STUDIES IN EMERGENT ORDER VOL 2 (2009): 1-18

Limits on the Communication of Knowledge in Human Organisations

Jacky Mallett

Introduction

Communication is ubiquitous. It is the invisible thread that binds all attempts at human organization together. We use communication to give orders, make suggestions, laboriously acquire and pass on knowledge; but rarely stop to think about the mechanics of the process of communication itself. All forms of co-operative human activity though, depend on communication through their success, or failure.

Hayek was notable among economists in considering the economic implications of the communication and possession of knowledge within society. He recognised the dependency of economies on the actions of multiple agents, none of which could possibly have complete knowledge of their mutual endeavour. He recognized in particular the impossibility of the condition that total knowledge of an economy could be known at a central point, the requirement that is implicit in the idea of centrally planned economies.

Knowledge as such, although very real, is hard to quantify. Its communication however is not. Constraints on communication in terms of bandwidth, the amount of data that can be transmitted; latency, the time it takes to transmit it; and message processing, the amount of effort or work receipt of the contents of the message incurs, have become well understood within the fields of real time data networks, and distributed computing. The successful construction and operation of complex, world spanning communication systems, the phone and data networks in particular, are the results of this understanding.

Communication in data networks is handled as a problem of message exchange between independent network nodes. Data networks are organised collections of nodes, controlling communication links that can now route traffic between hundreds of millions of end computers. The end computers in their turn run individual instances of distributed applications that co-ordinate their results over large numbers of computers by exchanging messages with each other, via the underlying communication network.

Communication within human society, can also be viewed as a problem of message exchange between independent agents, or people. Human society is organised as collections of people, working co-operatively to achieve co-ordinated results, using the exchange of messages between individuals to facilitate their activities. In many respects, with appropriate adjustments made for lower speeds and rather less reliability, the field of data communication networks provides a wealth of theory and experience for understanding the communication constraints that operate upon human organisations.

In this paper, we will discuss the effects of communication constraints on human social and economic organization, from the perspective of some fifty years of engineering and scientific endeavour in real time computer networks. From this perspective it is straightforward to prove that Hayek was correct about central planning, and that for large communicating systems, it is provably impossible to perform this function with anything approaching complete knowledge. Similar considerations of the arrangements of the communication paths between nodes, their topology as it is referred to in computer networking, play out in all forms of organized activity, with an opposing tension between the topology most suitable for efficient distribution of control or command information – the centralised, strictly hierarchical organisation - and that which provides the largest capacity for the distribution and sharing of information or knowledge, an organized partial mesh.

This allows us to add to Hayek's explanations of spontaneous or emergent order, an analytical explanation for some of their features in terms of the restrictions they place on the communications that can be performed by their constituent members, and the source of some of their advantages and disadvantages.

Information Capacity

Claude Shannon's paper, A Mathematical Theory of Communication[1], published in 1948, is generally credited with providing the foundations of the field of information theory and digital data communication. In it he formalized several key concepts critical to the development of communications theory, amongst them the idea of information as a single, unique message. This definition of information was important, since it introduced the idea that there is a quantitative difference between sending the same message to a number of recipients, or sending each of them a different message. Shannon's paper was a discussion of the mathematics of data transmission and receipt, with respect to how signals could be compressed in order to make communication more efficient. (At that time, low bandwidth and error prone communication links were the norm.) In this context the presence of repeated identical information allows for considerably more compression of the signal, and hence faster communication, than would otherwise be the case.[1]

Shannon quantified the amount of communication that could be achieved between a single sender and receiver in terms of the amount of information that could be transmitted on that channel during a specified time interval, i.e. the number of unique messages that could be sent. He showed that given the latency characteristics of a channel, where latency is the time taken to transmit the message, that there is a limit on the total amount of information that could be transmitted between two nodes in a given time period. Nodes may have more or less information than that available for transmission, but if one node tries to send more than the channel can provide the capacity for, then some messages must be dropped.

When larger numbers of nodes are connected together, other communication constraints emerge that are created by connectivity constraints between the nodes. Not all nodes may be directly connected to each other, and messages may have to be relayed by intermediary nodes. This reduces the total message carrying, or as we will term it, information capacity of the network. However, in real time networks with very large numbers of nodes it is hard to exactly characterise this impact. Recently however, an upper limit on the amount of information that can be transmitted by meshed communication networks has been derived as part of work in the sensor network field.

The idea that there are limits on the total amount of communication that can be performed by a set of nodes, is also the crux of Hayek's argument against central planning. Hayek presented it as an argument that economies are necessarily dependent on independent, imperfectly informed agents working on local knowledge[2]. It was apparent to Hayek that the form of conscious control represented by central planning was not capable of co-ordinating the day to day activities of individuals in something as complex as an economy. By recasting this as a problem of real time message transmission, receipt and processing, we can show that this is indeed the case.

