Self-Organising, Open and Cooperative P2P Societies – *From Tags to Networks*

David Hales
www.davidhales.com

Department of Computer Science
University of Bologna
Italy

Project funded by the Future and Emerging Technologies arm of the IST Programme

What's going on in the world? the answers are close to home!

Fatta la legge trovato l'inganno!

As soon as you make a law it will be broken!

Project funded by the Future and Emerging Technologies arm of the IST Programme

Apologies! - who is this guy who cant keep his mouth shut?

- Computer Science -> AI -> Agent-Based Social Simulation
- Brief background - PhD Essex 2001 (Jim Doran, Nigel Gilbert), Rome CNRS (Rosaria Conte, Cristiano Castlefranchi) a few months, Manchester CPM (Bruce Edmonds, Scott Moss) about 2 years
- Journal of Artificial Societies and Social Simulation (JASSS)
- ESOA - AAMAS - Franco Zabonelli

The big idea - society makes us nice!

- if I don't like the people I am with....
- I move to another group...
- if I am nasty...
- everyone leaves me alone...
- Hence Ostracism *via bounded* maximisation can deal with the single round PD in a P2P!
- It's a kind of group selection! - an unfashionable idea in biology (though that is changing) but that does not concern us.

DELIS

What's DELIS?

- Dynamically Evolving Large Scale Information Systems (DELIS)
- A four year EU funded Integrated Project (IP) of Framework Program 6 (FP6) within the Future and Emerging Technologies area (FET)
- 19 Partners across EU
- Bologna: Biologically and Socially inspired mechanisms (self-healing, scalable, robust)

DELIS

Talk Overview

- 1) Why study cooperation in P2P systems?
- 2) Some previous models of cooperation
- 3) Tags – a new novel mechanism
- 4) Translating tags into a P2P simulation

DELIS

1) Why Study Cooperation in P2P systems?

What's the big picture?
What's the big problem?
How do we solve it?

Project funded by the Future and Emerging Technologies arm of the IST Programme

DELIS

Why study cooperation in P2P?

We want to know how nodes (agents) can perform tasks involving:

- Coordination & Cooperation
- Specialisation & Self-Repair
- Scalability & Adapting to Change

WITHOUT centralised supervision and in a scalable way

DELIS

The Bigger Problem

- Often systems composed of agents with limited or faulty knowledge
- Agents may be malicious, deceptive, selfish or crazy (open systems and / or adaptive agents)
- Agents have limited resources
- How to design algorithms that allow agents to collectively emerge the desired properties under these difficult conditions?

DELIS

A Solution

- Required properties - a strong resemblance to those of “living” systems (organisms, groups, societies etc.)
- Historically studied within in the broad fields of Life and Social Sciences
- Theories & proposed mechanisms exist in various forms (including computer models!) Can we import some of these?

DELIS

Some Previous Models of Cooperation

The Prisoner’s Dilemma (PD) game

Ideas from Economics, Biology and Political Science

Project funded by the Future and Emerging Technologies arm of the IST Programme

DELIS

The Prisoner’s Dilemma

Two thieves are taken in. The police have little evidence. They interrogate them separately – each is offered a “deal”. If they give evidence against the other they get a lighter punishment (whatever the other does), otherwise they get some time in jail. If both keep quiet they get off lightly, if both talk then they both get put away for longer, but if one talks and the other stays silent then the “grass” walks free while the silent one goes away for an even longer time.

DELIS

The Prisoner's Dilemma

Given: $T > R > P > S$ and $2R > T + S$

Player 1 → / Player 2 ↓	C	D
C	R / R	T / S
D	S / T	P / P

DELIS

The Prisoner's Dilemma

- This is a "minimal form" of a "Commons Tragedy" (Hardin 1968).
- The "rational" game theoretic solution (the "Nash" equilibrium – is to defect)
- Selfish adaptive / evolutionary units would also tend to Nash (ESS)
- It is desirable for "societies" to maintain at least some level of cooperation in such situations and many seem to. But how?

