A New Approach to Modeling and Analyzing Security of Networked Systems

Gaofeng Da — UTSA

Maochao Xu — ISU

Shouhuai Xu — UTSA

HotSoS'14

Acknowledgement: ARO

The Problem: Quantitative Security Analysis of Networked Systems

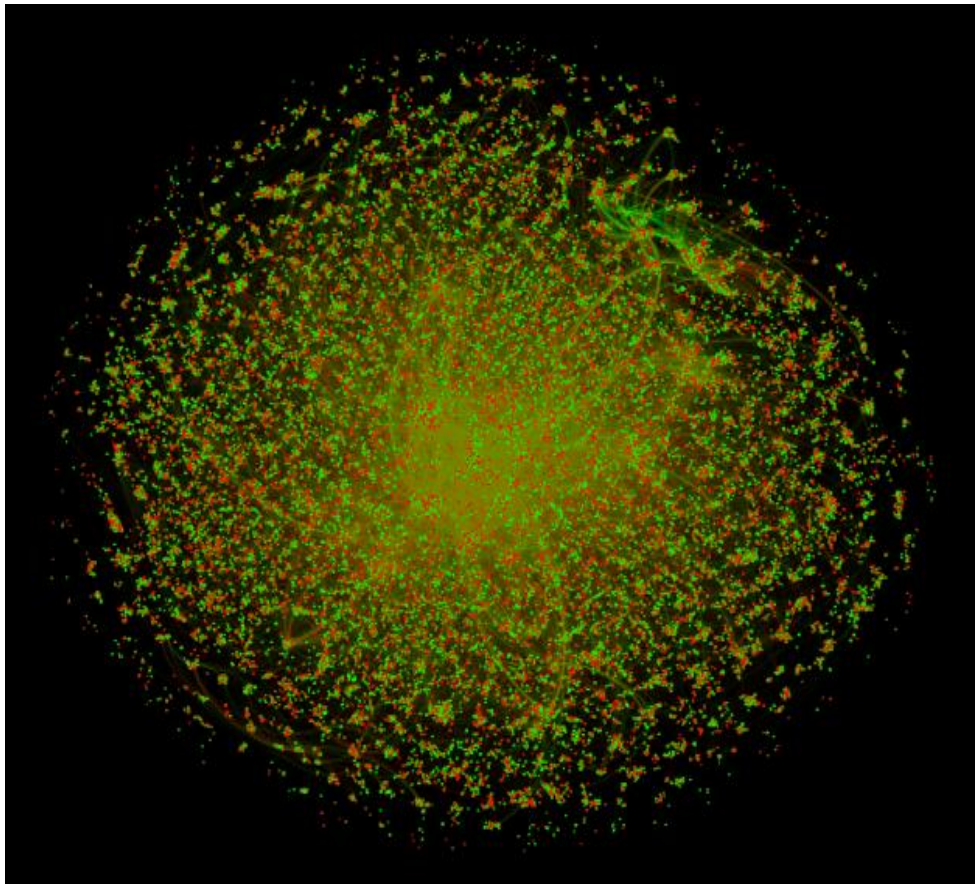

Green node: secure
Red node: compromised

- A problem we all want to solve, but we are far away from where we want to be.
- It cannot be bypassed!
- Very few (even early stage) results: extremely difficult in both modeling and analysis.
- But, the phenomenon is clear.

How Can We Model The Phenomenon

(i.e., evolution of security state of each computer)?