To demonstrate a simple example of these issues, consider a typical 1 hour project meeting for a large project co-ordinating many people's efforts towards a shared goal. Each "node" present in the meeting both increases the amount of potential knowledge about the project that can be contributed to the meeting, but also reduces the average amount of time available for that contribution. Nor is this solely a property of the broadcast communication method being used to communicate at the meeting, it is similarly a fundamental property of the physical topology used to connect the nodes, in conjunction with the real time constraint on communication imposed by the meeting's length.

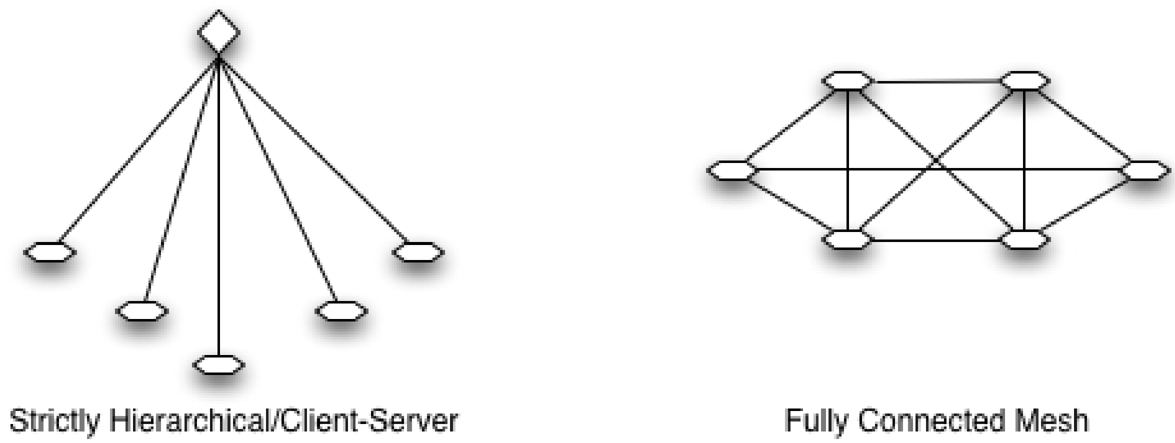

Figure 1. Communication Topologies

Consider the case for networks where all nodes are in direct communication with each other, in network parlance connected by a single hop, as shown in Figure 1. There are two distinctive topologies: the client-server or strictly hierarchical, and the fully connected mesh topology where each node has connections to every other node in the group. In a client-server topology, all nodes in the network communicate directly with a central node, so that any communication between clients also has to pass through the server, even if the central node has no need for the information in the message. The upper limit on the total amount of communication that can be performed by the network is thus constrained by the communication capacity of the central node. This is directly analogous to the topology of central planning, and the bottleneck on economic organisation that it can create.

In a fully connected mesh topology however, the communication capacity of the group of nodes is a function of N, the number of nodes in the network, since each node is connected directly to every other node. The upper limit on communications is the number of nodes that each node can connect to. However, direct messages between pairs of nodes can occur in parallel, unlike the client-server topology. This topological limit on one-hop connectivity is something we will refer to as the group size limit.

The significance of the group size limit is most apparent when group sizes are constrained to small numbers of nodes, such as in human organisations, than for distributed computer applications. Specialised equipment now allows some computer based client-server applications to scale to millions of nodes, albeit only for applications, such as credit card transactions, where client-server communication is well defined, extremely simple and there is little or no client-to-client messaging. In activities that require a continuous exchange of information in order to achieve co-operative results, achievable human 1-hop group sizes generally seem to range up to about 10. Scaling above that size requires that communications between individuals that are not directly connected must be relayed by intermediaries that are.

Both mesh and hierarchical topologies can be extended to scale beyond single groups, in terms of providing connectivity. That is, they can both provide message routes that provide a communication path between all nodes in the network. What neither can do is allow all nodes to communicate with each other simultaneously. The existence of a message route between nodes, and the ability of that route to deliver messages in real time are two very different things. So differences in information capacity that occur when the same number of nodes are arranged in different topologies are of interest, since this implies that they place a fundamental constraint on the amount of messages, or communicated knowledge in some sense, that the topology can accommodate.

Theoretical understanding of the behaviour of traffic within large packet switched networks has lagged behind the ability to build them. In particular, the limits on the amount of traffic that can be carried within the network have not been well defined.[2] One result on capacity limits on information capacity of large, partial mesh networks, has recently been provided as a result of work in the field of sensor networks by Gupta and Kumar[2], and refined by Scaglione and Servetto[3]. Since Scaglione's proof is extremely elegant, and also illustrative of the set of issues being discussed, we will reproduce it near verbatim changing only the descriptive node, to people.