DELIS

Maintaining Cooperation in the PD

- Binding Agreements (3'rd party enforcement) – expensive, complex, tends to centralisation (Thomas Hobbes 1660)
- Repeated Interactions so can punish defectors – requires enough repeated interactions and "good guys" at the start (Axelrod 1984)
- Fixed spatial relationships – lattice or fixed networks – not good with dynamic networks (Nowak & May 1992)
- **Tags – scalable, single round, simple (Holland 1993, Riolo 1997, Hales 2000)**

DELIS

Tags – New and Novel Mechanism for Cooperation

A little detail on a previous tag model Hales (2000, 2004).

Project funded by the Future and Emerging Technologies arm of the IST Programme

What are Tags?

- Visible and changeable markers attached to agents (e.g. dress style, accent, hair-style)
- If agents preferentially mix with those sharing same tags
- Distinct groups are formed - By excluding those without the same tags
- By changing tags agents move between groups
- Membership of some groups may be more desirable than others

Evolving Tags

If we assume (evolutionary process):

- Strategies and tags of agents obtaining high credit tend to get copied
- Periodically agents randomly mutate tag and strategy bits
- Result is all defection – since a defector never gets less credit from an interaction than its partner (ESS and Nash)

Evolving Tags

- But if we bias partner selection to those with matching tags (if any exist)
- We get unstable yet high levels of cooperation
- A dynamic group formation and dissolution process
- Tags mutate and are copied like strategies (but with a *higher mutation rate*)

Agents - a Tag and a PD strategy

Tag = (say) Some Integer

Game interaction between those with same tag (if possible)

How Tags Work

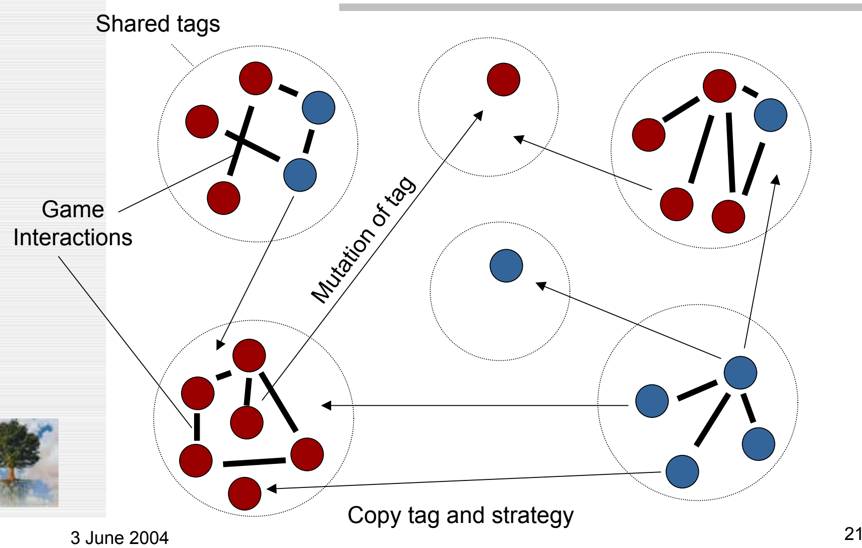

Visualising the Process (Hales 2000)

