Open Research Online

The Open University's repository of research publications and other research outputs

Towards the evolution of social structure

Journal Item

How to cite:

Edmonds, Bruce; Norling, Emma and Hales, David (2009). Towards the evolution of social structure. Computational & Mathematical Organization Theory (CMOT), 15(2) pp. 78–94.

For guidance on citations see [FAQs](#).

© 2008 Springer Science+Business Media, LLC

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Version: Accepted Manuscript

Link(s) to article on publisher's website:

<http://dx.doi.org/doi:10.1007/s10588-008-9052-0>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Towards the Evolution of Social Structure

Bruce Edmonds*, Emma Norling* and David Hales⁺

*Centre for Policy Modelling

Manchester Metropolitan University

⁺Department of Computer Science

University of Bologna

Motivation

To what extent can social structure result from evolutionary processes, as opposed to deliberately organised by a collection of intelligent individuals? That is, could social organisations (which are usefully identifiable as an entity in their own right) result from a mechanism which only relies on the processes of random variation and selective reproduction. Or, on the other hand, is it necessary for the individuals to have acquired (through evolution or design) sufficient cognitive abilities to deliberately organise themselves into such entities (using planning, reasoning, deliberate experiment, anticipatory learning or the like).

The question is important, for the answer (w.r.t. any particular system) makes a difference. For example, in the former (evolutionary) case the abilities of the individuals can co-evolve with the social structures that they can (unthinkingly) sustain so that some of their features (e.g. a propensity towards cooperation or a distrust of strangers) could have evolved to reinforce or build upon the existing social structures and so boot-strap the complex society we inhabit. If this is not the case then we must understand the particular social structure as essentially intended institutions built *subsequent* to the abilities developed in other circumstances. This might be important in either understanding observed societies or producing/managing new ones (including artificial ones, such as Multi-agent systems).

In a sense this question is the complementary half of the Machiavellian Intelligence theses which postulates that our intelligence evolved (partly) because it provided an ability to manipulate social structures and situations to our (individual) advantage. It can be seen as a part of the, more general, Social Intelligence Hypothesis [2], which is that suggestion that our intelligence gives us evolutionary advantage via the social structures it enables. An example of this is the ability to imitate [5] which may allow

groups of people to develop a culture of skills, rules and traditions that equip that group to inhabit specialist ecological niches [20].

Of course evidence has the primary role as far as observed social structures are concerned. For example, the apes can not sustain as sophisticated social structures as humans, so the present abilities of humans that go beyond those of apes must be crucial for their production and/or maintenance. However we do not have strong evidence about how our present abilities developed over biological evolution, there might have been a whole succession of structure then ability then structure etc. which resulted in the current position. Clearly new social structures and institutions have arisen over a historical timescale and thus it seems clear that some of our social structures rely upon our present abilities.

On the other hand, some social structures clearly do not require complex cognition by its individuals. Slime molds organise themselves (under the right conditions) into what can be meaningfully called a single entity (a fruiting structure). Since they, as individuals, have no (or only simple) information processing abilities, some social structure does not require complex cognition.

Thus the question should be rephrased into: ***which** mechanisms might result in **what** social structures under **which** conditions?*

Computational simulation/systems can help answer this question by establishing some possibilities in terms of the causal connection between mechanisms, conditions and resulting structures. That is, by starting to map out the possible triples of: (mechanism, conditions, resulting structure). This is not an easy task, for simulations can be deceptive in terms of the robustness of their results against “small” changes in their set-up and, given their intrinsic complexity, we can easily be mistaken as to our interpretation of how and why they produce the results they do. At the end of each example model described below I will attempt to summarise the (mechanism, conditions, resulting structure) triple. This is not aiming to be a survey of all possible such triples for a range of conditions and mechanisms – that would take a whole book or thesis (hint) to do, but an indication of the possibility of such a catalogue. I am not of the view that it will be possible to more than roughly outline the conditions under which every mechanism works – it probably will always be possible to find a new

alteration which would invalidate any specific set of necessary and sufficient conditions.

This paper looks at one stream of computational experiments that start to address this project – they examine how a particular class of mechanism (tag-related mechanisms) can produce group-like structures for cooperative behaviour. In this sense this paper is a mini-survey and prospective concerning these models and their results – models and results in which I have played only a part, the main work being done by others including: David Hales, Rick Riolo and Emma Norling.

