrees tends to decline faster to zero as the degree increases, therefore the characteristic *heavy-tail* of the power-law distribution disappears. This is exemplified by the different scaling of the horizontal axes in figure 8.2 for power-law and exponential structures, respectively. As the figure 8.2 shows, in the exponential networks the frequency of nodes with a certain degree tends to decline to zero before the degree exceeds 100, whereas degrees up to the maximum of 400 still can occur in the power-law structures .

To further illustrate the qualitatively different network patterns shown in figures 8.2 and 8.3, we replicated the simulations for a set of conditions that allows easy graphical representation of the network topology ($N = 200$, $M_c = 150$, $C = 150$, $M_o = 10$, $Q = 1$, $ME = explicit$). Figure 8.4 shows three different networks obtained by the LO-model for three different levels of $\frac{E}{B}$. The graphs confirm visually the three qualitatively different regimes. The star-like structure in the leftmost figure

is exemplified by two big clusters that each center around a few highly central nodes. The power-law pattern (middle figure) shows a few highly connected nodes in the center of the graph, with layers of nodes surrounding the center that are increasingly less connected as they are more distant from the center. Finally, the exponential regime (rightmost figure) exhibits a more modest level of centralization than the power law graph, visually represented by a less clearly accentuated center-periphery structure. We used the layout proposed by Kamada and Kawai [127] to draw the graph.

To summarize, the above figures show that the local interaction model can generate not only power-law structures, like many other models in the literature, but it also produces central-periphery networks and small world networks with exponential degree distribution. To explore model robustness, we repeat in figure 8.3 the experiment for different memory size M_c, and capacity C. The figure shows that the same qualitatively different structures appear than in the previous experiments. We only graph the *cumulative* degree distribution $P_c(k)$ in figure 8.3 and the following figures, because this is a better statistical estimator for small samples than the degree distribution.

To what extent do the results of the local interaction model correspond to structures generated by a model that is specifically designed to reproduce complex network structures? To answer this question, we used the model proposed by Krapivsky et al. [143] who extended the original preferential attachment algorithm of Barabási and Albert [22]. Preferential attachment can be formulated as the assumption that the likelihood of receiving a new relationship increases with the node's connectivity degree k_s. Formally, $\pi[k_s(t)] = \frac{k_s(t)}{\sum k_s(t)}$. In [143] Krapivsky et al. extended the preferential attachment to be non-linear, that is, $\pi[k_s(t)] = \frac{k_s(t)^\alpha}{\sum k_s(t)^\alpha}$. It turns out that this model generates three distinct regimes depending on α that clearly correspond with the qualitatively different patterns generated by the LO-model:

1. Linear case ($\alpha = 1.0$): this is the original Barabási-Albert model, which produces the well-known power-law degree distribution $P(k) \sim k^{-\gamma}$, where $\gamma = 3$.

2. Sub-linear case ($\alpha < 1.0$): the degree distribution is an stretched exponential of the form $P(k) \sim k^{-\alpha} e^{-Ak^{1-\alpha}}$. However, when α tends to 1.0 the potential part dominates the exponential becoming very similar to a power-law distribution with an exponent $\gamma > 3$.

3. Super-linear case ($\alpha > 1.0$): In this regime there is no analytical solution, but for $\alpha > 2.0$ a *winner takes all* phenomenon emerges, such that almost all the nodes connect to a single node. 8.2) for the LO-model.

Figure 8.5 illustrates those different regimes. The displayed data has been generated using the original Barabási-Albert algorithm with a couple of modifications:

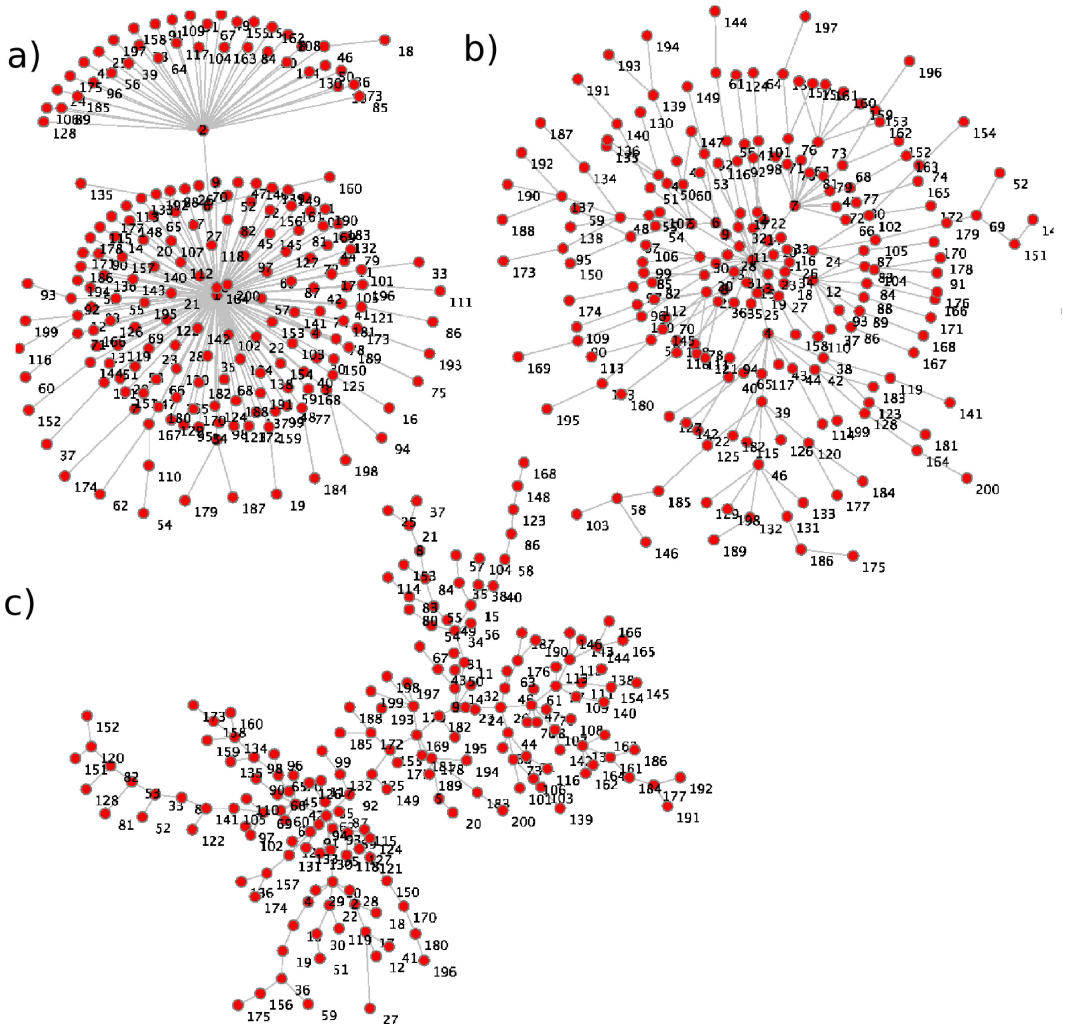

Figure 8.4: Network obtained by the model LO. Parameters: $N = 200$, $M_c = 150$, $C = 150$, $M_o = 10$, $Q = 1$, $ME = explicit$. The graph a is for $\frac{E}{B} = 0$ and corresponds to a star-like network. The graph b is for $\frac{E}{B} = \frac{2}{16}$ and corresponds to a power-like network. The graph c is for $\frac{E}{B} = \frac{8}{16}$ and corresponds to a exponential network.

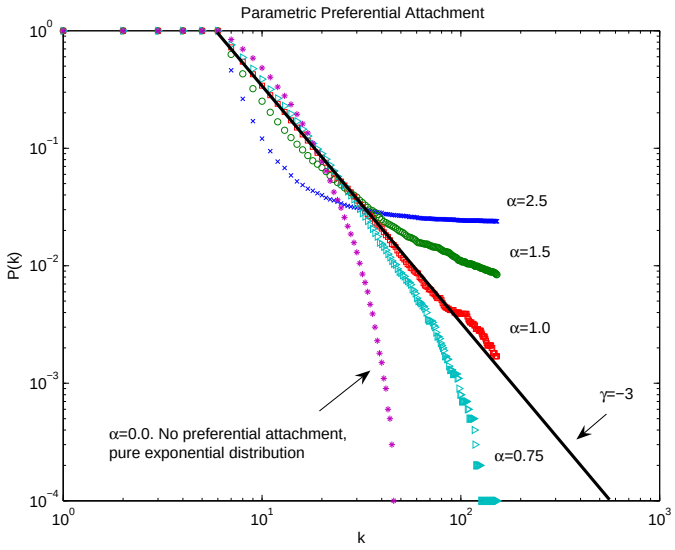

Figure 8.5: Cumulative degree distribution of non-linear preferential attachment

1) replacing the linear preferential attachment by the non-linear tuned by α, and 2) setting a maximum degree in order to match the capacity of our *LO* model.

We can observe that the non-linear preferential attachment yields different network structures that matches those obtained with the *LO* model. For $\alpha > 1.0$ we obtain a *star-like* structure. For $\alpha = 1.0$ we obtain a power-law. For $\alpha > 0.75$ we obtain power-law with an exponential influence. And, for $\alpha < 0.75$ the exponential part clearly dominates the power-law one. In the last case there is no preferential attachment. In this case, the probability of node i being chosen by a new node entering the network is $P(i) = \frac{1}{n}$, where n is the current number of nodes. Without preferential attachment the degree distribution is exponential.

We now turn to a closer inspection of the conditions that shape the topological features of the emergent networks in our simulations. Figures 8.2 and 8.3 show the simultaneous effects of two conditions, the cost to benefit ratio in exchanges in the support game $\frac{E}{B}$ and the variation between explicit and implicit knowledge exchange. Moreover, between these two figures there is the effect of memory size M_c and capacity C. The results demonstrate that particularly the ratio $\frac{E}{B}$ as one aspect of the harshness of the exchange game strongly affects the macro-structure that arises from agents' micro-behavior. At the same time, network structures seem to be more robust against variation in agents' learning capacity. We discuss effects of both conditions in turn, starting with effects of the cost to benefit ratio.

First, we analyze the cost to benefit ratio effect. Thus, we focus on the left sub-figure of figure 8.3. A star-like structure appears when actors face no social

dilemma or a very mild one, since the cost of giving support is very low compared to the benefit. Thus, we are in a situation of mild harshness of the game ($\frac{E}{B} = [0..\frac{3}{16}]$). When costs increase and, consequently, the game becomes harsher ($\frac{E}{B} = [\frac{4}{16}..\frac{7}{16}]$), the star-like structure is replaced by a power-law. This in turn is replaced by an exponential structure as soon as the cost to benefit ratio increases up to ($\frac{E}{B} \geq \frac{8}{16}$).

Why do we see these effects? Let us first consider the case of zero costs. In this situation, all agents get sooner or later acquainted with the most attractive system members and many of them actually exchange cooperatively with highly attractive agents. The reason is that due to zero costs these agents can benefit even from exchanges with the least attractive system members and thus do not reject them. As a consequence, the most attractive members of the population soon become known throughout the network as the most desirable interaction partners. They receive many proposals for interaction which are accepted until these agents exhaust their interaction-capacity. In the course of this learning process, highly attractive members tend to gradually optimize their network so that eventually the core arises in which highly attractive agents mainly relate to each other. This leaves little room for less attractive agents to relate to the core. Many less attractive members are frustrated in their search for further improvement by the frequent rejections they received. The graphs in the left column of figure 8.2 demonstrate a paradoxical consequence of these dynamics: despite the low costs of exchange, a large number of agents has a very sparse network in this condition with only ten or less exchange partners. This combination of large numbers of sparse ego networks with a highly connected core generates the star-like pattern.

