Stealthy Epidemics: Modeling Worm Attacks against CPS

Aron Laszka[1]

in collaboration with Nika Haghtalab[2], Ariel Procaccia[2], Yevgeniy Vorobeychik[1], and Xenofon Koutsoukos[1]

[1]Vanderbilt University, [2]Carnegie Mellon University

Motivation

* Highly sensitive systems, such as CPS for critical infrastructure, are usually supposed to be secured by the "air gap"
* However, computer worms that propagate over local networks and removable drives may infect even these systems
  * e.g., Stuxnet infected Iranian nuclear facilities

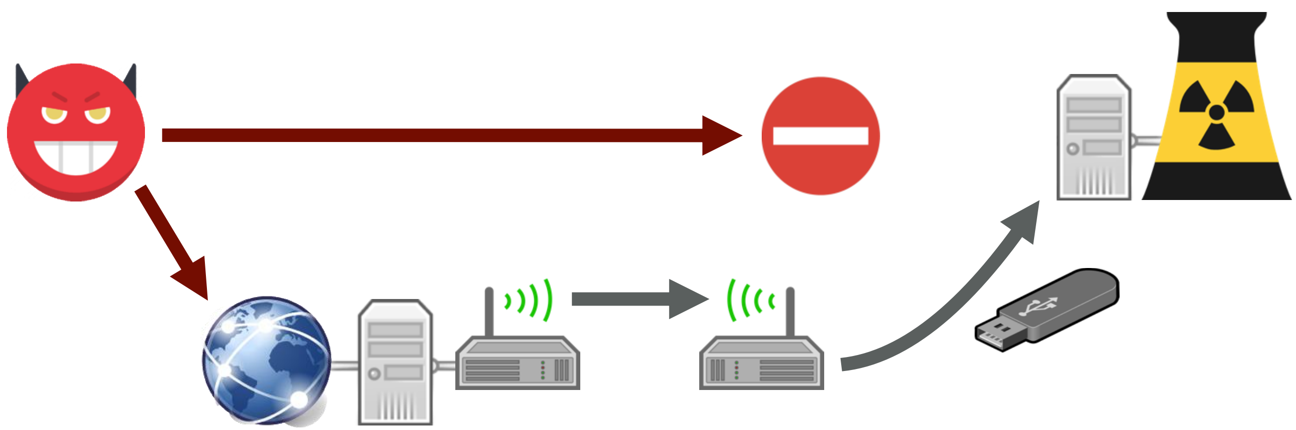

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Examples of Worm-Based Attacks #1

* Stuxnet worm
  * targeted Iranian uranium enrichment facilities
  * initially sent to companies working on industrial control systems in Iran
  * propagated over local area networks and removable drives
  * drastically reduced the lifetime and reportedly ruined almost one-fifth of Iran's nuclear centrifuges

http://www.businessinsider.com/stuxnet-was-far-more-dangerous-than-previous-thought-2013-11

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Examples of Worm-Based Attacks #2

* Shamoon worm
  * targeted energy companies in the Middle East, including Saudi Aramco and Qatar's RasGas
  * initially deployed on an Internet connected computer at Saudi Aramco
  * removed and overwrote information on hard drives
  * incapacitated 30,000 to 55,000 workstations at Saudi Aramco

http://www.bbc.com/news/technology-19293797

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Resilience to Worm-Based Attacks

- To stop a worm, we can
  - create antivirus signatures
  - patch vulnerabilities
  - …
- However, before we can implement these countermeasures, we first have to **detect the worm**
- Furthermore, it is imperative that we detect the worm **in time**
  - worm detection and alerting operators take some time
  - implementing countermeasures takes some time
- ➢ Attack-resilience depends on the **timely detection of worms**

Previous Work on Modeling Worms

- Mostly based on epidemic and influence maximization models
  - primarily concerned with **steady** or **equilibrium states**
- Generally, they do not consider the detection problem
  - in practice, a worm can be eradicated once it has been discovered
  - steady or equilibrium state might not be reached by the **time of detection**
- More importantly, they do not consider targeted attacks
  - usual assumption is that the worm is trying to infect as many computers as possible
  - targeted worms may try to be **stealthy** to avoid early detection

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Non-Targeted Worm Example: Code Red (2001)

