


# Worm vs Alert


Jared Saia


Joint with: James Aspnes, Bruce Kapron,  
David Kempe, Valerie King, Navin Rustagi,  
Amitabh Trehan, and Vishal Sanwalani


# Worm vs Alert


- Worm and Alert are fighting for control of a network
- Initially: one node infected, no nodes are alerted,  $\gamma$  fraction of the nodes are **detector** nodes
- Each alerted node sends out  $\alpha$  alerts per round, each infected node sends out  $\beta$  worms per round. When detectors receive a worm, they become alerted.
- Alerts can only spread through sparse alert network














# Comparison

- Advantage worm
  - head start
  - omniscient, except detector location
  - unconstrained by alert network
- Advantage alert
  - hidden detector nodes

# Analysis

- Just solve some DFQs to find who wins, right? **NO!**
- Q: Is there an alert strategy that saves almost all nodes, no matter what strategy the worm uses

# Answer

- Yes! provided that alert network has good expansion properties
- Strategy for alert is simple: each alerted node sends out  $\alpha$  alerts to randomly selected neighbors each round
- All intelligence is in the alert network

# Expansion

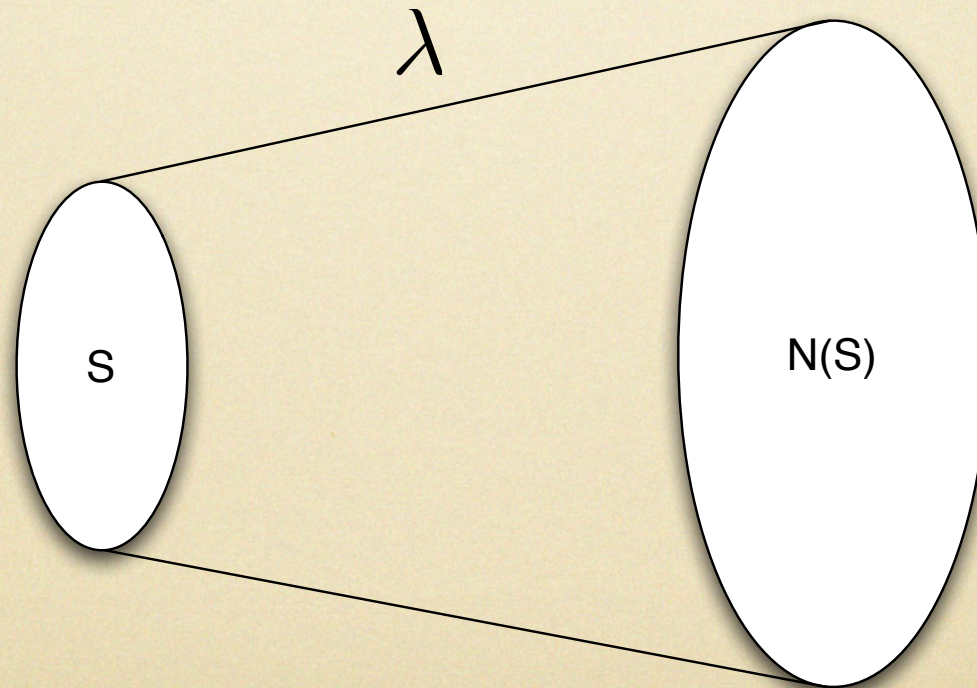
- A graph has expansion factor  $\lambda$  if for every vertex set  $S$  which is “not too large”:

$$|N(S)| \geq \lambda|S|$$

- Where  $N(S)$  is the set of neighbors of  $S$

# Expansion

Alert network is  $d$ -regular graph with expansion  $\lambda$




# Theorem

- Let  $r = \alpha/\beta$  and  $\gamma > 0$
- If 
$$\frac{r}{1 - \gamma} > \frac{2d}{\lambda}$$
- Then only  $o(1)$  fraction of nodes infected with probability  $1 - o(1)$


# Good News

- Even if alert spreads at same rate as worm, alert can save almost all nodes as  $n$  gets large
- True even though: 1) worm has head start, 2) worm can spread however it wants, 3) worm knows alert strategy

# Bad News



# Good News



$$\beta = 1$$

$$n = 10^6$$

$$d = 100$$

# Outcomes


- "Worm versus alert: Who wins in a battle for control of a large-scale network?" by James Aspnes, Navin Rustagi and Jared Saia, *In Principles Of Distributed Systems (OPODIS)*, 2007 ``A beautiful study of a problem with lots of appeal''


# Problem: False Alerts





- Vigilante: No False Alerts; poor coverage
- Most Systems: Some False Alerts; better coverage
- Goal: Algorithms that gracefully handle false alerts














# Good News

- With TTL fields on alerts:
- Each false alert spreads to at most polylog nodes
- Can still “catch” any fast worm
- Fast: spreads in logarithmic time

# Theorem

- If  $1 + \alpha\lambda/(2d) \geq 2 + 2\beta(1 - \gamma)$
- Then only  $o(1)$  fraction of nodes infected with probability  $1 - o(1)$
- **and** any false alert propagates to only polylog nodes

# Bad News

- With TTL fields:
- Provably can't catch a slow, stealthy worm
- Problem: slow worm infects one node at a time; if that node is a detector, quickly builds firewall around the alerts.

# Big Question

- Is there something smarter than TTL that will
  - catch a slow worm; and
  - quickly squelch false alerts???