CSC200: Lecture 34

Allan Borodin

Annoucements and todays agenda

- Announcements
  1. This weeks office hours Wed noon-2; best to email to schedule meeting
  2. First 3 questions for A4 have been posted.
  3. Some students are trying to survey how much interest there is in the course CSC375 (the enriched version of CSC373) in the hope that it will be offered again. If interested in particpating in this survey (even if you are not currently interested in taking CSC375) please look at: https://csc.cdf.toronto.edu/mybb/showthread.php?tid=12601

- Todays agenda
  1. Review the threshold model for influence spread and how it is induced by perceived payoffs
  2. Viewing a social network as a platform for advertising (of products, ideas, etc), how do we select a small set of iinitial adopters?
  3. The stochastic linear threshold and independent cascade models of influence spread.
  4. This leads to the introduction of monotone submodular functions and an algorithm for computing an approximately optimal set of initial adopters.

The uniform threshold model for spread/diffusion

- We recall that the main emphasis in Chapter 19 is on a very simple process of diffusion where each person has 2 alternative decisions:
  1. stay with a current "product" B
  2. or switch to a (new) product A.
- Some number of individuals are enticed (at time $t = 0$) to adopt a new product A.
- We assume all other individuals in the network are initially using a different product B (or no product).
- The diffusion model is that every node has a threshold q (in absolute or relative terms) for how many of its neighbors must have adopted product A before it adopts A.

Uniform threshold model (continued)

- If at some time t, the threshold for a node v has been achieved, then by time time $t + 1$, v will adopt product A.
- If the threshold has not been reached then v decides not to adopt A at this time.

Note

Although it is not explicitly stated, the initial adopters never reverse their adoption.

- To reason about a threshold model, we assumed that for every edge (v, w) in the network
  - There is payoff a to v and w if both v and w have adopted product A.
  - There is payoff b to v and w if both v and w have adopted product B.
  - A zero payoff when v and w do not currently utilize the same product.

Payoffs induce threshold

- Suppose node v has not yet adopted A at time t, but a fraction p of the $d(v)$ neighbors of v have already adopted A, then:
  - By switching, the payoff to v is $p \times d(v) \times a$.
  - By staying with B, v has payoff $(1 - p) \times d(v) \times b$.

- Thus node v will switch to A if

$$p \times d(v) \times a \geq (1 - p) \times d(v) \times b$$

  (and for simplicity we say v switches when payoffs are equal).

- This is then equivalent to saying that v will switch whenever p is at least $\frac{b}{a+b} = q$ which is then the relative threshold.

- That is, whenever there is at least a (threshold) fraction q of the neighbours of node v that have adopted A, then v will also adopt A.

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- A previous example (in lecture 33) showed a complete cascade where all nodes eventually adopt A.
- In this example, "tightly-knit communities" block the spread.

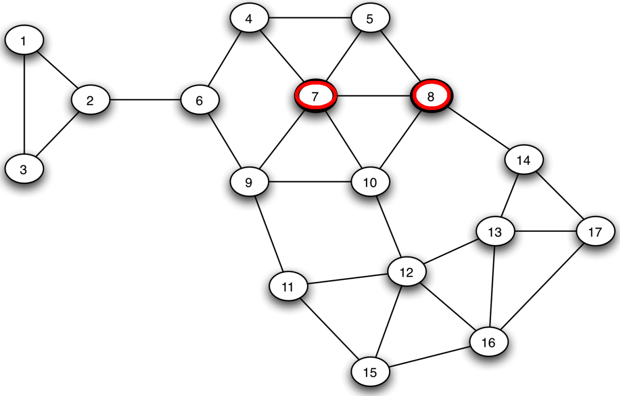

$t = 0$

[Fig 19.3, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- A previous example (in lecture 33) showed a complete cascade where all nodes eventually adopt A.
- In this example, "tightly-knit communities" block the spread.

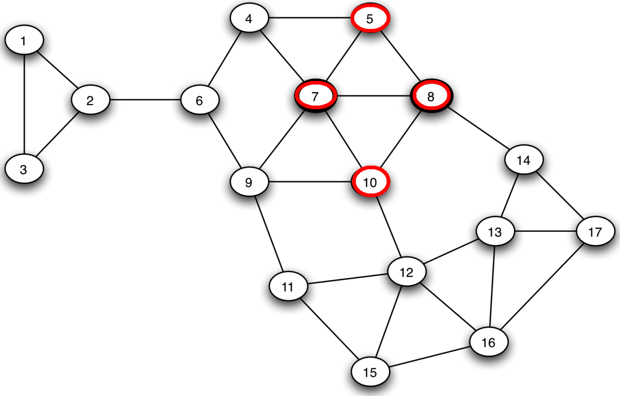

$t = 1$

[Fig 19.3, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- A previous example (in lecture 33) showed a complete cascade where all nodes eventually adopt A.
- In this example, "tightly-knit communities" block the spread.

