Rationality and Power in ICT: Addressing the gap in the middle

David Hales

Department of Computer Science, University of Bologna, Mura Anteo Zamboni 7, 40127 Bologna, Italy {hales}@cs.unibo.it

Summary. We are interested in designing better distributed software by explicitly modelling realistic user rationalities and the social, administrative and organisational structures that they inhabit and construct. Increasingly, software systems exist in networks that span multiple users and organisations. The components of such systems are loosely coupled and dynamically changing. In order to design such systems plausible models of user behaviour and social and administrative structures are required. Distributed systems design has tended adopt extreme models such as unconditional cooperation or rational action for user models and highly centralised or completely decentralised administrative structures. We aim to address the "gap in the middle" between these extremes by adapting and applying novel models derived from disciples within complexity science research in order to design better distributed ICT.

1 Introduction

Socially intelligent systems dynamically self-organise structures that coordinate individuals in a socially beneficial way such that the functionality of the social whole is greater than the sum of its individual parts. Examples of such systems are found in biological and social systems. They possess many desirable properties that designers wish to import into the emerging area of self-adaptive and self-organising ICT. But how can such systems be designed and engineered in a principled way? What are the current roadblocks in this endeavour? What kinds of research lines can be brought together to support each other and overcome such roadblocks?

ICT in the form of massive and open distributed software systems, that operate over networks, require large populations of processing nodes to dynamically coordinate and cooperate to achieve their goals. In the last few years a number of such systems have been successfully deployed[1]. They often

[1] For example the BitTorrent system allows users to share media content such as movies and music. The Skype system provides voice over IP (VoIP) telephony services.

exploit emergent properties and provide highly robust services under conditions of high flux (where nodes regularly join and leave the system) and in the presence of selfish or malicious behaviour.

To construct these kinds of systems effectively designers require, either implicitly or explicitly, models of user behaviour (user models) and of the social and administrative structures that relate users to processing nodes (social or administrative models). Many existing systems are designed and implemented with implicit versions of these models that are hard-coded into the implementation. For example, many P2P systems assume all nodes are essentially equal in terms of their administrative model (they are peers) and follow a naive user model in which cooperation or altruism is assumed by default.

We aim to explore these two dimension by explicitly specifying and modelling a set of plausible user models and social / administrative models. We will evaluate the design implications for constructing socially intelligent systems by applying these models over a range of application domains. The ultimate outcome will be a set of results which indicate what kinds of distributed design patterns are appropriate under different conditions of user behaviour and social / administrative settings.

2 User Models

Recently it has become apparent that to build robustness systems require incentive structures that discourage anti-social, selfish or malicious behaviour. In order to design and test effective incentive mechanisms it is necessary to select appropriate user models. It is fair to say that at the current time it is not widely agreed what kinds of user models are appropriate in given domains and contexts. One way to address this is through empirical measurement and analysis of the behaviour of systems in the wild in order to understand what users actually do. Alternatively, existing models from economics, game theory, evolution and social learning have been proposed. It is important to realise that any implementation (or protocol) is heavily influenced by the user model selected.

In this project we intend to explicitly specify a set of plausible user models drawing on and developing recent innovations from complexity science, social simulation, economics, empirical and experimental research. We will evaluate the effect of applying these user models to the design of effective distributed protocols. For example, a protocol design may perform efficiently under the assumption of an economically rational user model, from game theory say, but poorly when a boundedly rational social learning model is assumed, from, say, evolutionary economics. This kind of knowledge is essential in selecting and designing appropriate protocols for deployment in given domains and contexts.

Additionally we will address the issue of heterogeneous and dynamic rationalities. It is often assumed, again often implicitly, that a single user model

can characterise the behaviour of an entire population such that all users behave in the same way. It increasingly appears that this is unlikely. Empirical studies of P2P systems have shown how a small number of highly altruistic nodes support a large number of selfish nodes for example. A further aspect we will address is the idea that a single user may change or switch between different modes of behaviour depending on the context. For example, a user may be willing to behave altruistically if local resources (say bandwidth) are plentiful but may become selfish during shortages.

