Competitive Network Formation (draft)

Phil Jones, 2008

interstar@gmail.com

Abstract

Preliminary work is presented on research into "competitive network growth" models. In these models, networks grow through the addition of new links between existing nodes. However, they are distinguished from other growth models by the fact that links may fall into one of several "types"; and we study the larger-scale effects of different rules for choosing between the types.

These models are inspired by competition between rival communication protocols or other standards for which a "network effect" may be thought to occur. In such situations, when faced with a choice of which standard to adopt, an agent is influenced by the perceived popularity of the standard (either globally or within his own social network).

Introduction

We are all increasingly familiar with a proliferation of new communication products such as instant messaging, next generation phone services and web sites to help us keep in touch with old school friends. The internet provides a substrate on which it is relatively simple to innovate new kinds of communication and new virtual communities; and such services are appearing at an increasing rate.

Users faced with choosing between different flavours of communication will, naturally be influenced by the intrinsic properties of the protocol or the service. But they may also be influenced by the social environment. People hear about a new communication product via their friends. They judge which is most popular and which is most likely to "win" any potential "platform war" before making their decision.

More formally, economists have used the terms "network effect" or "network externality" to describe situations where a product's desirability to a potential buyer is affected by the number of other existing owners.

In a useful summarization of the field (Besen, 1999) Stanley Besen encapsulates what we think we know of the rules of engagement as follows :

- "Network effects may outweigh preferences for intrinsic product characteristics."
- There is "a tendency for a single technology to dominate".
- "When consumers choose sequentially, "stranding" can occur." Meaning that those who adopt one standard too early may suffer exclusion if the alternative takes off.
- "Network industries may exhibit path dependence, so that the behavior of early adopters may have a disproportionate influence on the equilibrium outcome."
- "Expectations may be critical to the final equilibrium because users must often choose among technologies before those technologies have reached their ultimate network size."
- "Lock-in may occur on the "wrong" technology because if, for whatever reason, the wrong technology is chosen, it may be difficult to achieve the coordinated movement of large numbers of users required for the "right" technology to become the standard."

All of which leads to the following strategic advice for those who are in the business of providing a new service or communication standard and wish theirs to thrive : to get in earlier than rivals; to try to sign up "complements" (in the sense that providers of films or television programmes are complements of a new video-disk standard); to pre-announce coming technologies before they are available so as to create the expectations that your standard will dominate; and to commit "to low future prices in order to assure adopters that they will not be stranded on a small network."

Besen also points out that organizations do not always engage in competition but may sometimes co-operate within an agreed standard, especially when the customer may so fear the cost of becoming stranded with a losing standard that she chooses not to buy this category of product at all.

Can this situation be modelled? And what might we hope to gain from doing so? In general, advocates of Agent Based Computational Economics such as Leigh Tesfatsion (Tesfatsion, 2005) have argued that modelling can account for real world problems and situations that other mathematical idealization miss.

Models can also throw light on outstanding controversies. To take one example, the economist S. J. Liebowitz, a skeptic about the overall significance of network effects in the market, writes :

> “First of all, the extent (and symmetry) of network effects may be much more limited than is commonly assumed. For example, in the case of spreadsheets and word processors, it may be quite important for a small group of collaborators to use identical software so as to be perfectly compatible with each other. Similarly, compatibility may be important for employees within a firm. But compatibility with the rest of the world may be relatively unimportant, unimportant enough to be overwhelmed by differences in preferences, so that multiple networks could survive. Networks that serve niche markets well (such as wordprocessors specializing in mathematical notation), might not be significantly disadvantaged by network effects.” (Liebowitz, S.J., 1998)

While he undoubtedly has a point; there are niches for specialist products; we may read this in a different way : as an intuition about how far the influence of merely local network effects will range.

Often the most striking results in agent-based research come from discovering situations that surprise our intuitions, such the early classic study by Schelling (Schelling, 1971) which upset our intuitions by showing that the cumulative effects of apparently innocuous, small racial preferences, were far from trivial.

Another area of interest we may have is how the competition between the network types affects the structure of the network. Are networks grown in a competitive environment the same “shape” as other networks? Does a competition of types leave a mark or are early adopters of the "losing" type merely equivalent to other late adopters of the dominating type?

Model

A preliminary model has been created using the Erlang programming language. Erlang is a functional programming language designed for large, reliable telecommunication systems that are distributed across a network of processing nodes. (Armstrong, 1996) As such, it supports features which may be interesting to those engaged in agent-based simulation: parallel processes with asynchronous message passing are part of the core language. While the experiments described in this paper have been written and tested on a single processor machine, the decision to use Erlang was made, in part, so that these simulations could be scaled up to run on multi-processor computers in the future. The source code is available to be downloaded from (Jones, 2008)

In this first model, a network is produced by progressively adding connections between pairs of nodes chosen at random. There is no attempt to impose any more interesting or significant structure on the population. All nodes have an equal chance of being chosen; the result is therefore a random graph. Arcs are directional in that A->B and B->A are considered to be two distinct connections.