↓↑ successful attack ↓↑ unsuccessful attack ✖ curing

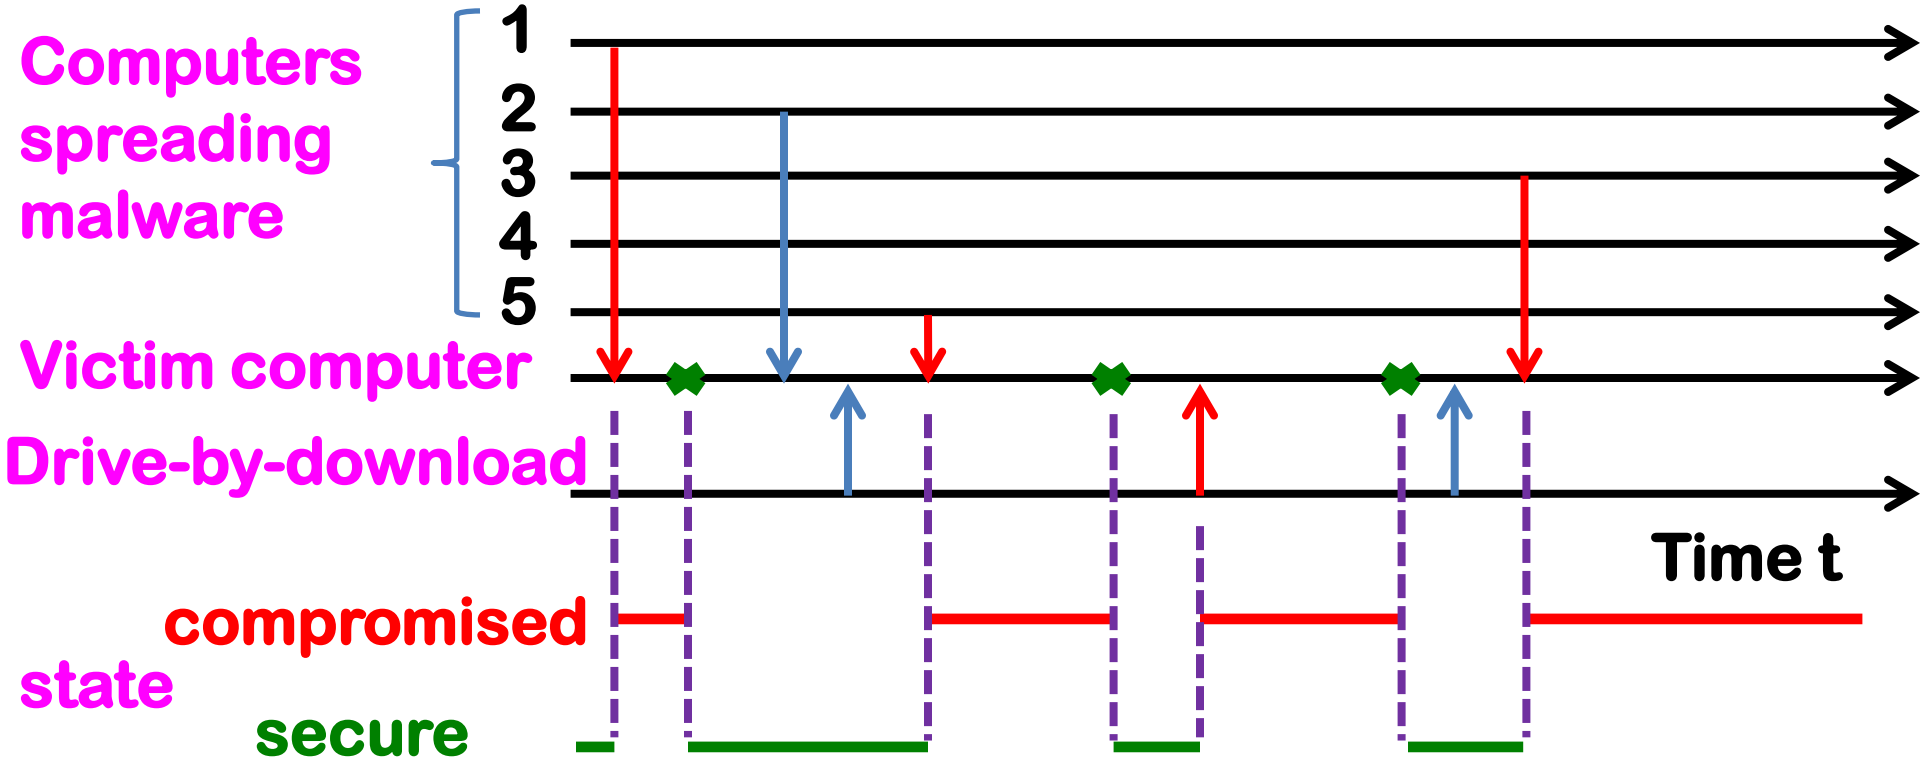

- ❑ Attack inter-arrival time is non-exponential.
- ❑ Stochastic process is non-Markovian.
- ❑ Attack (power) can be dependent/adaptive.

A Killer Problem

Internet or cyberspace

abstraction

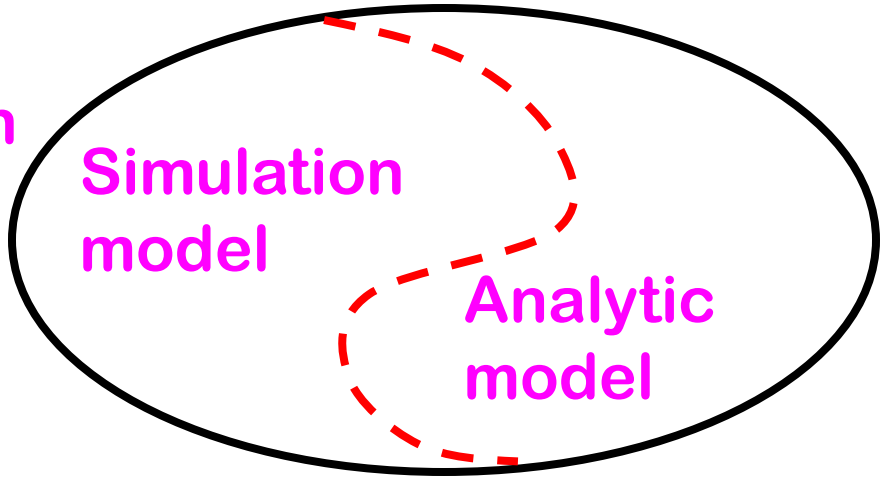

Where is the boundary?

We wish to know for what kind of security analysis purpose we can or cannot build analytically tractable models.

What Is This Paper About?

- The first model that can accommodate a certain degree of *adaptive attacks*.
  - Simple-minded model: 2^n-dimension
  - Our model: n-dimension (approximations/bounds)
- Two security metrics for individual nodes/computers:
  - Time-to-compromise: not necessarily steady-state
  - Steady-state compromise probability: base for
    - Byzantine Agreement assumption (1/3 compr.)
    - Selecting threshold for threshold cryptography
    - Risk management: node compromise probability

Roadmap

- The new approach and the resulting model
- Analyzing the model
  - Analyzing time-to-compromise
  - Analyzing steady-state-compromise probability
- Related work
- Conclusion and future researc directions

The *New* Approach

- What classes of attacks the model accommodates?
  - Push-based attacks in networks: malware spreading
  - Pull-based attacks in networks: drive-by-download
- Push-based attacks induce an (attacker, victim) relation, which induces a graph network G=(V,E)
  - V = node set: computers (or other resolutions)
  - E = edge set: (attacker, victim) relation
  - Node v's neighbors in G formulate *local environment*
- Pull-based attacks is modeled via *global environment*

The *New* Approach (cont.)

- What classes of defense the model accommodates?
  - Preventive defense:
    - Network- and/or host-based Firewall/IDS
    - Filtering recognized attacks
    - *Capability against push-based attacks: c_1*
    - *Capability against pull-based attacks: c_2*
  - Reactive defense:
    - Detecting and curing successful attacks; e.g., anti-malware tools
    - Capability in detecting and curing attack: r.v. $R_{v,i}$

The *New* Approach (cont.)

- **Push-based attacks can be captured by:**
  - ***Magnitude: power/capability of attack ($X^{(1)}$)***
  - **Attack inter-arrival time ($Y^{(1)}$)**
  - ***Success condition: attack magnitude ≥ threshold c_1***
- **Pull-based attacks can be captured by:**
  - ***Magnitude: power/capability of attack ($X^{(2)}$)***
  - **Attack inter-arrival time ($Y^{(2)}$)**
  - ***Success condition: attack magnitude ≥ threshold c_2***