About Tags

“Tags” are a socially distinguishable mark or signal, for example wearing a smart suit [17]. The important thing is that tags are not “hardwired” to any particular behavioural trait so that even if a particular tag comes to be associated with some behaviour (clerical garb is worn by trustworthy individuals) the connection is fallible – someone new can always adopt the tag regardless of whether they exhibit the associated behaviour (e.g. a con man pretending to be a respectable businessman). A biological example is where a wasp effectively advertises its ability to sting via its yellow and black stripes and thus reduces danger to itself by warning off others. The stripes are the tag of the species – although they may have originated as the result of a random mutation it persists because of the evolutionary advantage they confer. However, other species may evolve to use the same tag without the ability to sting (e.g. the hoverfly) to exploit the signal that the tag has come to be associated with. Thus although tags can be initially meaningless, they can allow for the recognition of classes of individual, albeit in a fallible manner. Thus displaying and recognising tags can be seen as a minimal ability that allows for a new set of social structures to arise, namely the appearance of clusters that can be sensibly thought of as “groups”. For cooperative behaviour, the basic rule is to preferentially interact with those with the same or similar tags.

Example 1 – integral tags on the one-shot PD game

I start with the simplest model that I know of, which is the one described in [11]. Here there is a population of individuals simulated over a number of discrete generations. Each generation, each individual is paired with a fixed number, n , of

others with each of whom it plays a one-shot prisoners' dilemma game (PD) according to its strategy flag, which is fixed as cooperate or defect. At the end of each generation each individual's score is totalled and individuals are propagated into the next generation in accordance to this score (higher scoring individuals are propagated more, on average, than lower scoring ones). There is also a small probability of a progeny's strategy "mutating" between cooperation and defection.

If the model were only as so far described the result is that quickly defectors would come to dominate the population. However the ability to display and recognise tags is now added to the mix and this radically changes the outcomes. With tags: each individual has a tag represented by an integer from a certain range; progeny of individuals have the parent's tag (with a small probability of mutation to a random integer); thus the prevalence of tags as well as the strategies are determined by an evolutionary process. The exploitation of the tags comes from the basic rule that individuals only "pair" with those with the same tag if that is possible (otherwise completely randomly). More information about this model is given in the appendix.

The dynamics observed in this model can be summarised as follows: eventually, by chance, a small group (usually a pair) of only cooperative individuals with the same tag arises; these are only paired with each other and hence score highly compared to the background population and hence these individuals are preferentially propagated into the next generations with the same tag and strategy; eventually an individual arises with the same tag but a defector; this individual does better than anyone and hence is propagated at the highest rate so the set of individuals with this tag becomes flooded with defectors; now the cooperative individuals in the group do very poorly compared to the background population and hence are selected out leaving only defectors with that tag; these defectors also do badly and are selected out. It makes sense to interpret all those with the same tag as a "tag group". This "life cycle" occurs repeatedly and in parallel – there are tag groups arising and dwindling all the time. This rising and falling of tag groups is illustrated in Figure 1.

Figure 1. The continual rise and fall of tag groups. Each line on the y-axis represents a unique tag and time is represented on the x-axis

If the rate at which cooperative groups arise is sufficiently high, the average time until a defector arises long enough, and the “killing effect” of a defector on a tag group fast enough, then the overall level of cooperation can be high, even given that cooperation is not an “evolutionary stable strategy”. The graph is not very informative - a high level of cooperation is quickly established thus it looks almost identical to Figure 1.

There have been many variations of this kind of model explored, including where tags are represented as a real number and there are degrees of tolerance, with probabilistic preferences for interaction based on tags, different selection and mutation regimes etc. The important property of these models is they show a way of maintaining the overall level of cooperation even when it is (a) possible for defectors (who cooperate with no one) to arise and (b) where it is locally and immediately beneficial for an individual to do so. Furthermore the variety of models in which such an effect occurs indicates its robustness. Of course if one builds in further protection against defectors then the tag groups can be extremely stable over long periods of time. However, this is not very surprising and depends upon such an enforcement being there – this is the sort of enforcement that might be possible to implement with more sophisticated cognitive machinery using punishment or forced ejection (tag-change).