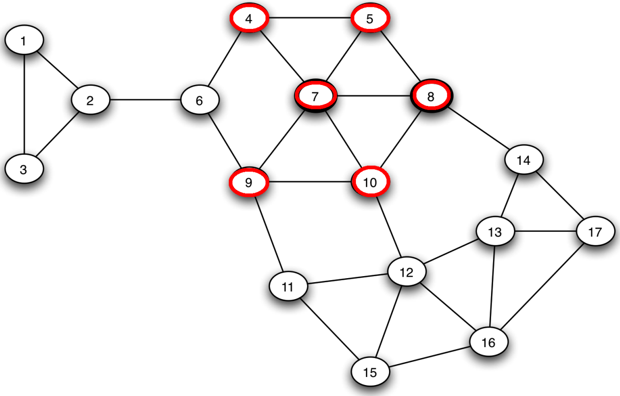

$t = 2$

[Fig 19.3, E&K]

Complete cascades vs tightly-knit communities (example: $a = 3$, $b = 2$, $q = 2/5$)

- A previous example (in lecture 33) showed a complete cascade where all nodes eventually adopt A.
- In this example, "tightly-knit communities" block the spread.

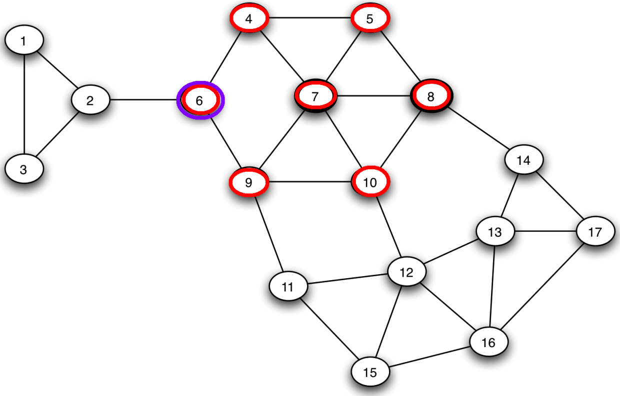

$t = 3$

[Fig 19.3, E&K]

Clusters vs complete cascades

- Suppose we have a network threshold spread model with threshold q, an initial set I of product A adopters and $V' = V - I$ is the set of nodes that are not initial adopters.

- Then we have the following (provable) intuitive result that characterizes when complete cascades will or will not form:

  - If V' contains a cluster C of density greater than $1 - q$, then the initial adopters will not cause a complete cascade. Furthermore, no node in C will adopt A.
  - If a threshold q spread network with initial set I of adopters does not cause a complete cascade, then the non initial adopters nodes V' must contain a cluster of density greater than $1 - q$.

When nodes have different thresholds

- As remarked before the assumption that all nodes have the same threshold is not essential.
- Consider a node v. Suppose now that for every adjacent edge (v, w), node v has payoff $a(v)$ (resp. $b(v)$) if both v and w have adopted product A (resp. B) and a zero payoff if v and w currently utilize different products.
- If node v has not yet adopted A at time t, but a fraction p of the $d(v)$ neighbours of v have already adopted A, then:
  - By switching, v has payoff $p \times d(v) \times a(v)$.
  - By staying with B, v has payoff $(1 - p) \times d(v) \times b(v)$.
- Thus node v will switch to A if

$$p \times d(v) \times a(v) \geq (1 - p) \times d(v) \times b(v).$$

- This is then equivalent to saying that v will switch whenever

$$p \geq \frac{b(v)}{a(v) + b(v)} = q(v)$$

which is then the threshold for node v.

