FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

University of Michigan - Ann Arbor

Defense Policies for Partially Observed Spreading Processes on Bayesian Attack Graphs

Erik Miehling
Mohammad Rasouli
Demosthenis Teneketzis

FORCES Meeting, Berkeley, CA — May 28-29, 2015

Berkeley UNIVERSITY OF CALIFORNIA

Massachusetts Institute of Technology

NSF

MICHIGAN

VANDERBILT UNIVERSITY

Motivation

- Three key factors in ***information security*** problems:

Confidentiality (C) — Ensuring data does not get into the wrong hands, that is, maintaining privacy

Integrity (I) — Maintaining accuracy and trustworthiness of information

Availability (A) — Ensuring that data is always available to trusted users

- We are interested in the problem of protecting specific, important resources
  - Closely related to confidentiality and integrity
  - Need to ensure key resources are still available while protecting assets

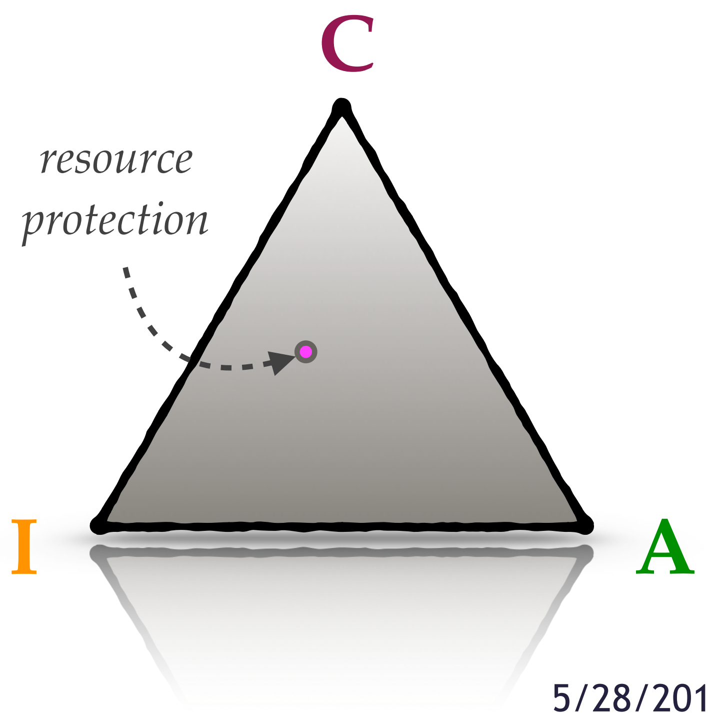

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

The Conflict Environment

- We consider a dynamic setting where a network is continually being subjected to attacks with the objective of compromising some ***target resources*** through ***exploits***
  - Resources that contain sensitive data
  - Resources that, when compromised, give an attacker control of a critical part of the system, potentially with catastrophic consequences
- Aspects of our model
  - ***Progressive attacks*** — recent exploits build upon previous exploits, progressively degrading the system
  - ***Dynamic defense*** — defender is choosing the best action based on ***new*** information
  - ***Partial knowledge*** — the defender only possesses a *guess* of the current exploits
- The defender can ***control services*** in the network to prevent the attacker from reaching the target resources

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Attack Graphs

- Insufficient to look at single vulnerabilities when protecting a network
  - Attackers combine vulnerabilities to penetrate the network

- *Attack graphs* model how multiple vulnerabilities can be combined and exploited by an attacker
  - Explicitly takes into account *paths* that the attacker can take to reach the critical exploitation

Graph Theoretic Representation

- Consider a directed graph, denoted by $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$

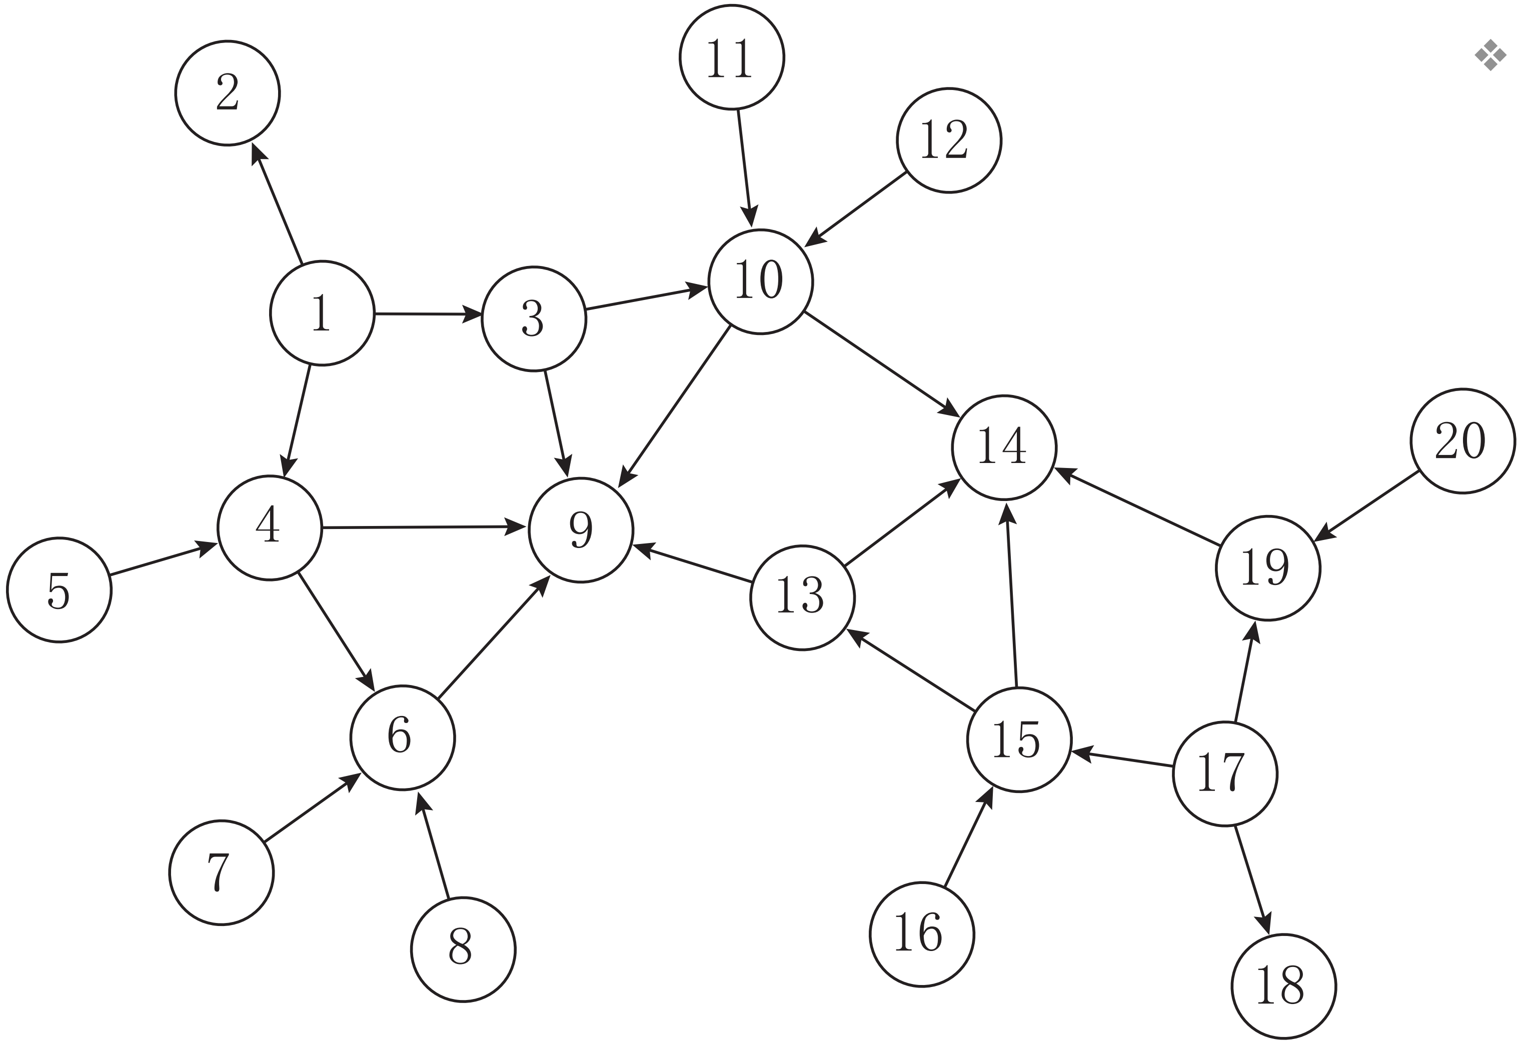

$$\mathcal{N}_R = \{1, 5, 7, 8, 11, 12, 16, 17, 20\}$$
$$\mathcal{N}_C = \{9, 14\} \subseteq \mathcal{N}_L = \{2, 9, 14, 18\}$$