Arcs also have one of two “types” which we can, for convenience, label type-1 and type-2. And when two nodes are connected there is a choice of which type to use. The decision is determined by the behavioural rules that are parameters to a particular run of the experiment.

In some rules, agents can be assumed to have a preference for one or the other type. However this is not to say that an agent will always get what it wants. Because two agents are involved in each decision, sometimes their preferences will conflict and the rules can be thought of as representing a kind of negotiation.

In this paper, the following rules for selecting arc-type are looked at :

- **Random** – either type is equally likely
- **Prefer Last Used** – each agent remembers the type that was used for its previous connection, and prefers to use that type again. If neither has a previous type, or there is a conflict, the type is again chosen at random.
- **Prefer Usual** – each agent remembers how many of each type it is already using, and prefers the one with the highest count . When there is a conflict then the type which is most used overall, wins. (For example, if A1 prefers type-1 due to already participating in 4 type-1 connections, whereas A2 prefers type-2 based on its single type-2 arc, then type-1 will be the winner.)
- **Neighbourhood Norm** – agents prefer the most popular type among their immediate neighbours. Note that this includes neighbours which are linked by arcs of either type. So that A1may be connected to A2 via a type-1 arc, but if A2 has more type-2 arcs than type-1, then her “recommendation” will still be for type-2. Again, if there is a conflict, then the type which is most commonly used wins.
- **Diffident** – as with Neighbourhood Norm, agents will prefer the type which is already most popular among their neighbours. However, if neither of the agents has at least two neighbours using a particular type, the pair will refuse to form a connection at all (analogous to two people deciding not to adopt any instant messaging technology if none of their friends are using instant messaging).

Results

To a population of 100 agents, initially disconnected, we added 1500 links sequentially. The decision as to the type of each new link depended on up-to-date state of the network. Each simulation was run 10 times and the results averaged.

Our chosen index was the proportion of the dominant (most frequently used) type within the whole population of types : $T_{max} / (T_1 + T_2)$ where T_{max} is the number of links having the dominant type and T_1 and T_2 are the number of type-1 and type-2 links respectively. This value is graphed below against the number of arcs (time).

The results are not, overly, surprising. We see that with two of the rules, the proportion tends towards 0.5, meaning that neither type is dominates the other. The slight preference for continuity by agents who try to reuse their previous link-type for the next (Prefer Last Used), is swamped by the randomness which is injected whenever there is a conflict.

For the other three rules, a clear winner emerges. This is most pronounced in the "Diffident" case, followed by the "Neighbourhood Norm" and finally, the "Prefer Usual" rule, which, while still having a clear winner, sustains a noticeable minority who "Think Different".

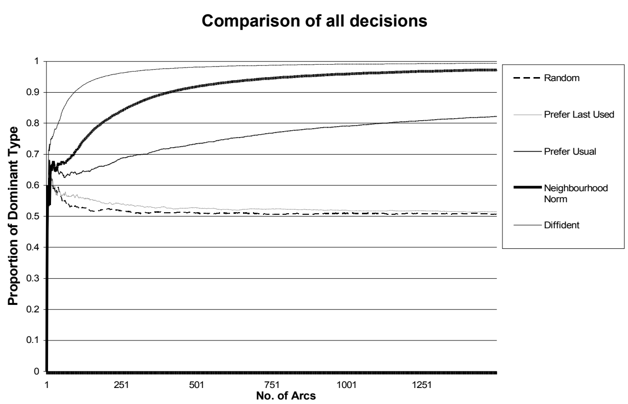

Conclusion

We are likely not surprised that the two rules allowing for social influence ("Diffident and Neighbourhood Norm") both give rise to a clear winner. But it is worth revisiting Leibowitz's intuition, above. This is not necessarily to criticize Leibowitz, as the phenomenon of diversity in the market clearly exists and requires explanation. Also the models presented here are simplistic in the extreme. It would be premature to infer too much. Nevertheless a straw-someone might have thought that the effects of local influence would not be so pronounced. The need for interoperability extends only so far.

In the real world, we look to very public proxies for information about the overall state of the standards war. What does the media say about Blu-Ray vs. HD-DVD? Which do the largest retailers commit to?

The simulations presented here have no such overview. Agents make their decisions affected only by immediate neighbours. And yet the result is still globally decisive. Consensus spreads through the whole population.

