

Self-Organization in Peer-to-Peer Systems

Hermann de Meer
Chair of Computer Networking and Communications
Faculty of Computer Science and Mathematics
University of Passau
demeer@fmi.uni-passau.de

Abstract Peer-to-Peer systems have often been advertised to as being “self-organizing”. This has become a somewhat elusive property with meaning and significance being far from clear. Answers sought in this research therefore focused on questions like: How well is self-organization understood in terms of peer-to-peer systems? How much self-organization is there really? How can, if at all, mechanisms of self-organization be exploited for control and management of peer-to-peer systems?

Self-organization is closely related to the phenomenon of structures that *emerge* out of activities by groups of lower components. The interactions among those components, thereby, are governed by rules purely based on local information. There is no external ordering influence that could impose the emergent properties that appear at the global level. Emerging structures are driven by positive feedback information that amplifies fluctuations. To balance systems, on the other hand, negative feedback information plays a crucial role for self-regulation. Examples of negative feedback would be depletion of resources due to saturation or competition. As a result, structure emerges as a function of an interaction between positive and negative feedback information. The emerging structure often exhibits so-called small-world characteristics where a few components have relations to a very high number of other components and many components have only very few relations to a few other components. Self-organizing systems operate in a *critical* state that allows them to respond adaptively, but not chaotically, to some stimuli. Typically, there exist several such semi-stable, critical states such that bifurcation, that is, state change, may occur in response to parameter variations. Multi-stability in critical states is crucial for adaptive, self-organizing systems. Self-organization arises in systems with multiple iterative interactions among components that behave according to some rules-of-thumb. This results in a recursive process with cascades of interacting mechanisms where complexity unfolds progressively.

Self-organization has to be distinguished from systems that are centralized and from systems for which there are blueprints, recipes or templates available. Here components are “unaware” of any global situation and complex group behaviour evolves from a component-level behaviour that is much less complex. The emerging structure is non-predictable and non-deterministic. Different group-level patterns could thus arise from identical conditions. Diversity of responses could result in qualitative and quantitative changes in groups while components do not change behaviour. This is part of stimulus-response patterns in critical state systems. The environment, thereby, acts as strong constraint that shapes the self-organization process.

Scalability, flexibility, evolveability and self-management without any central control are among the most prominent and attractive properties of self-organizing systems. This is particularly attractive from the perspective of the Internet, where limitations in scalability and flexibility have been increasingly manifesting themselves as major problems and challenges. This is in particular so due to prevalent client-server architectures with their implied limitations on the server-related resources on the one hand and plenty of unused resources at the client side on the other hand.

The recent emergence of peer-to-peer systems has countered that trend. Here interacting end systems (peers) jointly use resources in end systems in a more or less decentralized fashion.

Equal and autonomous participants cooperate in a form that has often been referred to as “self-organizing”. However there are several classes of peer-to-peer systems that all exhibit different properties, not the least with respect to the degree of self-organization involved. The property that can be distinguished most clearly is the degree of centralization involved in the different operations and types of peer-to-peer systems. There are resource mediation (where to find resources) and resource access (who may access resources, when and how) operations to be differentiated to what degree they are centralized or decentralized. The more decentralized, the more self-organizing peer-to-peer systems seem to be.

In addition to routing, searching for and accessing of resources, other features need to be taken into account as well. Those features entail reliability of the overall system where the individual nodes are autonomous and unreliable, efficiency of (network or enabling) resource usage, service discovery, or (topology maintaining) search, and identification and status control of neighbours within a peer-to-peer (overlay) system. The obvious question is, can self-organization be the key for those control and management tasks to be accomplished effectively?

While the degree of decentralization obviously can be taken as an indication peer-to-peer systems being self-organizing, it is firstly much less clear how definitions of other (core) properties of self-organization are met. And secondly, it remains an open question of what are the limits of the benefit in basing management and control of peer-to-peer systems on self-organization.

More specifically, what rules-of-thumb need to be defined on the component level to achieve a well-behaving emergent structure? Having simpler rules rather than complex holistic behaviour would be good for further automating management of peer-to-peer systems, but maybe difficult to find. Since emerging structures are non-deterministic, forecasting them based on local rules could be arbitrarily difficult. How could a system be steered towards an attractor that corresponds to the best possible fitness in a given condition? How to set the control cycles right, in particular, how to balance positive and negative feedback well? And these are only some of the challenges that need to be met.

While it seems obvious that current peer-to-peer systems, whether structured or unstructured, are further from being (fully) self-organizing as might have been anticipated (and as they are often misleadingly claimed to be), it remains to be seen what are the limits of self-organization itself.

What degree of complexity is generated by self-organization that is still of use for peer-to-peer systems and matches essential properties of peer to peer? How far can that complexity be self-controlled and self-managed, since automatic and decentralized control and management appear of utmost importance for current peer-to-peer systems? How well can self-organization be combined with alternative mechanisms that are still likely to be needed, like pattern formation and decision making?

Summary Peer-to-peer systems resemble some features of self-organization, like decentralization for resource mediation and access. However, peer-to-peer systems need much more control and management mechanisms in place with security, performance and perhaps anonymity in mind. It is much less clearer what role self-organization could play here and how. Self-organization, on the other hand, is characterized by emergent properties on the macro level based on simple rules-of-thumb that determine system behaviour on the micro level. While self-organization is always decentralized, it relies on cascades of interaction and recursion and on positive and negative feedback loops. It remains unclear whether such a non-linearity would be attractive for peer-to-peer systems and to what extent. Self-organization does not easily provide “control knobs” that can be used for tweaking peer-to-peer systems comfortably.