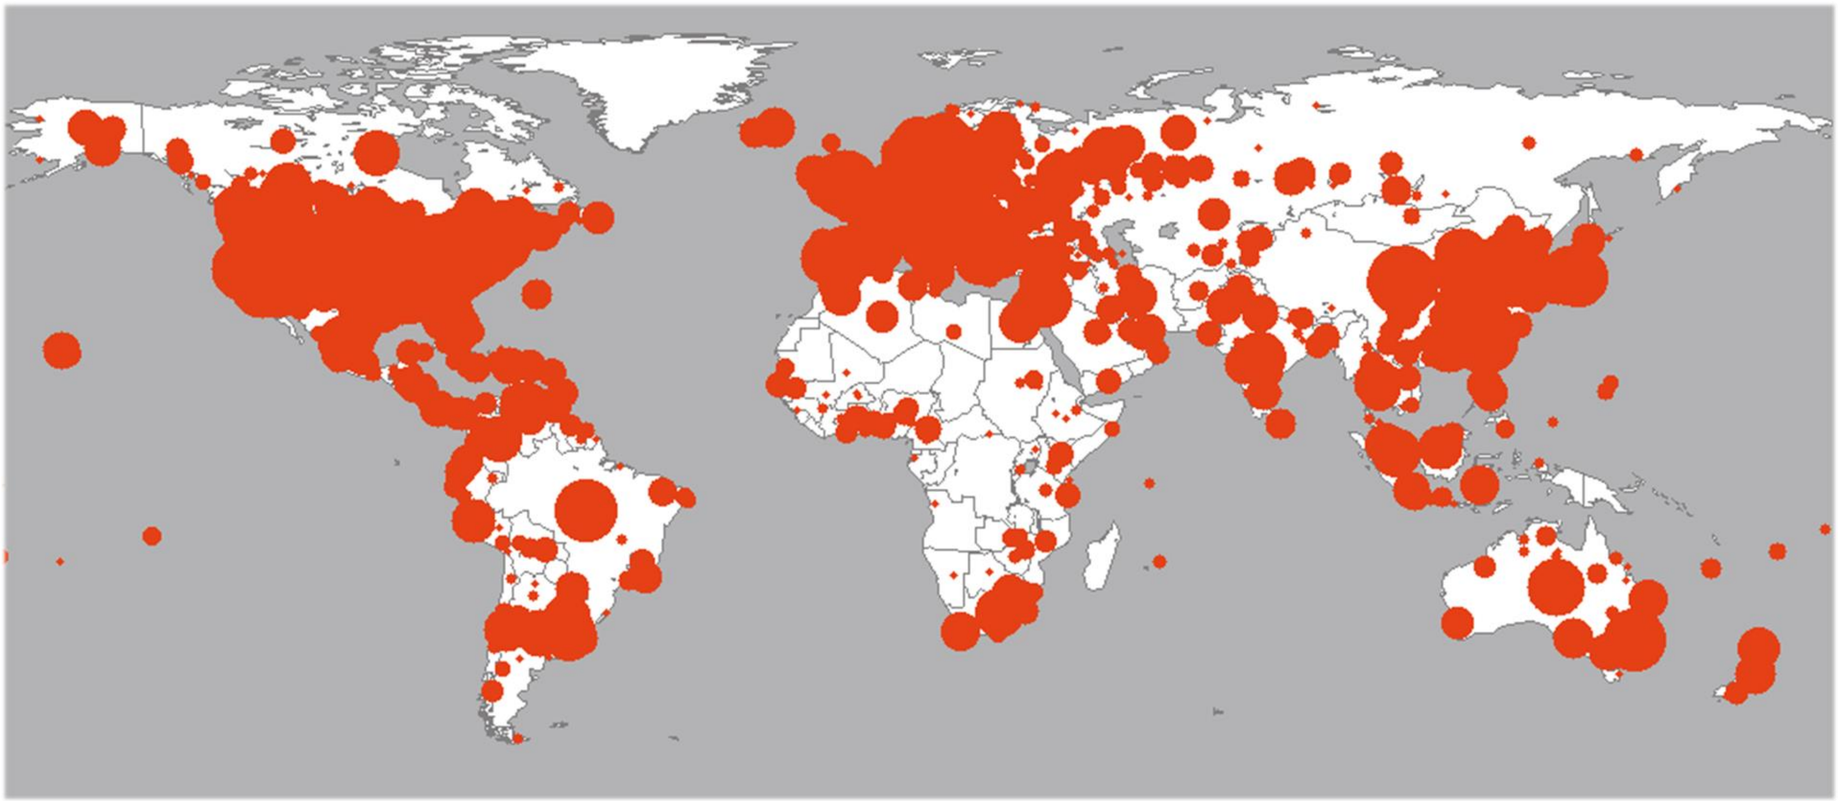

Targeted Worm Example: Flame (2012)

