An Analysis of Constructive Network Formation Models (Extended Abstract)

Gary Fredericks
University of South Carolina
Columbia, SC 29208
frederga@email.sc.edu

José M. Vidal
University of South Carolina
Columbia, SC 29208
vidal@sc.edu

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence—*Intelligent Agents, Multiagent Systems.*

General Terms

Algorithms, Economics

Keywords

Network formation

1. OVERVIEW

Network formation was originally studied by Jackson and Wolinsky [1]. We analyze a constructive model of network formation that is particularly suited to exhaustive computation.

Our model deals with a fixed-size network where each individual in the network wants to minimize its distance to the other individuals. Distance is measured as simply the number of hops between individuals, and is considered infinite if they are not connected. Building a link has a fixed cost (α) which is in the same units as the measure of distance. We assume an individual will want to build a link iff their decrease in total distance to the rest of the network is greater than or equal to the cost to them for the link.

The network formation process begins with an empty graph on N vertices and at each step a random not-yet-existing edge is added from the set deemed **feasible** by the payment rule. This is repeated until there are no more feasible edges to add. A **payment rule** is a rule for which vertices have to pay for the edge, how they split the cost, and how they decide whether or not to do it. The payment rule determines which edges can be added in a given situation. If a graph can be reached through a network formation process (i.e., can result from a series of feasible transitions), then we say it is **reachable**. The probability that a reachable graph will appear is its **reachability**. Finally, any reachable graph to which no further edges can be added (given a payment rule and α) is called a **sink-graph**. Since every network formation process will necessarily end with a sink-graph, our primary strategy for comparing payment rules in Section 3 is to compare the attributes of the sink-graphs for each payment rule.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

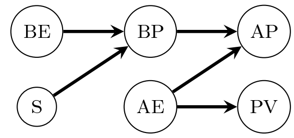

Figure 1: Partial ordering on rules.

We consider six different payment rules. Two natural choices are that the edge might be paid for by one of the two vertices it connects (S), or by both of them equally (BE). It also may be split equally among the whole network (AE). We may also modify the second two so that the cost is divided up fairly, in proportion to how much each vertex benefits from the new edge (BP and AP, respectively). In all the previous payment rules we assume the edge is not built unless all payers agree. Finally, since the AE rule turned out to be trivial (see Section 2), we added a modified payment rule that only requires a strict majority of the network to be in favor of building the edge (PV). In each of the global payment rules (AE, AP, and PV), we also assume that any edge that connects previously disconnected components can be added, for otherwise the process would never get past the empty graph.

2. THEORETICAL ANALYSIS

We have proven several interesting facts about these payment rules. We showed that the AE rule always deadlocks (and is therefore absent from Section 3). We then demonstrated a partial ordering of the six rules under a certain kind of subset relation. Finally, we proved the presence or absence of a certain extreme sink-graph under different rules.

The AE payment rule requires that for an edge to be built each node in the network should be willing to pay $\frac{\alpha}{N}$ for the edge, which means that their expected immediate value for it should be at least $\frac{\alpha}{N}$. However, we have proven that for any minimally-connected graph and not-yet-existing edge there must be a vertex that cannot benefit from adding that edge, and so will not be willing to pay $\frac{\alpha}{N}$. The result is that once the network formation process reaches a tree it cannot proceed further. Therefore AE is omitted in most of what follows.

We have also been able to prove the partial ordering shown in Figure 1. The partial order is a kind of subset relation. Given two payment rules X and Y, we say that $X \subseteq Y$ if for all α the feasibility of any transition under X implies its feasibility under Y. This implies that for a fixed α, any graph reachable under X is also reachable under Y.

All the arrows in Figure 1 have been proven, and all the pairs of payment rules that are incomparable in the diagram have been shown to be incomparable by example.

Finally, we have shown that a certain extreme type of graph is a sink graph for some of the rules but not others. We use the term "Lollipop Graph" to refer to any graph for which $N-1$ vertices form a clique, and the remaining vertex has only a single edge. This graph is the largest possible graph that can still be disconnected by removing a single edge, and so seems inefficient for most applications. Surprisingly, we have shown that for BE lollipop graphs are never sink-graphs, for PV they are never sink-graphs if $N > 5$, but for AP, BP, and S the lollipop graphs are sink-graphs for certain ranges of α. This is unexpected both because it is a sink-graph at all, and also because there are local and global payment rules in each category.

