Efficiency of Selfish Investments in Network Security

Libin Jiang, Venkat Anantharam and Jean Walrand
Dept. EECS, UC Berkeley

NetEcon'08 Workshop
Aug 22, 2008

Outline

- Introduction and Goal
- Price of Anarchy (POA) in one-shot game
  - Two models
- Using “Weighted POA” to Bound Cost Regions
- The “best equilibrium” in the repeated game

Introduction

- Roles of technology and incentives in network security
- Example: Virus propagation
  - Infected computers continue infecting others
- Users invest to reduce security risk
  - Money: Anti-virus software
  - Time: patching, scanning
- Positive externalities → under-invest

Basic model

- User *i*'s cost function: $g_i(\mathbf{x}) = f_i(\mathbf{x}) + c_i x_i$
  - $x_i \geq 0$: Investment: *strategy*
  - $f_i(\mathbf{x}) \geq 0$: Risk, decreasing and convex

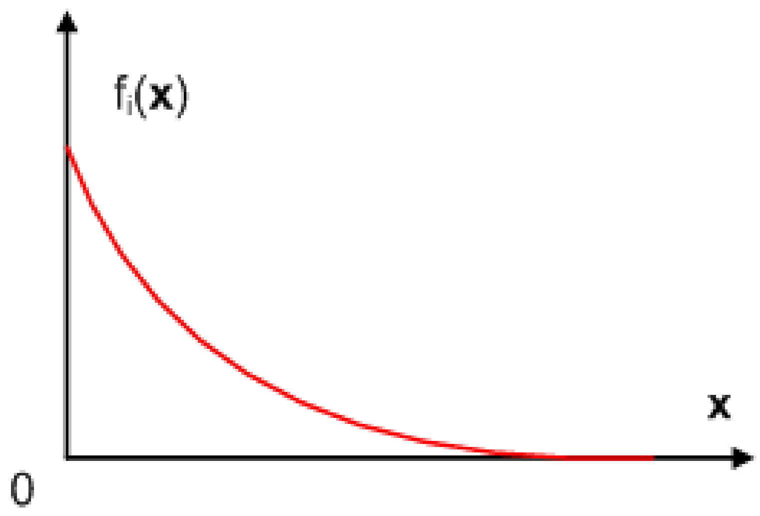

Goal

- How bad is selfish investment?
  - Is regulation necessary?
- Important factors that determine POA?
  - Network topology, players' inter-dependency
  - Players' heterogeneous cost functions
  - Strategic-form (one-shot) game or Repeated game

Outline

- Introduction & Goal
- Price of Anarchy in one-shot game
  - Two models
- Using “Weighted POA” to Bound Cost Regions
- The “best SPE” in the repeated game

POA

$$g_i(\mathbf{x}) = f_i(\mathbf{x}) + c_i x_i, i = 1, 2, \ldots, n$$

- $\bar{\mathbf{x}}$: a Nash Equilibrium (NE); $\mathbf{x}^*$: Social Optimum (SO). Find an upper bound of

$$\rho := \frac{\bar{G}}{G^*} = \frac{\sum_i g_i(\bar{\mathbf{x}})}{\sum_i g_i(\mathbf{x}^*)}$$

- Proposition 1:

$$\rho \leq \max\{1, \max_k\{(-\sum_i \frac{\partial f_i(\bar{\mathbf{x}})}{\partial x_k})/c_k\}\}$$

- Note that $(-\sum_i \frac{\partial f_i(\bar{\mathbf{x}})}{\partial x_k})$ is the marginal benefit to the overall security by increasing x_k at a NE; whereas c_k is the marginal cost

POA

- Specific forms of risk function $f_i(\mathbf{x})$, related to network topology, traffic rates, etc
  - Effective-investment (EI) Model
  - Bad-traffic (BT) model
- Use Prop. 1 to obtain explicit expressions of POA

