Diffusion of Innovations on Random Networks: Understanding the Chasm

Marc Lelarge

INRIA-ENS
Paris, France
marc.lelarge@ens.fr

Abstract. We analyze diffusion models on sparse random networks with neighborhood effects. We show how large cascades can be triggered by small initial shocks and compute critical parameters: contagion threshold for a random network, phase transition in the size of the cascade.

1 Introduction

In Crossing the Chasm [12], Moore begins with the diffusion of innovations theory from Everett Rogers [16], and argues there is a chasm between the early adopters of the product (the technology enthusiasts and visionaries) and the early majority (the pragmatists). According to Moore, the marketer should focus on one group of customers at a time, using each group as a base for marketing to the next group. The most difficult step is making the transition between visionaries (early adopters) and pragmatists (early majority). This is the chasm that he refers to.

In this paper, we analyze a simple model of diffusion with neighborhood effects on random networks and we show that it can explain this chasm. Most of the epidemic models [14], [15] consider a transmission mechanism which is independent of the local condition faced by the agents concerned. But if there is a factor of persuasion or coordination involved, relative considerations tend to be important in understanding whether some new behavior or belief is adopted [17].

We begin by discussing one of the most basic game-theoretic diffusion models proposed by Morris [13]. Consider a graph G in which the nodes are the individuals in the population and there is an edge (i, j) if i and j can interact with each other. Each node has a choice between two possible behaviors labeled A and B. On each edge (i, j), there is an incentive for i and j to have their behaviors match, which is modeled as the following coordination game parameterized by a real number $q \in (0, 1)$: if i and j choose A (resp. B), they each receive a payoff of q (resp. $(1 - q)$); if they choose opposite strategies, then they receive a payoff of 0. Then the total payoff of a player is the sum of the payoffs with each of her neighbors. If the degree of node i is d_i and S_i^B is her number of neighbors playing B, then the payoff to i from choosing A is $q(d_i - S_i^B)$ while the payoff from choosing B is $(1 - q)S_i^B$. Hence i should adopt B if $S_i^B > qd_i$ and A if $S_i^B \leq qd_i$. A number of qualitative insights can be derived from a diffusion model

C. Papadimitriou and S. Zhang (Eds.): WINE 2008, LNCS 5385, pp. 178–185, 2008.

even at this level of simplicity. Specifically, consider a network where all nodes initially play A. If a small number of nodes are forced to adopt strategy B (the seed) and we apply best-response updates to other nodes in the network, then these nodes will be repeatedly applying the following rule: switch to B if enough of your neighbors have already adopted B. There can be a cascading sequence of nodes switching to B such that a network-wide equilibrium is reached in the limit. This equilibrium may involve uniformity with all nodes adopting B or it may involve coexistence, with the nodes partitioned into a set adopting B and a set sticking to A. Morris [13] considers the case of infinite regular graph G and provides graph-theoretic characterizations for when these different types of equilibria arise.

Our work allows us to study rigorously an extension of this model, the symmetric threshold model, when the underlying network is a random network with given vertex degrees. We are able to characterize the relation between the network and the individual behavior. In particular, we compute the contagion threshold of the random network and validate a heuristic result of Watts [18]. We also show that there is a phase transition for the set of adopters at a critical value of the size of the initial seed. To the best of our knowledge, this result is new and our work is the first rigorous analysis of a general threshold model on a random network. Although random graphs are not considered to be highly realistic models of most real-world networks, they are often used as first approximation and are a natural first choice for a sparse interaction network in the absence of any known geometry of the problem.

In [4], the influence maximization problem is defined as follows: given a social network, find a small set of 'target' individuals so as to maximize the number of customers who will eventually purchase the product following the effect of word-of-mouth. Hardness results have been obtained in [8],[3] and there is a large literature on this topic. However, in most practical cases, the structure of the underlying network is not known and then one has to rely on distributional assumptions (like distribution of the degrees). Our model allows to answer the probabilistic version of the influence maximization problem, when the exact topology of the social network is not known.

The rest of the paper is organized as follows. In Section 2, we describe our model. Section 3 contains the main results in particular, the contagion threshold is computed and the phase transition phenomena is explained. Section 4 contains technical details and we conclude in Section 5.