- Nodes, $\mathcal{N}$, represent attributes

  $\mathcal{N}_R \subseteq \mathcal{N}$: root nodes

  - No prior exploit occurred
  - Outer layer of network (exposed to world)

  $\mathcal{N}_C \subseteq \mathcal{N}$: critical nodes

  - Deepest exploit level
  - Attacker is attempting to achieve one of the attributes

- Directed edges, $\mathcal{E}$, denote ***exploits*** (transitions between attributes)

Spreading Process

- The attacker’s behavior is assumed to follow a ***probabilistic spreading process*** (i.e. Bayesian attack graph)
- Each attribute (node) *i* can be in one of two states

  Disabled: $X_t^i = 0$ **Enabled:** $X_t^i = 1$

- *Infection seed and spread*: At each time *t*

  A. Each root attribute is enabled with probability α_i

  B. Infection spreads according to ``predecessor rules”

At time $t = \tau$:

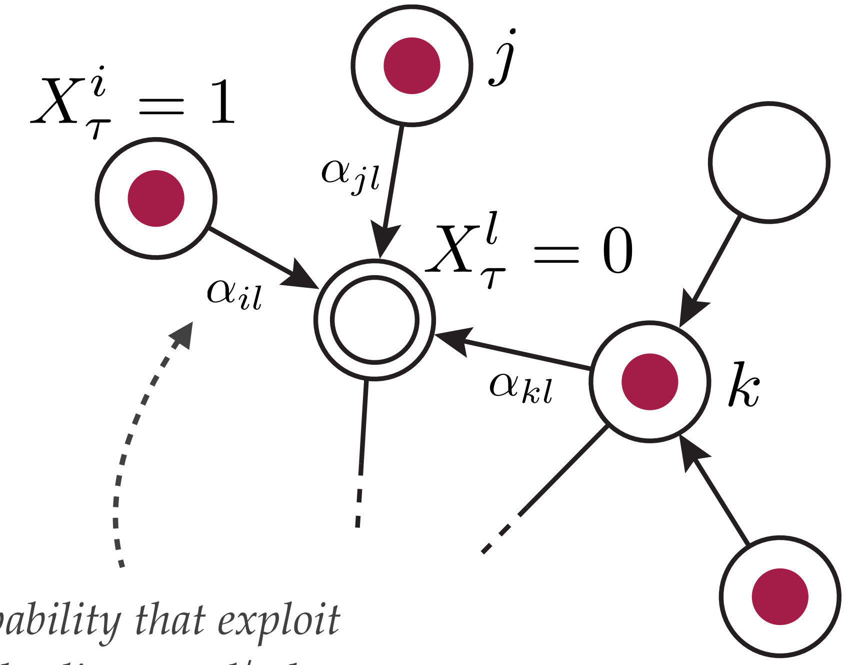

probability that exploit will be discovered/taken by the attacker (public knowledge)

Spreading Process – Predecessor Rules

- Each attribute (node) is one of two types
  - **AND** attribute
  - **OR** attribute
- The type of the attribute dictates the nature of the spreading process
- For **AND** attributes, e.g. node *l*

set of direct predecessors

$$P(X^l_{t+1} = 1 | X^l_t = 0, X_t) = \begin{cases} \prod_{p \in \bar{\mathcal{D}}_l} \alpha_{pl} & \text{if } \bigwedge_{p \in \bar{\mathcal{D}}_l} X^p_t = 1 \\ 0 & \text{otherwise} \end{cases}$$

- For **OR** attributes, e.g. node *k*

$$P(X^k_{t+1} = 1 | X^k_t = 0, X_t) = \begin{cases} 1 - \prod_{p \in \bar{\mathcal{D}}_k} (1 - \alpha_{pk}) & \text{if } \bigvee_{p \in \bar{\mathcal{D}}_k} X^p_t = 1 \\ 0 & \text{otherwise} \end{cases}$$

At time $t = \tau$:

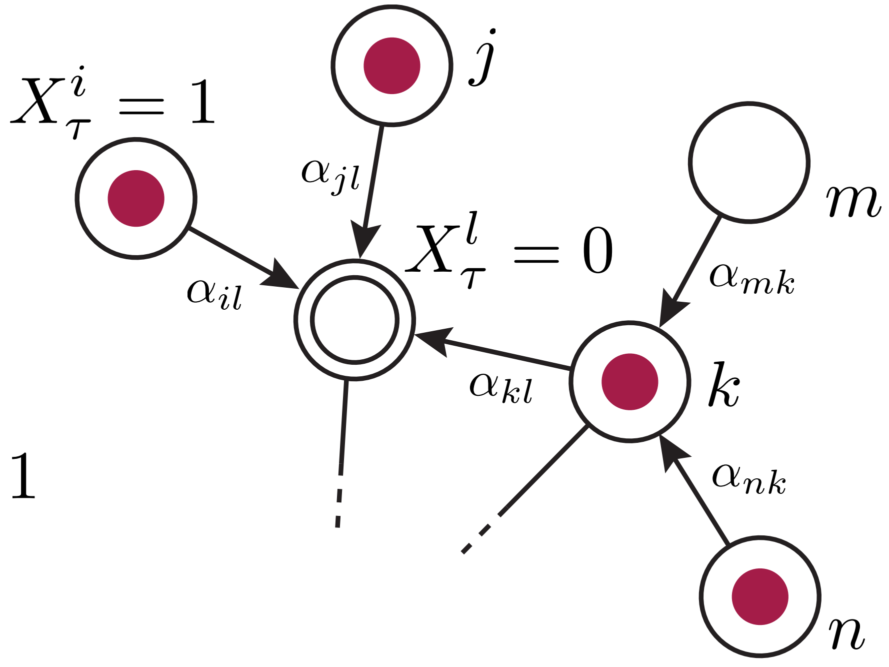

Defender’s Observations

- Defender only partially observes this process
  - The ***probability of detection*** at node i is β_i
- **Rationale**: defender may not known the full capability of the attacker at any given time
- Defender thus observes a subset of enabled attributes that have been discovered at each time-step

$$Y_t \in \{0,1\}^N$$

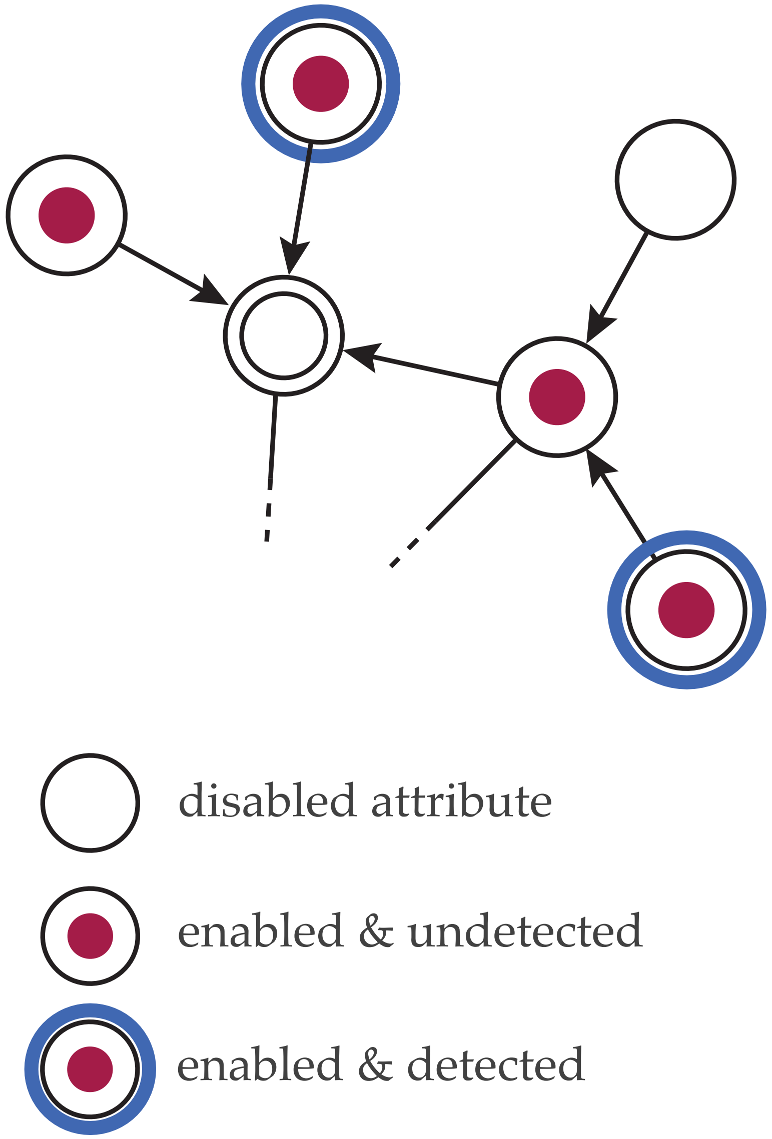

○ disabled attribute

◉ enabled & undetected

⦿ enabled & detected

Defense Actions

- We employ a ***moving target defense*** scheme, termed ***network hardening*** to protect against exploits
- Existence of exploits depend on protocols (services)
  - **Secure Shell (SSH)**
  - **File Transfer Protocol (FTP)**
  - **Port scanning**
  - **etc.**
- Defender can thus temporarily block or disable these services to stop the attacker from progressing