However, regardless of whether the tag cooperation is crisply defined (as in the model described above) or fuzzily defined using probabilities for interaction or tolerance

ranges and real number tags, the clusters of individuals with similar tags who preferentially interact with each other and display such characteristic “life-cycles” are usefully identifiable as “groups” in the same way that a connected cluster of cells might be identified as a person. However such groups have no structure beyond this clustering – they simply partition the population (possibly fuzzily).

Here the mechanism is that of that of simple crisp tag recognition where interaction is preferred with those of identical tags. The conditions for this to work are that new “clean” groups can form relatively quickly compared to the rate at which defectors appear with the same tag, then when they do that they effectively “kill” that group quickly. Another necessity is that the task is one where mutual cooperation is much better than mutual defection, but where a minority of defectors does much better still. The resulting structure is that of a collection rising and quickly falling groups so that at any moment there are many that are in cooperator-only groups. Often these groups are developing in parallel with each other.

Example 2 – single group formation with floating point tags

This is the model described in [21] and probed in [7]. Here there is a fixed population of individuals that are randomly paired a fixed number of times. They are each randomly paired a fixed number of times with other individuals to whom they donate 1 (and incur a lesser cost, c). The sum of their score determines their “fitness” which affects the rate with which they reproduce into the next generation.

Here, the tag behaviour of individuals is determined by two floating point numbers, both drawn from $[0, 1]$: the first is its tag value and the second its “tolerance”. An individual will cooperate with another it is paired with if the difference between its tag value and that of the individual it is paired with is less than or equal to ($\leq$) its tolerance value. Thus the tolerance value represents how cooperative an individual is: a tolerance value of 1 means that it will cooperate with every other it is paired with; a tolerance value of 0 means that it will only cooperate with another whose tag value is exactly the same as its own and no others. This mechanism is illustrated in Figure 2.

Figure 2. The tag/tolerance mechanism – the individual with the larger tolerance donates to the one with smaller tolerance but (in this case) not vice versa

The dynamics of this model are that cooperation is very quickly established and is extremely stable from there on with only very rare collapses of this group and the establishment of a new group. The cooperation is composed of only one dominant group at a time, and this group is mostly composed of exact tag-clones of each other who are forced to cooperate with each other due to the model set-up. Thus this model implements a sort of kin-selection, whereby individuals give to those directly related to them. If one changes this model so that complete defection is possible, so that cooperation only occurs if the difference between its tag value and that of the pairer is strictly less than ($<$) its tolerance value then the cooperation effect disappears. Thus the tolerance value here is not really important and only really has a function as a symmetry breaking mechanism (depending upon the selection mechanism implemented). For a complete analysis of this model see [7].

Thus the mechanism is similar to that of example 1, where groups of identical tagged individuals cooperate but with a possible “penumbra” of nearby hangers-on. Here the tolerance values do not play an essential role except sometimes as symmetry breakers (as documented in [7]). Similar conditions as example 1 hold, but in particular it is necessary that pure defectors (those that give to nobody under any circumstances) are not allowed. The structure that emerges is that of an extremely stable single dominant group with a core of identically-tagged individuals and a varying periphery of close individuals created by mutation. This group only very rarely gives way to another group.

Example 3 – evolution of “symbiosis” with floating point tags

The next model makes a small step towards adding structure with a group, by adding (and encouraging) mixtures of individuals with different skills within a group using tags. This is the biologically inspired model described in [8]. Here as well as tags, individuals have a *single* specific skill from a small range. “Nutrition” is provided to the environment in kinds that correspond to those skills – only individuals with the right skill can “harvest” each unit of nutrition. Individuals can store a certain amount of each kind of nutrition. However individuals need all kinds of nutrition to survive (expending some of each every generation) and more of all kinds to reproduce. The idea is that there is some reason why the individuals have had to specialise – clearly an individual who could harvest more than one type would out-compete those who rely on symbiosis. Thus this simulation might be interpreted as representing a situation where nutrition is difficult to come by (thus requiring specialise and expensive skills) rather than one where nutrition is plentiful.

Here individuals are randomly paired a number of times and may share any excess of any kind of nutrition they are storing with those with sufficiently similar tags to their own. As before the individuals do not directly gain anything for themselves by sharing (indeed they lose potential future survival and reproductive possibilities) so one might expect that, over evolution, more “selfish” individuals (those that share with few or none at all) would dominate.