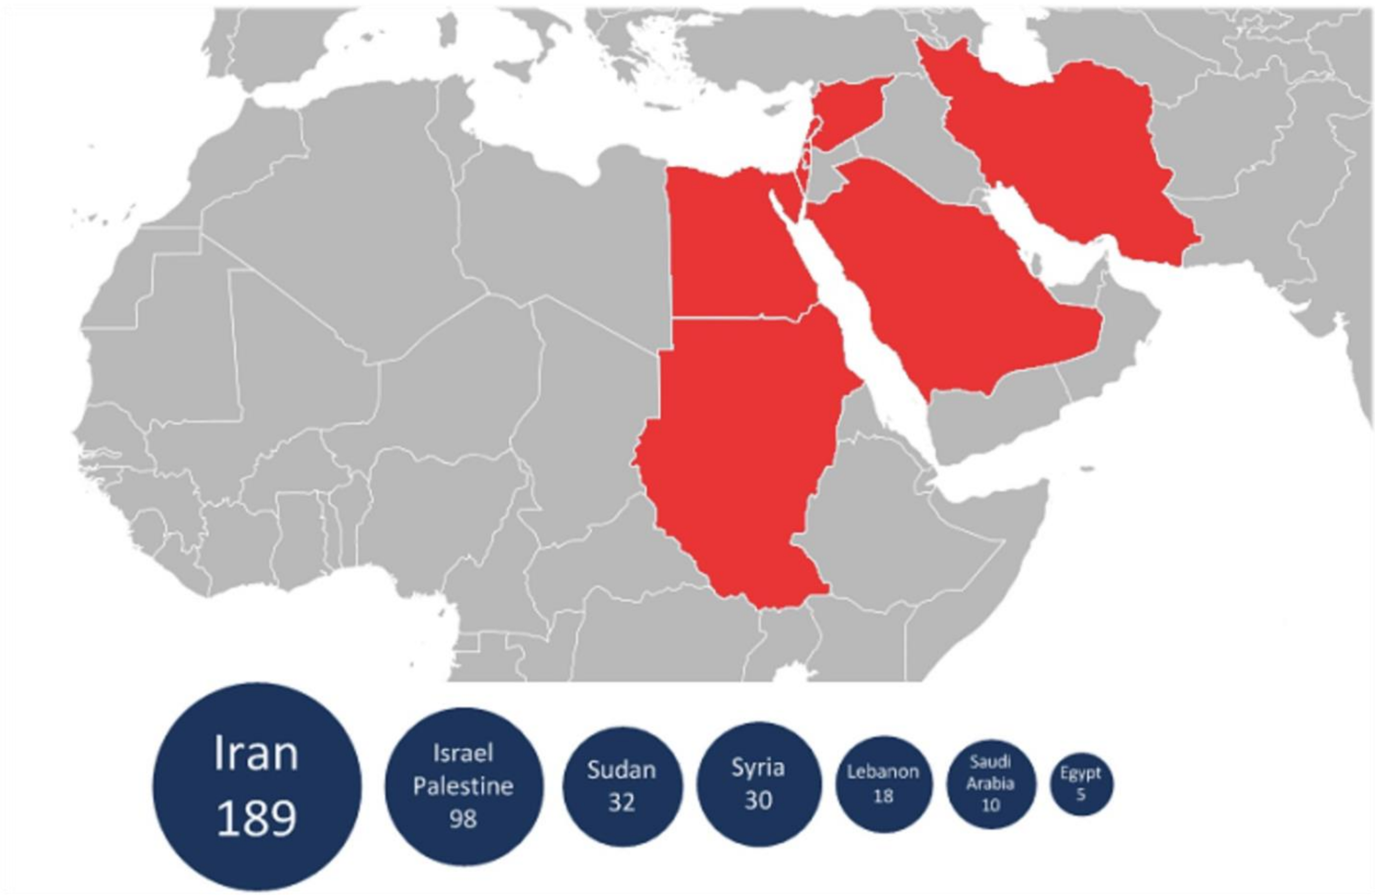

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Outline

* Model
  - Network model
  - Propagation model
  - Detection model
* Results
  * computing the probability of detection
  * optimal assignment of resources to detection
  - Non-strategic attacks
  - Strategic attacks

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Network Model

* Directed graph $G = (V, E)$
  * node = computer system (or tightly coupled group of computers that can be infected together)
  * edge = possible infections
    * e.g., local area connections, regularly shared removable drives
    * weight = probability of propagation

