10. Strategic knowledge sharing: a small-worlds perspective

Mike Metcalfe, ***School of Management, University of South Australia***

Abstract

This paper is about designing knowledge sharing in wicked systems. The perspective the paper takes is that of the self-organising 'small-worlds' phenomenon. Specifically, this paper argues that strategic knowledge sharing can be viewed as designing small-worlds networks so as to allow a wicked socio-technical system to self organise a coordinated strategic response to unpredictable environmental changes. The evidence used comes from the softer systems literature, biology (insect) literature and social-network literature.

Introduction

Centralised governance of effective knowledge sharing is very difficult in times of rapid change, especially for purposeful, information rich, socio-technical wicked systems. The lines of communication quickly become clogged, leaders suffer information overload and are unable to fully appreciate problems at the local level. Decentralisation of knowledge sharing runs the risk of causing local overload, with key information not being prioritised or depending on actors who only have experience at processing local problems. Alternatives such as 'middle-out' (Keen, 1999) have been suggested, where strategically informed middle level actors play a coordination role between the top and bottom level actors. This paper explores an alternative, using the small-worlds phenomenon, which is itself seen as a self-organisational response that enables actors in a wicked and dynamic socio-technical system to share knowledge and thereby generate an effective strategic response to environmental surprises.

For those who are concerned about the deep-rooted assumption that all socio-technical systems need a hierarchy to become organised, this paper can be seen as a small contribution to the anti 'hands on' top-down view of leadership, where a 'John Wayne' figure leads the herds of awestruck battlers through some life-threatening disaster. Rather, leadership is seen as a socio-technical system that is capable of allowing knowledgeable actors to interact strategically, as they see the situation, using their different experiences. To those who have some appreciation of the very limited impact even caring 'hierarchical leaders' can really have on the activities of any complex system such as regional government, this paper may provide some improved sense of the complexity of leadership in these dynamic situations. That said, this paper is not primarily about how to organise a response, but rather how to envisage a self-organising socio-technical wicked system. Examples of wicked problems in which this self-organisation design is believed to be required include broad area wildfires, rapidly evolving environmental disasters such as the one outlined below, blitzkrieg warfare, industry reorganisations and national IT policy in recent times.

Wicked problems and wicked systems

A wicked problem is defined by Rittel and Webber (1973) as one characterised by the following:

1. There is no definitive formulation.
2. Any solution is not true-or-false, but rather good-or-bad.
3. There is no immediate and ultimate test of any solution.
4. It is a 'one-shot operation' since there is no opportunity to learn by trial and error.
5. There is no enumerable (or an exhaustively describable) set of potential solutions, nor is there a well-described set of permissible operations that may be incorporated into the plan.
6. It is essentially unique.
7. It can be considered to be a symptom of another problem.
8. It results from a discrepancy that can be explained in numerous ways, and the choice of explanation determines the nature of the problem's resolution.
9. It does not allow the planner the right to be wrong.

Item 8, where the explanation of the problem is crucial to perceptions of how it can be resolved, is central to the design thinking in this paper. A wicked socio-technical system is, therefore, one that is made up of people, supported by technology, who appreciate that they are dealing mainly with these wicked problems. This includes most human organisations. Although not explicitly mentioned in the list above, these problems are also dynamic; they change over time due to a mixture of events, including new technology, new knowledge and possible shifts in participants' perspectives. The design task is, therefore, to allow actors in a wicked strategic problem situation to self-organise in order to produce what they see as an acceptable resolution.

This paper will discuss self-organisation and small-worlds phenomenon, providing two examples of wicked problems.

Self-organisation

The concept of 'self-organising systems' is in danger of losing its effectiveness and becoming as vague a term as 'general systems theory' and 'autopoesis', with their abstract talk of generic open and closed systems. For example, Georges and Romme (1995) define a self-organising system as one that is both open and closed; evoking the old debate about what constitutes a closed system. Mingers (1997) definition of *autopoesis* (self-reproducing) would appear to subsume self-organising systems. However, doing so may hide some of the advantages of using the perspective of self-organising systems for automatic knowledge sharing when wicked systems pose problems that need to be solved. In order to be able to reproduce, a system needs to be organised, and it may be hierarchical. A self-organised system, however, is one that does not need a hierarchy to respond to environmental surprises. It is the assumption of the need for a hierarchy to direct knowledge sharing that is being challenged here. Self-organisation is not being used here in the sense of the development of identity in a hostile environment, (such as the establishment of the early Christian Church, the labour movement or the feminist movement). Rather, this paper is concerned with how a wicked socio-technical system might be designed to share knowledge so as to provide an effective response to environmental surprises when there is no explicit internal hierarchy. Ideas about how these systems might be designed come from analogies with the world of insects. Some swarm, as in ant nests, and some bee nests have no boss, no corporate plan, and no strategic planner, but a higher level organisation has *emerged* that serves to enable the unsuspecting insects to make a strategic response to unpredictable large scale problems that suddenly impinge upon their world.

There is an extensive literature (e.g. Mingers, 1997) on self-reproducing, self-replicating, and similar systems. This paper will bypass revisiting these and merely synthesise from two other related areas. The first is the empirical scientific biological literature about what insect colonies do to share knowledge to provide an effective strategic response to problems. The second is the small-worlds literature, which has recently moved from the sociometric to the sciences, as more and more biological systems are seen to use the small-worlds structure to share knowledge. These will be discussed in terms of a story from the crisis management literature, which tells what actually happens in response to a rapidly changing, community-based problem, in particular when the strategic response has voided any pretence of controlled top-down knowledge sharing.

The insect literature

There has been a lot written about self-organisation in the biological and related sciences. Much of this literature is presented as a mathematical analysis of patterns that emerge, e.g. waves, sand dunes, tree structures and the markings on animals. Camazine et al. (2001), however, provide an empirically based explanation using insect systems. This paper's interpretation of what is meant by self-organisation draws heavily on this, and thus draws on analogies from the world of the insect nest. Camazine et al. (2001) observe that some complex actions emerge through simple interactions internal to the system, without intervention by external directing influences. More formally they define self-organisation as:

> … a process in which pattern at the global level of a system emerges solely from numerous interactions among the lower level components of the system. Moreover, the rules specifying interactions among the systems components are executed using only local information. [p8]

Camazine et al. (2001) do not accept that the queen insect in an ant's nest or beehive is somehow 'giving instructions' to the millions of insects who have never been near her. The term 'queen' is misleading; the term 'womb' would be more acceptable from a knowledge sharing perspective. Each individual ant or bee bases its behaviour on its perception of the position and behaviour of its nearest neighbour, rather than on knowledge of the global behaviour of the whole group. Local dynamic knowledge sharing is all that is present, yet the insects are able to make strategic responses to a global threat to the whole nest. A strategic response somehow 'emerges' from lower level actions, evidenced by the very existence of a nest that has specialised integrated operations. The individual ants, for example, are not even thinking about this higher order purpose; rather they are only concerned with their own small function in the nest. If this emerging strategy appears different to the actions of lower level activities, then the system may be described as complex. Individual ants forage for food, build the nest, care for the eggs and milk the queen, yet somehow these activities have become coordinated to produce a species that has survived, and very successfully, for millions of years.

Camazine et. al. (2001) summarise the now significant amount of empirical research that has been conducted on insect nests to better understand how a strategic response can emerge. For example, a few ants placed in a Petri dish were found to move sand around in a random fashion, achieving nothing. But when enough ants were added, the probability of the production of a randomly constructed shape that the ants recognised and would respond to, increased. The presence of these particular shapes then acted to suddenly start the ants working in a coherent manner, constructing recognisable elaborate structures. In another example, an ant's nest was deliberately damaged and metal plates used to divide the damaged area in such a way that knowledge sharing between the two damaged areas was impossible. The strategic response, the reconstruction, matched perfectly. When the dividing plate was removed, the rebuilt sections looked like one singular rebuilding exercise, perfectly orchestrated. In summarising this empirical literature, Camazine et al. (2001) identify a series of conditions necessary to enable the emergence of a knowledge sharing system from the insect activity that results in a coherent strategic response. These include the presence of:

1. group influence;
2. stigmergy;
3. decentralised control, dense heterarchies; and
4. dynamic knowledge sharing.