Metric 1: Time-To-Compromise

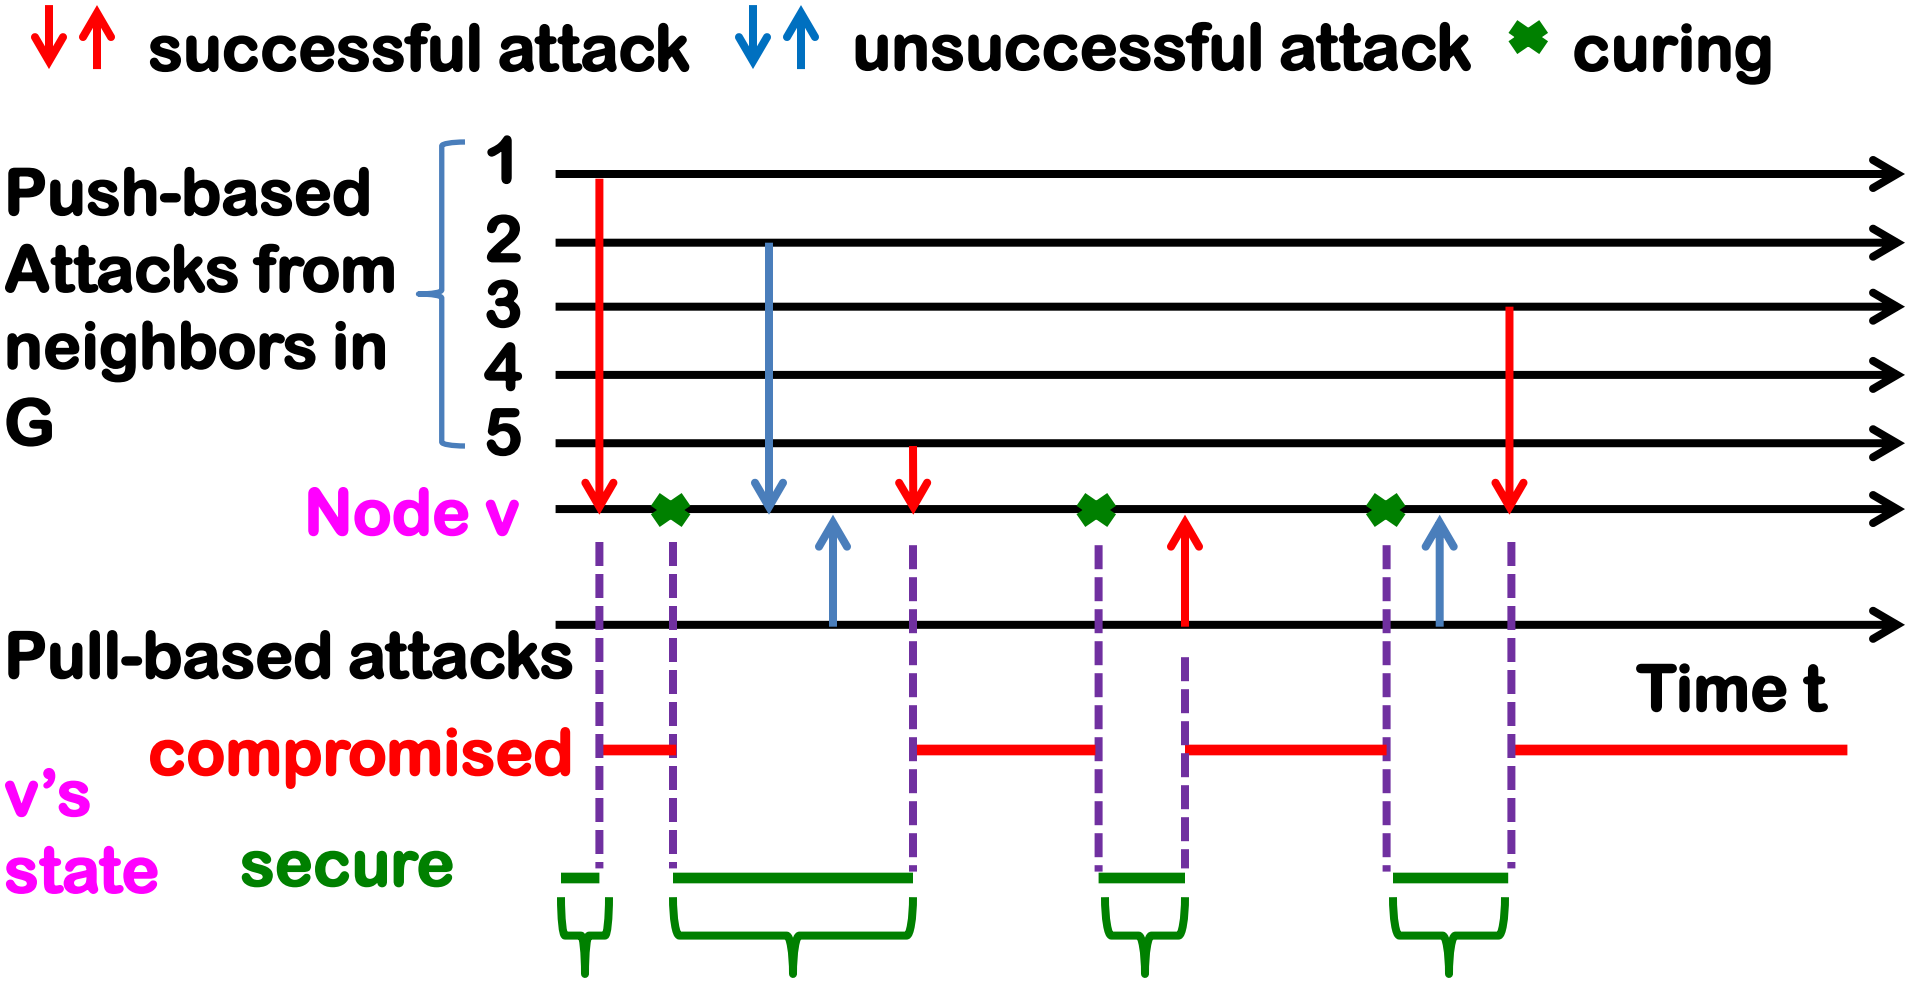

Time-to-compromise $T_{v,c=(c_1,c_2)}$: what is the distribution of this random variable?

Metric 2: Steady-State Comp. Prob.