We can speculate that that has something to do with the fact that the graph which is being built up has no more specific structure, and so, on average, distances between nodes are short. Were the graph to be formed by some other principle such as one leading to a lattice, the situation may be different.

The faster rise to domination of one type under "Diffident" may be explained as these agents being more reluctant to take a bet on a potentially unsuccessful standard. Adoption maybe therefore slower, but the winning tendency becomes obvious earlier.

Future Work

Simple as it is, there is more work to be done with this model. As of time of writing, time and processor power didn't permit the analysis of other properties of the networks such as applying clustering or the CNM algorithm (Newman, 2004) We would like to know if the networks grown here are the same shape as other random networks.

We plan to look at further variations on the rules: what happens if the links do have some intrinsic differences? If, for example, type-1 is just "better" (and recognized as such by the agents who have a bias in its favour). Is it possible, even so, to sometimes see the path-dependency and "lock-in" by the inferior type? Under what conditions?

What happens if the types themselves can affect their rate of adoption? If type-1 is just more likely to be copied from a neighbour than type-2, or type-2 presents less of a "diffidence" barrier than type-1?

Why, really, does the diffident rule seem to lead to a more decisive early win for one type? Does some kind of conservatism in agents actually improve their ability to agree on a common opinion?

A Curious Graph

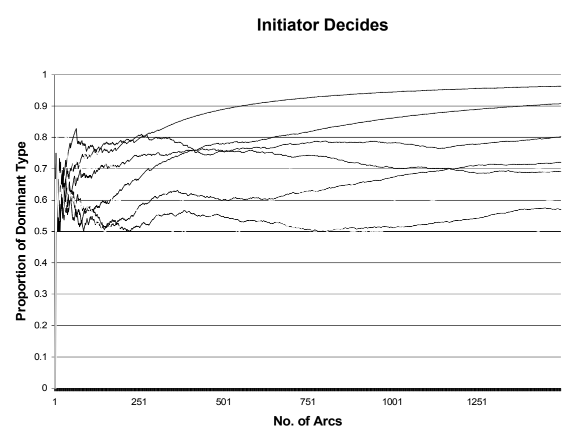

This graph shows the result of 10 runs of another rule we call “Initiator Decides”. This rule is similar to the “Prefer Last Used” rule in that agents try to use the same type that they used in their most recent previous connection. However, in this rule, one agent is considered to be the “initiator” of the connection, while the other is the recipient, and here the initiator always gets its way.

We can think of this as analogous to the case of certain communication and social networking services that grow by invitation. To be clear, in our model, agents were still chosen at random, and an agent did not have to be an existing user of a type in order to issue an invitation for it. So this is different from, say, the early days of Orkut. However, it is common in the real world that people join a social networking service, or start using a communication tool, not because they have sat down and rationally weighed up the alternatives or because they have polled their friends to find the most popular, but simply because one of their friends or colleagues asks them to sign up.

The result is quite different from the “Prefer Last Use” rule which tends towards a roughly equal balance of the types. (Largely because conflict there is, unrealistically, resolved by “tossing a coin”.) Instead, without conflict, each run of this model seems to find a stable, if slowly shifting, mix of the two types. However, different runs produce quite different outcomes. This rule can lead to domination of one type over another, it can lead to roughly similar quantities of the two types or it can lead to other, apparently stable, mixes. Here is some hint of a growth model which is susceptible to “path dependency”.

Further work is planned to look more closely at this “initiator decides” and other growth through invitation models.

References

Armstrong, J., R. Virding, C. Wikström, M. Williams. (1996) Concurrent Programming in Erlang, 2nd Edition. Prentice Hall,

Besen, Stanley M. Innovation, Competition, and the Theory of Network Externalities, http://www.econ.yale.edu/alumni/reunion99/besen.htm

Liebowitz, S. J. and Stephen E. Margolis, 1998, Network Externalities (Effects) in The New Palgrave's Dictionary of Economics and the Law, MacMillan, Available: http://www.utdallas.edu/~liebowit/palgrave/network.html

Jones, P. (2008). Erlang source-code for competitive network formation experiments. (http://www.nooranch.com/optimaes/compnet)

Newman, M. E. J. and M. Girvan. (2004) Finding and evaluating community structure in networks. Physical Review E, 69:026113

Schelling, T. C. (1971), Dynamic models of segregation, *Journal of Mathematical Sociology*, 1, 143-186.

Tesfatsion, Leigh (2005). Agent-Based Computational Economics : A Constructive Approach to Economic Theory, (forthcoming in L. Tesfatsion and K. L. Judd (editors), Handbook of Computational Economics, Volume 2:Agent-Based Computational Economics, but pre-print available http://www.econ.iastate.edu/tesfatsi/hbintlt.pdf)