To Inoculate or Not to Inoculate—That is the Game

Eric Brenner, Daoji Huang, Daniel Schroeder
Carleton College

0. Model Description

Inoculation Game: an Introduction

In the original inoculation game, each player is a node on a connected graph of n people. Each player chooses a strategy, either inoculating against an infection or not inoculating. Then, a single player is randomly chosen and infected. The infection spreads between uninoculated people who are connected, but does not pass through those who are inoculated. If an inoculated player is chosen to be randomly infected, the infection does not spread. For a node v, the **attack component** is all the nodes capable of infecting v.

Players are assigned a cost according to what decision they and the other players have made. The cost function for a node v is

$$Cost(v) = \begin{cases} C & \text{if inoculated} \\ L\dfrac{k}{n} & \text{if uninoculated.} \end{cases}$$

where C is the cost of choosing to inoculate, L is the cost of infection, and k is the size of the attack component of v.

The Modified Game—Inoculation Succeeds with a Probability

In reality, inoculations do not always succeed. As a result, in our modified version of the game, a new parameter is added: the probability p that an inoculation succeeds. Now, each player that decides to inoculate is successfully inoculated with probability p, and is left uninoculated otherwise. The spread of the inoculation now depends on which node is initially infected and on which nodes succeed or fail to inoculate.

Player costs are changed accordingly. Uninoculated node costs are now based on the *expected size* of the attack component. Inoculated players still pay the cost C of inoculating, but in the $(1 - p)$ chance that the inoculation fails, they are treated as uninoculated players and additionally pay the cost described above.

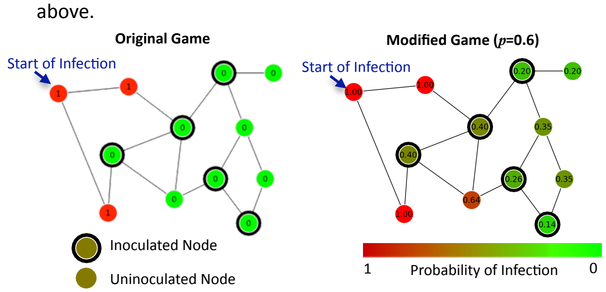

Comparing the spread of an infection in the two games (the numbers in the nodes denote the probability of getting infected)

Terminology

Best response dynamics is a process in which subsequent players are given the chance to change strategies, if doing so will lower the player's cost. A **Nash equilibrium** is a configuration of strategies in which no player will change strategies under best response. There may exist multiple Nashes, even for the same C and p values. An **optimal solution (OPT)** is a configuration of strategies that minimizes the **social cost**, the sum of the costs of the players. The **Price of Anarchy (PoA)** is the ratio of the largest social cost under any Nash over the social cost under an optimal solution.

1. When is this model nice to work with?

Players in an attack component share the same cost because if any single member gets infected, the others do as well. In **complete graphs** (graphs in which every pair of distinct vertices is connected by a unique edge), we find our model friendly to work with. Because nodes in an attack component all share the same cost, all uninoculated nodes have the cost function $L(n - mp)/n$ where m is the number of players inoculated and all inoculated nodes have the cost function $C + L(1 - p)(n + p - pm)/n$.

Given values of C, p, and n, we can calculate the number of inoculated nodes in the Nash and the number in the OPT. Furthermore, the number of inoculated nodes in the Nash never exceeds the number in the OPT. We demonstrate these below.

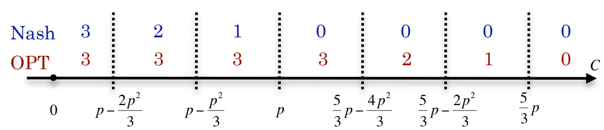

For the complete graph of 3 nodes, the number inoculating under Nash and OPT for $0 < p < 1/2$ as a function of C.

We also found that the PoA is maximized when $C = p$. We would like to investigate whether this is true for general complete graphs.

As in complete graphs, all inoculated nodes have the same cost and all uninoculated nodes have the same cost in cycle graphs with evenly spaced inoculations. Such cycle graphs have given us insight on proving the existence of Nash on cycle graphs for any C and p values. So far we have only proven the existence of Nash equilibria on complete graphs.

2. When is this model difficult to work with?

In the original model, [1] has shown using a potential function that a Nash always exists. However, best response does not always terminate in the modified game, so there is no such potential function based solely on best response.

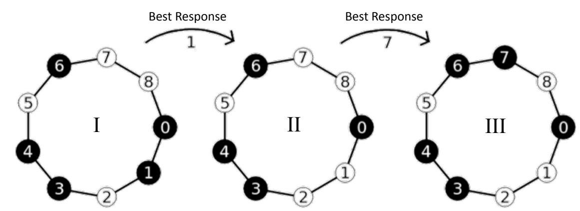

A case where best response reaches symmetrical states I and III; repeating this pattern will return to state I (C=0.291666, p=0.5).

At the moment, the algorithm we have for calculating the expected size of a player's attack component has complexity $O(2^m)$. As a result, it is difficult to analyze Nashes and OPT in general graphs.

3. General Graph Results

- If some player a inoculates as a result of best response dynamics, then the social cost decreases by at least as much as a's individual cost decreases.
- When $C > p$, the only Nash is for nobody to inoculate.
- Let E_{min} be the smallest possible attack component size of any node in the graph when all of the nodes choose to inoculate. If $C < pE_{min}/n$, then the only Nash is for everybody to inoculate.

4. References and Acknowledgements

We thank our advisors, Professors David Liben-Nowell, Tom Wexler, Alexa Sharp, and Kevin Woods for their guidance and suggestions. Special thanks to Mike Tie for saving us from a raging tornado.

This work was supported in part by NSF grant CCF-0728779 and by grants from Oberlin College and Carleton College.

[1]J.Aspnes, K. Chang, and A. Yampolskiy. Inoculation Strategies for Victims of Viruses and the Sum-of-Squares Partition Problem. In *Proc. 16th ACM-SIAM Annual Symposium on Discrete Algorithms (SODA)*, 2005