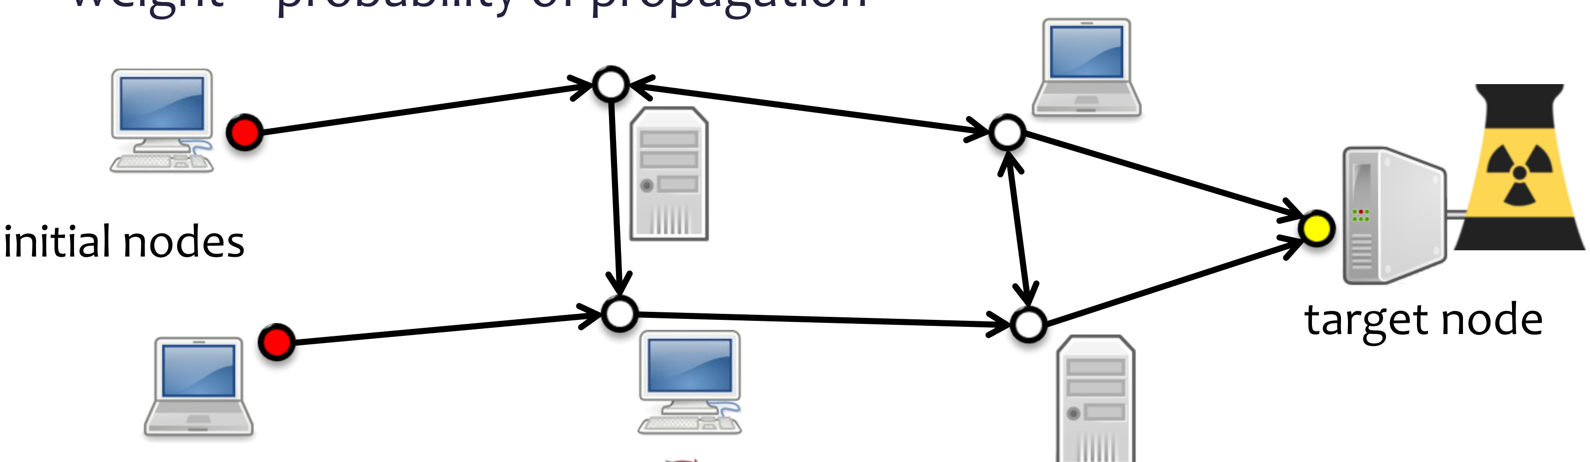

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Propagation Models

- Time
  - at the beginning, only the initial nodes are infected
  - in each time step, additional nodes may be infected
- **Independent cascades** model
  - nodes that were infected in the previous round may infect their neighbors
- **Repeated independent cascades** model
  - nodes that are infected may infect their neighbors

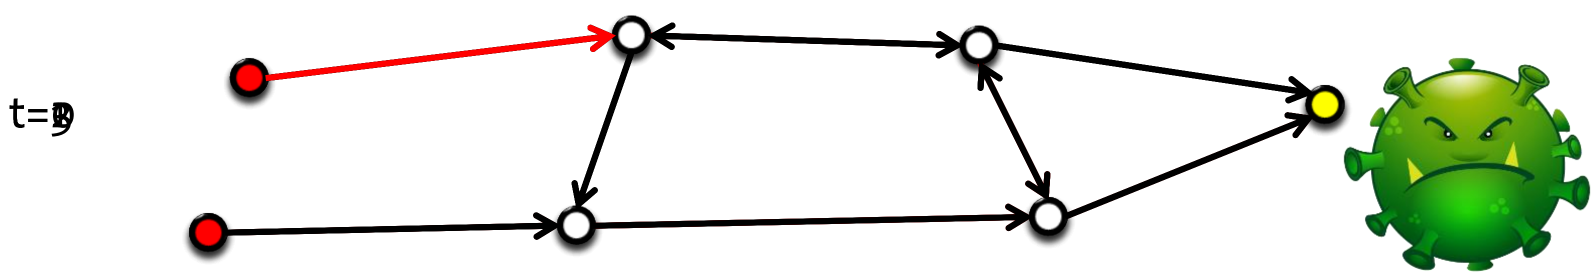

Monitored Nodes

- Monitored nodes
  - in order to detect worms, a defender monitors some nodes
    - e.g., performing thorough audits
  - since monitoring is costly, at most k nodes can be monitored
  - furthermore, the set of nodes that can be monitored is restricted
    - e.g., nodes that are not operated by the defender cannot be monitored
- Delayed detection
  - mitigation is successful if the worm reaches a monitored node m at least D_m time steps **before** it reaches the target
    (or if it never reaches the target)

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Problem Formulation

- Goal:

select a set of k monitored nodes M that maximizes the probability of detection $U(M)$

- Formulations
  - **non-strategic attacks**: fixed set of initial nodes
    - e.g., nodes that are connected to the Internet
  - **strategic attacks**: set of initial nodes is chosen by an attacker, who wants to minimize the probability of detection
    - set of possible initial nodes S is restricted (e.g., nodes that are connected to the Internet)