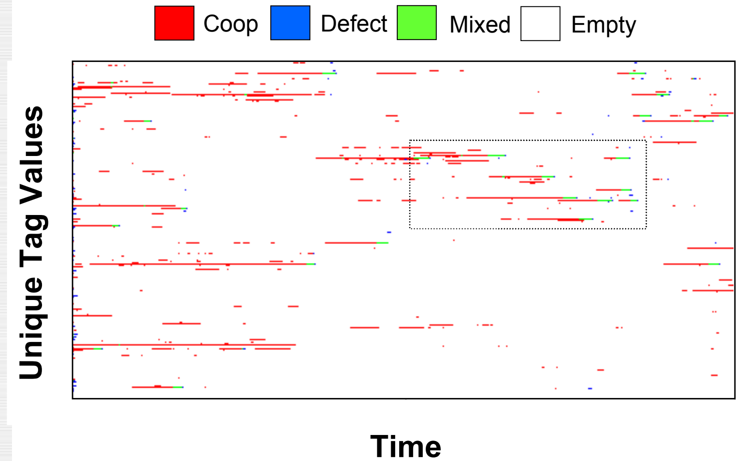

Visualising the Process

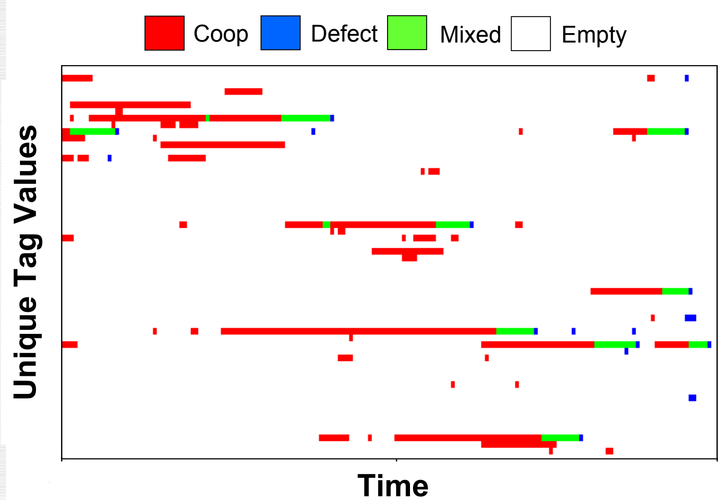

A Reverse Scaling Property

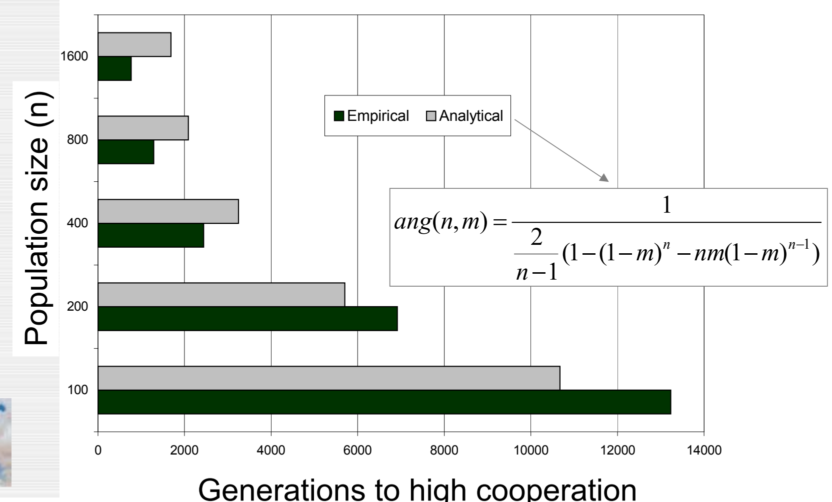

Recent finding (Hales 2004) – tag mutation rate needs to be higher

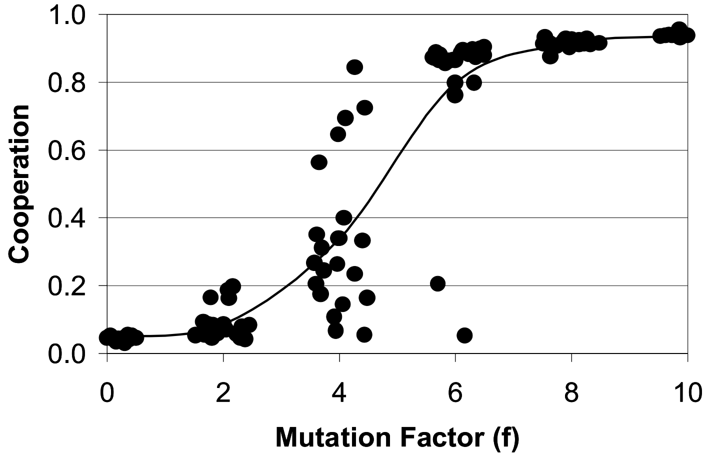

Translating Tags into a P2P Scenario

All well and good, but can these previous results be applied to something like looks more like: unstructured overlay networks with limited degree and open to free riders