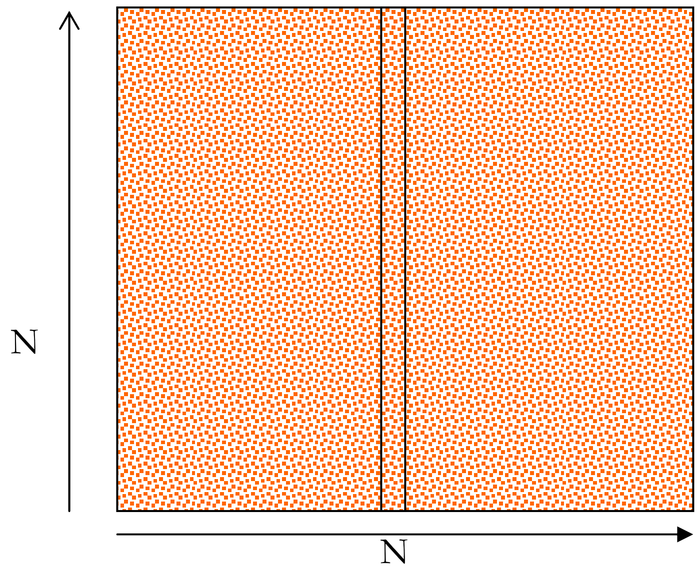

Figure 2: The Capacity of Mesh Networks

In Figure 2, N individuals are spread uniformly over a unit square, [0,1] x [0,1] (N large). Take a differential area of size Δ (Δ small). With high probability the approximate number of people in the strip, shown in the figure is NΔ. Since the total number of people is N, we must have N = NΔ x NΔ (because the total area of the unit square, is the product of the areas of two strips as shown in the figure, one horizontal and one vertical), and hence the number of people in a strip as shown is $N\Delta \approx O(\sqrt{N})$ for sufficiently large N.

In the original sensor network problem discussed by Scaglione, all the nodes in the network were required to receive information from all the other nodes on a continuous basis, and the question posed by the paper, to which the answer was clearly no, was is this in fact possible? In the example here, we will impose a less demanding requirement, that there is a single planner node somewhere in the network, that must receive all information from all the other nodes in order to create a perfectly informed plan. In order to reach the planner, the information must flow through the nodes in the strip. The question is, can the network in fact transport this much information?

The identical argument used by Scaglione applies here. The people present in this network have a limited number of other people they are each in direct contact with, we will call this L – their link capacity. Consequently, the maximum capacity of the strip cannot be greater than $O(L\sqrt{N})$. From the max flow/min cut theorem we know that the capacity of any network cut of this form is an upper bound on the network capacity, and so the total transport capacity of this network cannot be higher than O(L√N). In order for the central planner to have complete knowledge on which to base his plans, information must be obtained from every person in the network, so O(N) messages must be received by the planning node. But for large numbers of people this considerably exceeds the O(L√N) capacity of the network. So the central planner cannot have complete information of the network state on a continuous basis.

This is not to say that the network is incapable of obtaining information from every person, and transmitting it to a single point. The problem is that in real time, economic agents must be assumed to be continuously changing state and consequently producing information updates which have to be collected on a continuous basis. Once the size of the organisation has gone above the group size limit of its topology, this is no longer possible.

Central planning though, might still be regarded as an optimal choice, if it could be demonstrated that it allows more information about the system to be processed and intelligently acted upon than any alternative arrangement. Unfortunately, exactly the opposite is in fact the case.