When costs of the exchange are moderate relative to the benefit, a power-law structure emerges. In this situation, the most attractive agents are more likely to reject interaction proposals from less attractive members even before their interaction capacity is exhausted, or they may defect once an interaction has been established. As a consequence, knowledge about the most attractive agents spreads slower and to less recipients than in the zero-cost condition. Agents throughout the network hold more different perceptions of who may be a desirable interaction partner for them than in the zero-cost condition. The reason is that now there is no *absolute top* that is optimal for everybody. Because the social dilemma is harder, less attractive agents are more often exploited than in the zero cost condition. Accordingly, these agents become more conservative in their partner search. Exploration tends to diminish and the system ends sooner in a stable configuration than in the zero cost condition. This explains why the resulting degree distribution is *thicker* for intermediate degrees than in the star-like structure and *thinner* in the right tail. Now also agents with intermediate attractiveness are often chosen as interaction partners and there are less agents who become highly popular stars in the network.

Finally, when exchange costs are high, the conditions under which agents are

willing to cooperate are very restrictive. A focal agent is now only willing to cooperate when his interaction partner has very similar or higher attractiveness. As a consequence, the large majority of network members experiences that most of the interactions they try fail, or they get exploited in these interactions. Gradually most network members stick to those partners with whom they have positive experience and interact only with a small subset of all network members. Interaction networks tend to be less optimized than under milder social dilemma conditions. In particular, in many instances potential cooperation partner of the most attractive network members fail to find them before they stop exploration. Correspondingly, these most attractive members receive less choices than in milder circumstances. This explains why the resulting degree distribution is thinner in its right tail than the power law pattern that we found for moderate social dilemmas.

However, the topology of the simulated networks is not only affected by costs and benefits of the exchange but also by conditions that shape the efficiency of agents' learning process. Figures 8.2 and 8.3 show that under most conditions the network structures for implicit and explicit knowledge exchange are very similar, yet in some cases there are differences. For instance, in figure 8.3 we can observe that when $\frac{E}{B} = \frac{4}{16}$ explicit memory exchange ends up in a power-law distribution, while implicit memory exchange yields a star-like distribution.

To explore effects of memory exchange with more precision, we conducted a more fine-grained analysis of the interaction effect between cost to benefit ratio of and mode of memory exchange, upon network topology. We simulated both explicit and implicit memory exchange for cost-benefit ratios varying in the range $\frac{E}{B}$: $[0..\frac{14}{16}]$. Further parameters of the model are equal to those used for figure 8.3, that is: $N = 10000$, $M_c = 200$, $C = 150$ and $Q = 5$. To distinguish network structures quantitatively we use two different measures:

1. the determination coefficient of the cumulative potential regression r^2. This coefficient indicates how well the observed degree distribution can be fitted to the cumulative potential function $P(k) \sim k^{-\theta}$. A determination coefficient equal to or larger than $r^2 = 0.95$ implies a good fit of the regression (for $N = 1000$ we lower the acceptance threshold to 0.9, because for this small number of nodes statistics are very sensitive to random perturbations). Accordingly, the observed degree distribution can be classified as a potential distribution of exponent $-\theta$. Unfortunately, this measure is problematic when the statistics are poor. In this particular case the maximum degree is relatively small, which makes it difficult to discriminate between an exponential distribution and a power-law with a exponential cutoff. In order to overcome this problem rely on an additional measure, the *variance* of the degree distribution.

2. *variance* of degree $\frac{\langle k^2 \rangle}{\langle k \rangle^2}$. This ratio measures the variance between the real

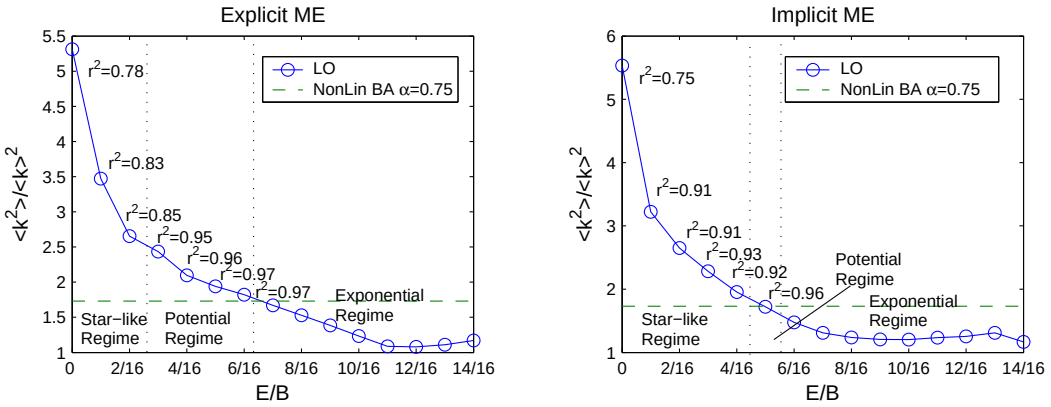

Figure 8.6: Classification in the star-like, power-law and exponential regimes of the networks generated by the LO model. Parameters: $N = 10000$, $M_c = 200$ and $C = 150$. The cost to benefit ratio $\frac{E}{B}$ of the support game varies from 0 to $\frac{14}{16}$. Left sub-figure: Explicit Memory Exchange. Right sub-figure: Implicit memory exchange. r^2 is the determination coefficient of the regression over the cumulative degree distribution. The dashed line is variance of degree for a network generated using the non-linear BA-model for α set to 0.75, with maximum degree set to 150, number of nodes set to 10000 and m set to 5, so then $\langle k \rangle = 10$. The dotted line identify the boundaries between regimes.

connectivity degree compared with the average degree. In the case of a random network this ratio would be close to 1.0, since nodes' degree do not deviate from the average. However, when *heavy-tails* are observed it means that there are nodes whose degree is much higher than the average degree. Therefore, this measure is very useful to discriminate between power-law and exponential structures. On the other hand this measure does not discriminate between power-law and star-like, since both distributions display *heavy-tails*. In that case, we rely on the regression's determination coefficient.

More precisely, to determine the type of distribution we use an algorithm that works as follows: If $\frac{\langle k^2_{LO} \rangle}{\langle k_{LO} \rangle^2} < \frac{\langle k^2_{Non-Lin-BA_{\alpha=0.75}} \rangle}{\langle k_{Non-Lin-BA_{\alpha=0.75}} \rangle^2}$ then the distribution is exponential. Otherwise, we check the determination coefficient of the potential regression r_{LO}. If $r^2_{LO} < 0.95$, then the distribution is star-like. Otherwise, it is power-law.

Figure 8.6 reveals profound interaction effects between the ratio of cost to benefits in the exchange game and the type of memory exchange on the qualitative structure of emergent networks, indicated by the different size of the regions of *star-like, potential and exponential degree distribution* in both sub-figures. Broadly, the

pattern that we find is that both higher cost to benefit ratios and the efficiency of memory exchange determine how difficult it is for agents to find suitable exchange partners. The higher this difficulty, the closer the network matches the exponential distribution. Conversely, the easier suitable exchange partners can be found, the more similar the patterns become to a star-like central-periphery structure. In between these two extremes, we find networks with power-law degree distributions. To illustrate these interaction effects we focus on two specific scenarios.

Consider the case where $\frac{E}{B} = \frac{6}{16}$. With explicit memory exchange (left part of figure 8.6) the LO-model generates a power law network, whereas we obtain an exponential structure under implicit memory exchange. Why this difference? In this condition, the costs of giving support are relatively high. Accordingly, agents experience often exploitation on their interactions, or rejection on their proposals to interact. As a consequence, agents' memory contains a lot of information about suboptimal or even non-profitable exchange partners. Implicit memory exchange exacerbates the difficulties of finding suitable partners in this condition, because agents get only randomly generated information from their interaction partners. Thus, with implicit memory exchange it is very likely that after an interaction, an agent receives a reference to at best a mediocre third party, so that the process of finding attractive partners is slow and inefficient in comparison with explicit memory exchange. With explicit exchange, the references agents receive are filtered, so that bad entries are not propagated. As a result, highly attractive agents are found by a larger number of advice seekers so that relatively more highly attractive network members obtain a high degree. This explains why in this particular condition explicit memory exchange results in a power law network while implicit memory exchange yields an exponential distribution.

Effects of memory exchange are strikingly different when we look at the other end of the spectrum where the social dilemma actors face is mild. Consider the case $\frac{E}{B} = \frac{3}{16}$. Figure 8.6 shows that in this condition explicit memory exchange entails a power-law structure, while under implicit memory exchange the structure is closer to a star-like pattern than to a power-law network as can be observed in figure 8.2 left-bottom sub-figure. Notice that this is just the opposite of the effect that we observed before. This time, explicit memory exchange generates the pattern that corresponds to the less favorable social dilemma structure as compared to the result of implicit memory exchange, while in the previous case this was reversed. The reason is that under explicit memory exchange references to third parties are filtered. As a consequence, all agents recommend more or less the same very limited subset of potential interaction partners, who in turn — are not capable to process all interaction proposals they receive. The frequent occurrence of rejections introduces a large amount of noise in the system. By noise we mean here that many agents get distorted perceptions of the real attractiveness of the other network members so

Table 8.2: Regimes boundaries for different settings in terms of cost to benefit ratio ($\frac{E}{B}$). Settings: $M_o = 150$ and A) $N = 10000, M = 500, M_c = 400,\ Q = 5$, B) $N = 10000, M = 500, M_c = 400,\ Q = 10$, C) $N = 10000, M = 200, M_c = 150,\ Q = 5$ D) $N = 5000, M = 200, M_c = 150,\ Q = 5$ E) $N = 1000, M = 200, M_c = 100,\ Q = 5$

Setting	*Mem. Exch.*	*Star-like Reg.*	*Power-law Reg.*	*Exponential Reg.*
A	Explicit	$[0..\frac{5}{16}]$	$[\frac{6}{16}]$	$[\frac{7}{16}..\frac{14}{16}]$
A	Implicit	$[0..\frac{6}{16}]$	$[\frac{7}{16}]$	$[\frac{8}{16}..\frac{14}{16}]$
B	Explicit	$[0..\frac{3}{16}]$	$[\frac{4}{16}..\frac{6}{16}]$	$[\frac{7}{16}..\frac{14}{16}]$
B	Implicit	$[0..\frac{4}{16}]$	$[\frac{5}{16}]$	$[\frac{6}{16}..\frac{14}{16}]$
C	Explicit	$[0..\frac{2}{16}]$	$[\frac{3}{16}..\frac{6}{16}]$	$[\frac{7}{16}..\frac{14}{16}]$
C	Implicit	$[0..\frac{4}{16}]$	$[\frac{5}{16}]$	$[\frac{6}{16}..\frac{14}{16}]$
D	Explicit	$[0..\frac{2}{16}]$	$[\frac{3}{16}..\frac{7}{16}]$	$[\frac{8}{16}..\frac{14}{16}]$
D	Implicit	$[0..\frac{4}{16}]$	$[\frac{5}{16}]$	$[\frac{6}{16}..\frac{14}{16}]$
E	Explicit	$[0..\frac{2}{16}]$	$[\frac{3}{16}..\frac{4}{16}]$	$[\frac{5}{16}..\frac{14}{16}]$
E	Implicit	$[0..\frac{3}{16}]$	$[\frac{4}{16}]$	$[\frac{6}{16}..\frac{5}{16}]$

that they can not find the best available interaction partner. By contrast, in implicit knowledge exchange the references are not filtered, but random. Here, not all agents try to get connected to the same subset of stars. The subset of potential partners that are perceived as highly attractive is wider, so that less interaction requests are rejected and more new interactions take place. In the long run, this fosters the efficiency of agents' learning process and highly attractive agents can be found and accessed by a more network members. This, in turn, drives the system more towards a star-like structure than under the explicit memory exchange mechanism, where agents sooner stop exploring due to frequent rejections.

So far, we have shown the effect of the cost to benefit ratio $\frac{E}{B}$ and the memory exchange process ME for a particular setting of the model, concretely, when number of agents N is 10000, the memory size M_c is 200, and the capacity C is 150. In order to test consistency of the model we conducted the same analysis for other settings, changing parameters N, M_c, C and Q. The results are summarized in the table 8.2.