Project funded by the Future and Emerging Technologies arm of the IST Programme

A P2P Scenario

Consider a P2P:

- Assume nodes maintain some max. no. of links
- Node neighbours can be thought of as a group
- Nodes may be good guys, share resources with neighbours, or free-ride, using neighbours resources but not sharing theirs (PD)
- Sharing / free-riding is a Strategy
- The neighbour links (as a whole) a kind of “tag” (if clustering high enough)

A P2P Scenario

- Represent the P2P as a undirected graph
- Assume nodes are selfish and periodically:
  - Play PD with RND selected neighbour
  - Compare performance to some randomly selected other node
  - If other node is doing better copy its neighbourhood and strategy
  - Mutate strategies and neighbourhood.

Initial thoughts and questions

- For tag-like dynamics high clustering would appear to be required (groups required)
- Will dynamic nature of the scenario support this?
- Can cooperation be maintained without it?
- We might start simulations of the model with high clustering initially (say small world or lattice) and compare that to random networks
- Many schemes of “neighbourhood copying and mutation” are possible which to use?
- What kind of topologies emerge over time?

Design Decisions

- Mutation of neighbourhood = replace all neighbours with a single neighbour chosen at random from the population
- Mutation on strategy = flip the strategy
- Node j copying a more successful node i = replace i neighbourhood with j’s U j itself
- When maximum degree of node is exceeded throw away a randomly chosen link
- Payoffs as before: T=1.9, R=1, P=d, S=d

Social Climbing, Ostracism, Replication

Before

After

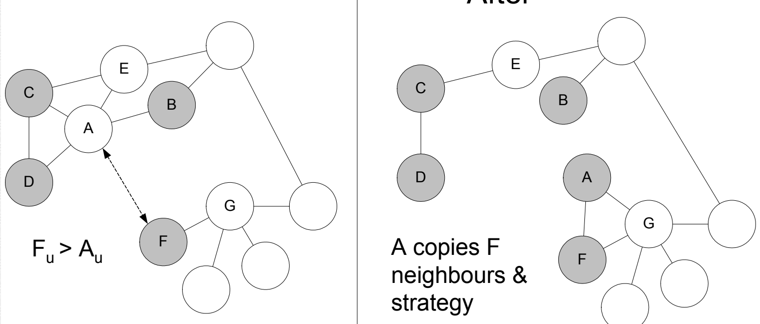

Where A_u = average utility of node A

In his case mutation has not changed anything

Mutation on the Neighbourhood

Before

After

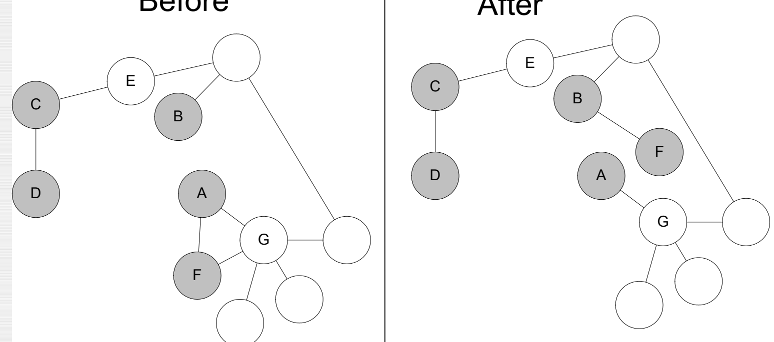

Mutation applied to F’s neighbourhood

F is wired to a randomly selected node (B)