Defense Actions

- Suppose there are a set of M services $\{u^1, \ldots, u^M\}$
- Taking action u^m corresponds to disabling service m
  - u^m disables a subset of the attributes $\mathcal{W}_{u^m}$

$$X^i = 0,\ i \in \mathcal{W}_{u^m}$$

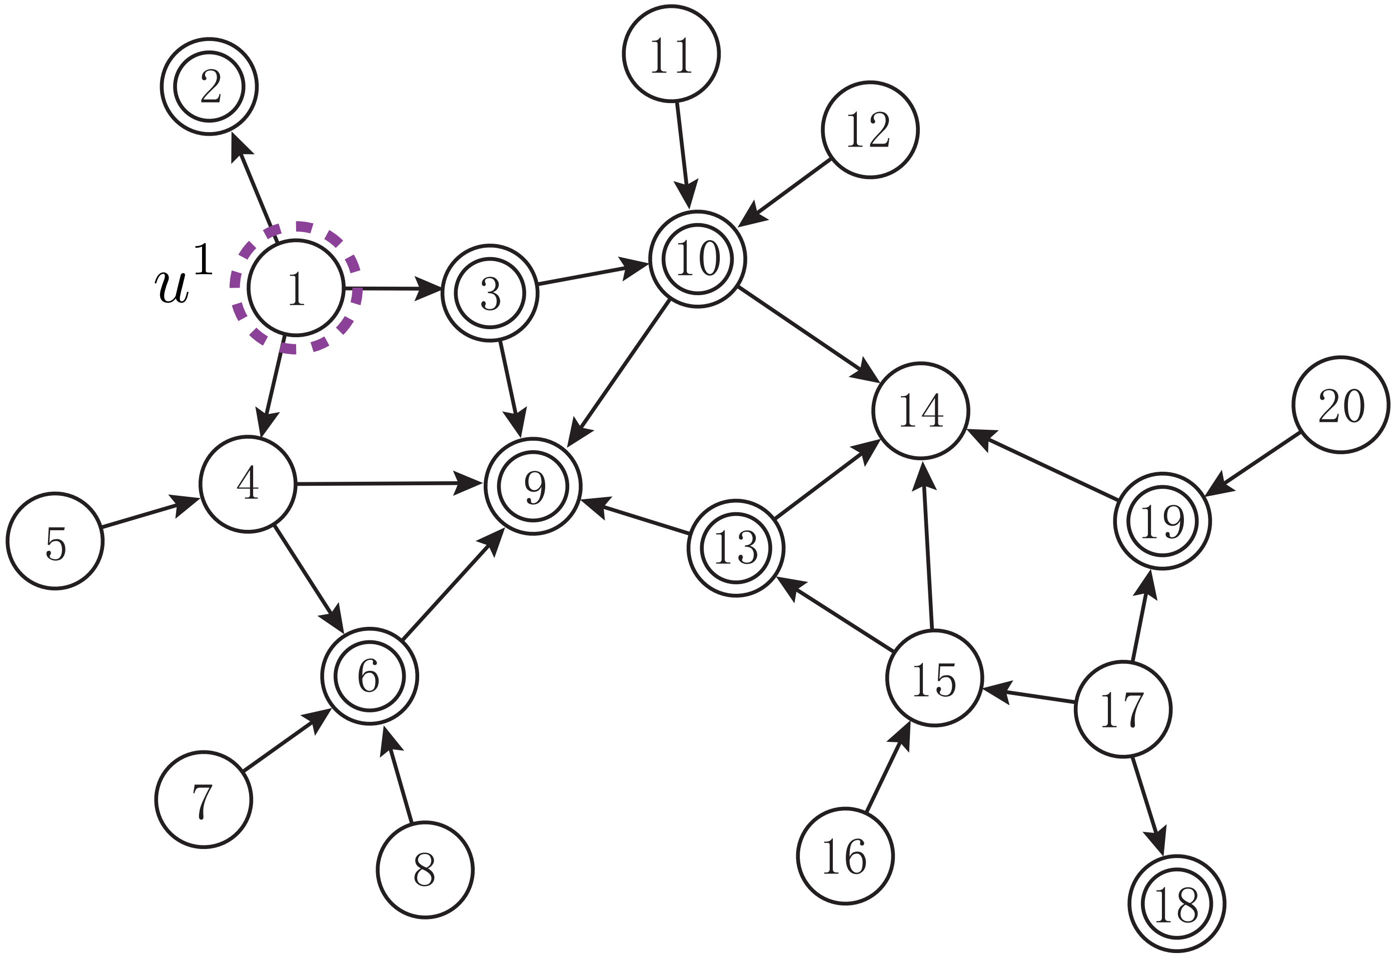

$$\mathcal{W}_{u^1} = \{1\}$$

- Action at time t

$$u_t \in \mathcal{U} = \wp(\{u^1, \ldots, u^M\})$$

- Assume that all root attributes are covered by at least one service

Defense Actions

- Suppose there are a set of M services $\{u^1, \ldots, u^M\}$
- Taking action u^m corresponds to disabling service m
  - u^m disables a subset of the attributes $\mathcal{W}_{u^m}$