Table 8.2 shows how despite having different population size N, memory size M_c, capacity C and initiation capacity Q all the analyzed settings are consistent with previous simulation runs. All settings display the three regimes, mainly distinguished in terms of cost to benefit ratio or *support game harshness*. Furthermore, the kind of memory exchange used affects the boundaries and the size of the power-law regime in a similar way as for setting C, which was the subject of the detailed analysis. Therefore, the detailed analysis carried out for setting C is also applicable

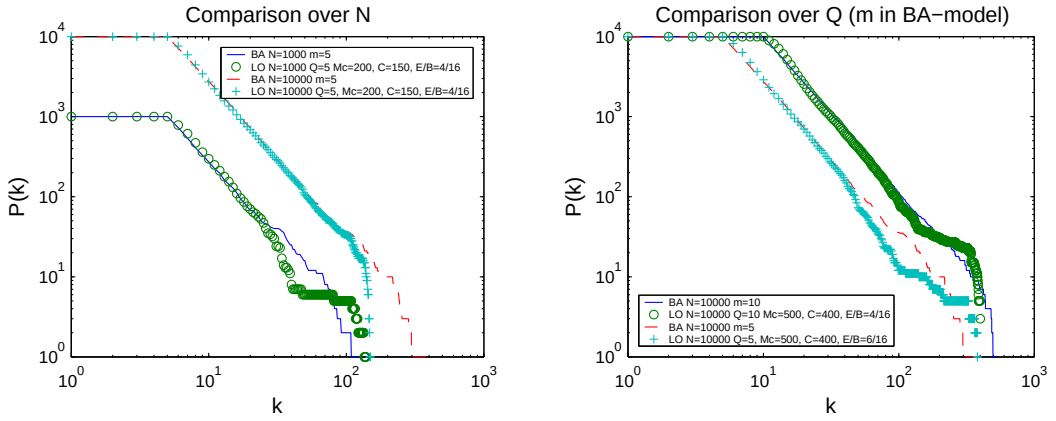

Figure 8.7: Comparison of distributions for the local optimization model. For the BA model we are using the original algorithm proposed by Barabási and Albert. In the left sub-figure we compare the graphs generated by both models for different networks size N. In the right sub-figure we compare the graphs for different initiation-capacity Q, which is the number of link that a node can deploy in the BA-model, that is to say, m. For the LO-model the memory exchange is always implicit, and the M_o is 150. See the legend for the rest of the parameters. Notice that the LO-model contains a capacity constraint C, which is set at $C = 150$ and $C = 400$ for left and right sub-figure, respectively.

to other settings.

Finally, we wish to test explicitly to what extent our LO-model can match results obtained from the original BA-model. For this purpose, we compared the results of both models with regard to simultaneous effects of the population size N and the initiation-capacity Q. We have shown the effect of the population size already in the previous results but we did not compare the obtained degree distributions for the two models. Figure 8.7 shows that the networks generated by the LO-model scale similarly to those obtained from the BA-model, both for the effect of N and the effect of Q on degree frequency. Moreover, the networks for both models seem to be very similar.

8.4 Discussion

We have argued that a sociologically plausible model of the micro-mechanism underlying complex network emergence needs to meet the following criteria. First, the locus of decision making about network changes should be at the level of individual actors. Second, individual decisions should be derived from the optimization of

individual goals based on bounded rational decision heuristics. Third, individual agents should use only local and imperfect knowledge of network characteristics to make decisions. Finally, the model should not trivialize conflicts between agents' interests with regard to the network changes they prefer. We proposed an agent-based computational model that satisfies these requirements, the local-interaction model (LO-model). The model describes network change as a consequence of a social exchange process in which agents differ in attractiveness, are initially unaware of their own and others' attractiveness, are free to change partners and acquire through interaction gradually local and partially distorted knowledge about others' attractiveness. In our model, agents seek to optimize their exchange relations based on backward-looking simple decision heuristics.

The model's assumptions rely on general theories of boundedly rational human decision making (e.g. [205, 213]) that have partially been validated in experimental research, but have not been specifically tested for individual decision making in making and breaking relations in social networks. While we believe that this is a significant step forward with regard to previous models of complex network emergence, we readily admit that more careful empirical studies of social network dynamics is needed to test underlying micro assumptions. One approach in particular that promises to be fruitful for this is the technique developed by Snijders [206] of *actor oriented statistics*, which is specifically designed to disentangled based on longitudinal network data the various individual motives that may drive network change.

We conducted simulations to explore how social conditions identified by the local interaction model may shape the structure of emergent networks. First and foremost, we found that despite the minimal and sociologically plausible assumptions we made about agents' knowledge and cognition, our model can replicate the celebrated *small world* and *power-law* network structures that have recently received much attention in the literature, due to their high robustness and efficiency for information propagation. We believe that for social network analysis, this is a significant improvement compared to previous models based on the mechanisms of *preferential attachment* and *random rewiring*. These models needed to employ implausible assumptions of globally available knowledge about structural positions or failed to explicate the individual goals and cognitions that motivate actors' decisions to make or break ties.

The second main result presented in this chapter is that the topological structure of the emergent social networks depends heavily upon the harshness of the exchange game, in particular the ratio of costs to benefits in a social exchange, and conditions that shape the learning capacities of agents, in particular the accuracy of information they receive about attractive exchange partners from their network relations. We show that it depends on the combination of these conditions whether star-like, small

world or power-law network structures emerge. Broadly, the pattern that we find is that both higher cost to benefit ratios and the efficiency of memory exchange determine how difficult it is for agents to find suitable exchange partners. The higher this difficulty, the closer the network matches the exponential distribution. Conversely, the easier suitable exchange partners can be found, the more similar the patterns become to a star-like central-periphery structure. In between these two extremes, we find networks with power-law degree distributions.

Our results not only show that complex networks can in principle arise from sociologically plausible behaviors of individual agents, but we have also obtained substantive insights that could not be derived from previous models of complex network emergence. The main reason we could do this is that our model contains parameters that relate to the costs, benefits and risks of social exchange actions, something that was not present in the mechanistic microassumptions based on random rewiring or preferential attachment. Accordingly, we believe that our insights may be fruitful for fields of research such as the study of social support or of the emergence of trust. In a society that faces a high level of harshness, e.g. in terms of the scarcity of material resources or harsh climatic living conditions, social support is both particularly needed and prone to social dilemma problems. Our model suggests that this should also show up in the network structures that arise in such a society. We find in our simulations that — paradoxically — the harsher are the conditions — that is: the more costly it is for actors to support others — the sparser and the less optimized are the emergent structures of support exchange. Such a hypothesis might be tested in a cross-national comparison where the relationship between economic wealth and the structures of social networks in different countries or regions is addressed. With regard to reputation and trust we find a similar paradoxical consequence of harshness. In harsh conditions actors are particularly vulnerable to exploitation by untrustworthy interaction partners. However, according to our analysis it is exactly here where emergent networks are least efficient for spreading the reputational information that actors need to protect themselves against exploitation. This, in turn, may imply the testable hypothesis that in poorer societies people need more time to develop trust into new potential interaction partners (e.g. immigrants) because it takes longer before they have information available that allows to assess these strangers' reputation.

Our work also adds a new note to research studying the effects of social network structures on cooperation in social dilemmas. Cohen et al. [52] have argued that it is mainly the stability of interaction structures in social networks and not so much their clustering that is needed to sustain cooperation. However, in their work they have treated interaction structures as an exogenous condition. We show instead how in the search for both attractive and cooperative partners agents may self-organize networks that are clustered and stable and sustain cooperative behavior

in social dilemmas. Emergent complex networks may be important for cooperation in a heterogeneous setting, because they match those pairs of agents who find each other sufficiently attractive to cooperate with each other in an exchange. This result also resonates with another recent study by Eguíluz et al. [77], in which the authors emphasize that cooperative role models may drive the self-organization of complex networks that sustain cooperation in social dilemmas.

Finally, we assume that agents are benevolent in the sense that they always cooperate with exchange partners as long mutual cooperation is preferable to mutual defection. The latter assumption in particular neglects the problem of opportunistic behavior that much of the social dilemma literature [16] deals with. However, we wish to point out that we did not entirely trivialize the problem of opportunism. Exploitation and unilateral defection of agents is possible and does occur in our model, but more sophisticated decision makers might exhibit these behaviors even more than our benevolent agents. Based on strategies already explored by Axelrod, a possible model extension here could be that agents vary in the extent to which they try to *test the waters* with occasional defections and then continue to defect unless the exchange partners retaliate. This obviously adds an extra dimension to the difficulties of finding an appropriate partner in our exchange game. We expect that the basic conclusion analysis will remain the same also for this complication. The more agents follow opportunistic strategies, the more difficult it is for them to find suitable partners and the more the emergent network will then exhibit an exponential degree distribution rather than a star-like structure.

Notwithstanding the need to explore in future work alternative, perhaps more realistic sets of assumptions, we believe this chapter offers an explicit model of sociologically plausible micro-processes that can generate a range of qualitatively different complex network structures.

Part V

Conclusions and Annexes

Chapter 9

Conclusions

Besides the particular summaries that are located at the end of each chapter we also want to sum up briefly the content of this thesis as a whole.

Throughout the research presented in this thesis we have contributed with evidences that support the *leitmotif* that best summarizes our work: *structure matters*. The structure of the network formed by the relationships of the individuals is indeed a very important factor that governs the dynamics of complex social systems, both at the micro and macro level. Moreover, structure also contains underlying knowledge about both levels. Therefore, by analyzing the structure, useful information about both the individuals and the systems can be acquired.

The aim of this thesis was not to hackle a well defined problem and to offer a solution, but rather it was an attempt to contribute to a body of research that turns to the structure of the relationships in order to gain better understanding of a system. Thus, our research is not a methodology to design or engineer artificial societies but a compilation of evidences that point out towards the need to consider the structure as a key factor of the system's dynamics. Although we also provide tangible methods, i.e. the algorithms to extract knowledge from the structure of the social network, most of our research is descriptive.

However, one must take into account that understanding is a necessary condition prior designing or engineering. Therefore, we believe that our findings and results might prove useful for multi-agent system research on its path towards the design and building of *full-fledged* artificial societies. To that purpose not only the social dimension of agents must be taken into account, but it is necessary to carefully consider the underlying social structure of the system as well.

Although the domain we studied was focused on multi-agent systems, we believe that the conclusions drawn by this thesis are directly applicable to other artificial societies such as electronic communities and peer-to-peer systems. Furthermore, the results are also relevant to the area of research dealing with complex social systems.

Let us briefly review the contributions of this thesis in relation to its goals.

- In chapter 4 and 5 we presented two algorithms that used network topology as a source of knowledge. The first algorithm is able to create a ranking of agents by relevance that is a good approximation to reputation. Thus, reputation of an agent can be inferred just from its position in the social network. The algorithm does not require to have complete information about the social network, i.e its adjacency matrix. The process is distributed through all agents, so that only local information is required. Furthermore, the algorithm adapts dynamically to different network topologies. The second algorithm, presented in chapter 5, is intended to extract the community structure from a network. By doing so, the groups of agents that form communities and the relations between these communities are revealed. This algorithm aims at improving time efficiency so that the community structure of very large networks (hundreds of thousands nodes) can be retrieved in reasonable time. The high efficiency of the algorithm does not come at the expenses of its accuracy since the maximum modularity yielded by the algorithm is comparable to the best algorithms in the current literature. So, this algorithm is indeed a sound choice for medium and large [social] networks. In brief the algorithms presented in the first part of the thesis are able to extract social measures such as reputation, group membership and relationships between groups which are valuable informations to navigate and interact in new societies. Therefore, the analysis of the structure can be called upon as a means to reduce the inherent complexity and uncertainty of open systems.

- In chapter 6 and 7 we saw how certain structures can favor the emergence and stabilization of coordination at system-level. The same agents with the same strategies will have different outcomes depending on the way agents are interconnected. Thus, macro-behaviours such as agreeing on a convention do not depend only on the strategies, or micro-rules, followed by the agents but on the underlying pattern of interactions. We saw in chapter 6 how certain structures — complex networks — promote the emergence of a convention in a reasonable time. Properties found in complex networks, such as their short diameter and short average path length, played a key role in the time required for the system to converge to a convention. Complex networks are much more efficient (in time) than other structures such as regular networks. In fact, complex networks are as efficient as fully connected networks, which are clearly unrealistic structures for a society design, since they imply global information and connections without cost. Furthermore, in chapter 7, we showed that if the structure displays a high clustering coefficient, which is a characteristic of social networks, not only the time of establishment of a convention, but also which

convention is eventually established is also affected by structure. Clustered networks lead the agents to adopt socially efficient conventions which, in turn, are stable against invasion of sub-optimal conventions. And they are so under a much wider range of conditions than in a non-clustered network. Most interestingly, the structural properties that facilitate the coordination regime, both in terms of time efficiency and in terms of choosing the socially efficient convention, can be found in many empirical complex systems.

- In chapter 8, we presented a model for complex network formation, notably power-law, small-world and center-periphery networks. Unlike most models of complex networks emergence, our model does not require implausible assumptions like global knowledge about structural position. The model presented also makes explicit the motivations that the agents have in order to start or terminate relationships with other agents. Thus, we showed that self-interested agents, performing a local optimization process grounded in social theory, are able to self-organize and arrange themselves into the same complex structures that are found in many real complex systems across very different domains. These structures — complex networks — are not the result of a centralized design but rather the consequence of the interactions between autonomous agents. We provided a model that requires neither global and perfect information nor it assumes implausible behavioural assumption to explain complex networks formation.

The research presented in this thesis is supported by several peer-reviewed publications. We include below a selection of these publications. Table 9.1 summarizes the correspondence between publications and thesis' chapters. All publications can be found at the authors personal homepage.[1]

1. **Pujol, J.M.**, Sangüesa, R. and Delgado, J. *Extracting Reputation in Multi Agent Systems by means of Social Network Topology.* Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems AAMAS-02. Vol. 1, pp 467-474. Bologna, Italy (2002).

2. **Pujol, J.M.**, Sangüesa, R. and Delgado, J. *A Ranking Algorithm Based on Graph Topology to Generate Reputation or Relevance.* In Web Intelligence (Ning Zhong, Jiming Liu, and Yiyu Yao eds.), ch 18, pp 382-395, Springer Verlag (2003), ISBN: 3-540-44384-3.

3. Delgado, J., **Pujol, J.M.** and Sangüesa, R. *Emergence of Coordination in Scale-Free Networks.* In Web Intelligence and Agent Systems Journal 1 (2003), 131-138.

[1] http://www.lsi.upc.edu/∼jmpujol/publications.html

Thesis' Chapter	*Related papers*
Chapter 4: Extracting Reputation	1 and 2
Chapter 5: Finding Community Structure	7
Chapter 6: Emergence of Social Conventions	3
Chapter 7: Emergence of Efficient Conventions	4
Chapter 8: Complex Networks through Local Optimization	5 and 6

Table 9.1: Correspondence between publications and thesis' chapters

4. **Pujol, J.M.**, Delgado, J., Sangüesa, R., and Flache, A. *The Role of Clustering on the Emergence of Efficient Social Conventions*. Proceedings of the 19th International Joint Conference on Artificial Intelligence, IJCAI-05, (Edinburgh, Scotland), July 2005, pp 965-970, ISBN: 0-938075-93-4.

5. **Pujol, J.M.**, Flache, A., Delgado, J. and Sangüesa, R. *The Emergence of Complex Networks through Local Optimization*. In Proceedings of the XVI European Conference on Artificial Intelligence ECAI-04 (López de Mántaras,R. and Saitta,L. eds.,Valencia, Spain), pp. 48 - 52, IOS Press 2004.

6. **Pujol, J.M.**, Flache, A., Delgado, J. and Sangüesa, R. *How can Social Networks ever become Complex? Modelling the Emergence of Complex Networks from Local Social Exchanges*. Journal of Artificial Societies and Social Simulation 8(4) (2005).

7. **Pujol, J.M.**, Béjar, J. and Delgado, J. *Finding Community Structure through Random Walkers*. Submitted to to Physical Review E (2006).

The final conclusion not only of this thesis but also of the Ph.D. is best summarized by the quote we started this document with: *To travel hopefully is a better thing than to arrive, and the true success is to labour.* It would be difficult to describe the process so accurately using less words. The true meaning of my research was not as much the Ph.D. as it was the joy of playing. I have no words of my own to thank the wonderful partners I had in my travel, poorly summarized in this document. Thanks to all, and especially to Ramon, Jordi and Andreas.

Appendix A

Pseudo-Code of the Models

As Macy and Flache pointed out in [87], one of the most important practices in agent based modeling is to facilitate independent replication of the models (see section 2.4.1 for a complete list of best practices). To that end, we provide a brief but comprehensive sketch of the models we have presented in this thesis (chapters 6, 7 and 8).

For the sake of clarity we focus only on the core of our model; the agent behaviour mainly. We do not cover manipulation of data structures, graphical display or collecting results. Do not hesitate to contact the author to get the sources.

A.1 Emergence of Conventions

Declarations

```
define A 1  // action A
define B 2  // action B
define GSM 10  // update rule
define HCR 11  // update rule

int N; // number of agents
int Ms; // memory size
int Q; // number of neighbours average connectivity is Q*2
int update; // defines the update rule to be used (GSM or HCR)
int layout; // the underlying network topology (see section 6.2)
double Beta; // part of the GSM update rule (eq 6.3.1)
double[2][2] G; // payoff matrix (figure 6.3)

Graph network; // the network that contains the underlying topology
```

```
Agent {
  double[2] outcome; // acumulated outcome of action A and B
  int pointer; // next memory slot to be used
  int action; // the current action (A or B) played by the agent
}

Agent mpAgents = Agent[N];  // vector of agents

double[N][Ms][3] mpMemory; // agents memory,
                      // mpMemory[agent][memory_position][x]
                      // where x is
                      //         0: the action (casted to int)
                      //         1: the outcome
                      //         2: the time (casted to int)
                      // of the interaction agent 'agent' had stored
                      // in its memory at the position
                      // 'memory_position'
```

Initialization

```
// initialization of the model
procedure init() {

  Agents mpAgents = new Agent[N];
  for each agent in agents {
    mpAgents[agent].pointer=0;
    mpAgents[agent].outcome[A]=0;
    mpAgents[agent].outcome[B]=0;
    if (random()<0.5) mpAgents[agent].action=B;
    else mpAgents[agent].action=A;
  }

   // setting the agents' interactions
   network = init_network(layout);

}

// initialize the graph that contains the underlying pattern of interactions
// complete, regular, small-world or scale-free networks. See section  6.2
```

```
function init_network(layout);
```

Agent's Behaviour

```
procedure main() {
  init();
  time=1;
  stop=false;
  do {
    agent=random(1,N); // choose one agent at random;
    run(agent,time);  // agent's behaviour
    time=time+1;
    ratioA=ratioA();
    // evaluates the stop criterion, Tc is set to 90%
    if ((rationA>0.9) || (ratioA<0.1)) stop=true;
    else stop=false;
  } while(!stop);

} //end main

// agent's behaviour
procedure run(agent, time) {

  // a neighbour of agent is choosen at random
  partner = chooseNeighbour(agent);

  if (update==GSM) {
     // if the update rule is GSM
     updateActionGSM(agent);
  }
  else {
    // if the update rule is HCR

    // agent and partner play the coordination game G defined in figure 6.3
    [outcome1, outcome2, action1, action2] = playGame(agent,partner);

    updateMemory(agent,outcome1,action1);
    updateMemory(partner,outcome2,action2);

    // action update rule Highest Cumulative Reward (see section 6.3.2)
```

```
    updateActionHCR(agent);
    updateActionHCR(partner);
  }

}

// functions required for the HCR case

// agents play the game defined in figure 6.3
function [outcome1, outcome2, action1, action2] = playGame(agent1, agent2) {
   action1 = mpAgents[agent1].action;
   action2 = mpAgents[agent2].action;

   outcome1=G[action1][action2];
   outcome2=G[action2][action1];
}

// updates the agent's memory with the last interaction
function updateMemory(agent, outcome, action) {
  if (mMemory[agent][mpAgents[agent].pointer][2]>=0) {
    // the memory slot has been used, so the accumulated payoff needs
    // to be updated (we need to substract the outcome of the interaction
    // to be forgotten needs)
    mpAgents[agent].outcome[mMemory[agent][mpAgents[agent].pointer][0]]-=
      mMemory[agent][mpAgents[agent].pointer][1];
  }

  // now the memory slot mpAgents[agent].pointer is free, and can be
  // updated with the current interaction
  mMemory[agent][mpAgents[agent].pointer][0]=action;
  mMemory[agent][mpAgents[agent].pointer][1]=outcome;
  mMemory[agent][mpAgents[agent].pointer][2]=time;
  // the accumulated outcome needs to be updated
  mpAgents[agent].pointer[mMemory[agent][mpAgents[agent].pointer][0]]+=
    mMemory[agent][mpAgents[agent].pointer[1];
  // the pointer has to be increased
  mpAgents[agent].pointer=(mpAgents[agent].pointer+1)%Ms;

}
```

```
// action update rule Highest Cumulative Reward (see section 6.3.2)
function updateActionHCR(agent) {
  // if the acumulated outcome of action A is bigger than the accumulated
  //outcomes of B then next action will be A, and vice versa
  if (mpAgents[agent].outcomes[A]>mpAgents[agent].outcomes[B])
    mpAgents[agent].action=A;
  else if (mpAgents[agent].outcomes[A]<mpAgents[agent].outcomes[B])
    mpAgents[agent].action=B;
}

// functions required for the GSM case

// action update rule Generalized Simple Majority (see section 6.3.1)
function updateActionGSM(agent) {
  //
  List l = getNeighbours(agent);
  neigh_in_A = 0;
  neigh = 0;
  for each element in l do {
    partner = element;
    if (mpAgents[partner].action==A) neigh_in_A++;
    neigh++;
 }

  // now evaluate the probability of replacing the agent's action
  // by its opposite, following eq. 6.3.1

  if (mpAgents[agent].action==A) neigh_in_not_S = neigh-neigh_in_A;
  else neigh_in_not_S = neigh_in_A;

  pchange = 1 / (1 + exp(2*Beta((2*neigh_in_not_S)/(neigh-1))));

  // now agent will update its state with probability pchange
  if (random()<pchange) mpAgents[agent].action=(mpAgents[agent].action+1)%2;

}

// common functions
```

```
// returns an int in the interval [i..f]
function random(int i, int f);
// returns the ratio of agents currently playing A
function ratioA();
// returns a neighbour of agent 'agent' picked at random
function chooseNeighbour(agent);
// returns a list of all neighbours of agent 'agent'
function getNeighbours(agent);
```

A.2 Emergence of Efficient Social Conventions

Declarations

```
define A 1 // action A
define B 2 // action B

int N; // number of agents
int Ms; // memory size
int rB;  // initial number of agents playing action B
int Q;  // number of neighbours average connectivity is Q*2
int layout; // the underlying network topology (see section 7.2.3)
double[2][2] G; // payoff matrix (figure 7.1)
double pImitation; // imitation propensity of action B

Graph network; // the network that contains the underlying topology

Agent {
  double[2] outcome;  // acumulated outcome of action A and B
  int pointer;  // next memory slot to be used
  int action;  // the current action (A or B) played by the agent
}

Agent mpAgents = Agent[N];  // vector of agents

double[N][Ms][3] mpMemory; // agents memory,
                      // mpMemory[agent][memory_position][x]
                      // where x is
                      //         0: the action (casted to int)
                      //         1: the outcome
                      //         2: the time (casted to int)
                      // of the interaction agent 'agent' had stored
                      // in its memory at the position
                      // 'memory_position'
```

Initialization

```
// initialization of the model
procedure init() {

  Agents mpAgents = new Agent[N];
  for each agent in agents {
    mpAgents[agent].pointer=0;
    mpAgents[agent].outcome[A]=0;
    mpAgents[agent].outcome[B]=0;
    if (random()<rB) mpAgents[agent].action=B;
    else mpAgents[agent].action=A;
  }

   // setting the agents' interactions
   network = init_network(layout);

}

// initialize the graph that contains the underlying pattern of interactions
// random, regular, small-world or scale-free graphs. See section
// 7.2.3
function init_network(layout);
```

Agent's Behaviour

```
procedure main() {
  init();
  time=1;
  stop=false;
  do {
    agent=random(1,N); // choose one agent at random;
    run(agent,time);  // agent's behaviour
    time=time+1;
    ratioA=ratioA();
    // evaluates the stop criterion, Tc is set to 99%
    if ((rationA>0.99) || (ratioA<0.01)) stop=true;
    else stop=false;
  } while(!stop);
```

```
} //end main

// agent's behaviour
procedure run(agent, time) {

  // a neighbour of agent is choosen at random
  partner = chooseNeighbour(agent);

  // agent and partner play the coordination game G defined in figure 7.1
  [outcome1, outcome2, action1, action2] = playGame(agent,partner);

  updateMemory(agent,outcome1,action1);
  updateMemory(partner,outcome2,action2);

  // action update rule Highest Cumulative Reward (see section 7.2.2)
  updateActionHCR(agent);
  updateActionHCR(partner);

  if (random()<pImitation)  {
    // with a probability pImitation it enters here
    if (mpAgents[agent].action=B || mpAgents[partner].action=B) {
       // at least one part played the 'contagious' action (B), so both
       // agents will play action B
       mpAgents[agent].action=B;
       mpAgents[partner].action=B;
    }
  }
}

// agents play the game defined in figure 7.1
function [outcome1, outcome2, action1, action2] = playGame(agent1, agent2) {
   action1 = mpAgents[agent1].action;
   action2 = mpAgents[agent2].action;

   outcome1=G[action1][action2];
   outcome2=G[action2][action1];
}

// updates the agent's memory with the last interaction
function updateMemory(agent, outcome, action) {
```

```
  if (mMemory[agent][mpAgents[agent].pointer][2]>=0) {
    // the memory slot has been used, so the accumulated payoff needs
    // to be updated (we need to substract the outcome of the interaction
    // to be forgotten needs)
    mpAgents[agent].outcome[mMemory[agent][mpAgents[agent].pointer][0]]-=
      mMemory[agent][mpAgents[agent].pointer][1];
  }

  // now the memory slot mpAgents[agent].pointer is free, and can be
  // updated with the current interaction
  mMemory[agent][mpAgents[agent].pointer][0]=action;
  mMemory[agent][mpAgents[agent].pointer][1]=outcome;
  mMemory[agent][mpAgents[agent].pointer][2]=time;
  // the accumulated outcome needs to be updated
  mpAgents[agent].pointer[mMemory[agent][mpAgents[agent].pointer][0]]+=
    mMemory[agent][mpAgents[agent].pointer[1];
  // the pointer has to be increased
  mpAgents[agent].pointer=(mpAgents[agent].pointer+1)%Ms;

}


// action update rule Highest Cumulative Reward (see section 7.2.2)
function updateActionHCR(agent) {
  // if the acumulated outcome of action A is bigger than the accumulated
  //outcomes of B then next action will be A, and vice versa
  if (mpAgents[agent].outcomes[A]>mpAgents[agent].outcomes[B])
    mpAgents[agent].action=A;
  else if (mpAgents[agent].outcomes[A]<mpAgents[agent].outcomes[B])
    mpAgents[agent].action=B;
}


// returns an int in the interval [i..f]
function random(int i, int f);
// returns the ratio of agents currently playing A
function ratioA();
// returns a neighbour of agent 'agent' picked at random
function chooseNeighbour(agent);
```

A.3 Complex Networks through Local Optimization

Declarations

```
int N;                  // number of agents
int Mc;                 // memory size
int Mo;                 // initial number of agents in memory
int Q;                  // number of interactions that can be initiated
int C;                  // interaction capacity
double B;               // Benefit
double E;               // Effort
ME={Explicit,Implicit} // Memory exchange

Agent {
int id;                // agent's id
double alpha;          // expertise of the agent
MemoryEntry memory;
}

MemoryEntry {
  int agent;   // agent in memory
  int expout;     // expected outcome of the agent
  int time;   // time of the last interaction.
                  //    if time<-1 agent belongs to the unknown set
                  //    if time==0 agent belongs to the known set
                  //    if time>0  an interaction occurred at that time
}
```

Initialization

```
// initialization of the model
procedure init() {

  Agents agents = new Agent[N];
  for each agent in agents {
    // init expertise of the agent in the [0.05..0.95] range
    // with a 10^-3 precision
    agent.alpha = (double)((random()*0.90+0.05)/1000);
    agent.memory = new MemoryEntry[Mc];
    for each me in agent.memory up to Mo {
      me.time=-1;
```

```
      me.expout=0.0;
      me.agent = random(N);
      // check that the me.agent is different than the agent itself
      // and it's not already in the agent's memory.
    }
  }
}
```

Agent's Behaviour

```
procedure main() {
  init();
  time=1;

  do {
    createPermutationOfAgents();
    do {
      agent=getAgentFromPermutation();
      run(agent,time);
      removeAgentFromPermutation(agent);
    } while(!isPermutationEmpty());
    time=time+1;
  } while(time<100);

} //end main

// agent's behaviour
procedure run(agent, time) {
  // chose a set of agents to interact with (<=Q)
  agentSet = chooseAgentsToInteract(agent);
  for each agentToInteract in agentSet {
   // is agentToInteract accepting the agent's proposal
   //for interaction?
   if (acceptInteractionProposal(agentToInteract,agent)) {
     [ag1_outcome,ag2_outcome,ag1_netbenefit,ag2_netbenefit]
                      = playGame(agent,agentToInteract);
      updateMemory(agent,agentToInteract,
                   ag1_outcome,ag1_netbenefit,time);
      updateMemory(agentToInteract,agent,
                   ag2_outcome,ag2_netbenefit,time);
```

```
      // if interaction is positive for both agents,
      //then exchange memories
      if (ag1_outcome>0 and ag2_outcome>0)
        exchangeMemories(agent,agentToInteract,time,outcomes);
   }
   else {
     interactionIsRefused(agent,agentToInteract);
   }
  }
}

// agent chooses a set of agents to interact with (up to Q).
// The set of agent is chosen at random or maximizing the
// expected outcome depending on the exploration
// probability
function chooseAgentsToInteract(agent) {
  agentSet = new List();

  // we have to calculate the exploration probability. So we
  // must know how many agents in agent's memory are known or
  // unknown (K and U set in section 8.1)
  unknown=0;
  known=0;
  for each mem in agent.memory {
    if (mem.agent>=0) {
      if (mem.time<0) unknown=unknown+1;
      else known=known+1;
    }
  }

  // equation 8.2.3
  explorationProb = (unknown / (unknown+known))^2.0;

  i=0;
  do {
    if (explorationProb < random()) {
      // agent is not exploring,
      mem = maximum(agent.memory);
      // mem is the memory entry with maximum expected outcome
      // (mem.expout) provided that mem.agent is not already
```

```
      // in agentSet
      agentSet.add(mem.agent);
    }
    else {
      // agent is exploring
      // choose one agent at random from agent's memory.
      // Check that mem.agent is not already in agentSet.
      mem = random(agent.memory);
      agentSet.add(mem.agent);
    }
    i=i+1;
  } while(i<Q);

  return agentSet;
}

// the agent must decide whether to accept the proposal
// made by initiatorAgent
function acceptInteractionProposal(agent,initiatorAgent) {
  mem=getMemoryEntry(agent,initiatorAgent,time);
  if (isNull(mem)) {
    if (getCurrentInteractions(agent,time)<C) return true;
    else return false;
  }
  else {
    if (mem.expout>=0.0
         and getCurrentInteractions(agent,time)<C) return true;
    else return false;
  }
}

// the agent updates its memory after the interaction
// equation 8.2.5
procedure updateMemory(agent, partner, outcome, netBenefit, time) {
  mem=getMemoryEntry(agent,partner);
  if (isNull(mem)) {
// the agent did not have partner in its memory
memToR=getLessAttractiveMemEntry(agent,time);
memToR.agent = partner;
memToR.time = time;
```

```
if (netBenefit>=0) memToR.expout = outcome;
else memToR.expout = -outcome;
  }
  else {
    // the agent did have partner in its memory
    if (netBenefit>=0) {
      // agent cooperated in the interaction with partner
      if (mem.time>0) mem.expout=(mem.expout+outcome)/2.0
      else mem.expout = outcome;
    }
    else {
      // agent defected in the interaction with partner
      mem.expout = -outcome;
    }
    // update the interaction time, if t>0 means that the agent
    // and partner have interacted, otherwise the knowledge
    // about partner comes from the memory exchange process
    mem.time=time;
  }
}

// returns the memory entry from agent's memory whose
// mem.expout (expected outcome) is minimum in
// absolute value. Provided that mem.time < time
function getLessAttractiveMemEntry(agent,time);

// remove the memory entry memToRemove from agent's memory
removeMemoryEntry(agent,memToRemove);

// add the memory entry memToAdd to agent's memory
addMemoryEntry(agent,partner,time,expoutPartner) {
  mem = getMemoryEntry(agent,partner);
  if (!isNull(mem)) {
    // partner was already in agent's memory
    if (mem.time<=0) {
      if (mem.time<0) mem.expout=expoutPartner;
      else if (mem.time==0)
            mem.expout=(mem.expout+expoutPartner)/2.0;
      mem.time=0;
    }
```

```
  }
  else {
    // partner was not in agent's memory
    memToR = getLessAttractiveMemEntry(agent,time);
    if (fabs(memToReplace.expout)<fabs(expoutPartner)) {
      // replace the old entry by the new one
      memToR.agent = partner;
      memToR.expout = expoutPartner;
      memToR.time = 0;
    }
  }
}

// agents exchange information about other agents
procedure exchangeMemories(agent, partner, time) {
  if (ME==Explicit) {
    memFromAgent = chooseMEExplicit(agent);
    memFromPartner = chooseMEExplicit(partner);
  }
  else {
    memFromAgent = chooseMEImplicit(agent,time);
    memFromPartner = chooseMEImplicit(partner,time);
  }

  addMemoryEntry(agent,memFromPartner.agent,time,
                 memFromPartner.expout);
  addMemoryEntry(partner,memFromAgent.agent,time,
                 memFromAgent.expout);
}

// the agent reduces the expected outcome after
// partner's rejection to interact
procedure interactionIsRefused(agent,partner) {
  mem = getMemoryEntry(agent,partner);
  mem.expout = (mem.expout+0.0)/2.0;
}

// return a memory entry whose mem.time is bigger than 0,
// which that at least one interaction between agent and mem.agent
// has occurred. The chosen memory entry is the maximum
```

```
// mem.expout in absolute value. In the case of  a clash the
// returned memory entry will be chosen at random among those with
// maximum expected outcome.
function chooseMEExplicit(agent);

// return a memory entry whose mem.time is equal to the current time,
// which means a current interaction of the agent. The memory entry
// is chosen at random
function chooseMEImplicit(agent, time);

// returns the current number of interactions of the agent, which is
// the number of memory entries in its memory with mem.time equal to
// the current time
function getCurrentInteractions(agent,time);

// play the game G
function playGame(agent1, agent2) {
  double Gij, Gji, Lij, Lji;
  double outcomeAg1, outcomeAg2;
  double netBenefitAg1, netBenefitAg2;

  Gij = (1.0-agent1.alpha)*(agent2.alpha)*B;
  Lij = (1.0-agent2.alpha)*(agent1.alpha)*E;

  Gji = (1.0-agent2.alpha)*(agent1.alpha)*B;
  Lji = (1.0-agent1.alpha)*(agent2.alpha)*mEffort;

  netBenefitAg1 = Gij-Lij;
  netBenefitAg2 = Gji-Lji;

  if (netBenefitAg1>=0.0 and netBenefitAg2>=0.0) {
    // both agents cooperate
    outcomeAg1=Gij-Lij;
    outcomeAg2=Gji-Lji;
  }
  else if (netBenefitAg1>=0.0 and netBenefitAg2<0.0) {
    // agent2 defects and agent1 cooperates
    outcomeAg1=-Lij;
    outcomeAg2=Gji;
  }
```

```
  else if (netBenefitAg1<0.0 and netBenefitAg2>=0.0) {
    // agent 1 defects and agent2 cooperates
    outcomeAg1=Gij;
    outcomeAg2=-Lji;
  }
  else {
    // both agents defect
    outcomeAg1=0.0;
    outcomeAg2=0.0;

  }

  return [outcomeAg1,outcomeAg2,netBenefitAg1,netBenefitAg2];
}

// create a random permutation; so agents are executed only once per
// simulation step and they the order is random
function createPermutationOfAgents();
// remove agent from the permutation
procedure removeAgentFromPermutation(agent)
// return when all the agents have been already chosen
function isPermutationEmpty();
// returns the memory entry form agent's memory
// corresponding to agent2 (agentTo=agent.mem.agent).
// If it does not exist return null
function getMemoryEntry(agent, agentTo);
```

Bibliography

[1] A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings Hawaii International Conference on System Sciences 33*, 2000.

[2] G. Abramson and M. Kuperman. Social games in a social network. *Physical Review E*, 63:090301, 2001.

[3] M. S. Ackerman and D. W. McDonald. Answer Garden 2: Merging Organizational Memory with collaborative help. In *Computer Supported Cooperative Work*, pages 97–105, 1996.

[4] L. Adamic. The small world web. In S. Abiteboul and A.-M. Vercoustre, editors, *Proc. 3rd European Conf. Research and Advanced Technology for Digital Libraries, ECDL*, number 1696. Springer-Verlag, 1999.

[5] L. Adamic and E. Adar. Friends and neighbors on the web. *Social Networks*, 25(3):211–230, 2003.

[6] L. Adamic and E. Adar. How to search a social network. *Social Networks*, 77(3):187–203, 2005.

[7] L. Adamic, R.M. Lukose, and B.A. Huberman. Local search in unstructured networks. *arxiv:cond-mat/0204181*, 2002.

[8] M. Ahuja and K. Carley. Network structure in virtual organizations. *Organization Science*, 10(6):741–757, 1999.

[9] W. Aiello, F. Chung, and L. Lu. A random graph model fro massive graphs. In *Proceedings of the 32nd ACM Symposium on the Theory of Computing*, pages 171–180. ACM, 2000.

[10] R. Albert and A.-L. Barabási. Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85:5234–5237, 2000.

[11] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Review of Modern Physics*, 74:47–97, 2002.

[12] R. Albert, H. Jeong, and A.-L. Barabási. The diameter of the world wide web. *Nature*, (401):130–131, 1999.

[13] R. Albert, H. Jeong, and A.-L. Barabási. The internet's achilles' heel: Error and attack tolerance in complex networks. *Nature*, 406:378–381, 2000.

[14] F. Amblard and G. Deffuant. The role of network topology on extremism propagation with the relative agreement opinion dynamics. *Physica A*, 343:725–738, 2004.

[15] A. Arenas, L. Danon, A. Díaz-Guilera, P. Gleiser, and R. Guimerà. Community analysis in social networks. *European Physics Journal B*, 38(2):373–380, 2004.

[16] R. Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.

[17] R. Axelrod. An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111, 1986.

[18] R. Axelrod. Advancing the art of simulation in the social sciences. *Complexity*, 3:193–199, 1997.

[19] R. Axelrod. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration.* Princeton University Press, Princeton, 1997.

[20] R. Axtell. The emergence of firms in a population of agents: Local increasing returns, unstable nash equilibria, and power law size distributions. Technical Report CSED N. 3, The Brookings Institution, 1999.

[21] A.-L. Barabási. *Linked: The New Science of Networks.* Perseus, New York, 2002.

[22] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[23] A.-L. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, (272):173–187, 1999.

[24] A.-L. Barabási, H. Jeong, E. Ravasz, Z. Néda, A. Schubert, and T. Vicsek. Evolution of the social network of scientific collaborations. *arxiv:cond-mat/0104162*, 2001.

[25] J.A. Barnes. Class and committee in a norwegian island parish. *Human Relations*, (7):39–58, 1954.

[26] J. Becker, B. Niehaves, and K. Klose. A framework for epistemological perspectives on simulation. *Journal of Artificial Societies and Social Simulation*, 8(4), 2005.

[27] J. Bendor and P. Swistak. The evolution of norms. *American Journal of Sociology*, 106(6):1493–1545, 2001.

[28] G. Bianconi and A.-L. Barabási. Competition and multiscaling in evolving networks. *Europhys. Lett.*, (54):436, 2001.

[29] P. Blau. *Exchange and Power in Social Life*. Wiley, New York, 1964.

[30] N. Boccara. *Modelling Complex Systems*. Springer-Verlag, 2004.

[31] M. Boguña, R. Pastor-Satorras, A. Díaz-Guilera, and A. Arenas. Models of social networks based on social distance attachment. *Physical Review E*, 70:056122, 2004.

[32] P. Bonacich. *Social Dilemmas. Theoretical Issues and Research Findings*, chapter Communication Networks and Collective Action, pages 225–245. Pergamon Press, Oxford, 1992.

[33] A.H. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence*, chapter An Analysis of Problems and Research in DAI, pages 3–35. Morgan Kaufmann, 1998.

[34] C. Borgs, J.T. Chayes, M. Mahdian, and A. Saberi. Exploring the community structure of newsgroups. In *Proceedings of KDD*, pages 783–787, 2004.

[35] R.A. Bourne, C.B. Excelente-Toledo, and N.R. Jennings. Run-time selection of coordination mechanisms in multi-agent systems. In *Proceedings of 14th European Conf. on Artificial Intelligence, ECAI 2000 (Werner, H. ed.)*, pages 348–352. IOS Press, 2000.

[36] P.S. Bradley and U.M. Fayyad. Refining initial points for k-means clustering. In *Proceedings of the 15th International Conference on Machine Learning*, pages 91–99, 1998.

[37] U. Brandes, M. Gaertler, and D. Wagner. Experiments on graph clustering algorithms. *Lecture Notes in Computer Science*, 2832:568–579, 2003.

[38] J.S. Breese, D. Heckermann, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the 14th conference on Uncertainty in Artificial Intelligence*, pages 43–52, 1998.

[39] D. Bu, Y. Zhao, L. Cai, H. Xue, X. Zhu, H. Lu, J. Zhang, S. Sun, L. Ling, N. Zhang, G. Li, and R. Chen. Topological structure analysis of the protein-protein interaction network in budding yeast. *Nucleic Acids Research*, 91(9):2443–2450, 2003.

[40] M. Buchanan. *Nexus: Small-Worlds and the Ground-breaking Science of Networks*. Norton, New York.

[41] R.S. Burt. *Brokerage and Closure: An Introduction to Social Capital*. Oxford University Press, 2005.

[42] V. Buskens. *Social Networks and Trust*. Kluwer Academic Publishers, 2002.

[43] K.M. Carley. Artificial intelligence within sociology. *Sociological methods and research*, 25(1):3–30, 1996.

[44] K.M. Carley. A comparison of artificial and human organizations. *Journal of Economic Behaviour and Organization*, 31(2):175–191, 1996.

[45] K.M. Carley. Computational organization science: A new frontier. *Proceedings of the National Academy of Science*, 99(3):7257–7262, 1999.

[46] K.M. Carley and M. Prietula. *Computational Organization Theory*. Lawrence Earlbaum Associates, 1994.

[47] J.M. Carlson and J. Doyle. Highly optimized tolerance: Robustness and design in complex systems. *Physical Review Letter*, 84(11):2529–2532, 2000.

[48] C. Castelfranchi. Modelling social action for ai agents. *Artificial Intelligence*, 103:157–182, 1998.

[49] E. Chou and A. Willsky. Analytical redundancy and the design of robust failure detection systems. *IEEE Transactions on Automatic Control*, 29(7):603–614, 1984.

[50] M.H. Coen. Non-deterministic social laws. In *Proceedings of AAAI'00*, pages 15–21. AAAI Press, 2000.

[51] M.D. Cohen, R.L. Riolo, and R. Axelrod. The emergence of social organization in the prisoners dilemma: How context-preservation and other factors promote cooperation. Technical Report 99-01-002, Santa Fe Institute, 1998.

[52] M.D. Cohen, R.L. Riolo, and R. Axelrod. The role of social structure in the maintenance of cooperative regimes. *Rationality and Society*, 13:5–32, 2001.

[53] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin. Resilience of the internet to random breakdowns. *Physical Review Letters*, (85):4626–4628, 2000.

[54] R. Conte and C. Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.

[55] R. Conte and C. Castelfranchi. *Understanding the Effects of Norms in Social Groups*. UCL Press, 1995.

[56] R. Conte, B. Edmonds, S. Moss, and R. Kaith Sawyer. Sociology and social theory in agent based simulations: A symposium. *Computational and Mathematical Organization Theory*, 7:183–205, 2001.

[57] R. Conte, R. Falcone, , and G. Sartor. Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, 7:1–15, 1999.

[58] N. Contractor, D. Zink, and M. Chan. IKNOW: A tool to assist and study the creation, maintenance, and dissolution of knowledge networks. In Toru Ishida, editor, *Community Computing and Support Systems*, number 1519 in Lecture Notes in Computer Science, pages 201–217, Berlin, 1998. Springer-Verlag.

[59] G. A. Cowan, D. Pines, and D. Meltzer. *Complexity: Metaphors, Models and Reality*. Addison Wesley, 1994.

[60] R. Cowan and Jonard N. Network structure and the diffusion of knowledge. *MERIT Research Memorandum*, 99-028, 1999.

[61] L. Danon, A. Díaz-Guilera, and A Arenas. Effect of size heterogeneity on community identification in complex networks. *arxiv:physics/0601144*, 2006.

[62] R.M. Dawes. Social dilemmas. *Annual Review of Psychology*, 31:169–193, 1980.

[63] J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141:171–185, 2002.

[64] J. Delgado, J.M. Pujol, and R. Sangüesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems Journal*, 1:131–138, 2003.

[65] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the Second ACM Conference on Electronic Commerce*, pages 150–157, 2000.

[66] F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7:69–79, 1999.

[67] P. Doreian. Event sequences as generators of social networks evolution. *Social Networks*, 24(2):5234–5237, 2002.

[68] P. Doreian, V. Batagelj, and A. Ferligoj. *Positional Analysis of Sociometric Data*, chapter Models and Methods in Social Network Analysis. Cambridge University Press, New York, 2005.

[69] P. Doreian and F.N. Stokman. *Evolution of Social Networks*, chapter The Dynamics and Evolution of Social Networks, pages 1–17. Gordon and Breach, New York, 1997.

[70] M. Dorigo, G. Di Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, 1999.

[71] S.N. Dorogovtsev and J.F.F Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW.* Oxford University Press, Oxford.

[72] S.N. Dorogovtsev and J.F.F Mendes. Scaling behaviour of developing and decaying networks. *Europhys. Lett.*, 52:33, 2000.

[73] S.N. Dorogovtsev, J.F.F. Mendes, and A.N. Samukhin. Structure of growing networks with preferential linking. *Phys. Rev. Lett.*, 85:4633, 2000.