Effective-investment (EI) Model

- WLOG, assume $g_i(\mathbf{x}) = f_i(\mathbf{x}) + x_i$, where
  $f_i(\mathbf{x}) = V_i(\sum_{j=1}^{n} \beta_{ji} x_j)$
  - $V_i(\cdot)$ is decreasing, non-negative and convex.
  - $\beta_{ji} \geq 0$ is the "relative importance" of the player *j*'s investment to player *i* . WLOG, assume $\beta_{ii} = 1, \forall i$
- Using Prop. 1, we have (tight bound)

$$\rho \leq Q = \max_k \{1 + \sum_{i: i \neq k} \beta_{ki}\}$$

User *k*'s "relative importance" to the society

EI Model: Dependency Graph and POA

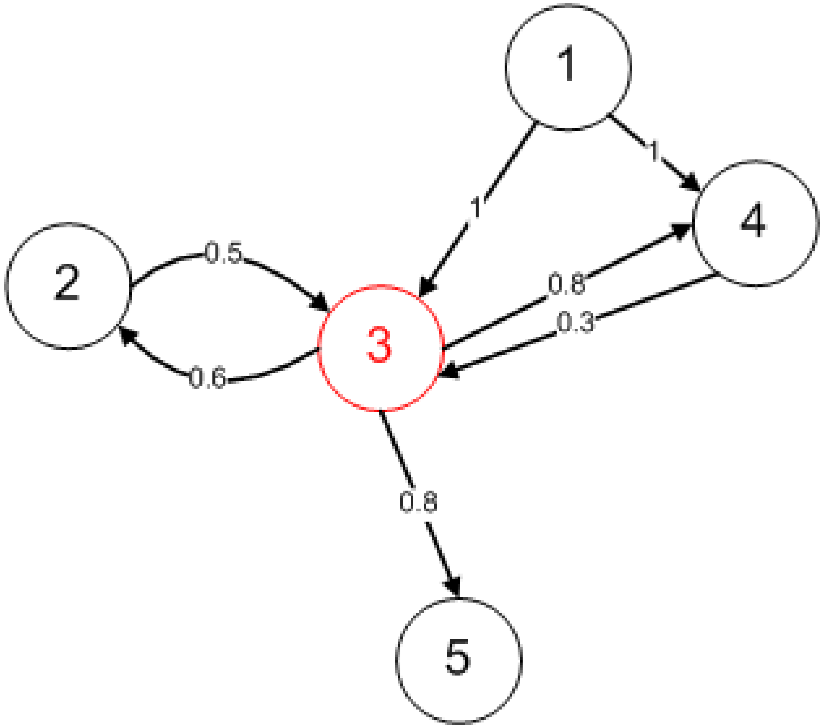

In this example, $\rho \leq 1 + 2.2 = 3.2$

If all players are equally important to each other, i.e., $\beta_{ki}=1, \forall\ k,i$, then POA is *n*, which is the number of players.

Bad-traffic (BT) model

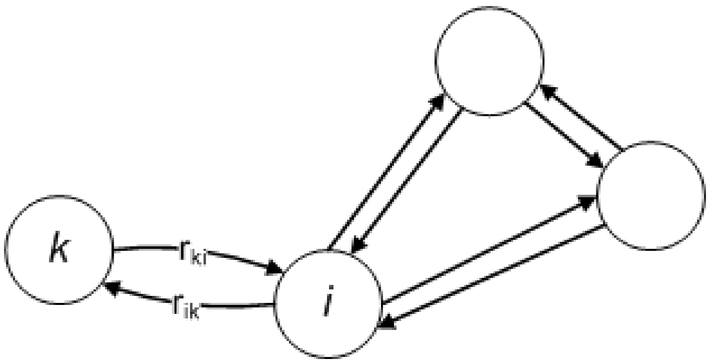

$$g_i(\mathbf{x}) = f_i(\mathbf{x}) + x_i$$

$$f_i(\mathbf{x}) = v_i \sum_{k \neq i} r_{ki} \phi_{k,i}(x_k, x_i)$$

- Assume “symmetric” firewall, i.e., $\phi_{k,i}(x_k, x_i) = \phi_{i,k}(x_i, x_k)$.
- Using Prop. 1, we have (tight bound)

$$\rho \leq Q = 1 + \max_{(i,j):i \neq j} \frac{v_i r_{ji}}{v_j r_{ij}}$$

- POA is the “maximum imbalance” of the network (+1).
- If each pair is “balanced”, POA=2.

Outline

- Introduction & Goal
- Price of Anarchy in one-shot game
  - Two models
- Using “Weighted POA” to Bound Cost Regions
- The “best SPE” in the repeated game

Weighted POA

- POA is 1-dimensional information
- "Weighted POA" Q_w: an upper bound of

$$\rho_{\mathbf{w}} := \frac{\bar{G}_{\mathbf{w}}}{G^*_{\mathbf{w}}} = \frac{\sum_i w_i \cdot g_i(\bar{\mathbf{x}})}{\sum_i w_i \cdot g_i(\mathbf{x}^*_{\mathbf{w}})} \le Q_{\mathbf{w}}$$

where $\mathbf{w} \in \mathcal{R}^n_{++}$. $\bar{\mathbf{x}}$ is a NE; and $\mathbf{x}^*_{\mathbf{w}}$ minimizes $G_{\mathbf{w}}(\mathbf{x}) := \sum_i w_i \cdot g_i(\mathbf{x})$