$$X^i = 0,\ i \in \mathcal{W}_{u^m}$$

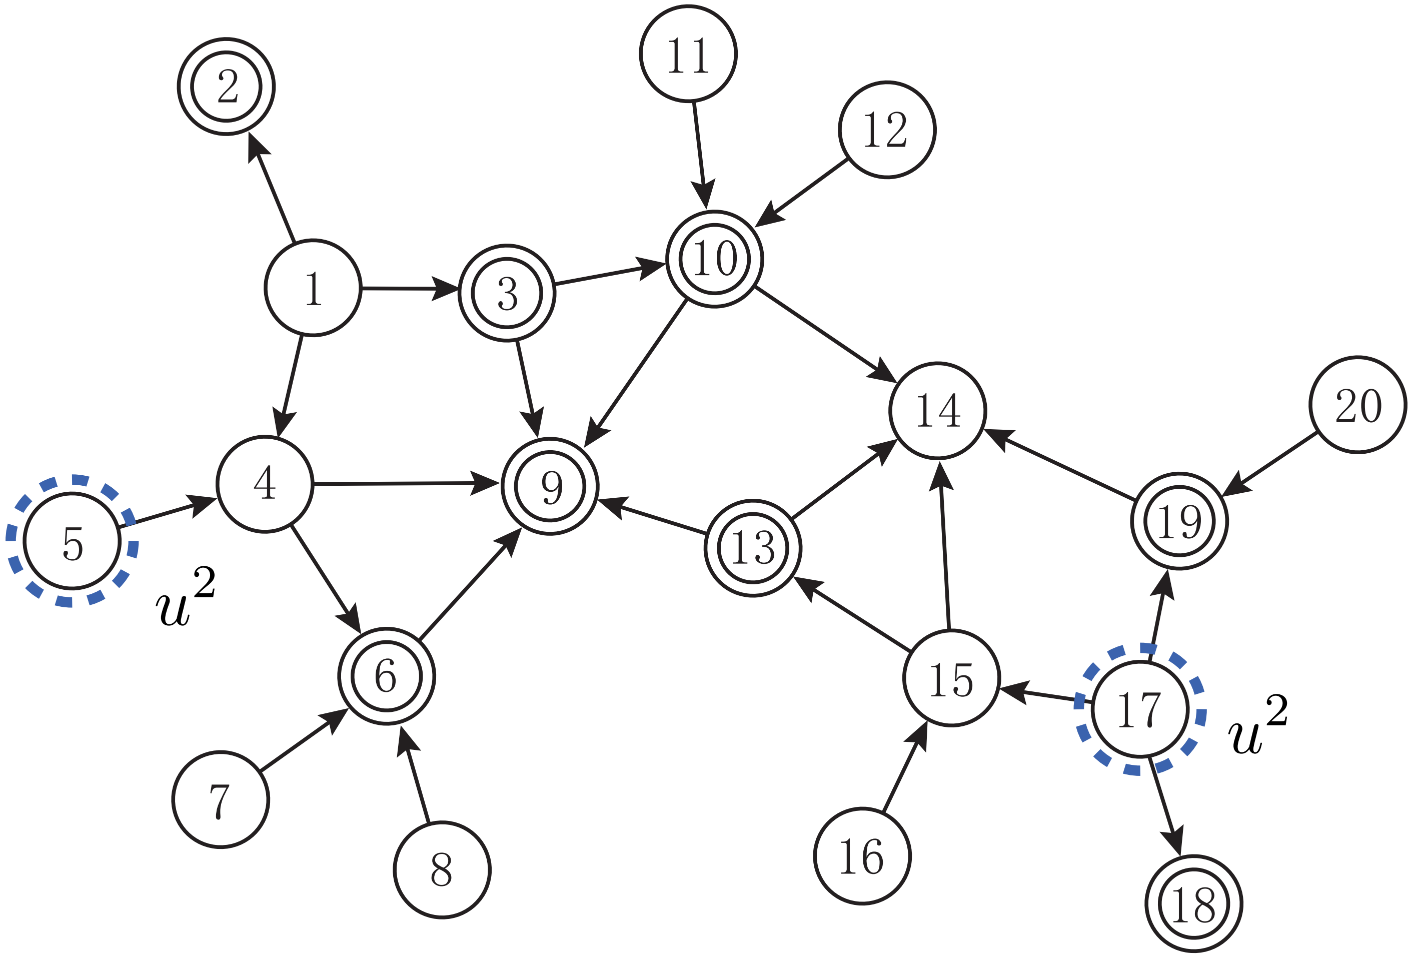

$$\mathcal{W}_{u^2} = \{5, 17\}$$

- Action at time t

$$u_t \in \mathcal{U} = \wp(\{u^1, \ldots, u^M\})$$

- Assume that all root attributes are covered by at least one service

Defense Actions

- Suppose there are a set of M services $\{u^1, \dots, u^M\}$
- Taking action u^m corresponds to disabling service m
  - u^m disables a subset of the attributes $\mathcal{W}_{u^m}$

$$X^i = 0, \; i \in \mathcal{W}_{u^m}$$

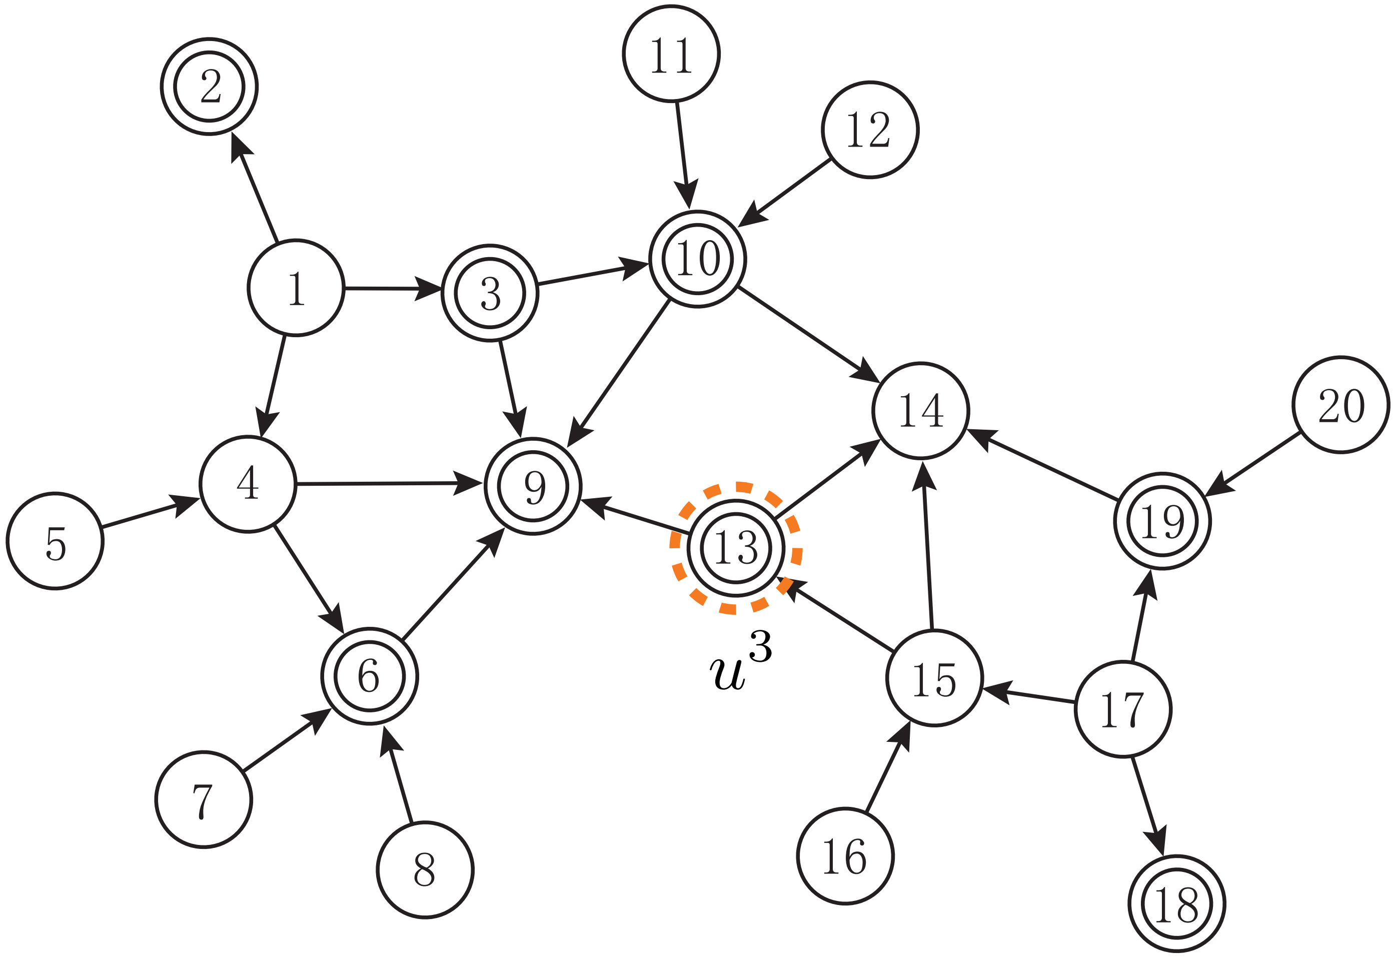

$$\mathcal{W}_{u^3} = \{13\}$$

- Action at time t

$$u_t \in \mathcal{U} = \wp(\{u^1, \dots, u^M\})$$

- Assume that all root attributes are covered by at least one service

Defense Actions

- Suppose there are a set of M services $\{u^1, \ldots, u^M\}$
- Taking action u^m corresponds to disabling service m
  - u^m disables a subset of the attributes $\mathcal{W}_{u^m}$

$$X^i = 0,\ i \in \mathcal{W}_{u^m}$$

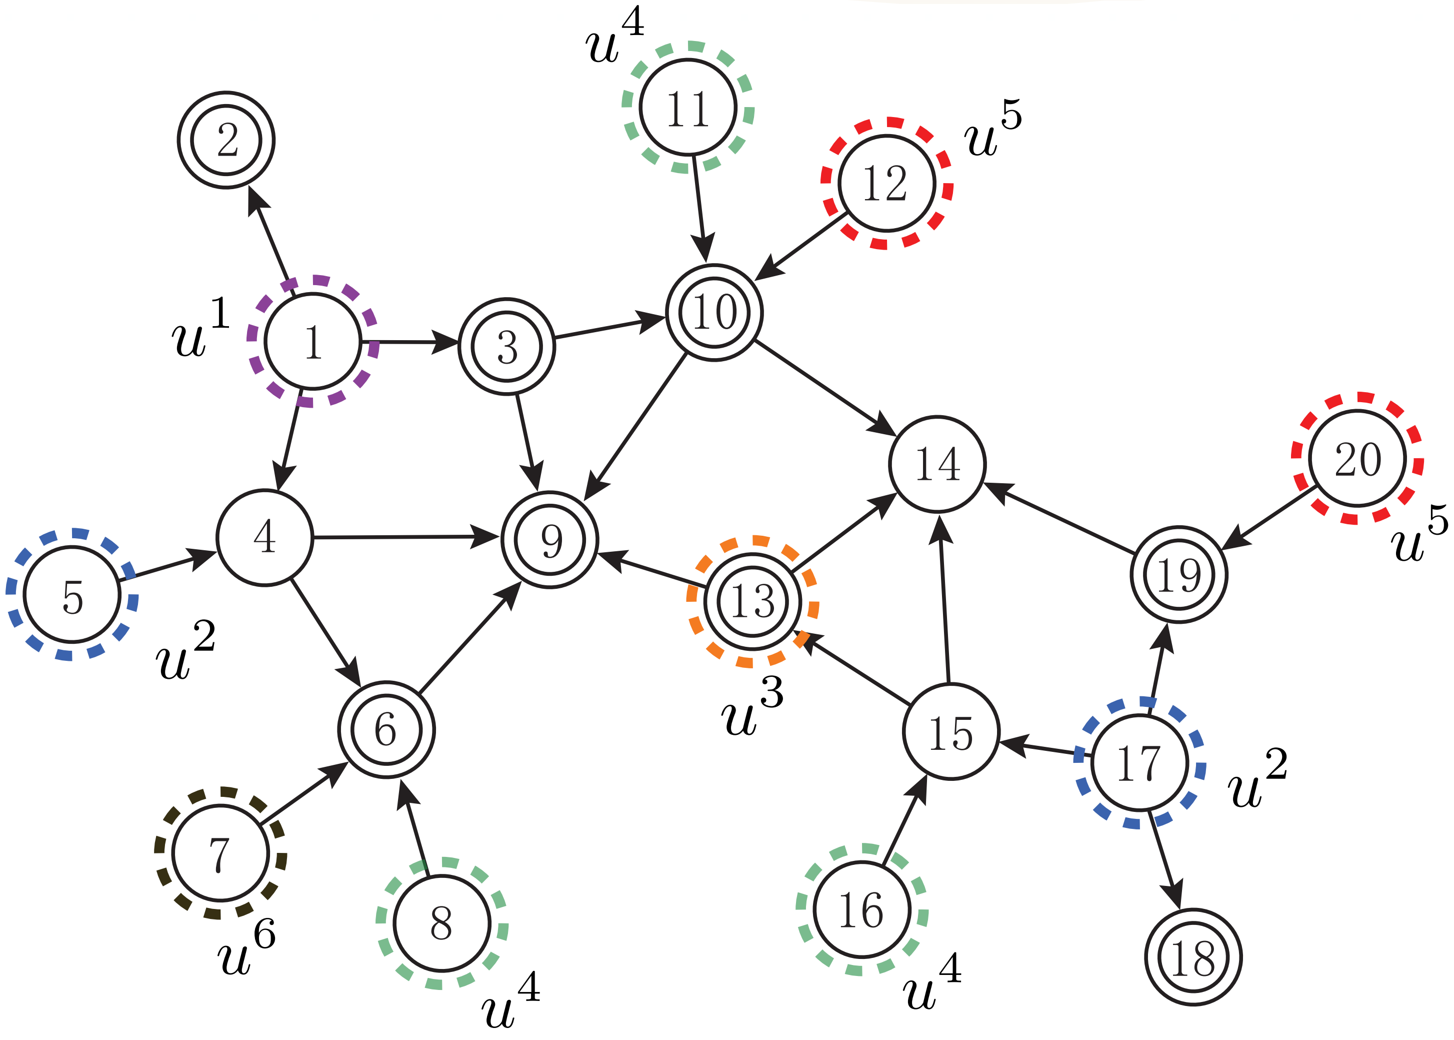

- Action at time t

$$u_t \in \mathcal{U} = \wp(\{u^1, \ldots, u^M\})$$

- Assume that all root attributes are covered by at least one service

Cost Function

- Cost of taking action $u \in \mathcal{U}$ in state $X \in \{0,1\}^N$

$$C(X,u) = \underbrace{C(X)}_{\textit{state cost}} + \underbrace{D(u)}_{\textit{availability cost}}$$

- **State cost,** $C(X)$**:** cost of being in a particular state
- **Availability cost,** $D(u)$**:** cost dependent upon how many resources the defense action renders unusable (due to the disabling of the service)
- The costs capture the confidentiality, integrity, and availability factors

Defender's Information States

- Define the history up to time *t* as $H_t = (\pi_0, u_1, y_1, u_2, y_2, \ldots, u_{t-1}, y_t)$
- We capture H_t by an ***information state*** $\pi_t = (\pi_t^1, \ldots, \pi_t^K) \in \Delta(\mathcal{X})$

$$\pi_t^i = P(X_t = x_i | H_t)$$

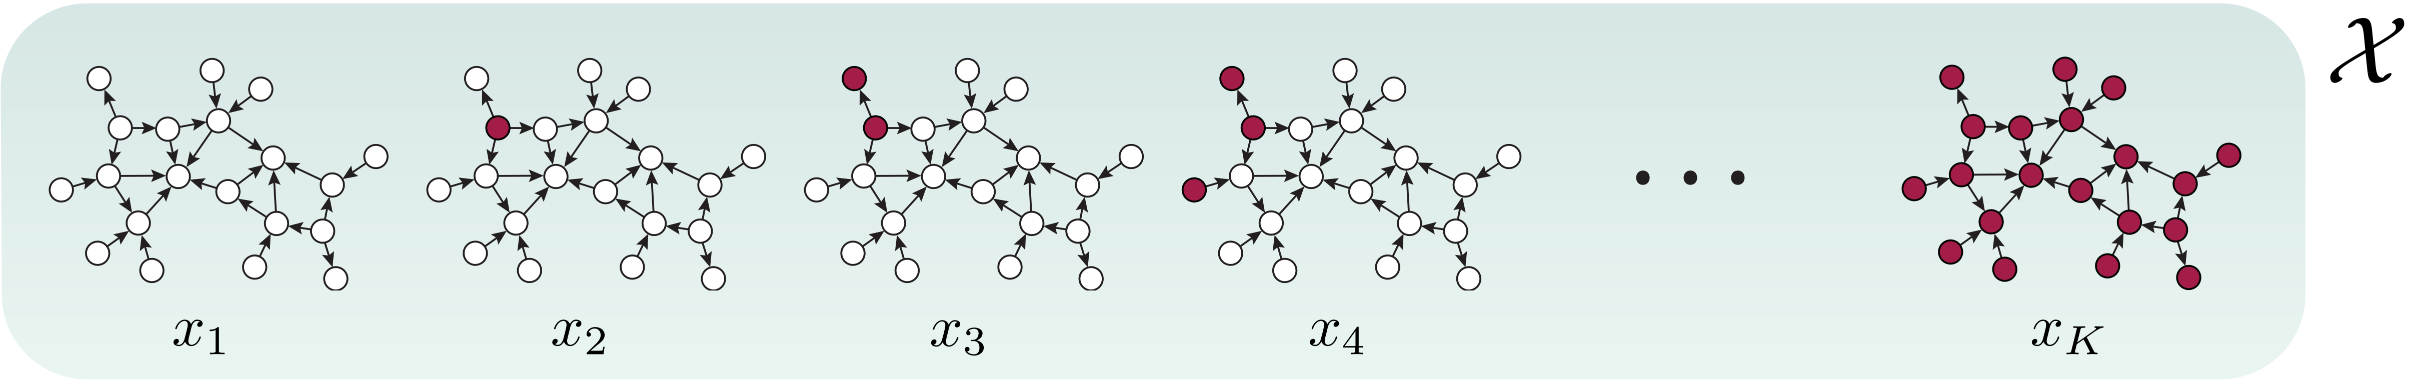

- Information state obeys the update rule $\mathcal{T} : \Delta(\mathcal{X}) \times \mathcal{Y} \times \mathcal{U} \rightarrow \Delta(\mathcal{X})$

$$\pi_{t+1} = \mathcal{T}(\pi_t, y_{t+1}, u_t)$$

Defender's Optimization Problem

- Choose a control policy $g : \Delta(\mathcal{X}) \to \mathcal{U}$, $g \in \mathcal{G}$ that solves

$$\min_{g \in \mathcal{G}} \mathbb{E} \left\{ \sum_{t=0}^{\infty} \rho^t C(\pi_t, g(\pi_t)) \Big| \pi_0 \right\}$$

$$\text{subject to } u_t = g(\pi_t)$$

$$\pi_{t+1} = T(\pi_t, y_{t+1}, u_t)$$

FORCES
FOUNDATIONS OF RESILIENT
CYBER-PHYSICAL SYSTEMS

Example

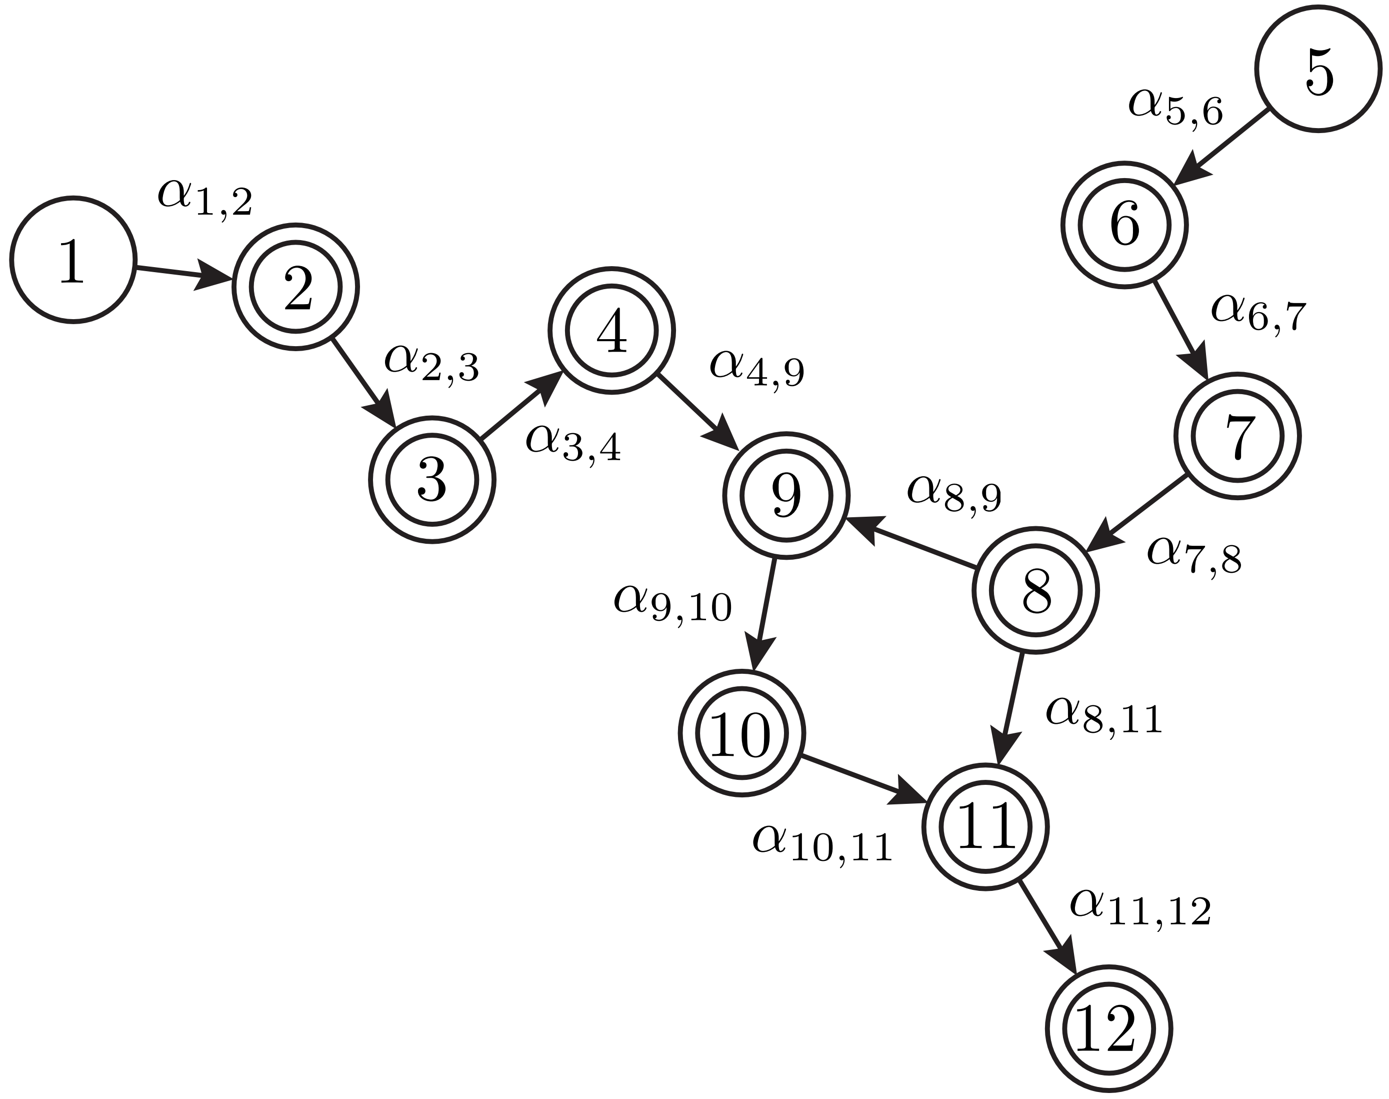

Attributes:

1. Vulnerability in WebDAV on machine 1
2. User access on machine 1
3. Heap corruption SSH on machine 1
4. Root access on machine 1
5. Buffer overflow on machine 2
6. Root access on machine 2
7. Squid portscan on machine 2
8. Network topology leakage from machine 2
9. Buffer overflow on machine 3
10. Root access on machine 3
11. Buffer overflow on machine 4
12. Root access on machine 4

Example - Countermeasures

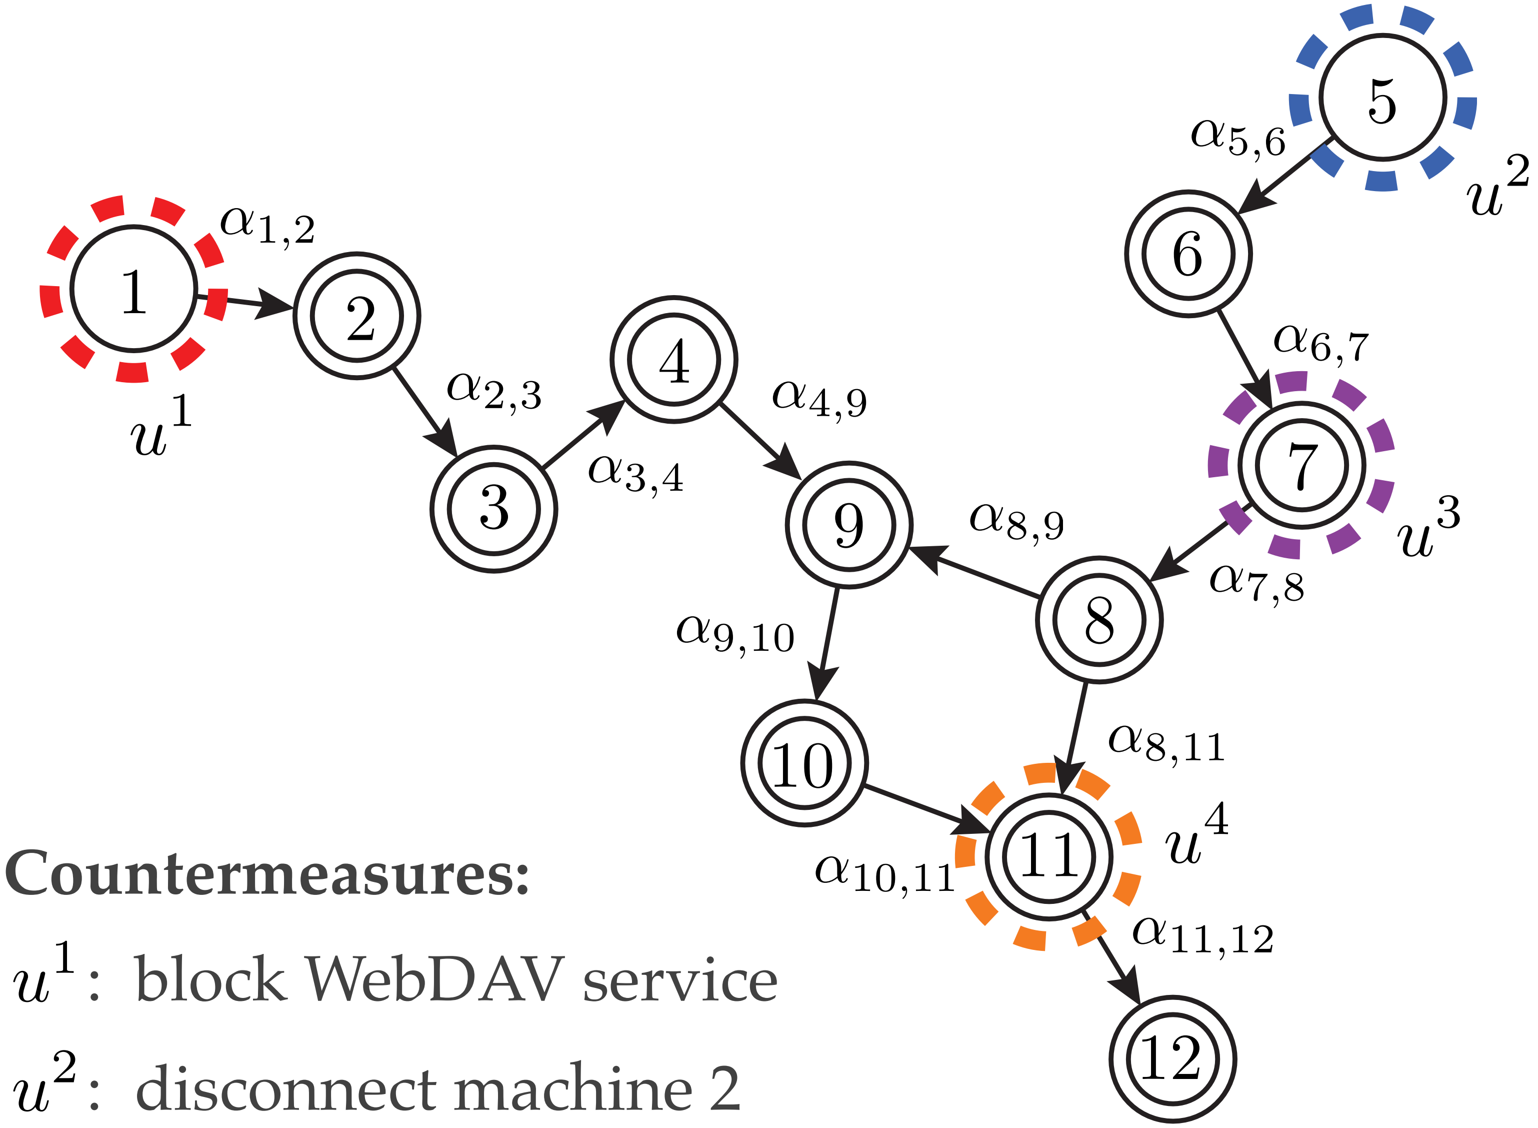

Attributes:

1. Vulnerability in WebDAV on machine 1
2. User access on machine 1
3. Heap corruption SSH on machine 1
4. Root access on machine 1
5. Buffer overflow on machine 2
6. Root access on machine 2
7. Squid portscan on machine 2
8. Network topology leakage from machine 2
9. Buffer overflow on machine 3
10. Root access on machine 3
11. Buffer overflow on machine 4
12. Root access on machine 4

Countermeasures:

u^1: block WebDAV service

u^2: disconnect machine 2

u^3: block port scanning

u^4: disconnect machine 4

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Future Work

- **Structural results**
  - Directed acyclic graphs give rise to a natural partial order
  - Can we use this to show threshold properties of the optimal policy?
    - If so, determining an approximately optimal policy would reduce to estimating these thresholds
- **Scaling the problem**
  - Exact POMDP solvers only capable of handling small examples
  - Realistic attack graphs are big…

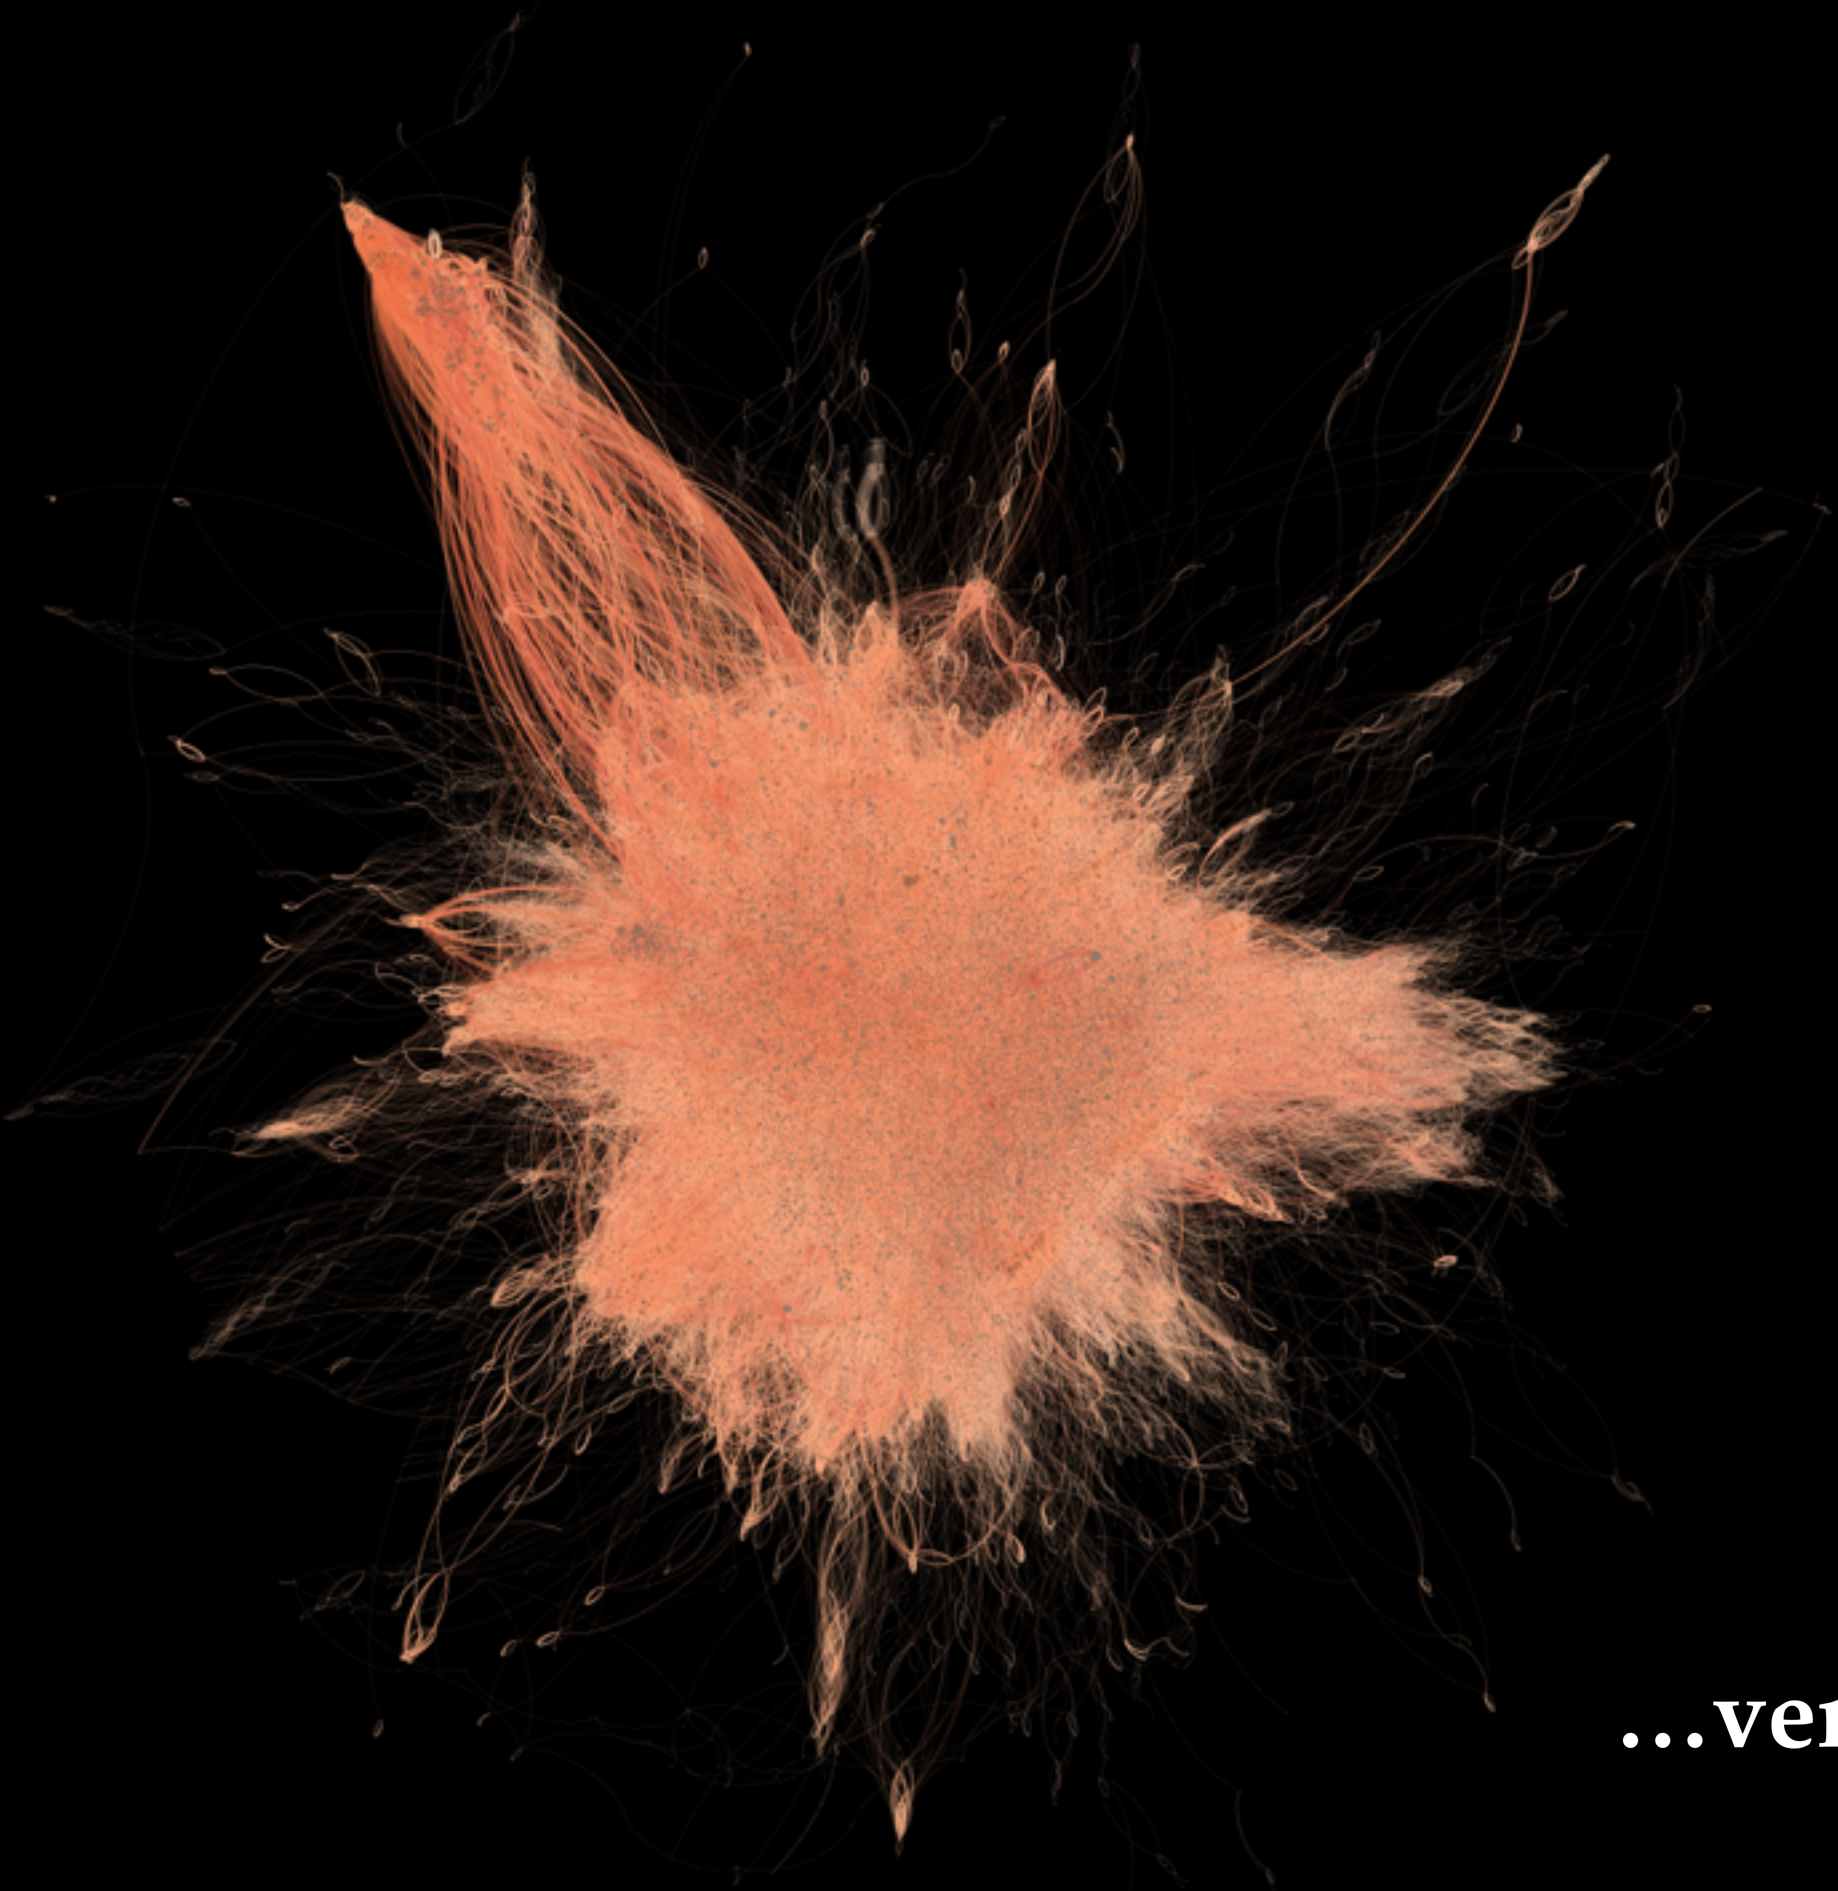

…very big.

Thank You!

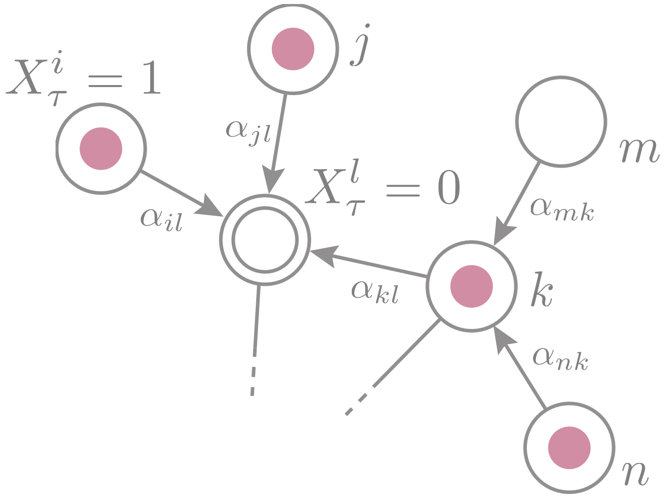

Questions?

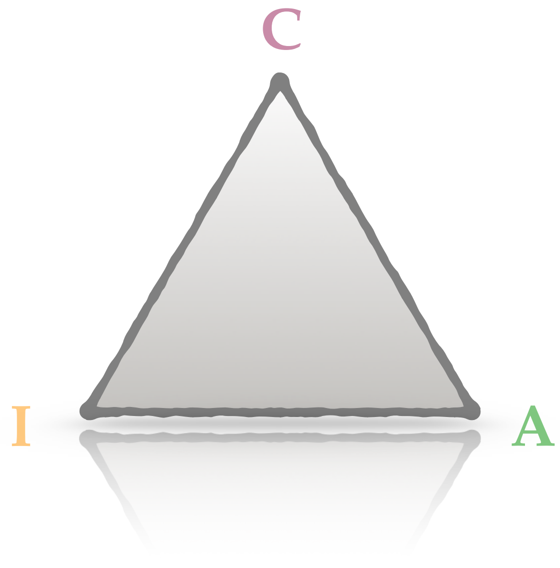

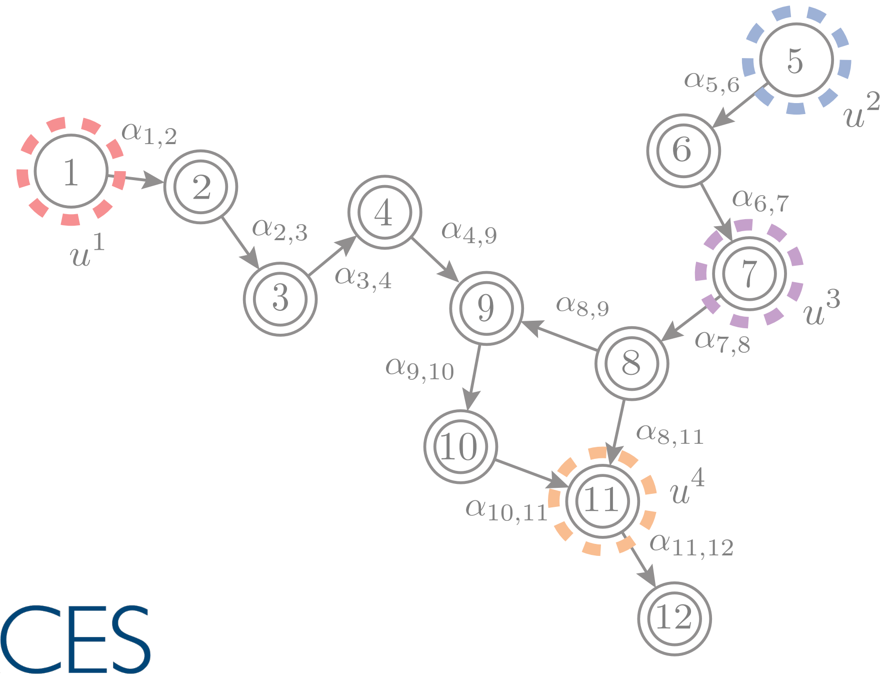

FORCES
FOUNDATIONS OF RESILIENT
CYBER-PHYSICAL SYSTEMS

Funding Acknowledgments

- Special thanks to the following funding sources

- **NSF —** Foundations Of Resilient CybEr-physical Systems (FORCES)

Grant: **CNS-1238962**

- **ARO MURI —** Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Building the Scientific Foundations

Grant: **W911NF-13-1-0421**

FORCES
FOUNDATIONS OF RESILIENT
CYBER-PHYSICAL SYSTEMS