The Simulation Cycle

LOOP some number of generations
  LOOP for each node (i) in the population N
    Select a game partner node (j) randomly from neighbour list
    Agent (i) and (j) invoke their strategies and get appropriate payoff
  END LOOP
  Select N/2 random pairs of agents (i, j) reproduce higher scoring agent
  Apply mutation to neighbour list and strategy of each reproduced agent with probability m
END LOOP

Parameters

- Vary N between 4,000..120,000
- Maximum degree 20
- Initial topology random graph
- Initial strategies all defection (not random)
- Mutation rate m = 0.001 (small) a previous
- Payoffs as before: T=1.9, R=1, P=d, S=d (where d is a small value)

Results

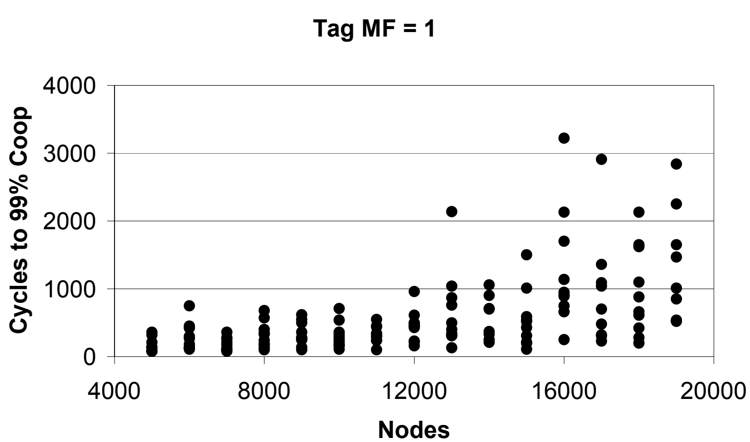

Results – increased mf=10

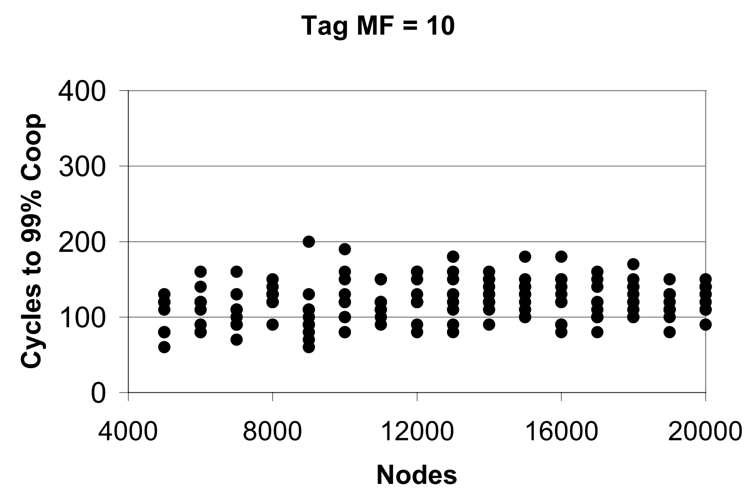

A few more nodes

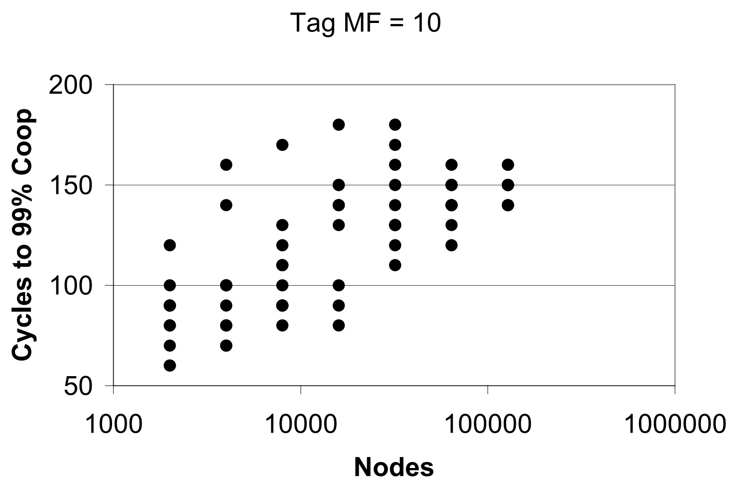

A typical run (10,000 nodes)