Selection Example

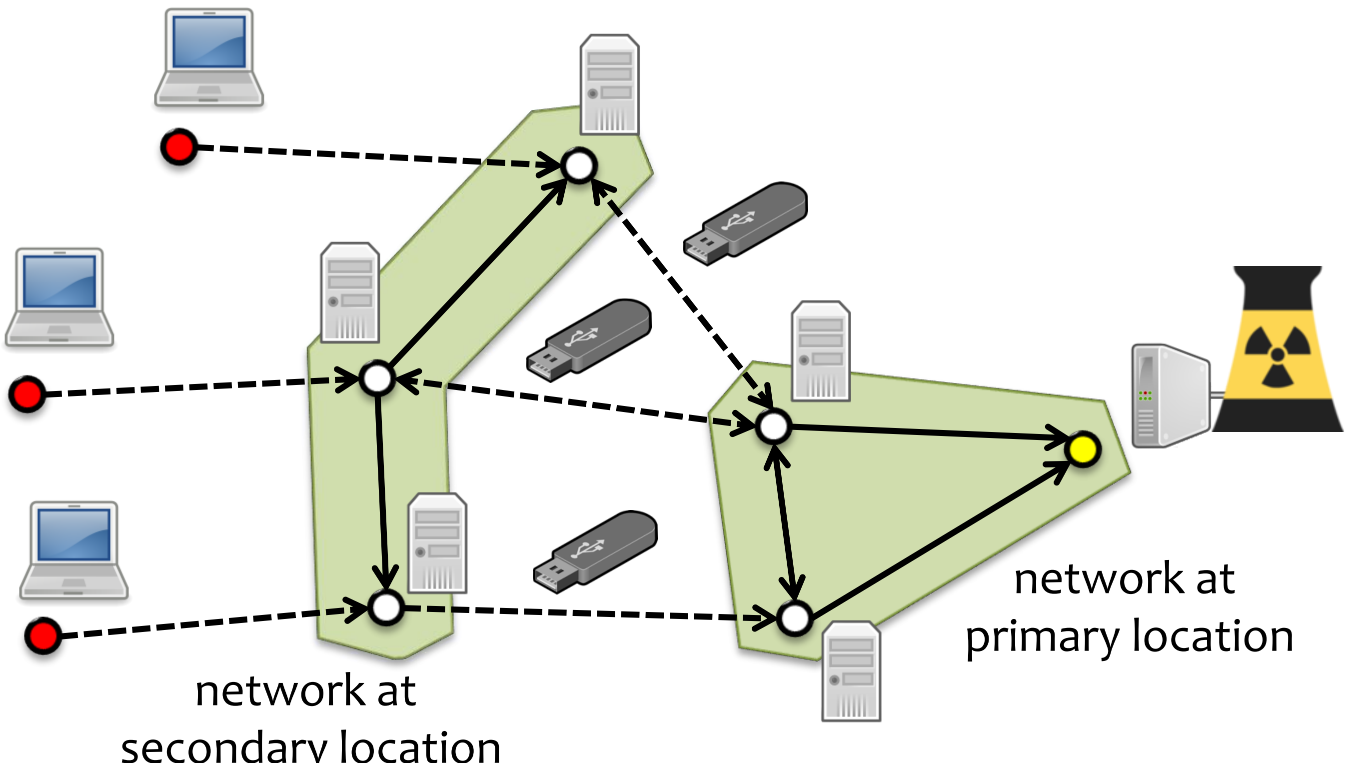

FORCES
FOUNDATIONS OF RESILIENT
CYBER-PHYSICAL SYSTEMS

Selection Example

* Monitoring budget: $k = 2$

set of possible monitored nodes

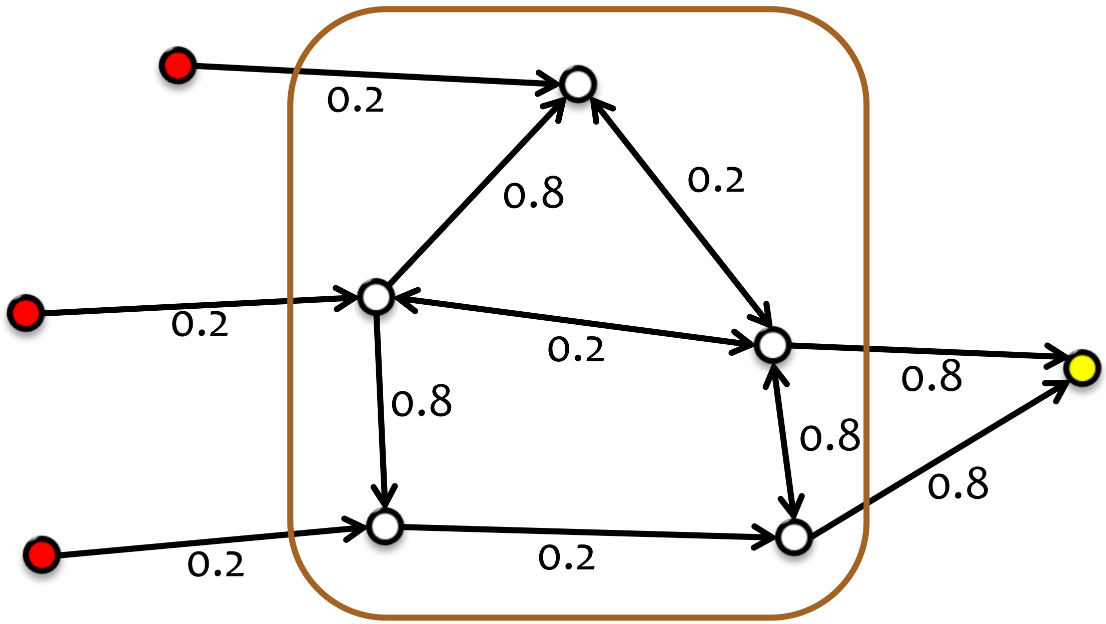

Selection Example

- Monitoring budget: $k = 2$
- Detection delay: $D = 2$

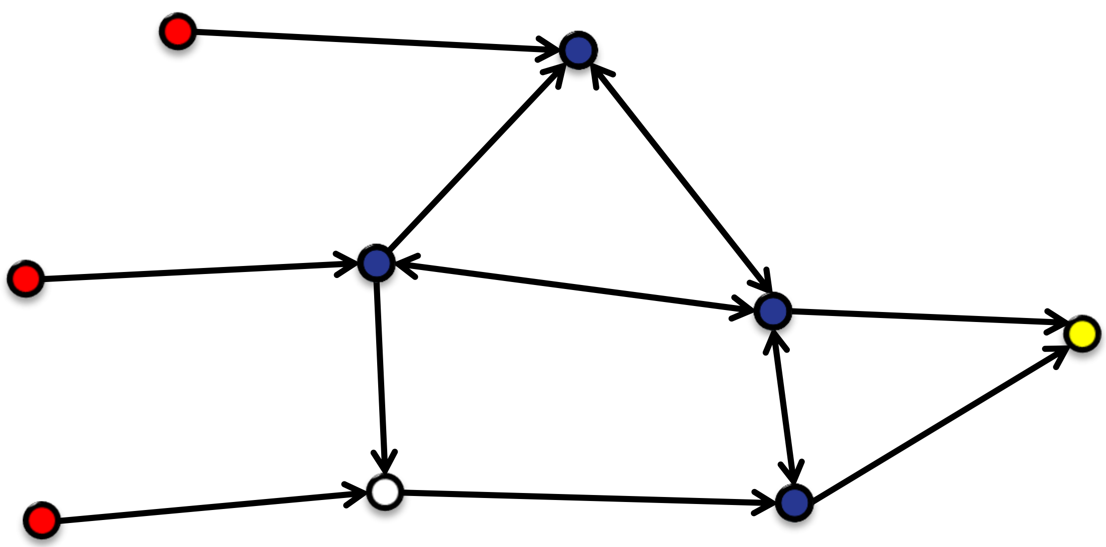

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Computing the Probability of Timely Detection

Computing the probability of detection $U(M)$ for a given set of monitored nodes M is a #P-hard problem.

  - #P is the set of counting problems associated with the decision problems in the set NP
- However, we can use simulations
  - error can be bounded using Hoeffding's inequality

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Optimal Monitoring against Non-Strategic Attacks

* Non-strategic = fixed set of initial nodes for the worm
* Computational complexity:

Finding a $(1 - 1/e + o(1))$-approximately optimal monitored set is NP-hard.

Optimal Monitoring against Non-Strategic Attacks

* Non-strategic = fixed set of initial nodes for the worm
* Computational complexity
* Approximation:

The probability $U(M)$ is a non-decreasing submodular set function of M.