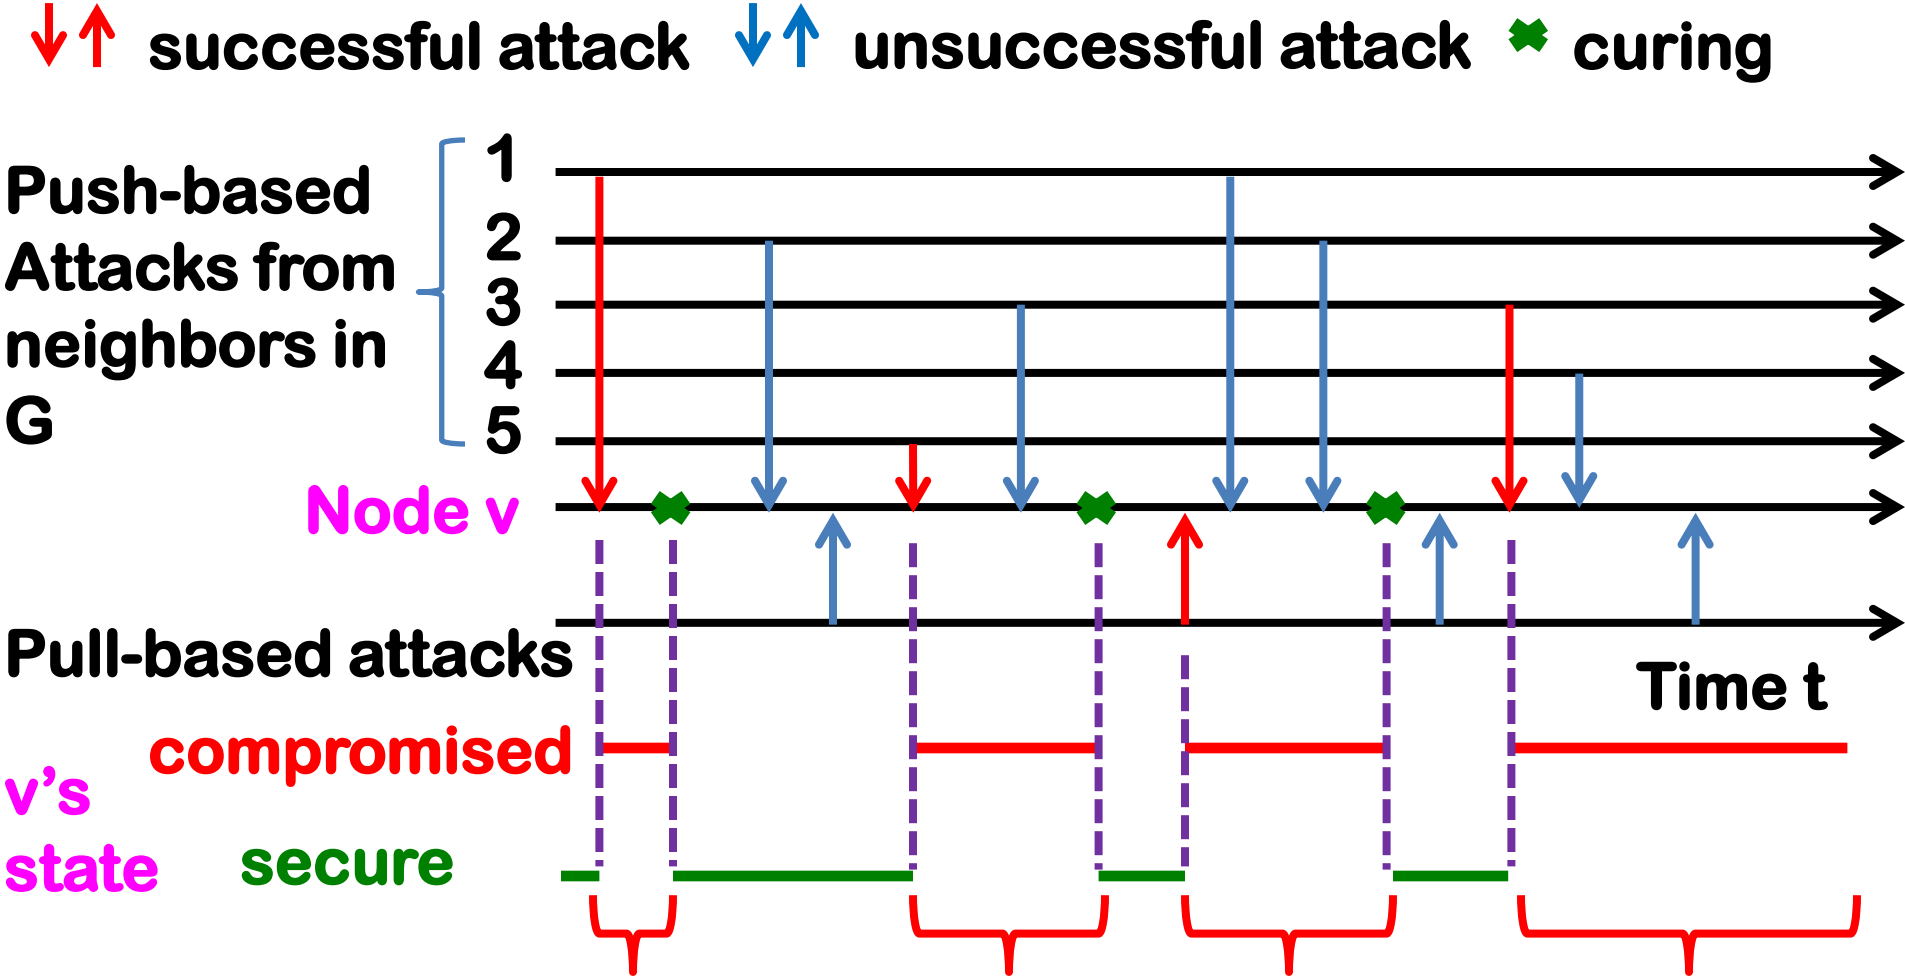

Steady-state compromise probability $p_{v,c=(c_1,c_2)}$: what is the portion of time during which v is compromised (length of red lines/total time)?

Formally Modeling Attacks

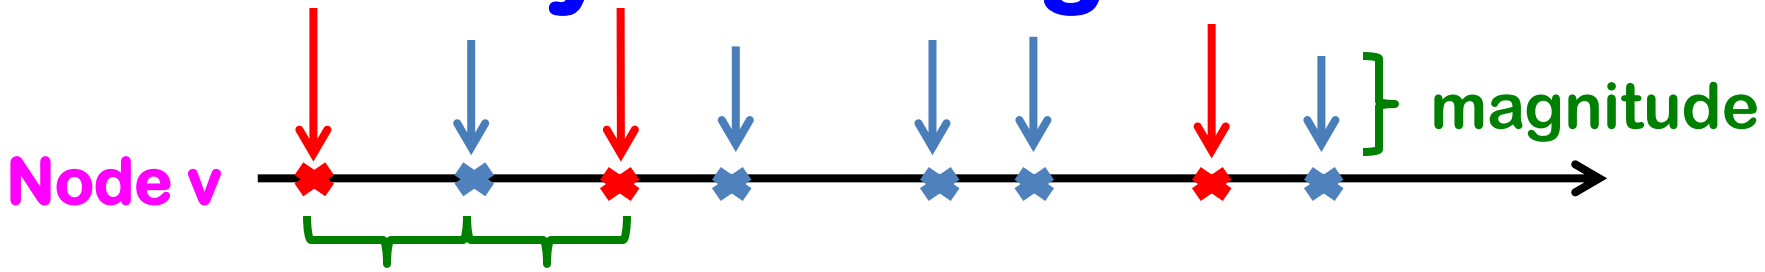

Push-based attacks against $v \in V$ formulate a point process:

$$\left\{ \left(X_i^{(1)}(J_v), Y_i^{(1)}(J_v) \right) \right\}, \quad i = 0, 1, 2 \ldots,$$

Attack ma… Attack inter-arrival time

- ✓ J_v: r… …ack environmen… (random variable)
- ✓ $X_i^{(1)}(J_v)$: random variable representing the magnitude (power) of the ith push-based attack against v; $X_0^{(1)}(J_v) = 0$
- ✓ $Y_i^{(1)}(J_v)$: random variable representing the time interval between the $(i-1)$th and the ith push-based attacks against v; $Y_0^{(1)}(J_v) = 0$.

Formally Modeling Attacks (cont.)

Pull-based attacks against $v \in V$ formulate a point process:

$$\left\{ \left(X_i^{(2)}(\Theta_v), Y_i^{(2)}(\Theta_v) \right) \right\}, i = 1, 2 \ldots,$$

- ✓ Θ er pace)
- ✓ $X_i^{(2)}(\Theta_v)$: random variable representing the magnitude (power) of the ith pull-based attack against v; $X_i^{(2)}(\Theta_v) = 0$
- ✓ $Y_i^{(2)}(\Theta_v)$: random variable representing the inter-arrival time between the $(i-1)$th and the ith pull-based attacks against v; $Y_i^{(2)}(\Theta_v) = 0$

Attack magnitude **Attack inter-arrival time**

Because of reactive defense, a compromised node v becomes secure after a random time $R_{v,i}$ for the ith time.