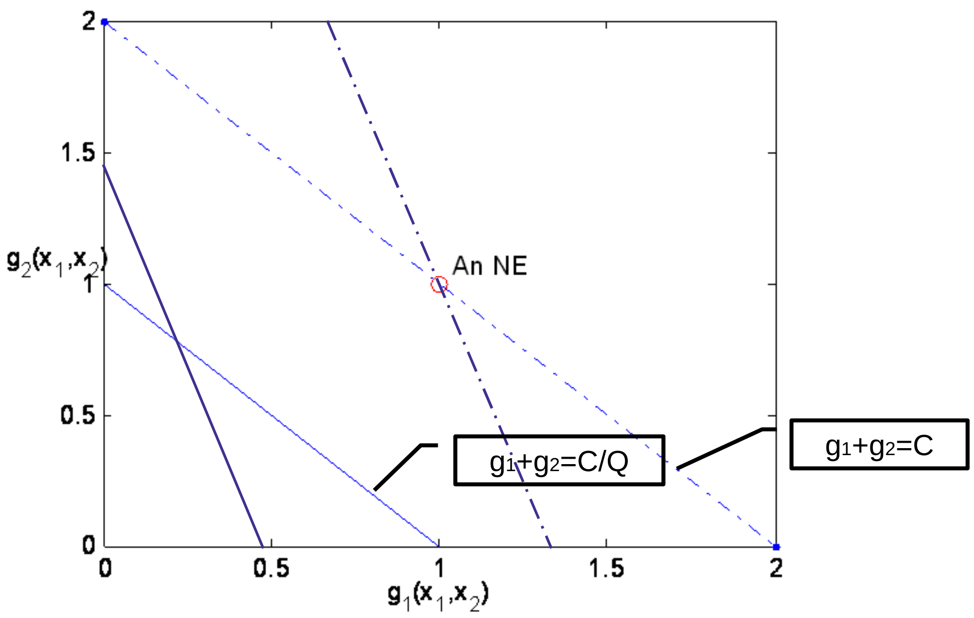

Weighted POA

- Define a modified game where the cost function of player *i* is $w_i g_i(\mathbf{x})$, then the POA in the game is $Q_{\mathbf{w}}$
  - EI model:

$$\rho_{\mathbf{w}} \le Q_{\mathbf{w}} = \max_{k}\{1 + \frac{\sum_{i:i\neq k} w_i \beta_{ki}}{w_k}\}$$

  - BT model:

$$\rho_{\mathbf{w}} \le Q_{\mathbf{w}} = 1 + \max_{(i,j):i\neq j} \frac{w_i v_i r_{ji}}{w_j v_j r_{ij}}$$

Weighted POA: Example

Two agents, and $\beta_{11} = \beta_{22} = 1, \beta_{12} = \beta_{21} = 0.2$.

$g_i(\mathbf{x}) = (1 - \sum_{j=1}^{2} \beta_{ji} x_i)_+ + x_i$, for $i = 1, 2$. It is easy to verify that $\bar{x}_i = 0, i = 1, 2$ is an NE, and $g_1(\bar{\mathbf{x}}) = g_2(\bar{\mathbf{x}}) = 1$.