Group influence

Camazine et al. (2001) do not clearly label this attribute of a self-organising group; rather they sum it up as 'I do what you do'. The idea starts with noticing that members of a group copy or mimic those around them; they are influenced by the actions of others. Children do what their parents do, artisans learn from their masters, business schools teach the 'echo of lies' of how management is done, and when at work we learn a corporate culture, we become team players. We learn the preferred way of doing things if we want to 'get along'. Examples of our compliance to our local group norms include our dress, religion, food and ethics. However, we can from time to time insert some small minor variation based on experiences we have had elsewhere. This is analogous to our genetic make-up; we are only minor variants of our parents, but we are variants. An invention, a new recipe or a clothes fashion change are examples of an individual changing a group's behaviour, but if we are honest, one person usually makes very little difference to the generic behaviours of a community. This 'get along, go along' behaviour seems related to our very strong 'inclusion' needs; we need to belong to a group. Horses are trained by threats to exclude them from the herd, which is far more sustaining as a threat than physical pain. Arguably, the worst punishment we inflict on other humans is solitary confinement. The need to belong is seen as an explanation of why herd species and insect colonies are influenced as they evidently are by the behaviour of the whole group; expressed as 'I do what you do'.

Being influenced by the behaviours of others, especially those immediately around us, is central to self-organisation. An insect seems to be born perceiving that the world will be intimately integrated with what the insects colony around are doing. An ant will merely do what the ants immediately around her do, using whatever genetically received devices she has at her disposal. More empirical evidence of this from insect research includes the behaviour of fire flies. When swarmed, fire flies, with their flashing tails, will all end up synchronising their flashes. The fire flies will alter their flash time and speed under the influence of the group. Infectious yawns, synchronised reproductive cycles, synchronous breathing and 'mobbing' are all examples of human group behaviour that influence individual behaviour. Wilson (1983), giving the example of a librarian thinking about the demand for books, uses the term 'cognitive authority' to identify who of those around us we choose to mimic. In an insect colony it is assumed the individual insects can only choose to mimic, to listen to, those immediately around them. Modern people, who have access to the media, books and different corporate cultures, and are able to travel to numerous different communities, have a much wider choice of cognitive influences to mimic.

In this social setting, knowledge sharing between those in immediate contact is expected to have already largely occurred. When a crisis occurs, more than one insect knows the same things.

Stigmergy

Camazine et al. (2001) are very cautious about the idea that colonies of insects carry in their heads a detailed recipe or fully laid out blueprint of what, for example, a nest should look like; a detailed vision of what the finished construction should do and be. This is justified with the empirical evidence for how nests respond to different physical situations. The insects build allowing for the physical conditions encountered, so every nest is slightly different. Yet overall common design features are observable. This is not attributed to the insects' knowledge sharing, but to their merely responding with a set of alternatives.

Stigmergy is a term attributed by Camazine et al. (2001) to Grasse. It refers to the mechanism whereby a swarm insect (such as an ant or bee) is stimulated to work constructively towards a common purpose by the presence of work in progress. The half completed work of other similar insects is recognised as an 'event' that induces automatic *responses* from those that see and recognise it. For example, an ant may see a pair of pillars and respond by building an arch between them, without having communicated directly with the earlier builders. This is an indirect form of knowledge sharing; an event is driving asynchronous knowledge sharing. The human equivalent may be the response of rescuers when a building is seen to collapse or a child is seen to be treated badly. In place of stigmergy, Michener uses the expression, 'indirect social interactions'. In systems management this may be called asynchronous knowledge, or sharing through design. For insects, it may be the result of the quantity of pheromones on the half finished building works, or it may be the physical shape that acts as the asynchronous stimulus.

It is possible to appreciate the importance of asynchronous knowledge sharing to the running of a complex system (one that has emergent properties), such as an ant nest, by drawing on the analogy of a modern corporation. The existence of multinational corporations has been attributed to asynchronous knowledge sharing of faxes, (e)mail and web pages. The size of organisation that can be controlled by means of oral knowledge sharing, is restricted. In large organisations, while oral synchronous knowledge sharing remains very important, time zones, legal records and very detailed specification require 'written', asynchronous knowledge sharing. Nations that have developed joint synchronous and asynchronous knowledge sharing have been dominant in economic and scientific terms. At a more modest level (in insects) Camazine et al. (2001) are suggesting a more subtle form of asynchronous social interaction, and thus motivation in the presence of half completed tunnels, pheromone paths and other work in progress.

Given the centrality of purposeful activity to systems thinking and design (Ulrich, 2002; Checkland, 2000), it seems necessary to mention that purposeful activity is presented by Checkland and Ulrich as an emergent property from the large, self-conscious human brain. This is thought to enable us to appreciate the purpose (drivers) of our actions and stand outside ourselves. But should insects be thought to be engaged in purposeful activity, rather than operating like parts of an alarm clock? Are insects living out genetic drivers to bring up young and continue the gene pool? Surely any human self-organising system would need to anticipate that the participants would be able to ask themselves why they should act. Moreover, in human systems, language could be used to provide a driver to act. Therefore, in human self-organising systems it may be necessary to emphasise why people should act if it is not otherwise obvious to them that they should. However, Camazine et al. (2001) understate the influence of purposeful activity, even those of genetic survival, for insects compared to group interaction influences of humans. They place much more emphasis on the insects' response than to their driving forces for achievement. The genetic drivers of gene survival are not emphasised, possibly because most of the insects never see or come in contact with the young.

Decentralised control and dense heterarchies

The decentralised control attribute identified by Camazine et al. (2001) is defined as a particular 'architecture of information flow'. Each insect responds to other insects immediately around it to learn what is to be done, rather than from messages from well-informed individuals (leaders) in the upper echelons of a control hierarchy. The organisation chart is one of small clusters of interacting insects responding to one stimulus, such as a half built arch in one part of the nest, or a food retrieval clique at another location in the nest. There is no tree of hierarchical knowledge flowing up and down; rather the structure is more a series of independent clusters of workers who, ninety percent of the time, only communicate directly with the other members of their cluster (described as cliques, or *small-worlds*). Only when they are unable to solve a problem with local knowledge sharing will they venture out to ask another cluster. The dense heterarchies attribute reinforces the image of a series of separate yet connected small clusters, each focusing on different but loosely interconnected tasks. Heterarchies are inter-independent groups; they are neither hierarchical and nor totally independent clusters. This raises concerns about how a strategic response from these roughly independent responding clusters is possible. The small-world literature may help here.

Small-worlds

The previous section briefly introduced some of the findings from empirical biological research as presented by Camazine et al. (2001). The next thread of the synthesis presented in this paper is that derived from the sociometric literature. In order to develop some appreciation of the knowledge sharing system considered so central to insect nest life, it is perhaps necessary to discuss this literature, or at least one part of it: the small-worlds literature.

The small-worlds phenomenon emerged as a result of Stanley Milgram's experiments (Milgram, 1967), and was later captured in the play and film 'Six Degrees Of Separation' (Watts, 1999). The idea proposed is that it appears that any two people picked at random are connectable via a chain of, on average, six intermediate acquaintances. This seems counter-intuitive given that, for most of us, frequent direct two-way conversation only occurs with fewer than 20 people; our small- world cluster. In sociometric network terms, this suggests an overall population network that can neither be described as 'everyone knows everyone else', nor, at the other extreme, one where local clusters of socially interactive persons have no means of contacting other clusters. The reality is a mix of the two, with imperfect knowledge shared between clusters. In rather simplistic terms, a small world network can be illustrated as in Figure 10.1. If one cluster were pressed to send a message to another, it would be possible to find a 'weak link' between the clusters; someone who can carry a message between them. In Figure 10.1, if A wants to send a message to F, whom he does not know, then first he would ask friends in cluster 1. B says she knows someone, C, in cluster 2, who may be able to pass the message on. When C gets the message, she asks her friends, and D suggests E, who does know F.

Solid lines represent knowledge sharing between people who talk to each other very frequently and dotted lines are between people who talk very infrequently.

Figure 10.1. Small-worlds.

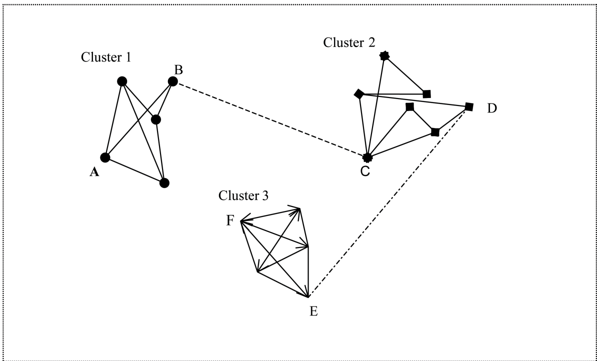

The small-worlds phenomenon provides a way of seeing knowledge sharing between small groups of ants working one particular project cluster, and occasional sharing with other groups of ants working within a different project cluster. It is perhaps predictable from the knowledge-sharing theory literature (Hare, 1976), which highlights that we can only have direct two-way knowledge sharing with a limited number of people. This is due to the exponential growth of knowledge sharing channels as the number of people involved increases. When three people wish to communicate with each other freely, there are only three knowledge sharing channels that need to be kept open (A to/from B, A to/from C, and B to/from C). For four people there are six, and for five there are 10 channels that need to be serviced. For people, that may mean exchanging pleasantries, as well as being able physically to get to and from the others at the same time and in the same place. Having to service a lot of channels becomes time consuming.

So with an ants' nest it is possible to imagine a situation where an ant responds to the ants immediately around it, obeying self-organisation driver number one (I do what you do), and joins in doing whatever they are doing; for example building a new passageway. When a problem arises with the harmony of this activity, no individual in the ant's immediate cluster knows what to do. One of them communicates to an ant nearby that was not involved in the passage building, an ant from another cluster it only knows weakly. The weakly known stranger may communicate that it is very busy collecting food. This then stimulates the passage builders to start collecting food. When a crisis occurs in this task, the ants look around for previously weak messages about other tasks.

The small-worlds research (Killworth and Bernard, 1979; Watts, 1999; Matsuo et al., 2001; Richardson and Lissack, 2001; Buchanan, 2002) extends the social network research (e.g. Mizruchi, 1994; Scott, 1996; Durrington, 2000; Cross et al., 2002) by suggesting what network shapes have self-organised in human groups, in the natural environment, in written communication and in biological systems. Management literature increasingly regards knowledge management as a social networking problem. Hansen (1999), Roubelat (2000) and Reagans and McEvily(2003) have studied management issues related to knowledge sharing using weak links.

To summarise, the discussion above suggests that small-worlds networks allow for effective knowledge sharing both in times of routine and when a strategic response is required. This further suggests that anyone responsible for designing the knowledge-sharing network in their organisation might use this lens to evaluate their communications systems.

Examples

Two very simple examples are now discussed. The first is an account of the response to a community level problem, with special emphasis on knowledge-sharing issues. Hopefully, the analogy to ants' nests and small-worlds phenomena is apparent.

Comfort (1994) argues that the citizens' response to a oil spill near Pittsburgh in 1988 was a self-organised one, as the situation developed too rapidly and was too complicated for a simple top-down leadership response. She reports that the crisis began with a four million gallon diesel fuel tank collapsing. This resulted in a seventeen-mile-long emulsified oil and water mixture flowing down and over the locks and dams of the Monongahela River, extending bank to bank. The river provided drinking water for the Pittsburgh metropolitan region but the risk of damage to water filtration systems made the water authorities shut down the water intakes, resulting in a lack of water for either drinking or fire suppression. For two weeks alternative arrangements had to be organised, requiring the coordination of 25 different types of organisations – public, private, and community non-profit. The zoo, the fire service, medical services, the coastguard, hazard waste services, car washes, and bottled water companies all had to be coordinated.

One can easily imagine groups of concerned persons establishing informal clusters around their particular concern, or expertise. The bottled water people may be one cluster, the fire services another and so on. Most of their knowledge sharing would be within their cluster, perhaps on a one-to-one basis. Every now and then one of these clusters would need information from another concerned cluster; for example, the bottled water people might need an estimate of how long the crisis would last, or need to know how to get access to extra transport, or bottle manufacturing facilities. They would then use their 'weak link' to make contact with another cluster, as they would need to keep an overall appreciation of what was going on. The whole system would only work if there were both the locally knowledgeable clusters and the presence of weak but effective inter-cluster links.

One knowledge-sharing centre handled an estimated 37 000 thousand incoming and outgoing messages during the crisis; averaging 154 per hour, 24 hours a day, seven days a week. This would have been a fraction of the knowledge sharing involved. Comfort (1994) emphasises the need for a dynamic decentralised information system, one that provided up-to-date local and overall information as the situation changed, able to record messages asynchronously and then supply relevant messages to inquirers at a later time. The danger was that critical information would not be stored and located effectively and so not be correctly identified due to the sheer mass of messages generated. Achieving this was not possible through one hierarchical knowledge sharing hub. A self-organisation or small-world knowledge-sharing system was required; one that needed to use human memory and awareness.

Another simple but familiar wicked problem example may help. A university is made up of numerous groups undertaking research in their own discipline area. This typically involves small groups ranging in size from perhaps only two to laboratories containing 10 to 20 or more members. These groups know much the same 'stuff', the discipline-specific research methods, the literature and the worldwide experts in their field. View these research groups as small-worlds knowledge clusters. The strategic imperative, or common purpose some of these wicked system clusters may appreciate is the need for multi-discipline research to provide a comprehensive research effort to deal with wicked problems such as poverty, terrorism or natural disaster response. This common purpose may spread through the weak inter-cluster links. The Research Dean may also further encourage this concern by allocating increased resources to multi-discipline research solutions. Each cluster has specialist knowledge and any excessive attempt to insist all its members spend a significant proportion of their time getting to know other cluster's research in detail may distract them from developing their own knowledge. However, these clusters do need to be 'weak linked' both with each other and with clusters knowledgeable about research resources. These weak links will need to be synchronous and asynchronous (stigmergy), using web pages, internal research newspapers, publication listings, signage, question-asking software like 'askme.com', web publishing of seminar PowerPoint slides, financial rewards and Listserv public acknowledgement of achievements. All of these are examples of the asynchronous (stigmergy) weak linking. Telephone lists by knowledge area, cross discipline coffee groups and conferences are examples of synchronous weak linking. The role of the Research Dean is merely to provide effective responses to those that do multi-discipline research compared to those that do not, and to encourage weak linking (not strong linking) between groups that would not normally even appreciate each others existence. Given the common purpose and the presence of weak links, the self-organisation perspective anticipates that members of the clusters will knowledge-share and self-organise an appropriate response for the improvement of the knowledge holding of the entire university.

Implications and conclusion

There is not much new about many of these activities, perhaps because weak linking across clusters is naturally efficient, an unappreciated theory in use (Argyris and Schon, 1978). However, this paper has attempted to make this theory explicit, and provide a clearer picture that makes sense of the phenomena of interest. This, it is hoped, will make the governance of strategic knowledge sharing more explicit. Given the complexity of wicked problems and the creativity needed to respond to them, hierarchical control of either participant's actions or of their knowledge sharing is considered naïve. The exact opposite of a 'need to know' knowledge sharing policy is required. Participants are to be encouraged to decide for themselves what they need to know and to be aware where they can get that information easily, as in the oil spill example.

This paper has argued that strategic knowledge sharing can be viewed as a task in designing small-worlds networks so as to allow a self-organised strategic response to wicked problems. Knowledgeable clusters and synchronous or asynchronous weak links can both identify wicked problems and respond strategically. A designer of this sort of network needs to encourage knowledgeable clusters that are only weak linked together, whether synchronously, asynchronously or both. Perhaps, for a commercial organisation, the designer may allocate resources to encourage a particular common purpose. Two examples of wicked problems being handled by wicked systems were outlined. In one, the wicked problem was a rather obvious oil spill threatening many dimensions of a community. In the other, the wicked problem required innovation and creativity. It is argued here that neither should be managed in a hierarchical sense.

Future research may continue the work to make explicit the design of existing social networks in public and research communities, so as to better appreciate if they follow a small-world structure. The limits of synchronous and asynchronous 'weak links' may also warrant further investigation, as might the role of information and communication technology from this point of view. Other suggestions may include: how people communicate, and with whom, when involved in a crisis; and what information they need or could not obtain. How the self-organisation process works also needs to be made more explicit.