Metric 1: Time-To-Compromise

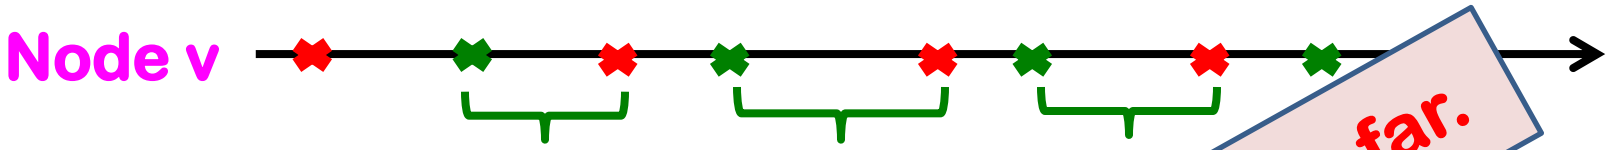

Given $(J_v, \Theta_v) = (r, \theta)$, let $N_r^{(1)}(t)$ be [illegible]nting process associated to sequence $\{Y_i^{(1)}(r)$ [illegible] let $N_\theta^{(2)}(t)$ be the counting processe associ[illegible]quence $\{Y_i^{(2)}(\theta),\ i \geq 0\}$

$$M_r^{(1)}(t) = \vee_{i=0}^{N_r^{(1)}(t)} X^{(1)}$$

$$M_\theta^{(2)}(t) = \vee^{N^{(2)}} (\theta).$$

$$T_{c_1}^{(1)}(r) = \inf\{t : M_r^{(1)}(t) > c_1\},$$

$$T_{c_2}^{(2)}(\theta) = \inf\{t : M_\theta^{(2)}(t) > c_2\}.$$

Good thing: no assumption is needed so far.

Time-To-Compromise

$$T_{v,\boldsymbol{c}} \equiv T_{\boldsymbol{c}}(J_v, \Theta_v) = T_{c_1}^{(1)}(J_v) \wedge T_{c_2}^{(2)}(\Theta_v),$$

$$T_{\boldsymbol{c}}(r, \theta) = T_{c_1}^{(1)}(r) \wedge T_{c_2}^{(2)}(\theta).$$

Key to describing adaptiveness

Metric 2: Steady-State Comp. Prob.

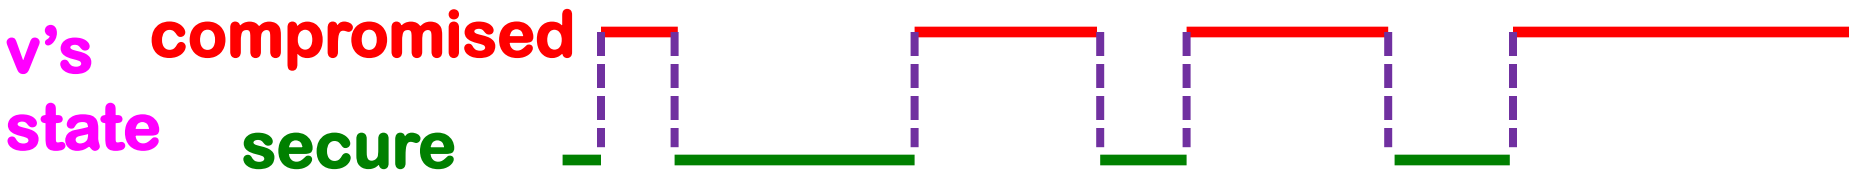

State of $v \in V$ alternates with time.
For given defense capabilities $\boldsymbol{c} = (c_1, c_2)$, the state of $v \in V$ at time t can be seen as Bernoulli random variable $X_{v,\boldsymbol{c}}(t)$, where $X_{v,\boldsymbol{c}}(t) = 1$ means v is compromised and $X_{v,\boldsymbol{c}}(t) = 0$ means v is secure. The probability that v is compromised at time t is:

$$p_{v,\boldsymbol{c}}(t) = \mathrm{P}(X_{v,\boldsymbol{c}}(t) = 1)$$

The probability v is compromised in steady state (if existing):

$$\boxed{p_{v,\boldsymbol{c}} = \lim_{t \to \infty} p_{v,\boldsymbol{c}}(t).}$$

Roadmap

- The new approach and the resulting model
- **Analyzing the model**
  - **Analyzing time-to-compromise**
  - Analyzing steady-state-compromise probability
- Related work
- Conclusion and future researc directions

Analyzing Time-To-Compromise

Assumption 1

(a) For any $v \in V$ and given local environment $J_v = r$,

(i) $\{X_i^{(1)}(r),\ i \geq 1\}$ is an independent sequence;

(ii) $\{Y_i^{(1)}(r),\ i > 1\}$ is an independent sequence;

- Unfortunately, we have to make "ugly" assumptions in order to attain analytic results.
- Still, a certain degree of adaptiveness of attacks can be accommodated or described under the assumptions.

(b) ...ner.

(c) ...ner.

push-based attacks and pull-based attacks) are independent of each other. J_v and Θ_v are independent each other, leading to $\{(X_i^{(1)}(J_v), Y_i^{(1}(J_v), i \geq 1\}$ and $\{(X_i^{(2)}(\Theta), Y_i^{(2)}(\Theta)), i \geq 1\}$ are independent of each other.

Adaptiveness Under Assumption 1

Adaptiveness: Push-based attack magnitude and attack inter-arrival time can be dependent upon each other.

another (possibly more) powerful push-based attack.

Distribution of Time-To-Compromise

Let $\pi_{v,r}$ be the probability mass function of J_v and $H_v(\cdot)$ be the distribution function of Θ_v. Under Assumption 1, distribution of $T_{v,\mathbf{c}}$ is:

$$q_{v,\mathbf{c}}(t) = \mathrm{P}(T_{v,\mathbf{c}} \le t) = 1 - \sum_{r=0}^{\mathrm{d}(v)} \pi_{v,r} \sum_{m=0}^{\infty} \prod_{i=1}^{m} F_{i,r}^{(1)}(c_1) \mathrm{P}\left(N_r^{(1)}(t) = m\right) \cdot \int_0^{\infty} \sum_{m=0}^{\infty} \prod_{i=1}^{m} F_{i,\theta}^{(2)}(c_2) \mathrm{P}\left(N_\theta^{(2)}(t) = m\right) dH_v(\theta).$$

- For general cases, we seek "bounds":
  - Upper bound of $q_{v,c}(t)$: $q_{v,c}(t)^+$
  - Expectation $E[T_{v,c}]$ and its lower bound $E[T_{v,c}]^-$
- Asymptotic results only for special case: $c_1, c_2 \to \infty$

General Cases: Upper Bound: $q_{v,c}(t)^+$

Assumption 2

(a) The same as Assumption 1(a).

(b) The same as Assumption 1(b).

(c) For any $v \in V$ and given environments $J_v = r$ and $\Theta_v = \theta$, $\left\{\left(X_i^{(1)}(r), Y_i^{(1)}(r)\right), i > 1\right\}$ are independently and identically

Much easier to compute $q_{v,c}(t)^+$ than to compute $q_{v,c}(t)$, because the former deals with random variables rather than stochastic processes.