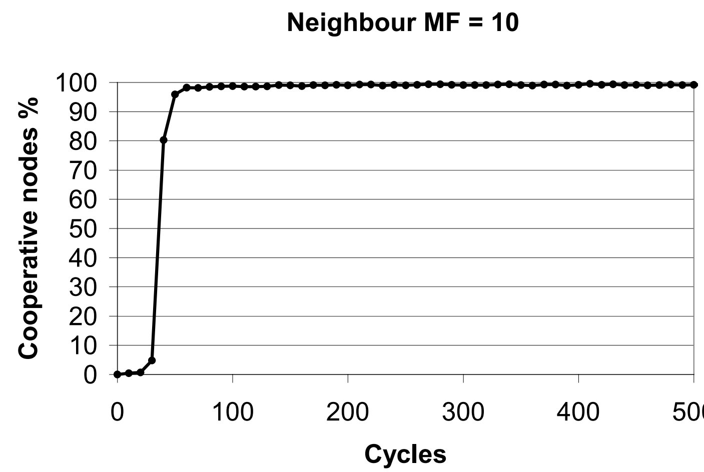

A 100 node example – after 500 generations

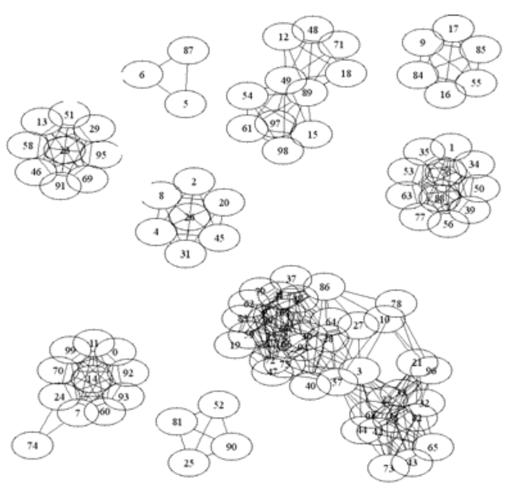

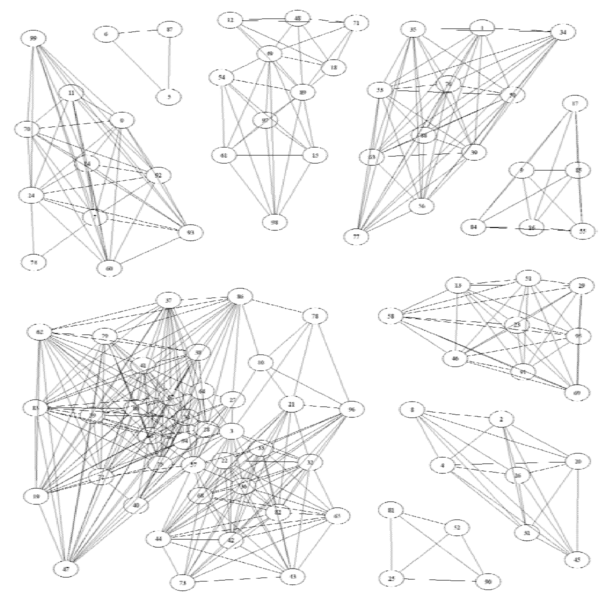

Topology Evolution – so far it seems….

- From ANY initial starting topology / strategy mix same outcome (tried random, lattice, small world, all nodes disconnected, all defect, random, all coop)
- Typically (very approx.) a max of n/10 unstable components exist at any one time which are highly internally connected (L not much more than 1 and C very high)
- But they are not of equal size
- Constantly reforming and changing due to mutation and replication
- Rough characterisation of disconnectedness = prob. that two random nodes are connected

Typical run, 200 nodes

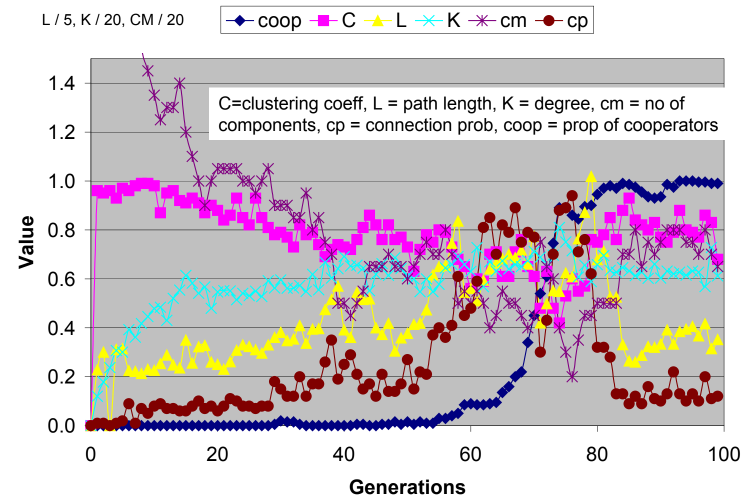

A message passing game

- Keep everything the same but change "game"
- A message passing game – select two nodes (i,j) randomly from G. i tries to send a message to j.
- Do a flood fill query from i to j.
- If a route of *cooperators* is found from i to j then i gets a "hit" (one point added to score)
- Only cooperators pass on a messages incurring a small cost for doing so, reducing score
- Hence defectors will do better than cooperators getting the same proportion of hits
- Tough task since need a route between specific nodes via a chain of coops only

Message Passing game - 200 nodes after 500 generations

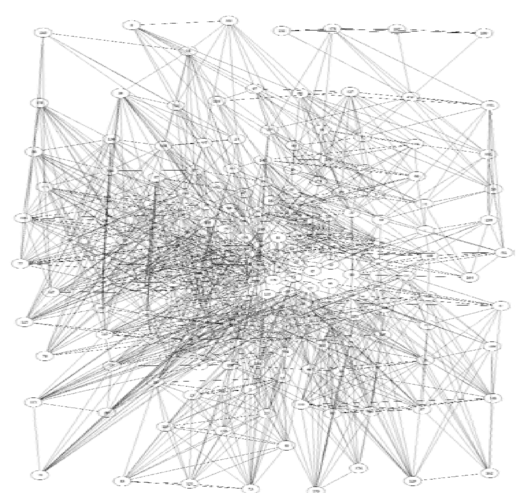

Message passing game - 200 nodes to 100 generations

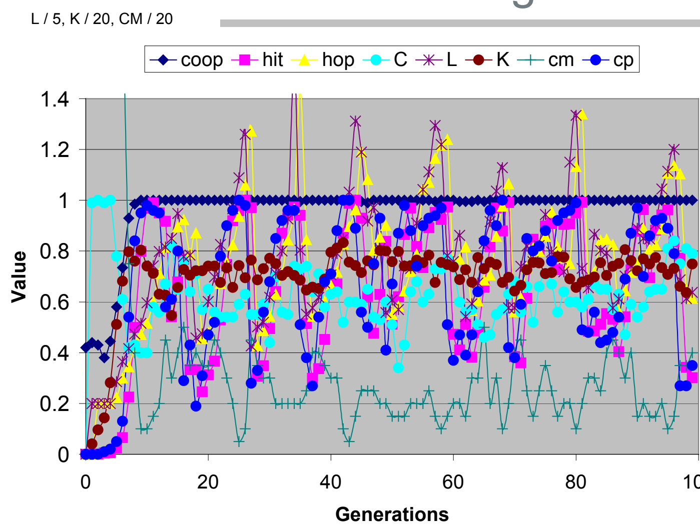

But its not as good as it seems...

