Social Networks and Social Choice **Lecture Date: September 30, 2010**

Lecture: Consensus, bribe and marketing

Lecturer: Elchanan Mossel *Scribes: A. Kaspersen*

In previous lectures we have been working with networks that were deterministic in the sense that given the model and initial values, the outcome could be determined by more or less complicated calculations. The models discussed in this lecture are of a different character since they involve randomness. Today's topics are:

- Randomized consensus protocols and the "voter model"
- Which voters to "buy in the voter model" and the viral marketing problem

1 Randomized consensus protocols

The **consensus problem** can be described as the problem of getting a group of n individuals in a network to agree on a value. Randomized consensus protocols are models that incorporate randomness in solving the consensus problem.

The framework we will consider is a social network represented as a graph $G = (V, E)$. We will now work in continuous time and in an asynchronous fashion. At time $t = 0$ each individual has the opinion $X_0(v)$, $v \in V$.

It is clearly in our interest to find "good" models that converge to consensus, but it is less obvious what a good model is. Intuitively we can list three properties that model should fulfill:

Property 1: Fairness with respect to alternatives

Property 2: Consensus is a fixed point

Property 3: Simple calculation/evaluation

Two different types of models will be discussed; **vertex models** and **edge models**:

Vertex models: In vertex models each vertex is updated in a *Poisson*(1) process, i.e. there is an expected waiting time of 1 after each update. At each update time, $X_t(v)$ is updated according to neighbors, own current opinion and a random factor. Vertex models focus on individual decision-making.

Edge models: Edge models are similar to vertex models with regard to the Poisson(1) process update times. At each update time, an edge (u, v) update both end points of the edge $X_u(t)$ and $X_V(t)$. This models decisions as results of interactions in the network.

A classic example of a consensus problem is the choice of location. A group of friends want to go see a movie, but there are two (or more general, several) movie theatres to choose from. Everybody prefers to see the movie with their friends, and the question to ask is if and how one can achieve a solution where everybody goes to the same location.

Figure 1: The choice of location problem

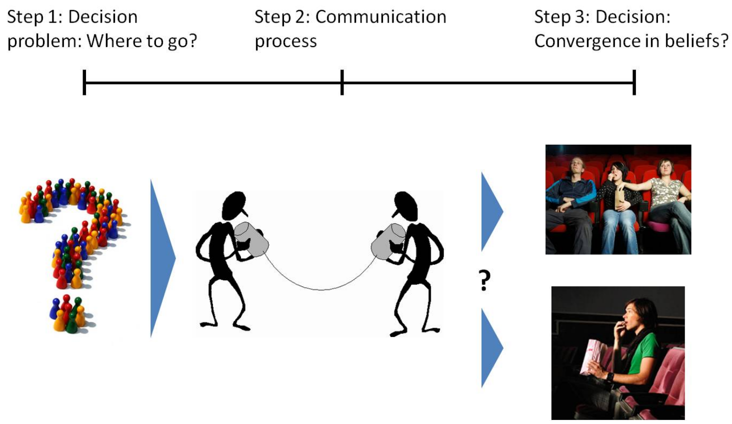

A lot of experimental work has been conducted to find out what people actually do to reach consensus. In an experiment by Latan and L'Herrou [9] human subjects in a lab were playing a conformity game in different network structure, i.e. each player is paid a fixed amount if he agrees with the majority of the players, but the players only have information about their neighborhood in the graph. The participants failed to win a bonus more than 25% of the time, even though at least half of the participants were guaranteed a win by the rules of the game.

For further reading about similar experiments, reference is made to [4] where a model for consensus with different payoffs is tested. [5] considers the coloring problem, where payoffs depend on finding the opposite action of the neighbors.

We will now try to model what has been discussed above. As usual we consider a graph where every player can see a subset of the other players. Each player can choose one of two labels *red* or *blue*. Every individual's state at time t depends on a) the state of its neighbors in the network, b) advice and c) randomness (flip the coin).

Figure 2: The voter model

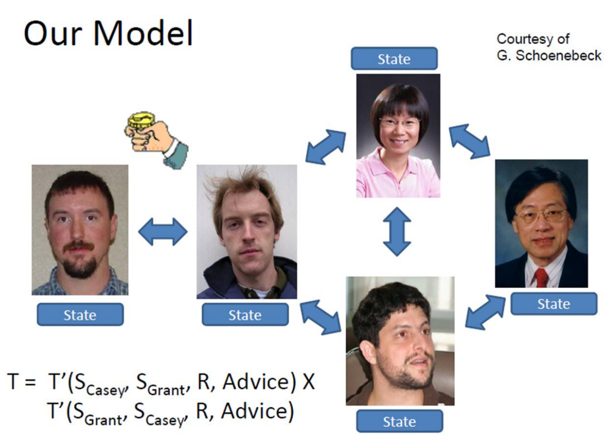

The two problems we will consider are the following:

- **The Coordination Problem** Do we arrive at consensus?
- **The Majority Coordination Problem** The coordination problem with the restriction that consensus must equal the majority of the original opinions.

Without further restrictions the solution to both problems would be trivial; the majority voting scheme. We therefore consider a graph where the network structure is only partially known to the players; everybody only has information about their neighborhood in the network. We also want the protocol to be symmetric with respect to the two states, and we should arrive to consensus naturally.

1.1 The coordination problem

We will first consider the coordination problem.

1.1.1 The voter model

A model that fulfills the abovementioned criteria is the **voter model**. The voter model is an edge model with the following updating rules at each update time

$$\text{If } \; X_v(t-) = X_u(t-) \; \text{ then } X_v(t) = X_u(t) = X_v(t-) = X_u(t-)$$

$$\text{If } \; X_v(t-) \neq X_u(t-) \; \text{ then } X_v(t) = \begin{cases} X_v(t-) \neq X_u(t) = X_u(t-) & \text{with probability } p = 1/2 \\ X_u(t) = X_u(t-) & \text{with probability } p = 1/4 \\ X_v(t) = X_v(t-) & \text{with probability } p = 1/4 \end{cases}.$$

This model was described by P. Clifford and A. Sudury [2] in the context of a spatial conflict where animals fight over territory (1973) and further analyzed by A.Holley and T.M. Liggett [3].

The model is fairly simple and has several nice properties: a) It is symmetric with respect to the states (red/blue) b) consensus is a fixed point and c) The number of reds (or blues) is random walk, since the number at each interaction either stays unchanged, or increases or decreases by one with equal probability. The random walk property allows us to use Markov chain theory.

Theorem 1.1. *The voter model converges to consensus and the convergence time is* $O(|V|^2)$.

Claim: Let opinions be in $\{+, -\}$ and let $X(t) = \sum_{v \in V} X_v(t)$. Then $X(t)$ is a martingale with respect to time. Since the only absorbing states are all 0 and all 1 we have $\lim_{t\to\infty} P(X(t) \in \{-n, n\}) = 1$. Furthermore, let T be the first time when $X(T) \in \{-n, n\}$, then we have that $P(x(T) = 1) = E(X(0) + 1)/2$. That is, if the initial configuration is 70% red and 30% blue, there is a 0.7 (0.3%) probability that beliefs will converge to red (blue).

We then consider $f'(t)$ where $f(t) = \text{var}(X_t)$ which is increasing and $\leq n^2$. We then have

$$\begin{aligned} f'(t) &= \lim_{h\to 0} \frac{E[X_{t+h}^2] - E(X_t^2)}{h} \\ &= \lim_{h\to 0} \frac{E[(X_{t+h} - X_t)^2]}{h} \\ &\geq \frac{4rh}{h} = 4r \times P(\text{ no convergence by time t}), \end{aligned}$$

where the first equality follows from the fact that martingale increments are orthogonal and the second follows since at a small time interval of time h there is at most one update and the probability that this update is between $-$ and $+$ is at least rh. From this it follows if that $P(4n^2/r) \leq 1/2$ is impossible, which gives the bound on convergence time.

For a complete proof see [11].

1.1.2 Choosing a leader

A different model of coordination is to choose a leader that decides what the group's opinion should be. Let each player choose at random how strong his opinion $S(v)$ is in the range $[n^{10}]$ say. We then let $X_v(0+) = X_v(0) \times S(v)$, that is every player has either a positive or a negative signal of strength $S(v)$. In each edge update, the player with the weaker signal copies the stronger. Then, given a unique strongest signal, $|X_v(t)| > |X_u(t)| \quad \forall\ u \in V,\ u \neq v$, consensus will converge to $sign(X_v(0))$. Note: If the strongest signal is not unique, but there are two or more players with equally strong signal strength and different opinions, and we assume that when two equally strong signals "meet", the players will choose label randomly, the two strongest signals will spread throughout the whole population until everybody is either + or - with the same signal strength. Then the dynamics will be equivalent to that of the voter model.

Theorem 1.2. *If there is a unique strongest signal in the "choosing a leader" model, the consensus belief will converge to that of the leader in the broadcast time of the network. If there is more than one strongest signal, the consensus model will follow the same dynamics as the voter model, and expected convergence time equals the broadcast time.*

Proof can be found in [11].

Table 1: Coordination models

Problem	Memory	Time	Required Advice
Voter Model	1	n^2	none
Greatest Element	$O(\log(n))$	broadcast	$\Theta(\log(\|V\|))$
Wait-and-See	Expected $O(1)$	O(broadcast)	$\Theta(\text{broadcast} \cdot \|E\|)$

1.2 The Majority Coordination Problem

We now turn the problem of reaching a consensus that reflects the original opinion of the majority. First we establish a basic, but easy result

Proposition 1.3. *The majority coordination problem cannot be solved with only one bit of memory*

Theorem 1.4. *The majority coordination problem can be solved with 2 bits of memory*

The model we will consider is one where in addition to opinion (red/blue), all players have strength of opinion (Strong/Weak). This turns out to be the second bit of memory needed to solve the problem. Furthermore, all players are strong initially. The updating rule goes in two steps:

- **Update color:** Strong influences weak, equal strength follows voter model
- **Update strength:** Two strong of different colors reduce both to weak, different strengths swap strenghs, two weak stays the same

For more details on the model and proof of Theorem 1.4, see [11].

Table 2: Majority coordination

	Memory	Time	Required Advice
[8]	1	impossible	
[1]	2	$< \infty$	none
Strong-Weak	2	$O(n^3)$	none
[6]	$O(\log(n))$	$O(n^7)$	$\|V\|$
Wait-and-see	expected $O(\log(\Delta))$	$O(d + log(n)) \cdot \log(n)$	Θ(broadcast $\cdot \|E\|$)

2 Which voters to buy - a different model

A natural question to ask, and one of high relevance in many applications, is which individuals in a group are the most critical in forming the opinion of the group. In other words, in a voter model, if you can change the opinions of k people from $-$ to $+$, who should you choose in order to maximize the probability that consensus converges to all $+$.

If we consider the voter model, it is easy to see that it does not matter which subset of nodes one chooses, since the random walk will be the same and thus the stationary distribution is the same for any possible state. The voter model is thus not a good model to study individual influences in a network.

We will therefore consider a different model. The model we will look at is a synchronous model where $X_v(t+1) = +/-$ with probabilities $\#\{w \in N(v) : X_w(t) = +/-\}$. We want to change $X_v(0)$ for k nodes to maximize the probability of final vote to be $+$. The question remains: Who should we choose?

Proposition 2.1. *In the synchronous model presented above, choosing the nodes of the opinion we want to change with the highest degrees maximizes the probability of changing the consensus outcome of the whole network.*

This follows from the fact that $\sum_v d(v)X_v(t)$ is a martingale.

2.1 Viral marketing and the Independent Cascade Model

It is well-known that word-of-mouth can be an effective marketing strategy. There is a lot of examples, especially from internet where a service like Hotmail spread “virally” by word-of-mouth rather than traditional advertising. These phenomena have led to the natural question of which individuals in a market are most likely to influence the rest of the market. The goal of **viral marketing** is to identify the influence of indivudials in a social network on the rest of the network, and thereby be able to target a smaller subset of “trendsetters” in the network.

One model to study the viral marketing problem is the independent cascade model. As usual we operate on a social network graph $G = (V, E)$, where each node can be in two states, deactivated or activated. Initially all nodes are deactivated. If a node u is activated, it gets one chance to activate each neighbor. The probability that a neighbor v of u is activated, is $p_{u,v}$. We then define an “influence function”, $F : 2^V \to [0, n]^{|V|}$, such that $F(S) := E(\#$ *of infected individuals if we initially activate subset* $S)$. The viral marketing problem is then the optimization problem $\max_{|S|=k} F(S)$ where a number of k individuals is chosen initially.

Submodularity The general influence maximization problem is *np*-hard, but it turns out that the independent cascade model belongs to a subclass of problems for which there exist an easy approximation algorithm, namely the sub-modular function.