[74] J. Duch and A. Arenas. Community identification using extremal optimization. *Phys. Rev. E*, 72:027104, 2005.

[75] N. Eagle and A. Pentland. Social serendipity: Mobilizing social software. *Pervasive Computing*, 2(4):28–34, 2005.

[76] H. Ebel, J. Davidsen, and S. Bornholdt. Dynamics of social networks. *Complexity*, 8(2):24–27, 2002.

[77] V. M. Eguíluz, M. G. Zimmermann, C.J. Cela-Conde, and M. San Miguel. Cooperation and emergence of role differentiation in the dynamics of social networks. *American Journal of Sociology*, 110:977–1008, 2005.

[78] J.M. Epstein. *Handbook on Computational Economics*, chapter Remarks on the Foundations of Agent-Based Generative Social Science. North Holland Press.

[79] J.M. Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999.

[80] J.M. Epstein and R. Axtell. *Growing Artificial Societies - Social Science from the Bottom Up.* MIT Press, 1996.

[81] P. Erdös and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.

[82] M. Esteva, J.-A. Rodriguez-Aguilar, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specifications of electronic institutions. *Lecture Notes in Computer Science*, 1991:126–147, 2001.

[83] R. Ferrer and R.V. Solé. *Statistical Physics of Complex Networks, Lecture Notes in Physics*, chapter Optimization in Complex Networks, pages 114–125. Springer, 2003.

[84] A. Flache. Individual risk preferences and collective outcomes in the evolution of exchange networks. *Rationality and Society*, 13(3):304–348, 2001.

[85] A. Flache and R. Hegselmann. Altruism vs. self-interest in social support. computer simulations of social support networks in cellular worlds. *Advances in Group Processes*, 16:61–97, 1999.

[86] A. Flache and R. Hegselmann. Rationality vs. learning in the evolution of solidarity networks: A theoretical comparison. *Computational and Mathematical Organization Theory*, 5(2):97–127, 1999.

[87] A. Flache and M.W. Macy. Social life from the bottom up: Agent modeling and the new sociology. In *Social Dynamics: Interaction, Reflexivity and Emergence. Proceedings of the Agent 2004 Conference*, 2004.

[88] G.W. Flake, R.E. Tarjan, and K. Tsioutsiouliklis. Graph clustering and minimum cut trees. *Internet Mathematics*, 1(4):385–408, 1993.

[89] F. Flentge, D. Polani, and T. Uthmann. Modelling the emergence of possession of norms using memes. *Journal of Artificial Societies and Social Simulation*, 4(4), 2001.

[90] E. Forgy. Cluster analysis of multivariate data: Efficiency versus interpretability of classifications. *Biometrics*, 21:768, 1965.

[91] U. Frank and K.G. Troitzsch. Epistemological perspectives on simulation. *Journal of Artificial Societies and Social Simulation*, 8(4), 2005.

[92] L.C. Freeman. A set of measures of centreality based on betweenness. *Sociometry*, 40:35–41, 1977.

[93] E.J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10:173–200, 2001.

[94] K. Fullam, T.B. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K. Suzanne-Barber, J.S. Rosenschein, L. Vercouter, and M. Voss. A specification of the agent reputation and trust (art) testbed: experimentation and competition for trust in agent societies. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems AAMAS-05*, pages 512–518, 2005.

[95] D. Gambetta. Can we trust trust? Oxford, 1990. Basil Blackwell.

[96] D. Gibson, J.M. Kleinberg, and P. Raghavan. Inferring web communities from link topology. In *Proceedings of the ACM Symposium on Hypertext and Hypermedia*, pages 225–234, 1998.

[97] A. Giddens. *The Constitution of Society.* University of California Press, 1984.

[98] N. Gilbert. Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy*, 15:1–14, 1990.

[99] N. Gilbert and R. Conte. *Artificial Societies: the Computer Simulation of Social Life.* UCL Press, 1995.

[100] N. Gilbert and K. G. Troitzsch. *Simulation for the Social Scientist.* Open University Press, 1998.

[101] M. Girvan and M.E.J. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Science*, 99(12):7821–7826, 2002.

[102] S. Golder and B.A. Huberman. The structure of collaborative tagging systems. *arxiv.org/cs.DL/0508082*, 2005.

[103] S. González-Bailón. *Mapping Civil Societies: a Network Analysis of the emerging forms of social capital on the web.* PhD thesis, Department of Sociology, Oxford University, 2006.

[104] R.V. Gould. Collective action and network structure. *American Sociological Review*, 58:182–196, 1993.

[105] R.V. Gould. Trade cohesion, class unity and urban insurrection: Artisanal activism in the paris commune. *American Journal of Sociology*, 98:721–754, 1993.

[106] M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.

[107] M. Granovetter. *Getting a Job*. Harvard University Press, Cambridge, MA, 1974.

[108] R. Guimerà and L.A.N. Amaral. Functional cartography of complex metabolic networks. *Nature*, 433:895–900, 2005.

[109] D. Gusfield. Partition-distance: a problem and class of perfect graphs arising in clustering. *Inform. Process. Lett.*, 3:159–164, 2002.

[110] M. Gustafsson, M. Hörnquist, and A. Lombardi. Comparison and validation of community structures in complex networks. *Physica A*, In Press, 2006.

[111] D. Hales and B. Edmonds. Applying a socially-inspired technique (tags) to improve cooperation in p2p networks. *IEEE Transactions in Systems, Man and Cybernetics*, 35(3):385–395, 2005.

[112] D. Harel and Y. Koren. On clustering using random walks. *Lecture Notes in Computer Science*, 2245:18–41, 2001.

[113] J. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press, 1977.

[114] R. Hegselmann. *Experimental ethics*, chapter A Computer Simulation of Classes, Cliques and Solidarity, pages 298–320. De Gruyter, Berlin, 1998.

[115] C.E. Hewitt. Open information systems semantics for distributed artificial intelligence. *Artificial Intelligence*, 47:79–106, 1991.

[116] J. Holland. The effect of labels (tags) on social interactions. Technical Report 93-10-064, Santa Fe Institute, 1993.

[117] P. Holme and B.J. Kim. Vertex overload breakdown in evolving networks. *Physical Review E*, (65):066109, 2002.

[118] G.C. Homans. *Social Behavior. Its Elementary Forms*. Harcourt Brace Jovanovich, New York, 1974.

[119] B.A. Huberman. *The Laws of the Web: Patters in the Ecology of Information*. MIT Press, 2001.

[120] B.A. Huberman and N.S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, (90):7716–7718, 1993.

[121] A.K. Jain and R.C. Dubes. *Algorithms for Clustering Data*. Prentice-Hall, 1988.

[122] A.K. Jain, M.N. Murty, and P.J. Flynn. Data clustering: A review. *ACM Computing Surveys*, 31(3):264–323, 1999.

[123] M. Janssen. Evolution of cooperation when feedback to reputation scores is voluntary. *Journal of Artificial Societies and Social Simulation*, 9(1), 2006.

[124] N. R. Jennings and J. R. Campos. Towards a social level characterization of socially responsible agents. *Proceedings on Software Engineering*, 144:11–25, 1997.

[125] H. Jeong, S.P. Mason, Z.N. Oltvai, and A.-L. Barabási. Lethality and centrality in protein networks. *Nature*, 411:41, 2001.

[126] H. Jeong, B. Tombor, R. Albert, Z.N. Oltvai, and A.-L. Barabási. The large-scale organization of metabolic networks. *Nature*, 407:651, 2000.

[127] R. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Inform. Process. Lett.*, 31:7–15, 1989.

[128] S. Kamvar, M. Scholsser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th International World Wide Web Conference*, pages 640–651, 2003.

[129] R. Kannan, S. Vampala, and A. Vetta. On clustering - good, bad and spectral. In *Foundations of Computer Science*, pages 367–378, 2000.

[130] K. Karvonen, L. Cardholm, and S. Karlsson. In *Proceedings of the Fifth Nordic Workshop on Secure IT Systems, NordSec-2000.*

[131] I. Kassabalidis, M. A. ElSharkawi, R. J. Marks, P. Arabshahi, and A. A. Gray. Swarm intelligence for routing in communication networks. *Globecom*, 6, 2001.

[132] S. Kauffman. *Origins of Order.* Oxford University Press, New York, 1993.

[133] H. Kautz, B. Selman, and M. Shah. The hidden web. *AI Magazine*, 18(2):27–36, 1997.

[134] K. Kelly. *Out of Control: The New Biology of Machines, Social Systems and the Economic World.* Perseus Books Group, 1995.

[135] J.E. Kittock. Emergent conventions and the structure of multi-agent systems. In D. Stein L. Nadel, editor, *In Proc. of the 1993 Santa Fe Inst. Complex Systems Summer School*, pages 384–389. Addison-Wesley, 1993.

[136] J.E. Kittock. The impact of locality and authority on emergent conventions: Initial observations. In *Proceedings of AIII'94*, pages 420–425. AIII Press, 1994.

[137] J.M. Kleinberg. Authoritative sources in a hyperlinked environment. Technical Report RJ 10076, IBM, 1997.

[138] J.M. Kleinberg, S. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. The web as a graph. In Lecture Notes on Computer Science, editor, *Proceedings of the International Conference on Combinatorics and Computing*, number 1627, pages 1–18. Springer, 1999.

[139] K. Klemm and V.M. Eguíluz. Highly clustered scale-free networks. *arxiv:cond-mat/0107606*, 2001.

[140] B. Kogut. The network as knowledge: Generative rules and the emergence of structure. *Strategic Management Journal*, 21(3):405–425, 2000.

[141] A. Komter. Reciprocity as a principle of exclusion. *Sociology*, 30(2):229–316, 1996.

[142] P.L. Krapivsky and S. Redner. Organization of growing random networks. *Physical Review E*, 63:066123014.

[143] P.L. Krapivsky, S. Redner, and L. Layvraz. Connectivity of growing random networks. *Physical Review Letters*, 85:4629–4632, 2000.

[144] D. Kreps and R. Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1998.

[145] S. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the web graph. In *Proceedings of the 42th Annual IEEE Symposium on the Foundations of Computer Science*, pages 57–65, 2000.

[146] S. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Trawling the web for emerging cybercommunities. In *Proceedings of the 8th World Wide Web Conference*, pages 403–415, 1999.

[147] Y. Lazebnik. Can a biologist fix a radio? or, what i learned while studying apoptosis. *Cancer Cell*, 2:179–182, 2002.

[148] R. Lempel and S. Moran. The stochastic approach for link-structure analysis SALSA and the TCK effect. In *Proceedings of the 9th World Wide Web Conference*, 2000.

[149] D.K. Lewis. *Convention: A Philosophical Study.* Harvard Univ. Press, 1969.

[150] F. Liljeros, C.R. Edling, L.A.N. Amaral, H.E. Stanley, and Y. Aberg. Sexual networks implications for the transmission of sexually transmitted infections. *Nature*, 411:907, 2001.

[151] N. Lin. *Social Capital: A theory of social structure and action.* Cambridge University Press, Cambridge, 2001.

[152] S.M. Lindenberg. *Handbook of Sociological Theory*, chapter Social Rationality versus Rational Egoism, pages 635–668. Kluwer, New York, 2001.

[153] E.D. Lumer and G. Nicolis. Synchronous vs. asynchronous dynamics in spatially distributed systems. *Physica D*, (71):440–452, 2000.

[154] Y. Luo and Ch. Shi. An emergence approach to behaviour convention in agent group through propagation of plans. In *Proceedings of the 2nd International Conference on Multi-Agent Systems*, pages 189–195. AAAI Press, 1996.

[155] M.W. Macy. *International Encyclopedia of the Social and Behavioral Sciences*, chapter Social Simulation. Elsevier, 2002.

[156] M.W. Macy and A. Flache. Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences*, 99(10):7229–7236, 2002.

[157] M.W. Macy and Y. Sato. Trust, cooperation, and market formation in the u.s. and japan. *Proceedings of the National Academy of Sciences*, 99:7214–7220, 2002.