For any ε, $\delta > 0$, a greedy algorithm running in time poly($|V|$, $1/\varepsilon$, $\ln(1/\delta)$) returns a set M such that with probability $1 - \delta$,

$$U(M) \geq (1 - 1/e)\, U(OPT) - \varepsilon.$$

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Numerical Results for Non-Strategic Attacks

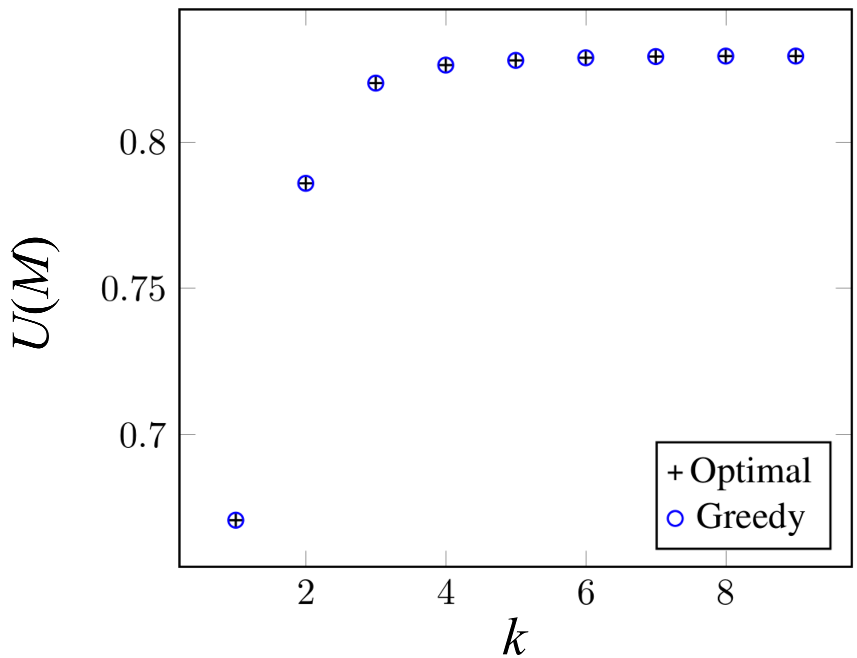

B-A graphs with 3 node clique and 3 edges per new node.

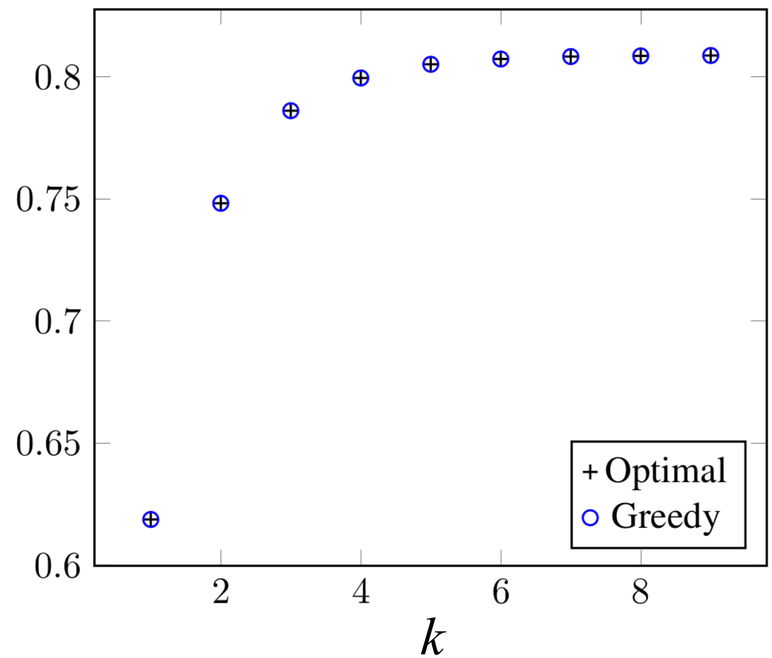

E-R graphs with 0.5 edge presence probability.

Randomly generated graphs with 100 nodes, 5 randomly chosen initial nodes, 10 randomly chosen possible monitored nodes, 1 randomly chosen target node, all edges having propagation probability 0.5, independent cascades propagation model, and 1 time step detection delay. Values are averages taken over 10 graphs.

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Optimal Monitoring against Strategic Attacks

- Strategic attacks = worst-case set of initial nodes for the worm
- Computational complexity:

For any ε, finding a set M of size at most $(1 - \varepsilon) \ln(|S|)$ such that

$$U(M) \,/\, U(OPT) > 0$$

is NP-hard.