2 Model

2.1 The Configuration Model

In this section, we define our random graph model which is standard in the literature on random graphs [2]. Let $n \in \mathbb{N}$ and let $(d_i)_1^n = (d_i^{(n)})_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even. We define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n, (d_i)_1^n)$ by the configuration model [2]: take a set of d_i half-edges for each vertex i and combine the half-edge into pairs by a uniformly random matching of the set of all half-edges. Conditioned on the multigraph $G^*(n,(d_i)_1^n)$ being a simple graph, we obtain a uniformly distributed random graph with the given degree sequence, which we denote by $G(n,(d_i)_1^n)$.

We will let $n \to \infty$ and assume that we are given $(d_i)_1^n$ satisfying the following regularity conditions [11]:

Condition 1. *For each n, $(d_i)_1^n = (d_i^{(n)})_1^n$ is a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even and, for some probability distribution $(p_r)_{r=0}^\infty$ independent of n,*

(i) $\#\{i : d_i = r\}/n \to p_r$ *for every* $r \geq 0$ *as* $n \to \infty$*;*
(ii) $\lambda := \sum_r r p_r \in (0,\infty)$*;*
(iii) $\sum_{i=1}^n d_i/n \to \lambda$ *as* $n \to \infty$*;*
(iv) $\sum_i d_i^2 = O(n)$.

In words, (p_r) describes the distribution of the degrees, λ is the average mean degree in the graph, condition (iii) ensures that the number of edges divided by n tends to the average degree divided by 2. The technical condition (iv) is required to transfer the results from $G^*(n,(d_i)_1^n)$ to $G(n,(d_i)_1^n)$ [5].

The results of this work can be applied to some other random graphs models too by conditioning on the vertex degrees. For example, for the Erdös-Rényi graph $G(n,p)$ with $np \to \lambda \in (0,\infty)$, the assumptions hold with p_r the distribution of a Poisson random variable with mean λ.

We consider asymptotics as $n \to \infty$ and say that an event holds w.h.p. (with high probability) if it holds with probability tending to 1 as $n \to \infty$.

2.2 Symmetric Threshold Model

The contagion model of [13] is the simplest model for cascading behavior in a social network: people switch to a new behavior when a certain threshold fraction of neighbors have already switched. Our symmetric threshold model generalizes this model by allowing the threshold fraction be a random variable with distribution depending on the degree of the node and which are independent among nodes. This is to account for our lack of knowledge of the exact threshold value of each individual. Formally, we define for each $d \in \mathbb{N}$, a sequence of i.i.d. random variables in $\mathbb{N}$ denoted by $(K(d), K_i(d))_{i=1}^\infty$. The threshold associated to node i is $K_i(d_i)$ where d_i is the degree of node i.

Now the progressive dynamics of the behavior operates as follows: some set of nodes S starts out adopting the new behavior B; all other nodes start out adopting A. We will say that a node is active if it is following B. Time operates in discrete steps $t = 1, 2, 3, \dots$. At a given time t, any inactive node i becomes active if its fraction of active neighbors exceeds its threshold $K_i(d_i) + 1$. This in turn may cause others nodes to become active leading to potentially cascading adoption of behavior B. We will suppose that $K_i(1) = 0$ for all i, so that any leaf of the network is active as soon as its parent becomes active.

It is easy to see that the final set of active nodes (after n time steps if the network is of size n) only depends on the initial set S (and not on the order of the activations) and can be obtained as follows: set $X_i = 1$ for all i in the set of initial adopters. Then as long as there exists i such that $\sum_{j\sim i} X_j > K_i(d_i)$, set $X_i = 1$. When this algorithm finishes, the final state of node i is represented by X_i: $X_i = 1$ if node i is active and $X_i = 0$ otherwise. It is easily seen that the linear threshold model [9] is covered by our framework (see [10] for a proof).