Figure 3: The independent cascade model

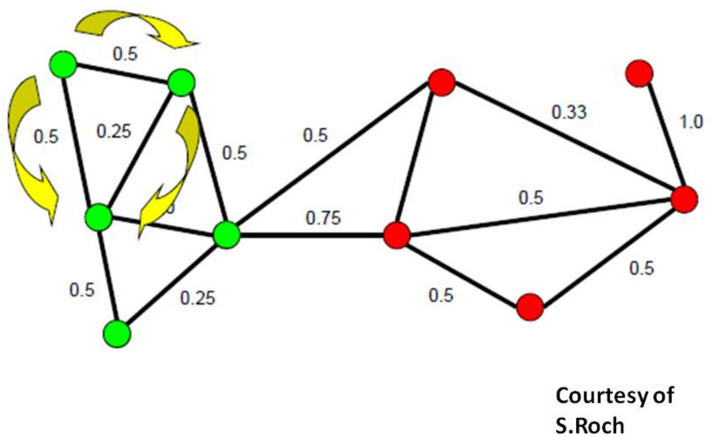

Definition 1. *A set function $f : V \to R$ is* ***submodular*** *if $\forall\, A, B \in V$*

$$f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$$

We return to our independent cascade model and we consider four different initially activated sets, S_0, S_1, $S_0 \cup S_1$ and $S_0 \cap S_1$ and a, b, c and d the expected size of the marketed set given each of the initial four sets. Then

Claim 1. $c + d \leq a + b$

This can be shown by using randomness to decide which edges to copy. The idea is to consider the random induced graph $G' \subseteq G$ of infectious edges, and count the expected number of vertices connected to the initial four sets. For a complete proof see [7].

This means that the expected size of the infected set is a submodular function of the set.

Theorem 2.2. *Let $f : 2^{[n]} \to [0, 1]$ be monotone and submodular. We want to solve the maximization problem $\max_{|S|=k} f(S)$. Then the greedy algorithm provides a $(1 - 1/e - \epsilon)$ approximate solution to the problem.*

Proof: Let $S_1, \ldots, S_k$ be the sets chosen by the greedy algorithm, and denote the optimal set by S^*. Write $x_i = f(S_i) - f(S_{i-1})$. Then

$$f(S^*) \leq f(S_i \cup S^*) \leq f(S_i) + kx_{i+1}$$
$$x_{i+1} \geq \frac{f(S^*) - f(S_i)}{k}$$

Then by induction it follows that

$$f(S_i = f(S_{i_1} + x_i \geq f(S^*)(1 - (1 - 1/k)^i)$$

We then take $i = k$ and obtain the claim

$$f(S_k) \geq f(S^*)(1 - 1/e)$$

□

Most general case We now introduce a general model for influence. Consider a single effected vertex v with the set S of infected neighbors. The vertex v draws u_v uniform$(0, 1)$. At any stage v is infected if $F_v(S) \geq u_v$.

This motivates the general result

Theorem 2.3. *If all F_v are monotone and submodular, then the influence function σ is monotone and submodular.*

See [10] for proof.

References

[1] F. Benezit, P. Thiran, and M. Vetterli. Interval consensus: From quantized gossip to voting. *ICASSP 2009*, pages 3661–3664, 2009.

[2] P. Clifford and A. Sudury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.

[3] R. Holley and T. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3(4):643–663, 1975.

[4] M. Kearns, S. Judd, J. Tan, and J.Wortman. Behavioral experiments on biased voting in networks. *PNAS*, 2009.

[5] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.

[6] M. Kearns and J. Tan. Biased voting and the democratic primary problem. *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE08)*, pages 639–652, 2008.

[7] J. Kempe, D. Kleinberg and E. Tardos. Maximizing the spread of influence through a social network. *Proc. 9th Intl. Conf. on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

[8] M. Land and R. Belew. No perfect two-state cellular automata for density classification exists. *Phys. Rev. Lett.*, 74(25):5148–5150, 1995.

[9] B. Latan and T. L'Herrou. Spatial clustering in the conformity game: Dynamic social impact in electronic groups. *Journal of Personality and Social Psychology*, 70(6):1218–1230, 1996.

[10] E. Mossel and S. Roch. On the submodularity of influence in social networks. *eprint arXiv:math/0612046*, 2006.

[11] E. Mossel and G. Schoenebeck. Reaching consensus on social networks. *Proceedings of the The First Symposium on Innovations in Computer Science*, 2009.