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Optimal Monitoring against Strategic Attacks

- Strategic attacks = worst-case set of initial nodes for the worm
- Computational complexity
- Approximation:

For any $\varepsilon, \gamma, \delta > 0$, we can find a set M in time $\text{poly}(|V|, 1/\varepsilon, 1/\gamma, \ln(1/\delta))$ such that $|M| \leq |S|\, k \ln(1/\varepsilon)$ and with probability $1 - \delta$,

$$U(M) \geq (1 - 1/e)\, U(OPT) - \gamma.$$

- algorithm: iterate over the set of possible initial nodes, and for each node s, select $k \ln(1/\varepsilon)$ monitored nodes in a greedy manner supposing that the attacker will select $\{s\}$ as the set of initial nodes

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Numerical Results for Strategic Attacks

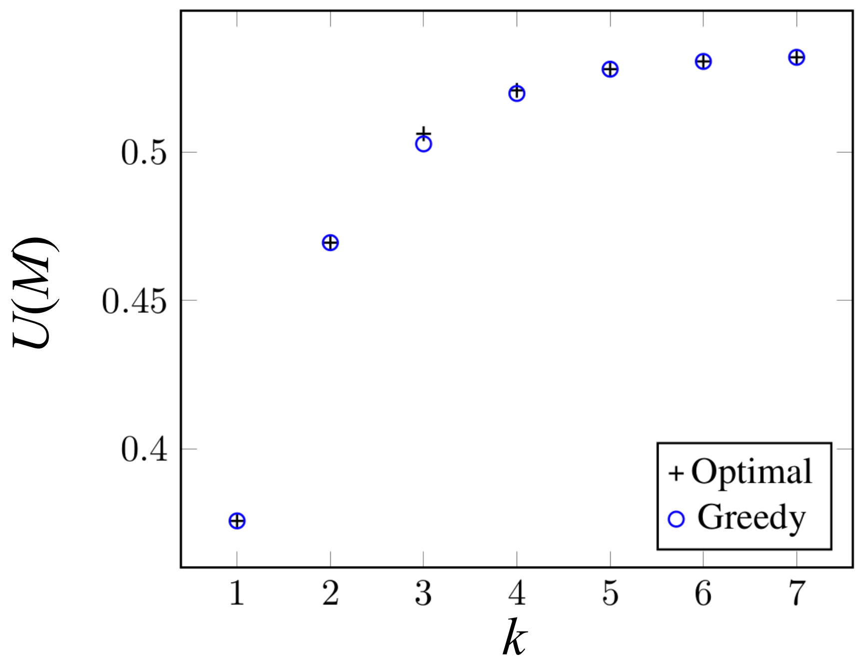

B-A graphs with 3 node clique and 3 edges per new node.

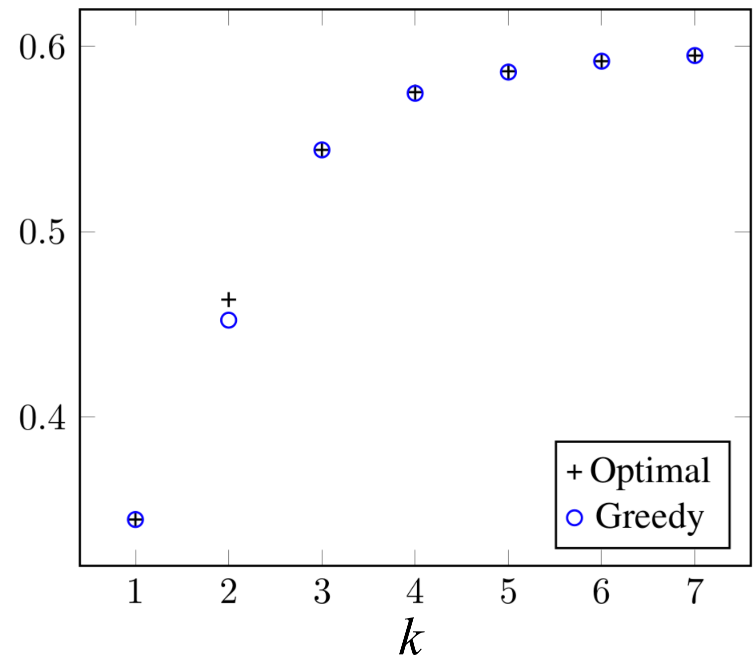

E-R graphs with 0.5 edge presence probability.

Randomly generated graphs with 100 nodes, 5 randomly chosen possible initial nodes, 10 randomly chosen possible monitored nodes, 1 randomly chosen target node, all edges having propagation probability 0.5, independent cascades propagation model, and 1 time step detection delay. Values are averages taken over 10 graphs.

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Conclusion

* Computer worms pose a serious threat to critical CPS
* In order to be resilient to such attacks, we have to be able to detect worms in time
* Selection of monitored nodes must be carefully planned
* Computational results
  * challenging, but can be solved
* Open problem: finding an optimal attack
  * NP-hard
  * but can we approximate it efficiently?

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Thank you for your attention!

Questions?

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS