Catch Me If You Can: Finding the Source of Infections in Temporal Networks

AAMAS 2026

Ben Bals
Centrum Wiskunde & Informatica and VU Amsterdam

Why Temporal?

Static SIR Model

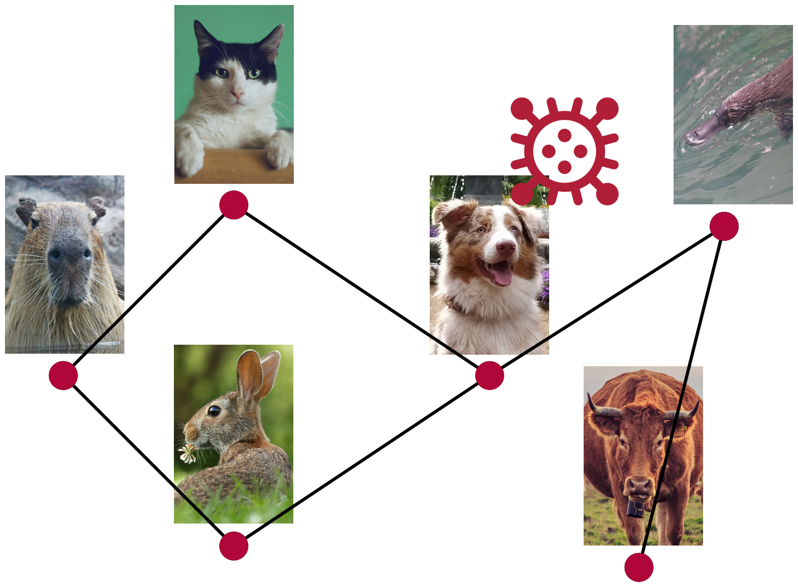

Temporal SIR Model

$\delta = 5$
$t = 8$

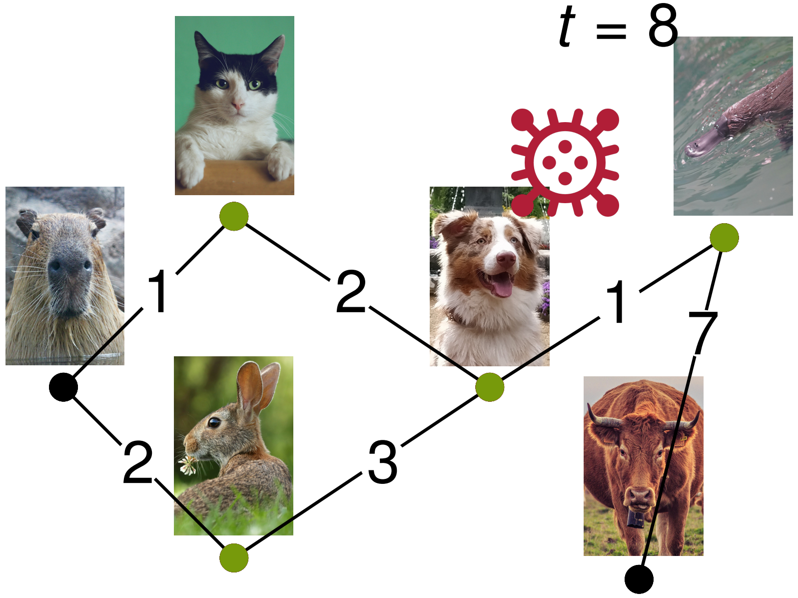

Formalizing Source Detection

Discoverer

Adversary

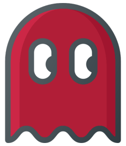

V, T_{max} ←

Pick node set V and $T_{max} \in \mathbb{N}$

— For rounds $i = 1, 2, \ldots$ —

Pick $v_i \in V$ to watch

$I_i(v_i)$ ←

Decide infection times and partner $I_i \subseteq V \times V \times [T_{max}]$

— Discoverer decides to end the game —

Suspect source $s \in V$

→ s

Pick edges E, labeling $\lambda \colon E \to [T_{max}]$, source s', seed time t_0

E, λ, s', t_0 ←

— End of game —

Adversary wins if:

- $s' \neq s$
- λ with seed infection at (s', t_0) is consistent with all $I_i(v_i)$

Source Detection—Results

Cost Measure
Number of infections (of all nodes) until the source is detected

	Trees		General	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Consistent				
Known	$\Omega(n \log n)$ wcp	$\mathcal{O}(n \log n)$ wcp	$\Omega(n\sqrt{n})$ wcp	$\mathcal{O}(n\sqrt{n})$ wcp
Unknown	$\Omega(n\sqrt{n})$ wcp	$\mathcal{O}(n\sqrt{n})$ wcp	$\Omega(n\sqrt{n})$ wcp	$\mathcal{O}(n\sqrt{n})$ wcp
Obliviously dynamic				
Known	$\Omega(n \log n)$ wcp	$\mathcal{O}(n \log n)^?$ wcp	$\Omega(n^2)$ wcp	n^2 det
Unknown			$\Omega(n^2)$ wcp	n^2 det

Open Question
What if the graph is a tree, but we don't know what it looks like?

wcp=with constant probability
?=when watching 2 nodes

Future Work

Novel field, plenty open questions

Probabilistic infections!

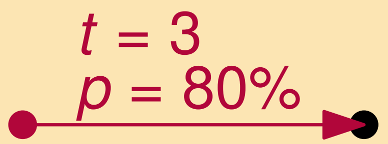

Fill our gaps

Open Question
What if the graph is a tree, but we don't know what it looks like?