3 Main Results

3.1 Contagion Threshold of a Random Graph

We consider the simple contagion model studied by Morris in [13] on a random graph, i.e. $K_i(d) = qd$ for all i. We define the contagion threshold of the graph to be the maximum q for which a single individual can trigger a global cascade, i.e. activate a strictly positive fraction of the total population, w.h.p. This notion is the natural extension of the contagion threshold defined in [13] for regular graphs.

Proposition 1. *The contagion threshold q_c is given by*

$$q_c = \sup \left\{ q : \sum_{1\leq s<1/q} s(s-1)p_s > \lambda \right\}.$$

This result is in accordance with the heuristic result of [18] (see in particular the cascade condition Eq. 5 in [18]) and is proved in Section 4. Figure 1 gives the contagion threshold as a function of λ, the mean degree of the graph.

Note that q is related to the quality of the new technology: the lower q is, the better the quality of the new technology is. In particular if $q < 1/2$, then technology B is better than technology A. Hence q_c can be interpreted as the minimal quality for technology B to get a non-negligible adoption with a finite initial seed of adopters.

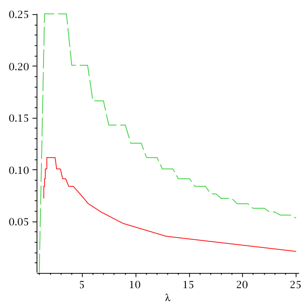

Fig. 1. q_c for the contagion model on a Poisson random graph (green dashed) and on a Power-law random graph (red) as a function of λ

3.2 Phase Transition in the Contagion Model

We now still consider the contagion model but for $q > q_c$. In this case, in order to trigger a large cascade, the set of initial adopters must be a non-negligible fraction of the total population. For simplicity, we assume that each node of the network is part of the initial set of adopters with probability α independently of everything else. In particular, the fraction of initial adopters is α and we now compute the final proportion of active nodes: $\Phi(\alpha) = \lim_{n\to\infty} \frac{\sum_{i=1}^n X_i}{n}$.

We need to introduce some notation first. For integers $\ell \geq 0$ and $0 \leq r \leq \ell$ let $b_{\ell r}$ denote the binomial probabilities $b_{\ell r}(p) := \binom{\ell}{r} p^r (1-p)^{\ell - r}$. We denote by D a random variable with distribution $\mathbb{P}(D = r) = p_r$. For $0 \leq p \leq 1$ we let D_p be the thinning of D obtained by taking D points and then randomly and independently keeping each of them with probability p: $\mathbb{P}(D_p = r) = \sum_{\ell=r}^{\infty} p_\ell b_{\ell r}(p)$. We now define $h(p) = \mathbb{E}[D_p \mathbb{1}(D_p \geq (1-q)D)]$.

The following proposition shows that the map $\alpha \mapsto \Phi(\alpha)$ exhibits point of discontinuity.

Proposition 2. *Consider a random graph such that $p_1 > 0$ and let $\tilde{p}$ be the largest local maximum point of $\psi(p) = h(p)/p^2$ in $(0,1)$. Then there is a phase transition at $\alpha_c = 1 - \frac{\lambda}{\psi(\tilde{p})}$: the function $\Phi(.)$ is discontinuous at α_c.*

Figure 2 shows an example of such a phase transition in the case of Poisson random graphs.

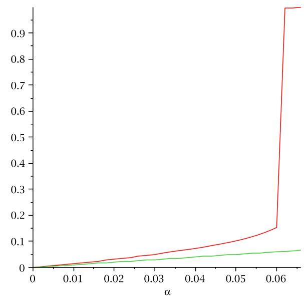

Fig. 2. Function $\Phi(\alpha)$ for the contagion model on a Poisson random graph with parameter $\lambda = 6$ and $q = 0.3$

Returning to the (probabilistic) influence maximization problem, our derivation of the function $\Phi(\alpha)$ is of crucial importance. In particular, the fact that this function is highly non-linear seems not to have been taken into account so far and will have a big impact on the optimal strategy. In the case where the marketer knows the degree of each individual (but not the underlying social network), our derivation of $\Phi(\alpha)$ will allow her to target her effort, by choosing the variable α.

3.3 Dynamic of the Epidemic

In previous section, α was related to the amount spent by the marketer and q corresponded to the quality of the new technology. We now consider that α is actually fixed and corresponds to the fraction of technology enthusiasts in the population. The rest of the population consists of pragmatists. Then the marketer's effort allows to increase the perceived quality by decreasing the value of q. It is easy to see that the phase transition described in previous section translates in a phase transition in the parameter q. Moreover, let consider the simple following dynamic of the epidemic: the edges of the active nodes become active(meaning that the end-point of the edge actually notices that his neighbor is active) at rate 1 (see [10] for more details) . Then Figure 3 shows the case where the real quality of the technology is $q = 0.3$. Without any marketing, a small fraction of the pragmatists adopt the new technology but with marketing, the diffusion is able to 'cross the chasm' and a large fraction of the population adopt the new technology.

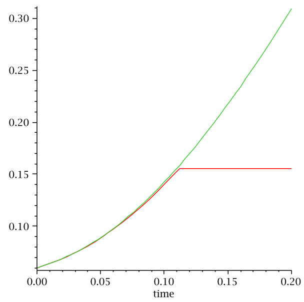

Fig. 3. Dynamic of the epidemic for Poisson random graph with $\lambda = 6$, $\alpha = 0.06$ and for $q = 0.29$ and $q = 0.3$

4 Exact Asymptotics

In this section, we state the theorem which is the corner stone of our work (see [10] for a proof). Recall that D_p is the thinning of D (defined in Section 3.2). We define the functions

$$h(p) := \mathbb{E}\left[D_p \mathbb{1}(D_p \geq D - K(D))\right], \tag{1}$$

$$h_1(p) := \mathbb{P}(D_p \geq D - K(D)). \tag{2}$$

Theorem 1. *Consider the graph $G(n, (d_i)_1^n)$ satisfying Condition 1 wehre each node is part of the initial set of adopters with probability α independently of everything else. Let $\hat{p} := \max\{p \in [0, 1] : (1 - \alpha)h(p) = \lambda p^2\}$.*

(i) If $(1 - \alpha)h(p) < \lambda p^2$ for all $p \in (0, 1]$, which is equivalent to $\hat{p} = 0$, then w.h.p. $\Phi(\alpha) = 1$.

(ii) If $(1-\alpha)h(p) \geq \lambda p^2$ *for some* $p \in (0,1]$*, which is equivalent to* $\hat{p} \in (0,1]$*, and further* $\hat{p}$ *is not a local maximum point of* $(1-\alpha)h(p) - \lambda p^2$*, then w.h.p.* $\Phi(\alpha) = 1 - (1-\alpha)h_1(\hat{p})$.

The proof of this Theorem consists in an extension of the work of Janson and Luczak [7] where the k-core problem is studied. Our model is related to the bootstrap percolation which is more or less the opposite of taking the k-core: with our notation, it consists in taking $K_i(d_i) = k$ a fix constant. For regular graphs (i.e. $d_i = d$ for all i), this process has been studied. Theorem 1 of [1] or Theorem 5.1 of [6] correspond exactly to our Theorem in the particular case of a d-regular graph, with fixed threshold.

Proof. of Proposition 1: Following the heuristic in [18], we introduce the following threshold: $K_i(d) = (d+1)\mathbb{1}(d \geq 1/q)$. In words, a node i becomes active if one of his neighbor is active and $d_i < 1/q$. Clearly the nodes that become active in this model need to have only one active neighbor in the original contagion model with parameter q. For any node i, let C_i denotes the final set of active nodes when starting with only i as active node. Clearly, if $j \in C_i$, then we have $C_i = C_j$. Now if we prove that $\Phi(0+) = \lim_{\alpha\to 0}\Phi(\alpha) - \alpha > 0$, then for any $\Phi(0+) > \alpha > 0$, at least one of the $n\alpha$ nodes say i in the initial set has activated at least $\frac{\Phi(0+)}{\alpha}n$ nodes. Hence we have $\#C_i/n \geq \frac{\Phi(0+)}{\alpha}$ and any point in C_i will activate at least the set C_i in the original contagion model. We now prove that for $q < q_c$, we have $\Phi(0+) > 0$. This will implies that the contagion threshold is larger than q_c. We have

$$\begin{aligned} h(p) &= \mathbb{E}[D_p\mathbb{1}(D \geq 1/q) + D\mathbb{1}(D < 1/q, D_p = D)] \\ &= \sum_{s\geq 1/q} spp_s + \sum_{s<1/q} sp_s p^s \\ &= p\left(\sum_{s\geq 1/q} sp_s + \sum_{s<1/q} sp_s p^{s-1}\right). \end{aligned}$$