Here, the tag behaviour of individuals is determined by two floating point numbers drawn from $[0, 1]$: the first is its tag value and the second its “tolerance”. An individual will share excess resources (in any store) with another it is paired with if the difference between its tag value and that of the individual it is paired with is strictly less than its tolerance value. This follows the tag structure of the model in Example 2 above but the strict inequality allows for completely selfish individuals which destroys the cooperation effect reported there. This mechanism is illustrated in Figure 2 above.

The resources in this model (of different kinds) are used up in sharing, living, and reproducing, thus the population size is controlled by the input of resources and the efficiency of the sharing. For individuals to survive and flourish they need to receive gifts of nutrition of the types they can not themselves harvest – thus the more types of

nutrition are necessary the more difficult it is to survive via cooperation (in the example below there are 3 kinds). More information about this model is given in the appendix.

Figure 3. The rise and fall of tags in the symbiosis model. The y-axis indicates the number of individuals holding a given tag value

The dynamics in this case are, in one sense, more complicated than in Example 1, although they follow the same basic outline. Clusters of cooperating tag groups do arise, punctuated by short periods of complete collapse (see Figure 3). Somehow, by chance, a small group of individuals covering the necessary range of skills and with compatible tag and tolerance values arises; these then are much more successful than the background population and quickly reproduce to form the dominant group; soon mutations arise with smaller tolerance values in one of the individuals with one of the skills – this one does better than the others (because it donates less) but if it reproduces too much its progeny starves itself of the necessary other nutrients because there are fewer of the others to donate to them. Thus as individuals mutate to ever more selfish individuals in a downwards arms-race the dynamics become more like those of Lotka-Volterra predator-prey dynamics. Eventually the dynamics, whose oscillations grow in amplitude, kills off one of necessary sub-groups and the

population collapses back to the near-zero state until a new group arises again by chance. This is illustrated in the figure immediately below.

Figure 4. The dynamics of the different skill groups: each line shows the number with each of the three skills in this example; the system starts (at the first arrow) in an unviable state; then a cluster of cooperating tag groups emerge (second arrow, the CE=coexistence phase); but by 250 (the third arrow) a predator-prey type dynamics has been established; which finally breaks down into another unviable state (fourth arrow and UV=unviable phase).

Due to the limited resources which force a kind of “winner takes all” dynamics between types of individual, only one cluster of tag groups occurs at any one time in this model – there is no parallelism between these clusters. Thus this model represents the cooperation that could occur within a single niche only. But it does show how cooperation can occur and be maintained with competing cooperating clusters. Unlike in the previous model the tag groups are not sharply defined but are fuzzy, with “hangers-on” at their edges who survive by receiving nutrition through a chain of intermediaries outwards. Thus although the clusterings are very definite when viewed in the results the “groups” are not crisply delineated in terms of their interactions.

The mechanism here is that of variable tags, but where no individuals are forced to cooperate just because they have the same tag. The conditions for this to work are that there is inherent specialisation of skills here, but where the results of these are all necessary for all individuals. Under these circumstances a single cluster of cooperating tag groups with different skills arise that fluctuate for a while due to population dynamics, but are slowly undermined by the arrival of (relative or

absolute) parasites which prey on the other groups, eventually causing their collapse. Only one such cluster can exist at a time in this model due to the winner-takes-all population dynamics, thus when this collapses there is a period before the cooperation restarts.

The extension of tags to local structure

The models above use a low information group membership indication (a single integer or floating point number) and simple rules determining whether an interaction occurs. This only supports fairly simple, unstructured groups. It can only encode crisp or fuzzy (respectively) group membership. Richer and more informative can be encoded in a social network structure, where the presence of a link (respectively absence) between two nodes indicates whether they might interact (respectively not interact). Thus, if one visualises the network, one may see the separate (or almost separate) groups with many internal connections but few (if any) connections between the groups occurring. Thus in the model below there are no tags other than the set of connections a node has – it is these connections which determine the structure.

Example 4 – job sharing using the SLAC algorithm

In this model there is a population of nodes of fixed number. Each of these can be connected to others (up to a maximum number) who they can interact with. Each node has a skill number (from a limited number of types indicated by an integer), a strategy bit (cooperate or defect) and a list of its connections. Each generation a fixed number of “jobs” are allocated to nodes at random, each of these jobs corresponds to one of the skill types. If a node has the right skill itself it does the job and gains a fixed reward. If not, it asks a fixed number of nodes it is connected to to do the job for them, if they are cooperative and have the right skill they do the job for the original node (thus incurring a cost themselves), the original node gets the reward if this happens. Thus a cooperative node does not directly benefit from doing work for others. A node's score is the total of its rewards minus the costs it incurred over that round of jobs. This task, called “SkillWorld”, is introduced in [9].

On top of this structure is an evolutionary algorithm based upon imitation and mutation. That is, nodes pick a random other node and compare their score with the score of that node. If it is doing better it drops its links, links to that node, and

imitates its connections and strategy (after each link is made if either exceeds its maximum number of connections is randomly drops one). Every now and then a node “resets” with a small probability, that is drops all of its nodes, and attaches to a random node. Also there is a small probability that a node changes its strategy bit. This is called the SLAC algorithm (Selfish Link-based Adaptation for Cooperation) [13]. There is more information about this model in the Appendix. The result of the this algorithm on the SkillWorld task is the rapid establishment of cooperation in the population which remains pretty stable and high from then on – see Figure 5.

Figure 5. Level of cooperation in the SLAC algorithm. The dark circles show the proportion of altruistic nodes. The grey crosses show the proportion of completed jobs

The sequence of network visualisations in Figure 6 shows the development of the network structure in this model. It starts with lots of disconnected nodes which quickly form up into clumps after 12 generations; at generation 117 a new defector arises (in the small circle); which then multiples into a bigger clump by generation 120 at the expense of the group co-operators it is connected to; but by generation 122 all the cooperative nodes that were connected to them have imitated ones that were not and the non co-operators have become isolated and disappear completely by the next generation.

Figure 6. The appearance then isolation of defectors, the defectors are circled in each “snapshot”

Thus here we see more structure resulting from the model as a result of an evolutionary mechanism. It turns out to be robust against many possible attacks against it. However the structures developed are not very sophisticated. It is not good for tasks which require competition between individuals or that need the network to be substantially connected (such as search via flood-fill queries in an P2P network).

In SLAC when a node imitates another it drops all its previous links – this corresponds to a probability of dropping existing nodes of 1. If one implements the model so that this probability can be set to any floating point number within $[0, 1]$, then one can “tune” the resulting network so that it stays connected and yet can maintain the connectedness of the whole – this variant is called SLACER [16]. For example setting this probability to 0.95 causes an almost imperceptible drop in the overall level of cooperation (see Figure 7) and yet results in an almost completely

connected network. Somehow the co-operators are able to “avoid” the defectors without breaking into lots of separate groups. I have not shown some example structures that this results in since it is very difficult to visually discern the structure.

Figure 7. Level of cooperation with SLACER $p=0.95$. The grey crosses show the proportion of completed jobs

Here the mechanism is the SLAC rewiring algorithm, allowing nodes to change who they connect to within a larger network. It works as long as it is possible for cooperating nodes to effectively escape the defectors and hence isolate them, a condition which does not occur when the network is densely connected into a single component. The structures that result are dependent upon the probability, p , of dropping old links when mutating or copying new ones. When $p=1$ the network breaks up into lots of smaller well connected groups containing all the necessary skills for a job. When p is just less than 1, the network is well connected, but retains enough “stretched” structure to enable the cooperators to avoid the defectors. When p is smaller the structure degenerates into a densely connected network of all defectors.

Example 5 – SLAC on unidirectional P2P file- sharing task

The final example given in this survey is that of a more realistic model of file-sharing via a P2P network. P2P networks are those without a centralised or client-server structure – every computer is both a client and a server. Each computer in the network “knows” of a limited number of others with which it interacts. To communicate with others further away the messages have to go through intermediary

computers. Well known P2P networks that operate on this basis work as a virtual layer over the internet include: e-donkey and BitTorrent.