3 Social Structure / Administrative Models

Nodes within a distributed system in the context of a given application or service domain have some administrative relationship. That is, nodes can be partitioned into subsets such that each subset is under the administrative control of some user (or organisation). For example, two administrative extremes can be contrasted: Centralized or closed systems and pure open P2P systems.

In centralized systems all nodes fall within a single administrative authority. This means that the deployment of protocols on, and user model of, all nodes are centrally controlled in some fashion. Consider for example a typical e-mail system within an organisation. Users within the organisation must register centrally and be bona fide members of the organisation. They must obey the local administrative rules of use and the mail server they use will be centrally administered.

Conversely consider pure P2P systems. Here each node is considered as an administrative domain in its own right. Deployment and behaviour of each node is under the control of a user at each node. There is no centralised administration or control. For example, a typical P2P file sharing system allows users to install and modify their clients and settings and decide what files to share or download. There is no check on user identities or centralised server administration.

These are two extremes, and often the reality of deployed systems lay somewhere between them. Interestingly however, a lot of protocol design tends to assume one of these extremes. Additionally, it is rare for models to focus on the possible dynamics of such administrative relationships. For example, where an existing authority might disappear or some subset of nodes may cede administrative control to another node spontaneously.

In this project we will explicitly specify a set of plausible administrative models and evaluate the effectiveness of protocol designs in relation to them. This work will draw on and develop recent models from organisational theory and complexity and social modelling in addition to emerging empirical analysis of administrative and organisational dynamics. Interestingly, recent models in organisation theory and economics have drawn on the dynamic network-centric view of organisational structures. We will also consider dynamic models and endogenously created and emergent administrative structures.

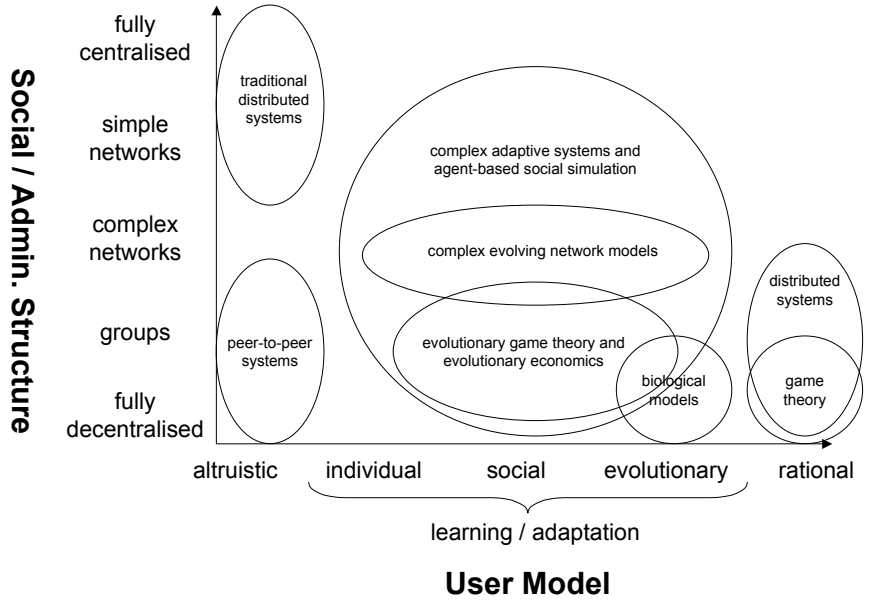

Fig. 1. A schematic showing the different research areas and their domain of focus on two axis of user model and social structure. Note that distributed systems research as tended to focus on extremes. We aim to focus on the "gap in the middle" centred on learning and adaptation and complex networks.

From a sociological perspective administrative models are essentially models of social structure and power relationships. Hitherto, protocol designers have tended to focus on two social extremes of either centralised dictatorship or distributed anarchy. We aim to explore the interesting and more plausible space between them - the "gap in the middle". Figure 1 shows a schematic of different research areas with their focus place on the two axis of user model and administrative / social model. This gives an idea of the kind of approaches we will be drawing on in relation to the gap we have discussed.