Proposition 1 (upper bound of $q_{v,c}(t)$)

Suppose Assumption 2 holds, and $Y^{(1)}(r)$ and $Y^{(2)}(\theta)$ have the NBU property for any given local environment $J_v = r$ and global environment $\Theta_v = \theta$. We have

$$q_{v,c}^+(t) = 1 - \sum_{r=0}^{\deg(v)} \pi_{v,r} [\bar{G}_r^{(1)}(t)]^{\bar{F}_r^{(1)}(c_1)} \int_0^\infty [\bar{G}_\theta^{(2)}(t)]^{\bar{F}_\theta^{(2)}(c_2)} \, \mathrm{d}H_v(\theta).$$

Adaptiveness under Proposition 1

The good.

+: But, the "memory" property of NBU adds another kind of adaptiveness/dependence to attack inter-arrival time:

$$\mathrm{P}\left(Y_{i+1}^{(1)}(r) > z_1 + z_2 | Y_{i+1}^{(1)}(r) > z_1\right) \leq \mathrm{P}\left(Y_{i}^{(1)}(r) > z_2\right).$$

The extra-waiting time for the $(i+1)$th push-based attack to arrive is, under the condition that the attack has not arrived after time z_1, shorter than the waiting time for the ith attack in the stochastic sense. Moreover, we have

$$\mathrm{P}\left(Y_{i+1}^{(1)}(r) > z_1 + z_2 | Y_{i+1}^{(1)}(r) > z_1\right) \leq \mathrm{P}\left(Y_{i+1}^{(1)}(r) > z_2\right).$$

Lower Bound of Expectation $E[T_{v,c}]^-$

Proposition 2 (lower bound $\mathrm{E}[T_{v,\boldsymbol{c}}]^-$)

Suppose Assumption 2 holds, and $Y^{(1)}(r)$ and $Y^{(2)}(\theta)$ have the NBUE property for any given local environment $J_v = r$ and global environment $\Theta_v = \theta$. We have

$$\mathrm{E}[T_{v,\boldsymbol{c}}]^- = \sum_{r=0}^{\deg(v)} \int_0^{\infty} \pi_{v,r} \left(\frac{\bar{F}_r^{(1)}(c_1)}{\mathrm{E}[Y^{(1)}(r)]} + \frac{\bar{F}_\theta^{(2)}(c_2)}{\mathrm{E}[Y^{(2)}(\theta)]} \right)^{-1} \mathrm{d}H_v(\theta).$$

- Similar to Proposition 1, NBUE accommodates another kind of adaptiveness/dependence.
- The bound is also easier to compute than $E[T_{v,c}]$ because it deals with random variables (not processes).
- Under NBU/NBUE, $E[T_{v,c}]$ exists (not necessarily true for heavy-tailed attack inter-arrival time).

Numerical Examples

Consider random node $v \in V$. Set

- ✓ $X^{(1)}(r) \sim Weibull(\alpha, 1/r)$, $Y^{(1)}(r) \sim Gamma(\beta, r)$, $\beta \geq 1$;
- ✓ $X^{(2)}(\theta) \sim Weibull(\gamma, 1/\theta)$, $Y^{(2)}(\theta) \sim Gamma(\lambda, \theta)$, $\lambda \geq 1$;
- ✓ $J_v \sim Binomial(p, \deg(v))$, $\Theta_v \sim Uni(a, b)$.

Note that for all r and θ, $Y^{(1)}(r)$ and $Y^{(2)}(\theta)$ have the NBU property because their shape parameters $\beta \geq 1$ and $\lambda \geq 1$.

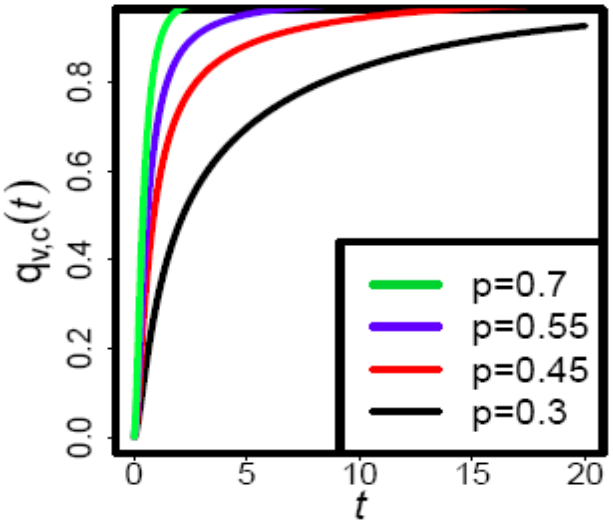

(a) $\deg(v) = 8, c_1 = c_2 = 3$

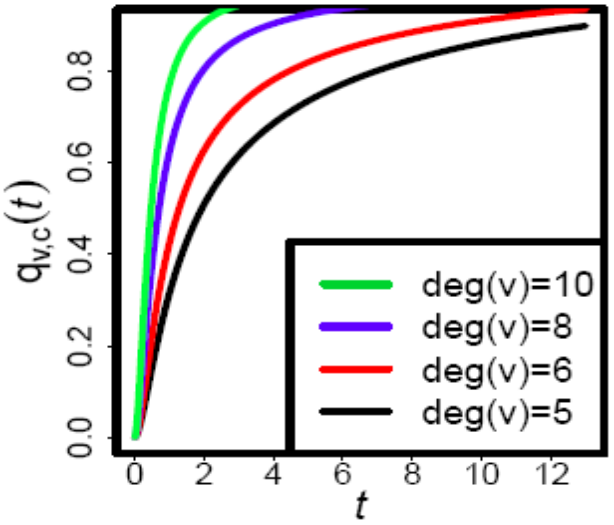

(b) $c_1 = c_2 = 3, p = .5$

effect of time ≈ effect of space (degree)?