This simulation is quite complex, so I do will not describe it in full here but rather sketch it. In this, nodes seek files by sending queries through the network using a “flood fill” algorithm – contacting other nodes they know – for a set number of “hops” (arc transitions). This is a decentralised network in that each node only “knows” about a limited number of other nodes and is not aware of the whole network (for example via access to a central server). To search the network a node sends its query to the nodes it knows and they pass it on to those they know and so on. This process continues until the queries have been passed on for a certain maximum number of hops at which point the relevant copy of the query “dies”.

If a node is currently sharing its files (it does not have to) and it happens to possess a requested file, that file is sent back to the originator of the query. Each node has a “satisfaction” level. When a node gets a file it does not have (and wants) its satisfaction level increases. Satisfaction decays exponentially at 10% per cycle. This means that nodes have to keep succeeding in their quest for files if they are to remain satisfied. If the satisfaction level of a node drops low enough then it will copy the connections and sharing strategy of a neighbouring node which is doing better (or with a low probability drop out of the network altogether and be replaced by a new node with a single random connection). This constitutes a social imitation process based on relative performance.

Unlike the network in the previous example, the network here is a directed one so that if node A sends queries to node B the reverse does not necessarily occur. Each node thus represents a person who is controlling some P2P software on the Internet. If the controller is unsatisfied with the number of files they are getting they may choose to imitate the way that another controller operates their software.

The general structure of the network that develops is illustrated in Figure 8 below. Due to the dynamics of the model, a core partition develops which is a single partition (that is there is a path from every node therein to any other following arrows from node to node) and the rest of the network, i.e. the periphery, mostly consists of branches that link into this core either directly or indirectly. There may also be one or more small isolated groups which quickly “die”. This is a result of the dynamic node

behaviour. If there were any nodes that were on a branch leading away from the core, these would not be viable in terms of file searching success and hence their satisfaction level would fall until they reset their connections to random ones thus “breaking” the branch. Of course, before this structure is established (and during transitory periods afterwards) different patterns may occur but the combination of core partition and periphery seems to be the “attractor” for the system.

Figure 8. An illustration of the typical network structure that results in the P2P model – there is a core partition that is totally connected and a collection of branches feeding into this. Arrows show the directions in which queries for files might be sent.

Although there is a lot of churn in the network in this model, it does seem to act so that the co-operators are better connected than the defectors. Table 1 shows some statistics for a typical run of this model over 600 simulation iterations of a 1000 iteration run. Uniformly the co-operators have a higher utility, average number of links and average centrality (an indicator of the nodes position in the group: 1 being the maximum and 0 the minimum indicating it is right on the periphery) than defectors and those in the core group also get higher measures than the equivalent kinds of non-core individuals.

Table 1. Average network properties by node type: network measures

Type	Average utility	Average number of links	Average centrality
<i>in-coop</i>	0.790	2.967	0.405
<i>out-coop</i>	0.512	2.500	0.309
<i>in-def</i>	0.373	2.005	0.269
<i>out-def</i>	0.324	1.492	0.189

However, when one looks at the patterns over time one does *not* see a good correlation of number of links or centrality score to utility (highest is 17% even over a variety of lags – see [6]). The measures for a single node are illustrated in Figure 9. This is because of the short term effects of defectors copying others with better links or more central contrasted to the inability of defectors to maintain such advantageous positions due to the co-operators they affect, relocating away. Thus the role of structure in this model is quite complex and certainly not a case of being able to effectively exclude defectors from one’s group.

Figure 9. The last 600 cycles of a simulation run for one particular node. The black line shows 1/3 of the number of outward links the node has. The grey line shows the measure of centrality: from 0 to 1 with 1 being the most central node. The light-grey line shows the level of satisfaction for the node when it is making its files available for sharing.

The mechanism is the same as the last example, but in a task which requires longer range cooperation and communication and with directional links. We did see cooperation here, but it is not clear whether this is due to the SLAC algorithm or the amount of “churn” it generates. A connected core+tails structure developed where the core provides most of the services to the rest and the tails use these. This structure seems to persist although its membership changes a lot over time.

Discussion

The co-evolution of structure and individuals

In the above we have seen how different “mechanisms” can result in different structures within an evolutionary process. What we do not see is the Darwinian evolution of these mechanisms so as to produce social structures useful to the

individuals who inhabit them (or otherwise). This would allow individuals to evolve a social “extended phenotype” along with their other characteristics.

Meaningful evolution of groups

In many of the examples above the structures that arise do undergo a sort of “life-cycle” and selection does indirectly and eventually occur upon individuals depending upon the viability of the structures they are in. However one is not justified in saying that the structures are evolving in any Darwinian sense. For the evolution of these structures more is necessary – new groups would have to be spawned from old groups and (at least some of) the characteristics of the old group would need to be transmitted to the new group it spawned. The transmission of these group characteristics would need to be very accurate, implying some sort of instruction or blueprint being transmitted, but with some small source of variation too.

However it is possible that the sort of mechanisms described in this paper could provide enough structure for such group evolution mechanisms to be implemented on top. Thus these sorts of systems might be a “stepping-stone” towards true group-level evolution. Even providing a low level of stability of groups might then make it possible for more sophisticated but culturally-driven structures to develop.

Conclusion

The broad conclusion of this work is that it is possible to encourage the “growth” of social structures. These structures are dynamic and fallible – it is always possible to circumvent and destroy them, but with suitable adaptive features built in they can be reasonably robust. In this paper we summarised work which sustains simple transitory groups for cooperation, clusters of related groups, isolated network clusters, and a core and periphery structure.

Although work of this kind can suggest possibilities of how individual abilities and social structure interrelate in observed societies (which may, or may not, turn out to be true), we anticipate that the real value of this kind of work is for informing the “design”, formation and management of complex, distributed IT systems. That such mechanisms are useful in such systems is already evident, for example the “reputation system” of eBay or the “tit-for-tat” systems in BitTorrent. Understanding some of the possibilities concerning what basic facilities/abilities/mechanisms enable the

development of “higher-level” social structures and institutions has the potential to multiply the usefulness of such systems. In Table 2 I attempt to summarise the mechanisms and resulting structures.

Table 2. A summary of the models

Model	Mechanism	Resulting Structure (under some conditions)
Integer Tags in PD games	Crisp Integer tags	Single-tag groups rising and falling in parallel
Floating tag model (Riolo et al model)	Floating tags forming single groups with a core of identical tags	A single tag-group with a core of identically tagged individuals
Symbiosis model	Floating tag and tolerances with a given specialisation of skills	Loose single tag group with sub-populations with periods of chaos
SkillWorld	SLAC/SLACER rewiring algorithm	High p : isolated totally connected groups. Lower p one totally connected group but somewhat locally separate
P2P	SLAC rewiring algorithm	A central partition and a periphery of connected trees of nodes

Of course, just because a system’s properties have been deliberately engineered into such systems does not mean that the dynamics that occur in such simulations will also occur in the target system. The continual investigation and re-modelling of such systems is necessary to understand them, given the ever-present possibility of emergence in such systems. However having a catalogue of some of the possible connections gives the best possible basis from which to do this.

Appendix

Support material associated with this paper is available as an online appendix at: <http://cfpm.org/~david/cmotappendix.pdf>

References

- [1] Axelrod, R. (1984) *The Evolution of Cooperation*, Basic Books, New York.
- [2] Barkow, J. H., Cosmides, L. & Tooby, J. (Eds.), (1992) *The adapted mind: evolutionary psychology and the generation of culture*. Oxford University Press.