4 Relationship to Mechanism Design

What we are proposing bears comparison with what has been termed mechanism design from game theory. However, here we intend to explore the two dimensions of user models and social / administrative structures that are plausible within open and real systems considering mechanism other than traditional markets or auctions. Hence we relax strong economic assumptions concerning rational action and associated solution concepts by considering a space of alternatives plausibly grounded in empirical reality.

5 Method

The complexity of the emergent behaviours we are interested in, comprising adaptive behaviour, social learning and administrative / social structures, dictates that often analytic solutions will not be possible. Hence we intend to make use of computer simulation and empirical analysis. Also, where feasible, we will implement prototype protocols and evaluate them on testing platforms (such as PlanetLab) and potentially by releasing them into the wild i.e. deployment over the Internet to real users.

6 Integrating ICT, user and administrative / social models

Historically, social simulation and complexity inspired modelling has focused on human behaviour and interactions. Even when those interactions are mediated by ICT, as they increasingly are, models tend focus on user behaviour and their emergent effects. ICT systems are rarely modelled in any detail. Conversely, ICT research, particularly in distributed systems, focuses on detailed measurement and modelling of technology with little emphasis on plausible user modelling. For example, many distributed systems protocol designers assume users can be modelled as either unconditionally cooperative or as rational actors using classical game theory / economics assumptions. Such user models are imported from other contexts and disciplines with little justification or plausibility. We believe this is partially a result of the cultural and academic divide between the human sciences and the technological sciences. What is required is to bring together these two areas to produce models with detailed and plausible ICT and user components. Both need to be given equal importance. This is essential when the systems under investigation result from tightly coupled interaction between users and ICT systems - so-called techno-social systems that is, social systems that are mediated and constructed with ICT.

7 Possible ICT application domains

However, it is not the case that user models and ICT models can be simply taken "off the shelf" and glued together. This is because the kinds of models needed, by definition, interact and affect each other. The way a user interacts with an ICT system influences performance and this in turn enables or constrains user behaviour. Hence users and ICT form a tightly coupled feedback loop. This means that plausible models have to be developed together in relation to a given application context or domain. Some possible application domains include:

- Peer-to-Peer middleware services in highly dynamic and open systems over the internet
- Disaster response scenarios where ICT can self-organise under high levels of failure and potentially extreme user behaviour
- Dynamic service orientated self-organsing ICT - augmenting and contributing to Services Science approaches

In order to address these issues we aim to bring together leading European researchers in social simulation [2] and distributed systems [1] to explore the design of next generation socially intelligent ICT systems.

8 Areas addressed

In the context of Objective IST-2007.8.4: FET proactive 4: Science of complex systems for socially intelligent ICT of FP7, the work discussed here covers the three main challenge areas in the following way:

- Theoretical and algorithmic foundations: by examining and applying novel user and social models to ICT between the traditional extremes we aim to develop simulation tools and design patterns that incorporate micro-maco and macro-micro feedback, emergent social structures and psychologically and socially plausible approaches.
- Data-driven simulation: We aim to draw on empirical measurement where appropriate to validate our models. Open and existing systems such as BitTorrent already have delivered valuable empirical data concerning how real users actually behave - how altruistic they are, their client preferences, levels of malicious attack etc. A key aspect of our approach is not to rely on overly simple models.
- Prediction and predictability: By developing plausible models of user behaviour and social structures we aim to give ICT designers tools to allow them to be able to have, at least, some level of predictability as to how their systems will behave "in the wild". Currently there is no way for designers to know how their systems will behave when they are massively deployed because it is not feasible to test a system with millions of users with actually realising it. Currently this is achieved via prolonged beta-testing but this can be highly costly and problematic.

References

1. Giovanna Di Marzo Serugendo et al (2007) Proceedings of The First IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2007), Boston, Massachusetts, USA, July 2007, IEEE Comp. Soc. Press.
2. The Journal of Artificial Societies and Social Simulation (JASSS), http://jasss.soc.surrey.ac.uk/JASSS.html.