3. STATISTICAL TESTS

We also investigate the payment rules statistically by exhaustively computing properties of all graphs on ten or fewer vertices, and try to identify patterns that seem most amenable to extrapolation to larger graphs. Most of our experimentation was done by treating the network formation process as an acyclic Markov chain, where the Markov states are graphs, the transitions represent adding an edge to the graph, and the transition probabilities are conditioned on α. We used the **nauty** program to populate our database with all 12,293,431 unlabeled graphs for $3 \leq N \leq 10$, and all 251,463,867 transitions between them. Unlabeled graphs were used as an optimization (there are over 35 trillion labeled graphs of the same orders), and extra calculation involving symmetries was necessary to ensure we were computing the same probabilities as would be obtained with labeled graphs.

We used dynamic programming to compute the feasibility thresholds for each transition and payment rule. Using these we then computed the reachability probabilities for each graph and payment rule, conditioned on α. Having this data, it was easy to identify sink-graphs, and to compute the expected value of various graph attributes given a payment rule and α. We compared the total count of reachable graphs for each payment rule as α varies, the expected connectivity and unfairness of the sink-graphs, and the probability that the sink-graph reached is regret-free.

The number of reachable graphs for each payment rule for $N = 1$ is plotted in Figure 2. The subset relationship described in Section 2 is evident in the dominance relationships between the lines in the plot. The line for AP dominates the BP line, which itself dominates both S and BE. The pairs of payment rules that are incomparable in Figure 1 have intersecting lines in Figure 2.

The connectivity of a graph is, formally, the minimum number of vertices that must be removed to disconnect it, so it measures the stability of the network against vertex failures. A disconnected graph has connectivity 0, a minimally connected graph (e.g., a tree) has connectivity 1, and a complete graph has connectivity $N-1$. When examining the expected connectivity of the sink-graphs for each payment rule, we found that (not surprisingly) the connectivity decreases as α increases (though there are curious slight exceptions for each payment rule). Furthermore we found that

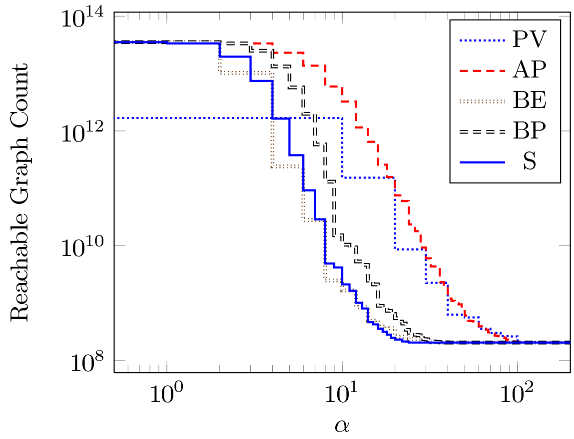

Figure 2: Number of reachable order-10 graphs for each payment rule as α changes. Since the x-axis is plotted logarithmically, the lines extend infinitely to the left; at $\alpha = 0$ all lines meet at the same value.

the S rule gives reliably less connectivity than the other four rules, which had mixed relationships amongst themselves.

We define unfairness as the maximum difference between any two vertices' total distances to the rest of the graph. When analyzing the expected unfairness of the sink-graphs for each payment rule, we found that the S rule was worst for nearly all α. The other local payment rules (BE and BP) were better for small α (< 6), while the global payment rules (AP and PV) were better for large α (> 20).

The final characteristic we analyzed was one we called "regret-free", which describes a graph for which all of the existing edges are still worth adding. That is, for all edges ij of the graph g, the transition from $(g - ij)$ to g is feasible. We computed the probability that the sink-graphs for each payment rule would be regret-free. We found that the probabilities are quite low for small α (in fact as the order increases, the minimum probability seems to approach zero for all payment rules), and slowly build up towards 1 as α increases (eventually the only reachable graphs are trees, which are trivially regret-free). The probabilities for the three local payment rules seem to rise much faster than the two global payment rules.

We also noted the strange fact that all of the regret-free sink-graphs for $N \leq 10$ for all of our payment rules have some symmetry in them (i.e., there is a non-trivial permutation of the vertices that produces the same graph), which is a relatively uncommon attribute in general.

4. CONCLUSION

Our work has so far dealt with only a small set of payment rules, and only networks with relatively few individuals. Future work could explore through sampling whether the statistical results in Section 3 apply to larger networks, and could try to characterize other payment rules or families of payment rules.

5. REFERENCES

[1] Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.