Tightness of Bounds

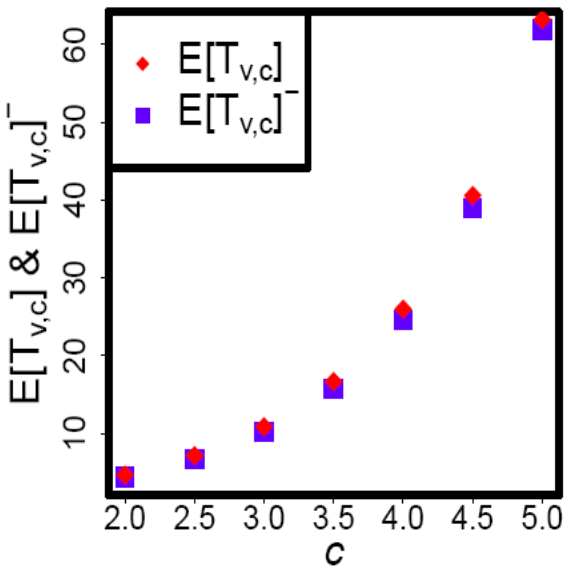

(a) $p = 0.2$

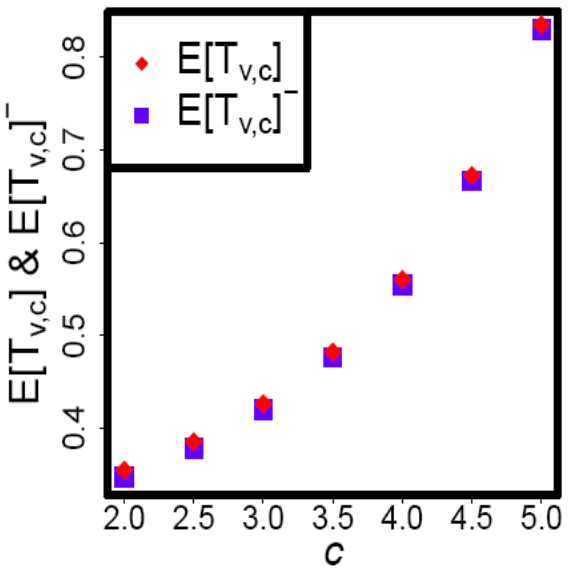

(b) $p = 0.8$

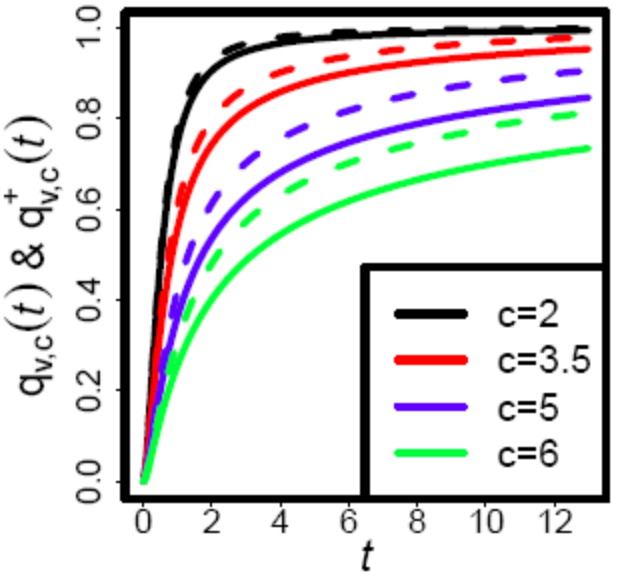

(c) $\deg(v) = 8, p = .5, c_1 = c_2$

Observation: bounds are tight when $c=c_1=c_2$ is small (≤ 2).

Implication: Time-to-compromise is relatively easy to quantify (via bounds) for badly defended networks (because the situation is too bad?).

Asymptotic $q_{v,c}(t)$ in Special Case

Proposition 4

Suppose Assumption 2 holds. Suppose μ_r, the mean of attack inter-arrival times $Y^{(1)}(r)$, and ν_θ, the mean of attack inter-arrival times $Y^{(2)}(\theta)$, are finite. As $c_1, c_2 \to \infty$, we have

$$q_{v,\boldsymbol{c}}(t) \sim \tilde{q}_{v,\boldsymbol{c}}(t) := \sum_{r=0}^{\deg(v)} \pi_{v,r}\left(1 - e^{-\bar{F}_r^{(1)}(c_1)t/\mu_r}\right) + \int_0^\infty \left[1 - e^{-\bar{F}_\theta^{(2)}(c_2)t/\nu_\theta}\right] \mathrm{d}H_v(\theta).$$

The adaptiveness accommodated by Proposition 4 is the same as the adaptiveness accommodated by Assumption 2 (and therefore slightly weaker than the adaptiveness accommodated by Assumption 1).

Asymptotic $q_{v,c}(t)$ in Special Case

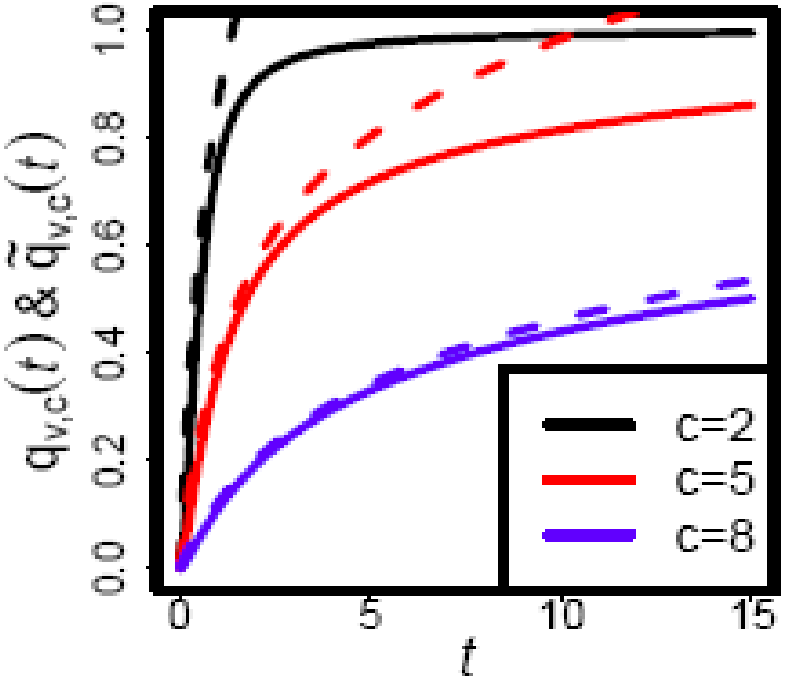

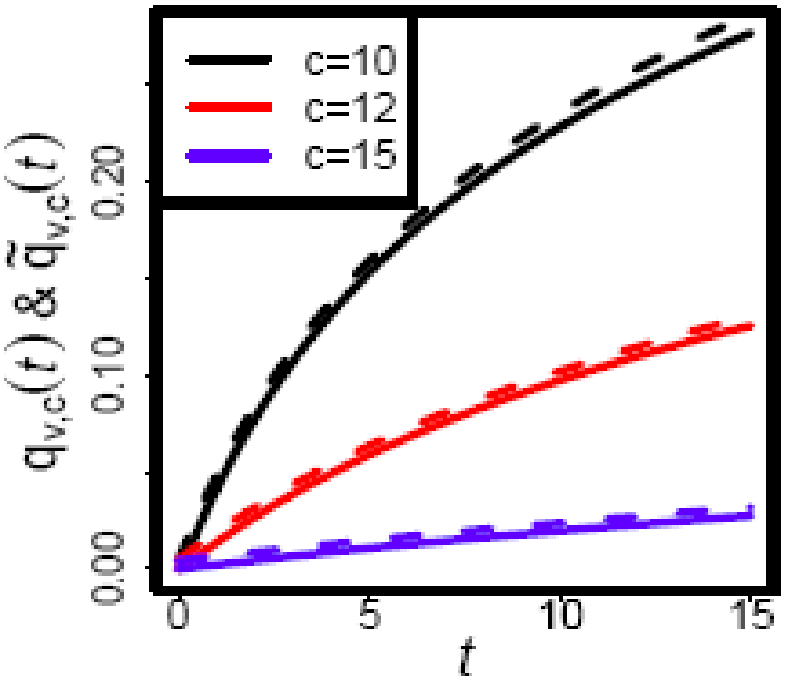

Asymptotic results are tights for $c = c_1 = c_2 \geq 12$ (or 8).