References

- Argyrios Deligkas, Michelle Döring, Eduard Eiben, Tiger-Lily Goldsmith, and George Skretas. Being an influencer is hard: The complexity of influence max- imization in temporal graphs with a fixed source. Information and Computation (2024)
- David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. KDD ’03 (2003), 137–146.
- Srijan Kumar, Chongyang Bai, VS Subrahmanian, and Jure Leskovec. Deception Detection in Group Video Conversations using Dynamic Interaction Net- works. In: ICWSM 2021. International AAAI Conference on Web and Social Media. 2021.
- Praneeth Netrapalli and Sujay Sanghavi. Learning the graph of epidemic cas- cades. ACM SIGMETRICS Performance Evaluation Review 40:1 (2012), 211–222
- Devavrat Shah and Tauhid Zaman. Rumors in a network: Who’s the culprit? IEEE Transactions on information theory 57:8 (2011), 5163–5181.

Credits

Icons: Flaticon.com

Cat: Photo by Manja Vitolic on Unsplash

Capybara: Photo by Dušan veverkolog on Unsplash

Rabbit: Photo by Gary Bendig on Unsplash

Cow: Photo by KAUE FONSECA on Unsplash

Platypus: Photo by Aaron De Wit on Unsplash

$\Omega(n^2)$ wcp Lower Bound

Setting: Obliviously dynamic source behavior, unknown static graph

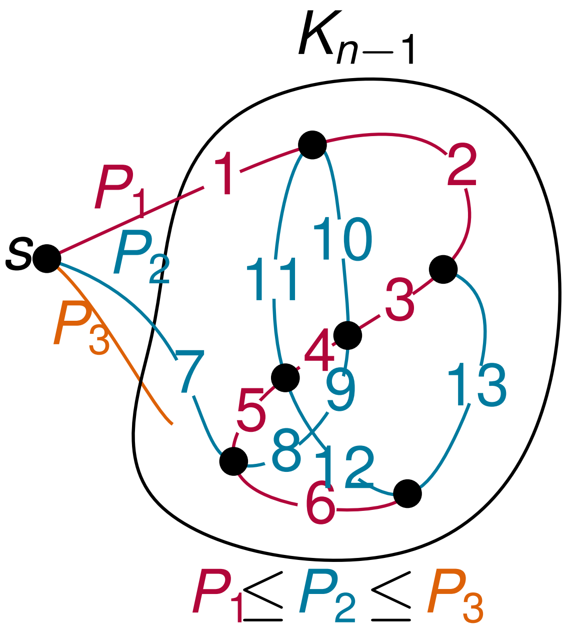

- Let n be odd
- Let $P_1, \ldots$ be $(n-1)/2$ Hamiltonian paths on the nodes $v_1, \ldots, v_{n-1}$
- If e is the i-th node on the j-th path, set $\lambda(e) := jn + i$
- In round x, infect s at xn-1
  - Thus the infection travels via P_x
  - Every node is infected in every round
- Discoverer must observe first or second node on the active part to find the source
- Proof intuition: observing a node only reveals its location on the current path, but not on any future paths
- We have probability at most 1/3 to discover the source in the first $n/2$ rounds (i.e., before $n^2/2$ infections)
- Generalize to arbitrary success probabilities

$O(\text{tw} \cdot n \log n)$ Source Detection

Setting: Consistent source, known static graph

1. Maintain a subtree of candidate nodes, and always compute a centroid as a balanced separator. Start with the whole graph as the candidates.
2. In each round: watch the current separator and one node in the current subtree picked uniformly at random.
3. If you receive information about the subtree to go into (either because the separator or the other node was infected), do so.
   - a. If the separator is infected, recurse into the side of the separation from which the infection originated.
   - b. If the randomly picked node is infected but not the separator, recurse into the side of the separation this node is a part of.

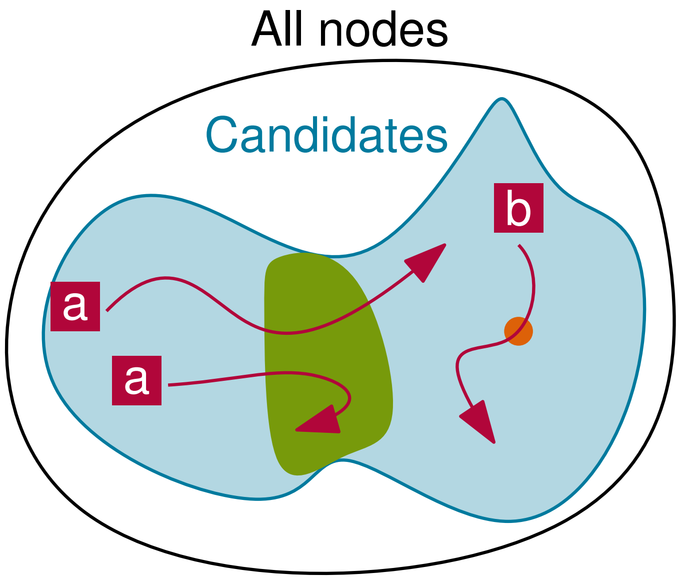

Balanced seperator of size tw

Proof idea

- Each round watches tw + 1 and must thus be expanded to tw + 1 individual rounds watching one node
- Each such phase will likely have a success in case a or b after a constant number of iterations
- After each phase, the number of candidates halves

.