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

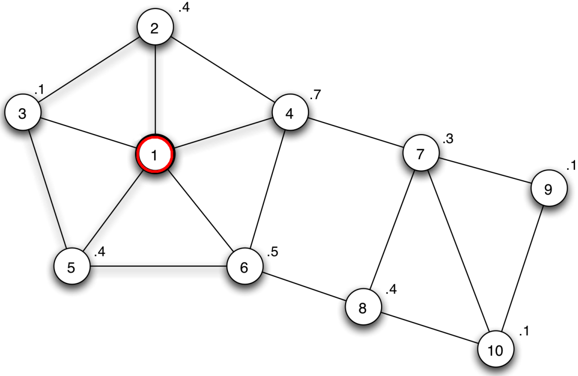

$t = 0$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

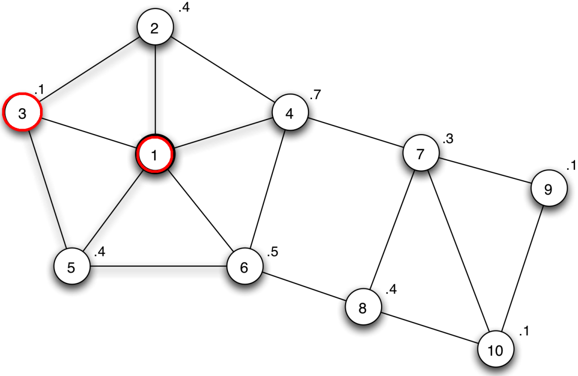

$t = 1$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

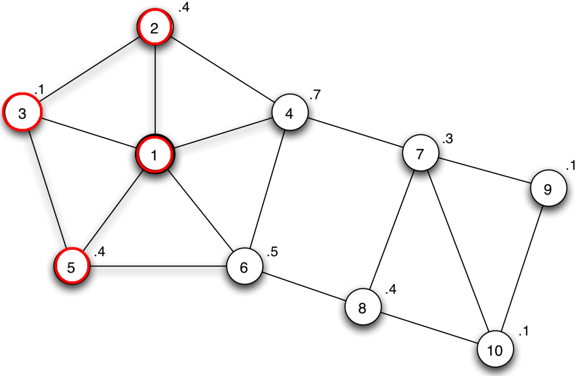

$t = 2$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

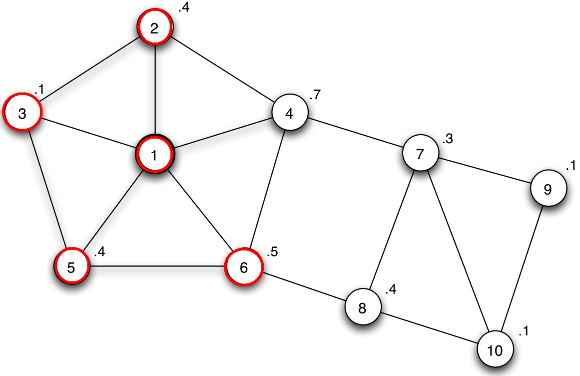

[Fig 19.13, E&K]

$t = 3$

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

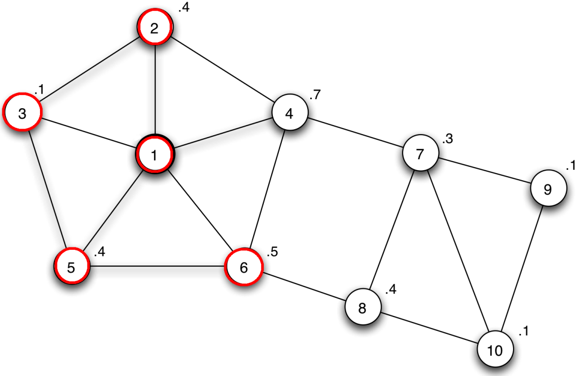

$t = 4$

[Fig 19.13, E&K]

Redefining blocking clusters

- A blocking cluster is now a set of nodes C such that every node $v \in C$ has more than a fraction $1 - q(v)$ of its adjacent nodes in C.
- It follows (as in the case of uniform thresholds) that a given set of adopters I in a network will not cause a complete cascade iff $V - I$ contains a blocking cluster C.

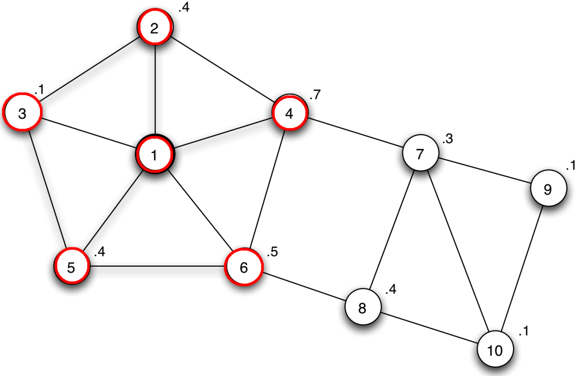

$t = 5$

[Fig 19.13, E&K]

Factors determining the rate and extent of diffusion in a social network

1. The structure of the network.

2. The relative payoffs vs costs for adopting a new product.
   - We haven't spoken of costs yet but we usually do have a cost for adopting a new product.
   - We can introduce such a cost into the model by saying that v will not adopt the new A unless

$$p \times d(v) \times a(v) \geq (1 - p) \times d(v) \times b(v) + cost$$

   - As in our discussion of population wide direct benefit effects, we can also add intrinsic values $v_A(v)$ and $v_B(v)$ to both sides of the above inequality to determine the threshold for v adopting A. Then v will switch to A if

$$p \times d(v) \times a(v) + v_A(v) \geq (1 - p) \times d(v) \times b(v) + v_B(v) + cost$$

Choosing a set influential initial adopters

- As an advertiser, manufacturer, politician, etc our goal is to influence as many people as possible. This raises an interesting computational question as to how to select the most influential nodes (within some budgetary constraint).

- That is, suppose we wish to spread a new technology and to do so we have money to influence some "small" set of initial adopters (e.g. by giving away the product or even paying people to adopt it).

- Even in this simple model of (non-competitive) influence spread, and even if we have complete knowledge of the social network, it is not at all clear how to chose an initial set of adopters so as to achieve the largest spread.

- Furthermore the spread process could be much more sophisticated.
  - For example, adoption by a node will likely be a random process (say adopting with some probability relative to the nodes threshold) and
  - Maybe the influence of neighbors first increases and then decreases over time

Choosing influential adopters continued

- To simplify the discussion, suppose we have funds/ability to influence k nodes to become initial adopters.
  - We can try all possible subsets of the entire $n = |V|$ nodes and for each such subset simulate the spread process.
  - But clearly as k gets larger, this becomes prohibitive especially for large networks.

- Earlier in the course, we mentioned that for many optimization problems (like the one being considered now), there is a widely held belief (with good supporting evidence from complexity theory):

Such "NP-hard problems" cannot be optimally solved in an efficient manner (and sometimes we cannot even get a good approximation to optimality).

Can we determine a "good" set of initial adopters?

- For even simple models of information spread as being discussed here, complexity theory (the P vs NP conjecture) argues that we cannot efficiently choose the best set of initial adopters.

- And even for our simple deterministic threshold process, there is a class of networks for which (assuming the complexity conjecture) it is not possible to obtain an approximation within a factor n^c for any $c < 1$.

- Instead we will identify properties of a spread process that will allow a good approximation: a good set of initial adopters that will do "almost as well" as the best set.

Computationally manageable influence maximization models; monotone submodular set functions

- Some spread models have the following nice properties.

Let $f(S)$ be size (or more generally a real value benefit since some nodes may be more valuable) of the final set S of adopters satisfying:

1. Monotonicity: $f(S) \leq f(T)$ if S is a subset of T
2. Submodularity: $f(S + v) - f(S) \geq f(T + v) - f(T)$ if S is a subset of T

- Where have we seen such functions before?
- The simple threshold examples considered thus far are monotone processes but are not submodular in general. Are these contrived worst case network examples?
- But some variants of the threshold model and related models do satisfy these properties. We consider two such stochastic models.

The St. Petersburg Paradox

- Monotone submodular set functions arise naturally in many contexts. One important context is in our utility for money. Although we don't try to specify our utility, it seems reasonable to believe that most of us have some diminishing marginal value for money. That is, once we have obtained some sufficiently large enough amount x of wealth, our utility for say $2x is not as much as twice our utility for $x.
- The St. Petersburg Paradox (attributed to Nicolas Bernoulli) is the following casino game: We ask a gambler to pay for the chance to participate in the
  1. We flip a fair coin and observe the first time t we get heads.
  2. We pay the gambler 2^t \$ (or really rubles)

  How much would you pay to take the gamble?
- If we consider the value for a set S of $|S|$ one dollar bills, then if our utility for money was a linear function $f(S) = c \cdot |S|$, then the expected utiltiy for the gamble is infinite; namely $\sum_{i=1\ldots} 2^{-i} 2^{i}$.
- One possible resolution (given by Daniel Bernoulli) is that we have a sharply decreasing utility for wealth.