Insight: $q_{v,c}(t)$ is relatively easy to quantify via asymptotic results (and bounds) for highly effectively (and badly) defended networks.

Roadmap

- The new approach and the resulting model
- Analyzing the model
  - Analyzing time-to-compromise
  - **Analyzing steady-state-compromise probability**
- Related work
- Conclusion and future researc directions

Recap: Steady-State Comp. Prob.

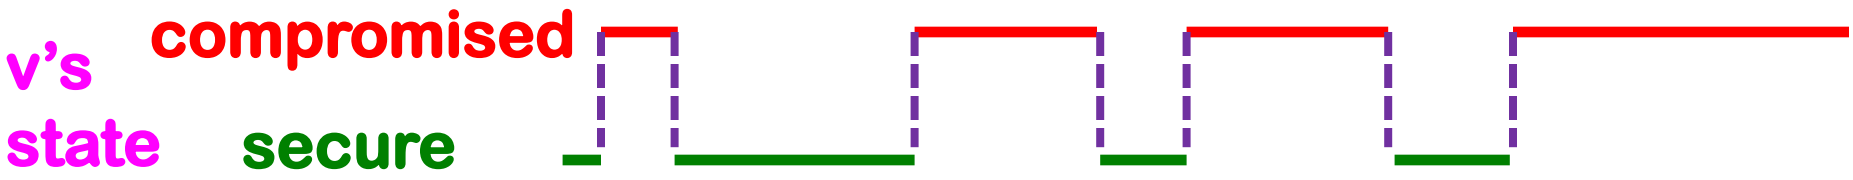

In order to get analytic results, we need to make restrictions (not very restrictive, though):

Under the condition that each compromise-and-recovery cycle has the same distribution after each recovery (i.e., reset to the secure state), $J_{v,i}$ for $i = 1, 2, \ldots$ are independently and identically distributed. Therefore, the stochastic process can be seen as an alternating renewal process with a sequence of vectors $(T_{v,\boldsymbol{c},i}, R_{v,i})$ for $i \geq 1$, which have the same distributions as $(T_{v,\boldsymbol{c}}, R_v)$.

Analyzing Steady-State Comp. Prob.

To attain $p_{v,\boldsymbol{c}}$ for $v \in V$, we need to solve the following system of equations:

$$p_{v,\boldsymbol{c}} \int_0^{\infty} \mathrm{E}\left[\left(F^{(1)}_{\sum_{u=1}^{n} a_{uv} X_u}(c_1)\right)^{N^{(1)}_{\sum_{u=1}^{n} a_{uv} X_u}(t)}\right] \cdot \mathrm{E}\left[\left(F^{(2)}_{\Theta_v}(c_2)\right)^{N^{(2)}_{\Theta_v}(t)}\right] \mathrm{d}t - (1 - p_{v,\boldsymbol{c}})\mathrm{E}[R_v] = 0, \quad v \in V.$$

n-dimension rather than 2^n-dimension

- In case the n-equation system is not numerically solvable, we offer the following bounds of $p_{v,c}$.
- The bounds can be used for decision-making.
- Numerically results show that sometimes (but not always) the upper bound is tight.
- How can we improve the upper bound?

Bounding Steady-State Comp. Prob.

The bunds require a further assumption (restriction):

Assumption 4

Suppose the following monotonicity holds:

(a) $X^{(1)}(r)$ is increasing in r in stochastic order (i.e., greater number of compromised neighbors implies greater magnitude of push-based attacks). $X^{(2)}(\theta)$ is increasing in θ in stochastic order (i.e., severer environment implies greater magnitude of pull-based attacks).

(b) $Y^{(1)}(r)$ is decreasing in r in stochastic order (i.e., greater number of compromised neighbors implies more frequent push-based attacks), and $Y^{(2)}(\theta)$ is decreasing in θ in stochastic sense (i.e., severer environment implies more frequent pull-based attacks).

Bounding Steady-State Comp. Prob.

Proposition 5 (bounds of $p_{v,\boldsymbol{c}}$)

Suppose Assumptions 2 and 4 hold. We have for all $v \in V$

$$\left(1+\frac{\mathrm{E}[Y^{(2)}(\bar{\theta}_v)]}{\bar{F}^{(2)}_{\bar{\theta}_v}(c_2)\mathrm{E}[R_v]}\right)^{-1} \leq p_{v,\boldsymbol{c}} \leq \left(1+(\mathrm{E}[R_v])^{-1}\cdot \int_0^\infty \mathrm{E}\left[(F^{(1)}_{\deg(v)}(c_1))^{N^{(1)}_{\deg(v)}(t)}\right]\mathrm{E}\left[(F^{(2)}_{\bar{\theta}_v}(c_2))^{N^{(2)}_{\bar{\theta}_v}(t)}\right]\,\mathrm{d}t\right)^{-1}.$$

Adaptiveness accommodated by Proposition 5:

Magnitudes: Since $X^{(1)}_i(r)$ is increasing in r for any $i \geq 1$ in stochastic order, we have

$$\mathrm{P}(X^{(1)}_{i+1}(J_v) > s' | X^{(1)}_i(J_v) > s) \geq \mathrm{P}(X^{(1)}_{i+1}(J_v) > s')$$

Inter-arrival times: Since $Y^{(1)}_i(r)$ is decreasing in r for any $i \geq 1$ in stochastic order, we have

$$\mathrm{P}(Y^{(1)}_{i+1}(J_v) \leq s' | Y^{(1)}_i(J_v) \leq s) \geq \mathrm{P}(Y^{(1)}_{i+1}(J_v) \leq s').$$

Related Work

- Fall into the big picture of *Cybersecurity Dynamics* (see poster & another talk tomorrow afternoon)
  - Anther approach to whole-system security modeling: attack graphs
- In terms of the means for achieving same/similar goal
  - Getting rid of exponential attack inter-arrival time [Internet Mathematics 2012]
  - Accommodating static dependence between random variables [Internet Mathematics 2014]

Conclusion and Future Work

- This paper: a first step towards ultimately modeling *adaptive attacks* (or dynamic dependence), while allowing analytic treatment.
- Future work (difficult to analyze):
  - Make J_v (local push-based attack environment), Θ_v (global pull-based attack environment) and $R_{v,i}$ (curing time) driven by stochastic processes.
  - Make c_1 (defense power against push-based attacks) and c_2 (defense power against pull-based attacks) r.v. or even driven by stochastic processes.
- Future work: orthogonal thread: obtaining parameters etc

Towards the Ultimate Goal

- How far can this approach go?
- What are the better ways to go?
- Where is the boundary between analytically tractable models and analytically intractable (simulation) models!

Enjoy exploring the unknown territory!