Figure 3: Network Capacity Limits vs Available Interconectivity

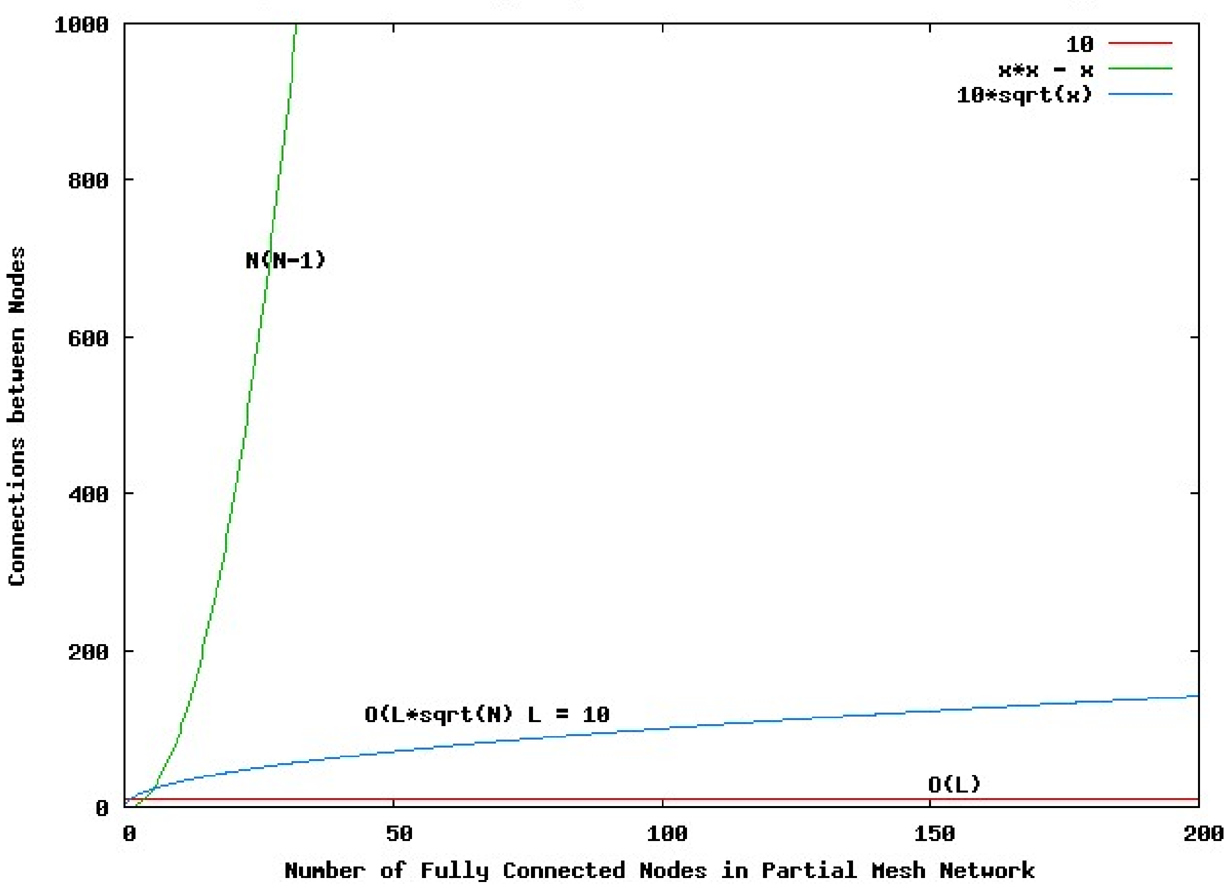

To illustrate this, consider Figure 3 where we can see the approximate network capacity obtained for two different topologies, the partial mesh analysed by Scaglione and Servetto, where all nodes are maximally connected to other nodes, and the strictly hierarchical topology where all nodes are connected through one single node. In this case the limit on direct node-to-node connectivity is 10. To give some idea of the rapid degeneration of the problem as the number of communicating nodes increases, the number of connections that would be required if every single node were connected to every other node, N(N-1) is also provided.

On the left hand side of the graph, we can see that for communication networks with small numbers of nodes, topology does not matter, the required information capacity is less than the theoretical maximum $O(L\sqrt{N})$. However, as the number of nodes exceeds the group size, the partial mesh clearly provides considerably greater capacity than the strictly hierarchical topology, which is always constrained by the capacity of its central node. For very large N though, even the partial mesh solution scales badly, as the extra information capacity added by additional nodes is negligible.

It is however, possible to improve on $O(L\sqrt{N})$ for large N, if the communication network is divided into groups of directly (ie. 1-hop), connected nodes, which are linked by overlapping membership between the groups. If we assume that the total number of nodes is divided into groups of size L-2, so allowing for inter-group connectivity with redundancy, we can approximate the information capacity of the resultant network as the sum of the individual information capacities of its' constituent groups, thus:

$$O((N/(L-2) * L * \sqrt{(L-2)})$$

shown in Figure 3, for a network of 1000 nodes, with a link limit of 10.

This is an interesting result, and bears examining. Why would introducing structure improve network information capacity? Going back to Scaglione's proof, and consider the connectivity problem for the entire network, if it is divided into slices of maximally linked nodes. In an unorganised network, messages must flow through a set of nodes in order to reach their destination. In the worst case, this is from one side to another, or through a series of slices. However if the network can be divided into a set of overlapping groups, where at least one member of each group is also a member of the overlap group, then routing distance is reduced to at most two hops. Since the primary loss of information capacity within networks arises from the need of intermediate nodes to route messages for other nodes in the network, minimising this loss increases the total available capacity to the network.

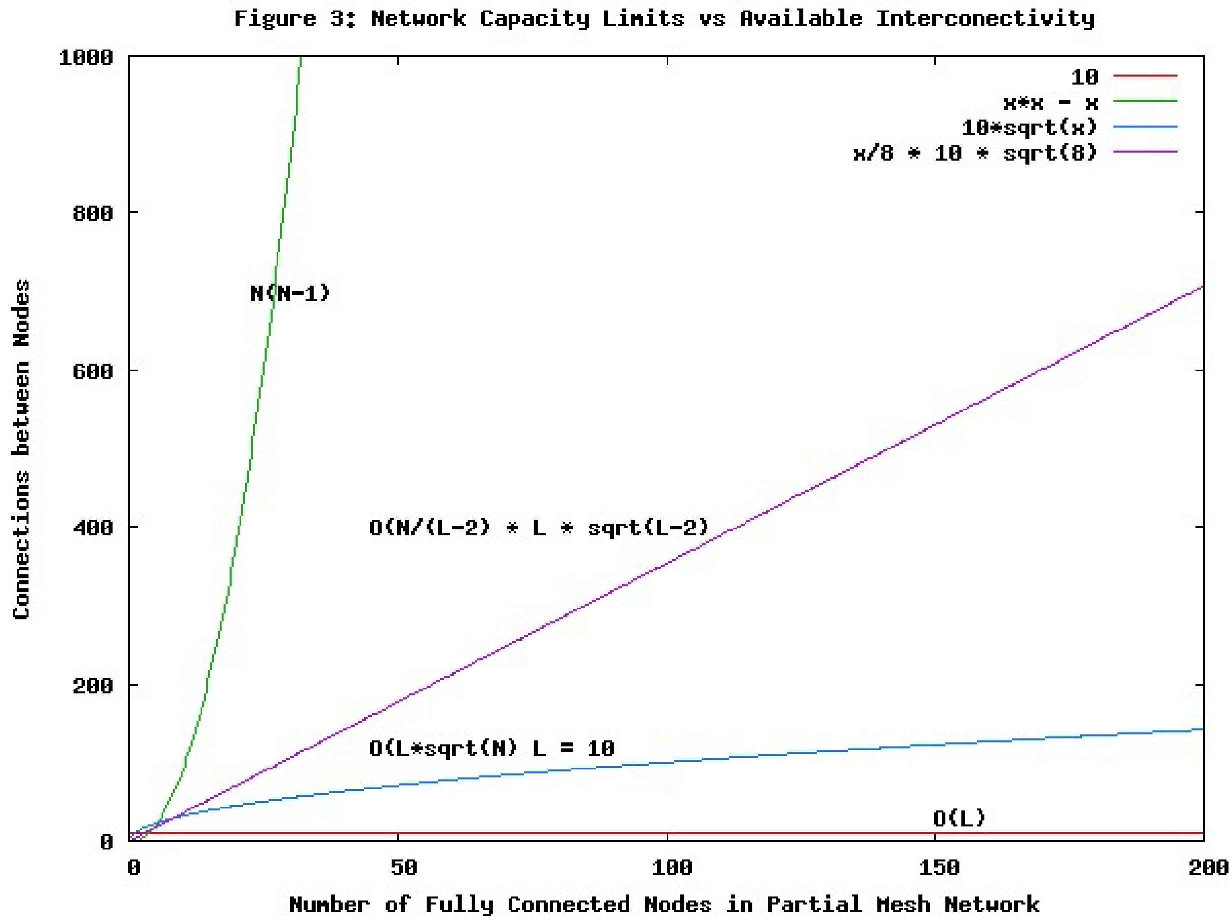

This communication capacity result for organised networks provides an alternative explanation for the fairly widespread phenomena of emergent orders in communicating organisations, to the topology growth explanation provided by Barabási and Albert[8]. In dynamic systems both effects probably play an important role, in that not only does organised distributed structure provide the most capacity, but also that there is a straightforward growth path to obtain it.

Taken together, these results provide a set of observations on the limits of information transfer within any communicating network, or organisation that are dependent on its topology. They are not strictly quantitative, significant details are being hidden behind big O notation, in particular more detailed calculation of the applicable communication latencies, which contribute to the value for L. Still the limits are relatively clear, in the presence of similar bandwidth and message processing rates, an organised mesh will be capable of more communication than a partial, unorganised mesh topology, and both considerably outperform a strictly hierarchical topology. However, in all cases, only for very small networks (or very limited data communication) can any topology accommodate all the communication that the nodes are capable of. Large networked systems will always be faced with the problem of adapting to limited communication space.

Mesh based organisations can take advantage of what for a hierarchically structured organisation is a large amount of wasted communication capacity between the client nodes. It bears repeating that communication is by definition real time. While a single central node is processing one set of messages, it is not capable of simultaneously handling other messages – they are queued waiting for attention, or discarded. Distributed topologies provide choice, if one link is congested, traffic can be directed towards an alternate route. If one node is overloaded, the information can potentially be processed by another. This principle applies at all levels of abstraction.

Information capacity cannot however, be the sole determinant of organisational topology, else it would be hard to see why any other topology than a partial mesh, or group based mesh would be used. The differences in information capacity shown above are very significant. Since historically, centrally planned and hierarchical organisations have from many perspectives dominated most of human development, there must be some other factor in play. It is certainly true that centralised systems are in many respects simpler to organise, and can work well with small numbers of nodes, which may and indeed often does lead people to believe that central planning is a good solution for large systems too. However the difference in information capacity between the two forms is several orders of magnitude. A clear historical preference for hierarchical organisation is hard to explain purely on the basis of simplicity.

The Influence of Latency on Communication Toplogy.