- Increased games to 25n per generation
- Start with random strategies (all def. no good)
- Does not appear to scale well (oscillations)
- More work needs to be done (only a few runs)
- A very tough test for scaling on this mechanism
- On reflection - surprising it did this well
- Naive translation for bio-like model (homogenous greedy bounded optimiser assumption - no cleaver nasty nodes - need to consider entire plausible space of possible node behaviours - not just the reduced space of my “carefully selected” genes) - ongoing discussions with Mark et al

Next steps

- Assume random selections from the population (will it work with net. generated selections - say using neighbours neighbours?)
- **Try more realistic task (file sharing) (Qixiang Sun & Hector Garcia-Molina 2004) (see next slides)**
- So far robustness tested as effect of mutation – static pop size – try drop or introduce lots of nodes at once (churning)
- Simplistically treats all neighbour links as “one chunk” rather than selectively removing links (eliminate comparison also? Vance Maverick’s idea) various schemes possible
- Translate model into PeerSim framework
- Some maths! help! Sociobiological models of group selection? (Wilson 1970’s)

Qixiang Sun & Hector Garcia-Molina 2004 - scenario

- Simplified form - (perhaps a bit too simplified!)
- P2P file sharing query answering task domain
- each node has variable giving proportion of capacity devoted to generating queries against answering them
- Each node has an answering power (probability of making a hit given any query) TTL’s etc.
- “Incentive-based” mechanisms
- Some results

some results

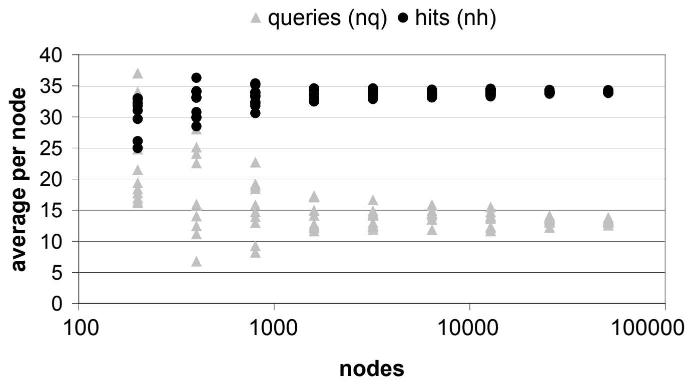

Results (averaged over cycle 40..50) for different network sizes (10 individual runs for each network size)

some results

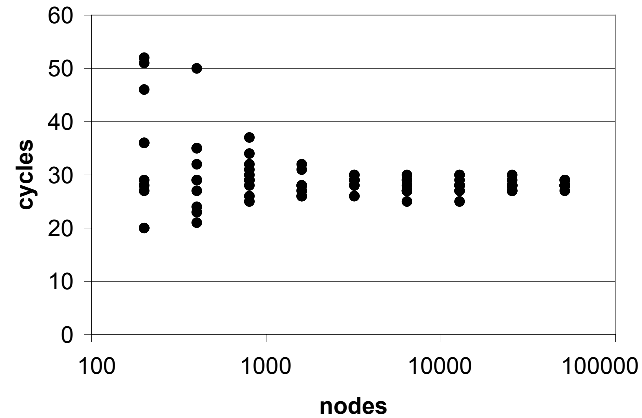

Cycles to high hit values (nh > 30) for different network sizes (10 runs each)

Conclusion

- Tag-like dynamics can be put into a network using simple rewiring rules
- Even simple rules appear flexible, able to create and maintain different topologies for different tasks
- Free-riding is minimised, even though node behaviour selfishly and have no knowledge of past interaction
- At least for close neighbour interaction the method scales well
- But much more analysis needs to be done and more realistic kinds of p2p task domain need to be tested - the more I learn the more nasty this problem is of course!
- Naive application of biological-type (or social) model needs sorting out - utility, copying of state (reproduction) space of behaviour etc.