[158] M.W. Macy and R. Willer. From factors to actors: Computational sociology and agent-based modelling. *Annual Review of Sociology*, 28:143–166, 2002.

[159] N. Mark. Beyond individual differences: Social differentiation from first principles. *American Sociological Review*, 63:309–330, 1998.

[160] G. Marwell, P. Oliver, and R. Prahl. Social networks and collective action: a theory of the critical mass. *American Journal of Sociology*, 94:502–534, 1988.

[161] P. Massa and P. Avesani. *On the Move to Meaningful Internet*, chapter Trust-Aware Collaborative Filtering for Recommender Systems. Springer-Verlag, 2004.

[162] E.M. Maximilien and M.P. Singh. Toward autonomic web services trust and selection. In *Proceedings of ICSOC*, pages 212–221, 2004.

[163] S. Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[164] H. Miller, P. Resnick, and R.J. Zeckhauser. Eliciting honest feedback in electronic markets. Technical Report RWP02-039-2002, KSG Working Paper, 2002.

[165] B. Misztal. *Trust in Modern Societies*. Polity Press, Cambridge, 1969.

[166] P.R. Monge and N. Contractor. *Theories of Communication Networks*. Oxford University Press, New York, 2003.

[167] M. Montaner, B. López, and J.L. De la Rosa. A taxonomy of recommender agents on the internet. *Artificial Intelligence Review*, 19(4):285–330, 2003.

[168] J.M. Montoya and R.V. Solé. Small world patterns in food webs. *Journal of Theoretical Biology*, 214:405–412, 2002.

[169] J. Moreno. *Who shall survive?* Beacon Press, New York, 1934.

[170] Y. Moses and M. Tennenholtz. Artificial social systems. Technical Report cs90-12, Weizmann Institute, 1990.

[171] S. Moss and B. Edmonds. Towards good social science. *Journal of Artificial Societies and Social Simulation*, 8(4), 2005.

[172] A. Newell. The knowledge level. *Artificial Intelligence*, 18:87–127, 1982.

[173] M.E.J. Newman. Scientific collaboration networks: I. network construction and fundamental results. *Phys. Rev. E*, 64:016131, 2000.

[174] M.E.J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89:208701, 2002.

[175] M.E.J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167–256, 2003.

[176] M.E.J. Newman. Fast algorithm for detecting community structure in networks. *Phys. Rev. E*, 69:066113, 2004.

[177] M.E.J. Newman. Modularity and community structure in networks. *arxiv:physics/0602124*, 2006.

[178] M.E.J. Newman, S. Forrest, and J. Balthrop. Email networks and the spread of computer viruses. *Phys. Rev. E*, 66:035101, 2002.

[179] M.E.J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113, 2004.

[180] M.E.J. Newman and J. Park. Why social networks are different from other types of networks. *Phys. Rev. E*, 68:036122, 2003.

[181] L.A.N. Nunes Amaral, A. Scala, M. Barthélémy, and Stanley H.E. Classes of small-world networks. *Proceedings of the National Academy of Sciences*, (97):11149–11152, 2000.

[182] G.M.P. O'Hare and N.R. Jennings. *Foundations of Distributed Artificial Intelligence*. John Wiley and Sons, 1996.

[183] L. Page, S. Brin, R. Motowani, and T. Winograd. Pagerank citation ranking: Bringing order to the web. Technical Report 1997-0072, Stanford Digital Library Working Paper, 1997.

[184] R. Pastor-Satorras, A. Vázquez, and A. Vespiagni. Dynamical and correlation properties of the internet. *Physical Review Letters*, 87:258701, 2001.

[185] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86:3200–3203, 2001.

[186] R. Pastor-Satorras and A. Vespignani. *Evolution and Structure of the Internet*. Cambridge, 2004.

[187] S.L. Pimm, J.H. Lawton, and J.E. Cohen. Food web patterns and their consequences. *Nature*, 350:669–674, 1991.

[188] J.M. Podolny and L. Page. Network forms of organizations. *Annual Review of Sociology*, 24:57–76, 1998.

[189] R.D. Putnam. *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster, New York, 2000.

[190] F.W. Rankin, J. van Huyck, and R.C. Battalio. Strategic similarity and emergent conventions: Evidence from similar stag hunt games. *Games and Economic Behavior*, (32):315–337, 2000.

[191] W. Raub and J. Weesie. Reputation and efficiency in social interaction: an example of network effects. *American Journal of Sociology*, 96:626–654, 1990.

[192] R.L. Riolo, M. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 411:441–443, 2001.

[193] G.L. Robins, P. Pattison, and J Woolcock. Small and other worlds: Global network structures from local processes. *American Journal of Sociology*, 96:626–654, 2005.

[194] J.-A. Rodriguez-Aguilar. *On the Design and Construction of Agent-mediated Electronic Institutions*. PhD thesis, Artificial Intelligence Research Institute, Spanish National Research Council, 2003.

[195] N.J. Saam and A. Harrer. Simulating norms, social inequality, and functional change in artificial societies. *Journal of Artificial Societies and Social Simulation*, 2(1), 1999.

[196] J. Sabater and C. Sierra. Reputation and social network analysis in multiagent systems. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 475–482, 2002.

[197] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1), 2005.

[198] R. Sangüesa and J.M. Pujol. Netexpert: A multiagent system for expertise location. In *Accepted to IJCAI'01 Workshop on Organizational Memories and Knowledge Management*, 2001.

[199] R. Sangüesa, A. Vázquez, and J. Vázquez. ACE: a multiagent recommender system using mixed collaborative and cognitive filtering. In *Proceedings of the WARS 2000, Workshop on Agent-based Recommender Systems*, 2000.

[200] J. Scott. *Social Network Analysis.* Sage Publications, 2000.

[201] Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agent systems. In B. Nebel C. Rich, W. Swartout, editor, *In Proceedings of Knowledge Representation and Reasoning*, pages 225–231, 1992.

[202] Y. Shoham and M. Tennenholtz. Co-learning and the evolution of social activity. Technical Report STAN-CS-TR-94-1511, Stanford University, Stanford CA (USA), 1994.

[203] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: Modeling, analysis and simulations. *Artificial Intelligence*, 94:139–166, 1997.

[204] H. D. Simon. Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering*, 2:135–148, 1991.

[205] H.A. Simon. *Models of Bounded Rationality.* MIT Press, 1982.

[206] T.A.B. Snijders. *Sociological Methodology*, chapter The Statistical Evaluation of Social Networks Dynamics, pages 361–395. Basil Blackwell, Boston and London, 2001.

[207] T.A.B Snijders. *Models for Longitudinal Network Data*, chapter Models and methods in social network analysis. Cambridge University Press, New York, 2005.

[208] T.A.B Snijders, C.E.G. Steglich, and M. Schweinberger. *To appear in Longitudinal Models in the Behavioral and Related Sciences*, chapter Modeling the co-evolution of networks and behavior. 2006.

[209] R.V. Solé, Smith R. Pastor-Satorras, R., and T.B. Kepler. A model of large-scale proteome evolution. *Advances in Complex Systems*, 5:43–54, 2002.

[210] C.E.G. Steglich, T.A.B. Snijders, and M. Pearson. Dynamics networks and behaviour: Separating selection from influence. *Submitted for publication*, 2004.

[211] F.N. Stokman and P. Doreian. *Evolution of Social Networks*, chapter Evolution of Social Networks: Processes and Principles, pages 233–250. Gordon and Breach, New York, 1997.

[212] M. Tennenholtz. Convention evolution in organization and markets. *Computational and Mathematical Organization Theory*, 2:261–283, 1996.

[213] P.M. Todd and G. Gigerenzer. Bounding rationality to the world. *Journal of Economic Psychology*, 24:143–165, 2003.

[214] E. Ullmann-Margalit. *The Emergence of Norms*. Oxford Univ. Press., Oxford, 1977.

[215] T.W. Valente. *Network Models and Methods for Studying the Diffusion of Innovations*, chapter Models and Methods in Social Network Analysis. Cambridge University Press, New York, 2005.

[216] J. van Huyck. *Handbook of Results in Experimental Economics (C. Plott and V. Smith, eds.)*, chapter Emergent Conventions in Evolutionary Games. North-Holland, 2001.

[217] A. Vázquez, I. Barrio, J. Vázquez-Salceda, J.M. Pujol, and R. Sangüesa. An agent-based collaboratory. In *Proceedings of ACAI2001 and EASS2001*, pages 135–142, 2001.

[218] J. Vázquez-Salceda. *The Role of Norms and Electronic Institutions in Multi-Agent Systems*. Springer, 2004.

[219] H.J.E. Verhagen. *Norm Autonomous Agents*. PhD thesis, Royal Institute of Technology and Stockholm University, 2000.

[220] A. Wagner. Robustness against mutations in genetic networks of yeast. *Nature Genetics*, 24(4):355–362, 2000.

[221] A. Walker and M. Wooldridge. Understanding the emergence of conventions in multi-agent systems. In V. Lesser, editor, *In Proceedings of 1st International Conference on Multi-Agent Systems*, pages 384–389. AAAI Press, 1995.

[222] G. Walker, B. Kogut, and W. Shah. Social capital, structural holes and the formation of an industry network. *Organization Science*, 8(2):109–125, 1997.

[223] T. Walsh. Search in a small world. In *In Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'99)*, volume 2, pages 1172–1177, San Francisco, July 1999.

[224] T. Walsh. Search on high degree graphs. In *In Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'01)*, pages 266–271, Seattle, August 2001.

[225] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.

[226] S. Wasserman and J Galaskiewicz. *Advances in Social Network Analysis*. Sage Publications, 1994.

[227] D.J. Watts. *The Dynamics of Networks Between Order and Randomness*. Princeton University Press, 1999.

[228] D.J Watts. Network dynamics and the small world phenomenon. *American Journal of Sociology*, 105(2):493–527, 1999.

[229] D.J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Science*, 99:5766–5771, 2002.

[230] D.J. Watts. *Six degrees: The Science of a Connected Age*. Norton, New York, 2003.

[231] D.J. Watts, P. Sheridan Dodds, and M.E.J. Newman. Identity and search in social networks. *arxiv:cond-mat/0205383*, 2002.

[232] D.J. Watts and S.H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393:440–442, 1998.

[233] M. Wellman and S.D. Berkowitz. *Social Structures: A Network Approach*. Cambridge University Press, Cambridge, 1988.

[234] E. Werner. *Foundations of Distributed Artificial Intelligence*, chapter Logical foundations of Distributed Artificial Intelligence. Wiley, 1996.

[235] M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002.

[236] M. Wooldridge and N.R. Jennings. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 10:115–152, 1995.

[237] F. Wu and B.A. Huberman. Social structure and opinion formation. *arxiv:cond-mat/0407252*, 2002.

[238] L. Xiong and L. Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):843–857, 2004.

[239] T. Yamagishi, K. S. Cook, and M. Watabe. Uncertainty, trust, and commitment formation in the united states and japan. *American Journal of Sociology*, 104:165–194, 1998.

[240] J.M. Yeomans. *Statistical Mechanics of Phase Transitions*. Oxford University Press, New-York, 1992.

[241] H.P. Young. Diffusion in social networks. Technical Report 2, Center on Social and Economics Dynamics, Brookings Institution, 1999.

[242] S. Younger. Reciprocity, normative reputation, and the development of mutual obligation in gift-giving societies. *Journal of Artificial Societies and Social Simulation*, 7(1), 2004.

[243] B. Yu and M.P. Singh. A social mechanism of reputation management in electronic communities. In *Proceedings of Fourth International Workshop on Cooperative Information Agents*, pages 154–165, 2000.

[244] B. Yu and M.P. Singh. Detecting deception in reputation management. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 73–80, 2003.

[245] G. Zacharias and P. Maes. Trust management through reputation mechanisms. *Applied Artificial Intelligence*, (14):881–907, 2000.

[246] W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33:452, 1977.