Latency is the time it takes to send a message within a communicating organisation. Strictly it should be regarded as having two components, the time taken to transmit the message, and the amount of computational effort it induces on receipt. The latter depends greatly on the application receiving the message, which might be required to perform complex, time-consuming calculations with the received data before it can process the next message. Outside of some specialised real time streaming applications, the time taken to physically transmit the message is now so small in the majority of computer networks, that transmission time itself is often disregarded.

Until relatively recently the situation for most human organisations was exactly reversed. Throughout most of history, transmission of human communications has been extremely time consuming. When messages have to be physically moved between people, or alternatively arrangements have to be made for individuals to meet, in order to exchange information, communication latency dominates as a constraint on any form of co-operative activity that requires the exchange of information.

Communication within distributed organisations is not solely dedicated to the exchange of information. It is also used for decision-making and for control. Centralised topologies are optimal for the quick distribution of a single set of instructions to all nodes, but not for receiving and processing information from those nodes. It may well seem highly desirable to maximise the amount of information available to all levels of decision making, but this again ignores the problems implicit in real time communication. If a decision has to be made by tomorrow, and it will take at least a week to receive a response to a query for information, then that decision must necessarily be made without that additional information.

This is where centrally organised systems can have a considerable advantage. In systems with long message communication times, and a need to make comparatively rapid decisions on an appropriate set of actions for the participating nodes; there may simply not be enough time for the distributed exchange of messages that allow such systems to enjoy a larger information capacity. As it is quicker with long transmission latencies both to issue orders from a central point, and to concentrate information there, a centralised topology is optimised for faster response, albeit a much less informed one. The longer the latency of the messages, the worse the communication problem will become though, until eventually, for large networks, central control also fails. In this circumstance, especially when there are local nodes significantly closer than the centre, centralised control is likely to break down in favour of local communication, as message communication becomes faster with the adjacent nodes. [3]

Since military activity depends for its success on solving real time communication, information gathering, and control and co-ordination problems, it provides many interesting examples of the communication challenges these problems create. Hayek (1944: 152) commented that the difference on commercial and military organisation was "a fundamental one between two irreconcilable forms of social organisation" in the Road to Serfdom, but did not particularly consider why military organisation should be so different. Certainly the persistence through history of certain organisational patterns within the military, such as the tiered, hierarchical group structure, is unlikely to be an accident.

In the context of military organisations, the clearly defined ranks and hierarchies of the military can be regarded as a solution to communication routing under conditions of sudden node failure. Network systems that use dynamic topologies, that is they can adjust the relationship of their communication links with each other, suffer from a problem of broadcast storms, where the instantaneous amount of communication required to inform nodes of the new topology can exceed the capacity available for communication. The ranks, and known reporting structures of the military provide a self-healing network topology which requires minimal communication overhead, in that nodes can recover their local part of the topology from purely local knowledge in order to create a global organisation with known information routeing paths, albeit extremely low capacity ones. This is critical, since if network connectivity is sufficiently damaged, not only is it impossible for nodes to co-ordinate their activities, but it may also be impossible to recreate the communication linkages that provided network connectivity in the first place. However, it can also create problems of overload and command failure.

There are many examples from military history that can serve to illustrate the real time nature of communication constraints, and the unfortunate choices that consequently confront centralised organisations. In 1893 at the height of British Naval power, the battleship HMS Victoria was rammed and sunk by the HMS Camperdown during a training exercise in calm waters on the Mediterranean. The incident resulted in the death of 358 sailors, including the Admiral in charge of the manoeuvre. At the time of the accident, the fleet was approaching the Tripoli coast in two parallel columns, and was ordered to execute a 180° degree turn, in order to reverse course. Rather than achieving this by turning outwards as was customary, the fleet was ordered to perform a manoeuvre whereby the ships turned inside of each other. Owing to a misinterpretation of received orders, and it should be said, a strict rule of obedience to those orders, the two lead ships turned onto collision courses with each other. [5]

Naval communication at the time was still performed by flags used to send messages between ships. This is a very low bandwidth medium, and over time, a system of codes had evolved which increased the amount of information that could be sent between ships. Groups of flags provided a code that referenced a pre-computed message from a codebook, essentially an early form of hash table. Although this allowed complex messages to be sent, sending, receiving and decoding was still time consuming and typically took several minutes.

In real time then, the fleet was already too close to the coast, when the order to turn was given, for any exchange of messages clarifying the orders to occur. In the time it would have taken to get a response, the fleet would have grounded itself on the Tripoli coast. In this example, and many others the emphasis in military training on blind obedience to orders can be seen as a reaction to communication constraints when operating under conditions of high message latency.[4]

The dilemma then that communication constraints create for communicating organisations, such as the Navy under conditions of slow message transmission, or high latency, is as follows. A strictly hierarchical topology can provide the optimal co-ordination of simultaneous orders that is necessary for the independent nodes, or ships to act together, since it optimises the transmission of information from a single point. However, if there is any miscommunication, misunderstanding or simple mistakes in the orders transmitted, then although local knowledge exists that this has occurred, there is not enough time to communicate this knowledge to the central point. As the subsequent inquiry revealed, the commander of the Camperdown was only too aware of what was potentially about to happen, but had to assume that the Admiral of the Fleet, on board of the Victoria knew what he was doing. Since the ships were so close, once the turn began, the laws of physics took over, just as the direct cause of the accident were the laws of information transfer.

Communication, Economics and Market based systems.

Can consideration of these various network effects that influence organisations provide possible analytic answers for the relative success of market based and democratic systems in human history? Hayek phrased his answer to these questions in terms of the importance of local knowledge, and the impossibility of central planning ever being able to provide this for large economic systems involving millions of agents. As we have shown, we can now support that answer by considering simply the constraints on the communication of knowledge within large systems of agents operating in a continuously changing real time environment. In such systems there will always be more information produced than the network is capable of transmitting to a single point for analysis, irrespective of the analytical tools available at that point.

From this perspective, following Hayek's example, we can analyse societies and economic organisation in terms of the ability of their underlying network structures to transmit information.

From a purely network perspective, we can regard money as a form of packetized information, and markets as hubs in a large scale distributed network that uses it to provide a continuous determination of relative supply and demand across society. Market based systems in this sense can be regarded as a distributed way of computing in real time, a local supply/demand equilibrium. It is not that the flow of information, in the form of money, (or indeed of market information), can be perfectly efficient in a market based system – it follows directly from the argument above on communication limits within large networks that neither the market itself, nor the market participants themselves can be – but simply that it is provably more efficient than any centrally planned alternative, and by several orders of magnitude.

Proponents of central planning, such as Lange, have often pointed to computer developments as the solution that will enable central planning to ultimately succeed. Leaving aside the communication problems we have so far discussed, there are more subtle issues with this proposal. Centrally organised systems have some advantages, they are conceptually easy to understand, and easy to create. Simplistically, only one piece of information needs to be communicated to each node in the network in order to create it – the identity of who they report to. In any situation where communication is limited, this is an important advantage, since communication of topological information to each node is an activity that can very easily exceed the information capacity of the network itself.[5] It may well be possible to establish a rigidly hierarchical control structure where it is not possible to create a more distributed one because there is not enough communication capacity to create it.

It could also be argued that there is nothing stopping a centrally planned organisation from deliberately creating a distributed communication structure. This might look identical to the similar structure that would have been created by an emergent process, but there would still be a significant difference, the structure chosen by the central planner would have been necessarily based on far less information, than that of the similar emergent structure.

Further once a communicating organisation has been established, it is unlikely to remain static. A hierarchical, centrally planned organisation can be regarded in some sense as the lowest achievable communication state for a set of co-operating nodes. If it is successful, it will act to increase the economic output and consequent complexity of its society, to the point where this exceeds the capacity of the central planers to completely control it in real time. Once this occurs the natural tendency for those at the periphery at the organisation will be to create local links in order to avoid the bottleneck at the centre. As much as latency considerations permit, the system will grow to a more distributed state. Absent coercive measures, a successful centrally planned economy is liable over long periods of time to sow the seeds of its own destruction, simply by building the network communication infrastructure that allows a more complex society to emerge.

The Age of Low Latency.

The phenomena of globalisation that has been widely noted in the last 20 years could be more accurately characterised as a world wide drop in communication latency to near zero. Although strictly it began with the introduction of the phone and cable communication systems in the 19th century, it is the widespread and extremely rapid deployment of the Internet, and mobile phone networks that has triggered world wide side effects. As we have seen, changes in latency can be peculiarly disruptive to communicating organisations, since they directly affect the topologies of organisations that rely on that transmission.

So it should perhaps not surprise us that we are living once again through a period of economic and social turmoil, as these effects play out in real time. Hayek's question from the end of another period of rapid communication changes in the 20th century remains to haunt us, "What is the problem we wish to solve when we try to construct a rational economic order?"[5] How can we control something that is provably uncontrollable? Should we even try?

The conflict between the hierarchical and distributed approaches is not a simple one to resolve, it is in some senses un-resolvable, since their advantages and disadvantages are often mirror images of each other. Hierarchical organisations have a single point of failure, distributed organisations can be robust to the point of indestructibility. Distributed organisations can process more information, and in theory at least, react more intelligently. However, there is no single point of control if circumstances require the organisation to react quickly; and they are often, and quite correctly, accused of being unable to easily reach decisions.

For much of human history, there has been little or no choice for societies operating above the group limit of organisation; communication latency considerations have forced them to use some form of more or less hierarchical control. With the introduction of the Internet and mobile communications, this is however no longer the case.

Emergent orders and social organisations that have evolved over centuries cannot be changed overnight, even if the communication conditions they are adapted too are changed underneath them. There is an often un-remarked persistence to the structure of human organisation, even through violent attempts to change it, that reflects the impossibility – in terms of communication overhead - of successfully transmitting wide spread topology changes within large networks. Local knowledge can allow a communicating organisation to be reconstructed, creating a new one requires that communication links are already present.

So the economic and political challenges of this century will be at least partially framed in the context of the adaptation of existing organisations, as well as the emergence of new ones, to the changes in communication conditions. We are already seeing this process with the emergence of Web based organisations such as Wikipedia, which have conceptually reversed the client-server topology to allow local knowledge from many sources to be concentrated at a single point. There will doubtless be many others, and given that low latency conditions in many respects favour distributed organisations able to access and combine local knowledge, it can be expected that there will be concerted attempts to leverage local knowledge rather than central control, creating tension within existing hierarchically controlled organisations.

These effects also raise many questions. What new topologies can be exploited, now that communication latency considerations no longer apply? Given that even with the enhancements that come from ubiquitous laptops and small addictive communication devices, the time taken to process messages has decreased far less than the time to communicate them, how can processing time be better distributed in organisations? How do we manage the resulting information overload that has been created, bearing in mind that one significant problem of partial mesh networks is that they lack the natural throttles on network overload that the strictly hierarchical network provides? In this context one advantage of a more distributed topology is that it can distribute message processing more evenly too – but this still leaves the problem of handling decision making better.

Hayek's arguments then, are perhaps even more important today, than when they were first published in the rubble of the 20th century. Central planning will always be a temptation in times of crisis, perhaps at least partly because our social instincts betray us. In the small co-operative groups that we have the most direct experience with, central control is more efficient, and the problems of low information flow that it creates are rarely experienced because these groups are too small to experience them. In the large and complex organisations and societies that we have constructed in the last century it has proved to be disastrous. Perhaps, a better appreciation of the communication issues within co-operating networks, their historical changes, and how they affect all levels of organisation, will help us navigate the latency changes of this century a little better than the previous one.

Acknowledgements: I would like to thank Dr. Mike Bove and the MIT Media Lab for support for the co-operative robotics research on which this paper is based; Sean Wheeler for first suggesting I should explore the works of F.A. Hayek; Dr Ari Benbasat for discussions on network capacity constraints in communicating systems; and Dr Lorin Wilde for patiently enduring and critiquing several discussions on network theory as applied to human organisation.

Notes

[1] For example, if it is known that the next part of a message consists of a repeating sequence, for example, 'ababababab'; then it can simply be suitably encoded as 5ab. Cfgyturerw on the other hand, would have to be sent in its entirety. Broadcast communication technologies such as Television for example, where the same message is sent to many participants, may be very efficient ways to communicate, but are extremely inefficient ways to distribute information.

[2] Wide area networks have been built more empirically than may be generally apparent, and could be regarded as an emergent order in their own right. Network operators typically monitor real time traffic behaviour, and add capacity when certain limits are reached; if the addition of new switches or links cause problems they are swiftly removed until the cause of the problem is understood.

[3] This has been noticed in the context of the decline of many empires by Turpin[7], who recorded an historical pattern of new empires arising on the borders of previous ones, rather than from the old centres.

[4] Modern armies now experience a different set of problems that can be broadly characterised as information overload. This is a direct result of a drop in communication latency allowing far more battlefield information to be provided than was hitherto the case. The attempts by military organisations which have evolved under conditions of long latency and limited communications capabilities, to adapt to conditions of low latency and information overload will be interesting to watch.

[5] This explains incidentally why it is critical to prepare disaster response organisations well ahead of time. Generally disasters dramatically reduce available network communications, whilst simultaneously increasing the amount of communication that is needed to co-ordinate a response, Consequently, it is important that disaster communication be pre-arranged and as much as possible pre-computed, in order to be as efficient and robust as possible; and to not require any extended communications to re-arrange the network's topology once the disaster has occurred. Temporal shifting of communication load, in the form of rules, procedures, and pre-arranged response plans, are another frequently seen solution to communication issues in human organisations.

References

Shannon, C.E. 1948. "A Mathematical Theory of Communication," *Bell System Technical Journal,* 27: 379-423.

Barabási, Albert-László and Réka Albert. 1999. "Emergence of Scaling in Random Networks," *Science,* 286:509-512.

Gupta P. and P. R. Kumar. 2000. "The Capacity of Wireless Networks," *IEEE Trans. Inform. Theory*, 46(2): 388–404.

Hayek, F.A. 1944. *The Road to Serfdom*, University of Chicago Press.

Hayek, F.A. 1945. "The Use of Knowledge in Society," *American Economic Review*, 35(4): 519-

530.

Hough, Richard. 1959. *Admirals in Collision*, NY: The Viking Press

Scaglione, Anna and Sergio Servetto. 2002. "On the Interdependence of Routing and Data Compression in Multi-Hop Sensor Networks," The Annual International Conference on Mobile Computing and Networking.

Turchin, Peter. 2002. *Historical Dynamics, Why States Rise and Fall*, Princeton: Princeton University Press.