Let $f(p) = \lambda p - \frac{h(p)}{p}$. The condition $\Phi(0+) > 0$ is equivalent to for $\epsilon > 0$ small enough $f(1-\epsilon) > 0$. We have $f(1-\epsilon) = \epsilon\left(-\lambda + \sum_{s<1/q} s(s-1)p_s\right) + o(\epsilon)$, which is the condition of the proposition. The proof that for $q > q_c$ a single active node cannot activate a positive fraction of the population is similar and omitted.

Proof. of Proposition 2: Let $f(p,\alpha) = \lambda p^2 - (1-\alpha)h(p)$. Note that $f(0,\alpha) = 0$ and $f(1,\alpha) = \alpha\lambda)$. Then we have $f(p,\alpha) \sim -(1-\alpha)p_1 p < 0$ as $p \to 0$ and $\alpha\lambda - f(p,\alpha) \sim (1-p)\left((1-\alpha)(\sum_{s<1/q} s(s-1)p_s - \lambda) - 2\alpha\lambda\right) < 0$ as $p \to 1$ and the result follows easily.

5 Conclusion

We proposed a simple model of diffusion with neighborhood effects which allows to explain the 'chasm'. We should emphasize that the random graph model considered eliminates a lot of the network structure from the problem (only the degree distributions are preserved). We expect that other local effects like clustering will have a significant impact on the diffusion. However our work shows that neighborhood effects 'alone' can explain the 'chasm' and we think that these effects will actually 'add up'. These issues are left for future research.

References

1. Balogh, J., Pittel, B.G.: Bootstrap percolation on the random regular graph. Random Structures Algorithms 30(1-2), 257–286 (2007)
2. Bollobás, B.: Random graphs, 2nd edn. Cambridge Studies in Advanced Mathematics, vol. 73. Cambridge University Press, Cambridge (2001)
3. Chen, N.: On the approximability of influence in social networks. In: SODA 2008 (2008)
4. Domingos, P., Richardson, M.: Mining the network value of customers. In: KDD 2001 (2001)
5. Janson, S.: The probability that a random multigraph is simple, arxiv.org:math/0609802 (2006)
6. Janson, S.: On percolation in random graphs with given vertex degrees, arxiv.org:0804.1656 (2008)
7. Janson, S., Luczak, M.J.: A simple solution to the k-core problem. Random Structures Algorithms 30(1-2), 50–62 (2007)
8. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: KDD 2003 (2003)
9. Kleinberg, J.: Cascading behavior in networks: Algorithmic and economic issues. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
10. Lelarge, M.: Diffusion and cascading behavior in random networks (in preparation)
11. Molloy, M., Reed, B.: A critical point for random graphs with a given degree sequence. Random Structures Algorithms 6(2-3), 161–179 (1995)
12. Moore, G.: Crossing the Chasm. HarperBusiness, New York (2002)
13. Morris, S.: Contagion. Rev. Econom. Stud. 67(1), 57–78 (2000)
14. Newman, M.E.J.: Spread of epidemic disease on networks. Phys. Rev. E 66(1), 016128 (2002)
15. Pastor-Satorras, R., Vespignani, A.: Evolution and Structure of the Internet: A Statistical Physics Approach. Cambridge University Press, New York (2004)
16. Rogers, E.M.: Diffusion of Innovations, 5th edn. Free Press (August 2003)
17. Vega-Redondo, F.: Complex Social Networks. Cambridge University Press, Cambridge (2007)
18. Watts, D.J.: A simple model of global cascades on random networks. Proc. Natl. Acad. Sci. USA 99(9), 5766–5771 (2002) (electronic)