- [3] Cohen, B. (2003) Incentives Build Robustness in BitTorrent. Presented at the 1st Workshop Economics of Peer-2-Peer Systems, 2003;
<http://www.sims.berkeley.edu/research/conferences/p2pecon/papers/s4-cohen.pdf>
- [4] Byrne, R. W. & Whiten, A. (eds.) (1989). Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans. Oxford Science Publications, OUP.
- [5] Dautenhahn, K. and Nehaniv, C. L. (eds.) (2002) *Imitation in Animals and Artifacts*. Cambridge, MA: MIT Press.
- [6] Edmonds, B. & Chattoe, E. (2005) When Simple Measures Fail: Characterising Social Networks Using Simulation. CPM Report 05-158, MMU. Presented at the Social Network Analysis: Advances and Empirical Applications Forum, Oxford, July 16-17 2005. <http://cfpm.org/cpmrep158.html>
- [7] Edmonds, B. & Hales, D. (2003) Replication, Replication and Replication - Some Hard Lessons from Model Alignment. *Journal of Artificial Societies and Social Simulation* 6(4). <http://jasss.soc.surrey.ac.uk/6/4/11>
- [8] Edmonds, B. (2006) The Emergence of Symbiotic Groups Resulting From Skill-Differentiation and Tags. *Journal of Artificial Societies and Social Simulation*, 9(1). (<http://jasss.soc.surrey.ac.uk/9/1/10.html>).
- [9] Hales, D. (2006) Emergent Group-Level Selection in a Peer-to-Peer Network. *Complexus* 2006;3:108-118
- [10] Hales, D. & Edmonds, B. (2003) Evolving Social Rationality for MAS using "Tags", In: Rosenchein, J.S., et al (ed.) *Proceedings of the 2nd International Conference on Autonomous Agents and Multiagent Systems*, Melbourne, July 2003 (AAMAS 2003), ACM Press, 497-503.
- [11] Hales, D. (2000), Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In Moss, S., Davidsson, P. (Eds.) *Multi-Agent-Based Simulation. Lecture Notes in Artificial Intelligence*, 1979:157-166. Berlin: Springer.
- [12] Hales, D. (2002) Wise-Up! - Smart Tag Pairing Evolves and Persists The Centre for Policy Modelling, Manchester Metropolitan University, UK. Discussion Paper CPM-02-90. (<http://cfpm.org/cpmrep90.html>)
- [13] Hales, D. (2004) From Selfish Nodes to Cooperative Networks – Emergent Link-based Incentives in Peer-to-Peer Networks. In *proceedings of The Fourth IEEE International Conference on Peer-to-Peer Computing (p2p2004)*, 25-27 August 2004, Zurich, Switzerland. IEEE Computer Society Press.
- [14] Hales, D. (2005) Change Your Tags Fast! - a necessary condition for cooperation? *Proceedings of the Workshop on Multi-Agents and Multi-Agent-Based Simulation (MABS 2004)*, LNAI 3415, Springer, 2005.
- [15] Hales, D. and Edmonds, B. (2005) Applying a socially-inspired technique (tags) to improve cooperation in P2P Networks, *IEEE Transactions in Systems, Man and Cybernetics*, 35:385-395.
- [16] Hales, D. and Arteconi, S. (2006) Article: SLACER: A Self-Organizing Protocol for Coordination in P2P Networks. *IEEE Intelligent Systems*, 21(2):29-35.

- [17] Holland, J. (1993), The Effect of Labels (Tags) on Social Interactions. SFI Working Paper 93-10-064. Santa Fe Institute, Santa Fe, NM. (available at: <http://www.santafe.edu/sfi/publications/wplist/1993>)
- [18] Margulis, L. and Sagan, D. (1986) *Microcosmos: Four Billion Years of Evolution from Our Microbial Ancestors*. Summit Books, New York.
- [19] Norling, E. and Edmonds, B. (2006) Why it is Better to be SLAC than Smart. 1st World Congress on Social Simulation (WCSS'06), Kyoto, Japan, August, 2006.
- [20] Reader, J. *Man on Earth*. London: Collins, 1988.
- [21] Riolo, R. L., Cohen, M. D. & Axelrod, R (2001), Evolution of cooperation without reciprocity. *Nature*, 411:441-443.
- [22] Roberts, G. and Sherratt, T. N. (2002) Similarity does not breed cooperation. *Nature* **418**:449-500.
- [23] Sigmund, K. and Nowak, M. A. (2001) Evolution - Tides of tolerance. *Nature* **414**:403.
- [24] Wilson, D.S. & Sober, E. (1994) Reintroducing group selection to the human behavioural sciences. *Behavioral and Brain Sciences* **17**(4): 585-654.

Author affiliations

- Bruce Edmonds: Centre for Policy Modelling, Manchester Metropolitan University, Aytoun Building, Aytoun Street, Manchester M1 3GH, United Kingdom, bruce@edmonds.name
- Emma Norling: Centre for Policy Modelling, Manchester Metropolitan University, Aytoun Building, Aytoun Street, Manchester M1 3GH, United Kingdom, norling@acm.org
- David Hales: The University of Bologna, Dept. of Computer Science, Mura Anteo Zamboni 7, 40127 Bologna, Italy. dave@davidhales.com