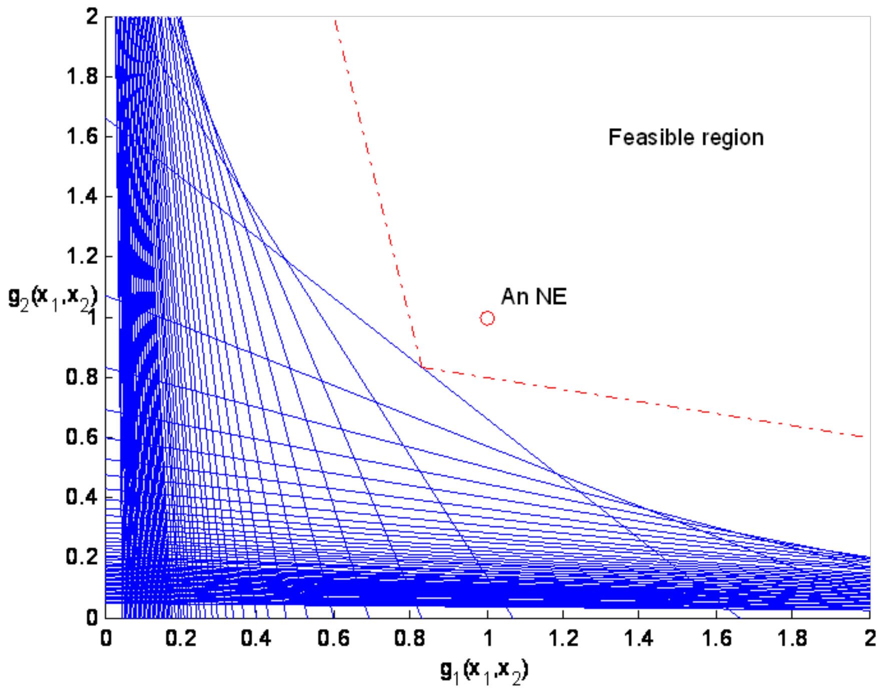

The bound applies to any non-negative, decreasing and convex functions $f_1(\cdot)$ and $f_2(\cdot)$.

Outline

- Introduction & Goal
- Price of Anarchy in one-shot game
  - Two models
- Using "Weighted POA" to Bound Cost Regions
- The "best SPE" in the repeated game

"Best" SPE in the Repeated Game

$$g_i(\mathbf{x}) = f_i(\mathbf{x}) + x_i, i = 1, 2, \ldots, n$$

- By *Folk Theorem*, any feasible vector $\mathbf{g} \leq \underline{\mathbf{g}}$ can be supported by a SPE
  - $\underline{\mathbf{g}}$: vector of *reservation costs* (max-min)
- "Best" SPE: SPE with the minimum social cost G_E

$$G_E = \min_{\mathbf{x}} \sum_i g_i(\mathbf{x})$$
$$s.t. \quad g_i(\mathbf{x}) \leq \underline{g_i}, i = 1, 2, \ldots, n$$

  - Compare to SO:

$$\gamma = G_E / G^*$$

Performance of the "Best" SPE

- SO is not always achievable by SPE

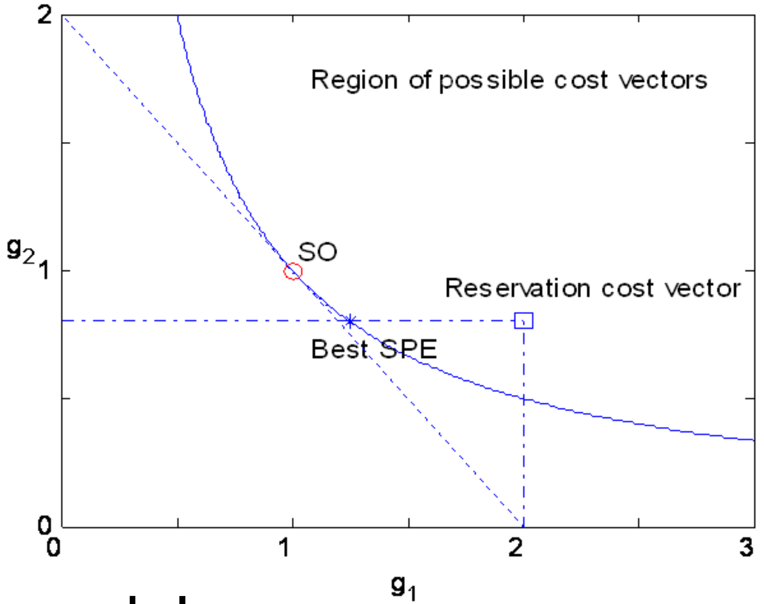

2-player game

- In the EI model

$$\gamma \leq \min\{\max_{i,j,k} \frac{\beta_{ik}}{\beta_{jk}}, Q\}, Q = \max_k\{1 + \sum_{i:i\neq k} \beta_{ki}\}$$

If $\beta_{ki} = 1, \forall k, i$ (symmetric), then $\rho_{max} = n$ but $\gamma = 1$: the best SPE is SO, n times better than NE.

Conclusion

- Positive externality
- In strategic-form game
  - POA can be very large and tends to increase with the network size, dependency, imbalance.
  - Severe efficiency problem
- In repeated game
  - Usually much better
  - SO possible if not conflicts individual interest
  - But need more communications and coordination
- Simple & flexible model. More general applications

How to improve security?

- *Stick and carrot*
  - Catch the (elusive) attackers…
  - Due care (with a regulator)
    - Each user must invest at least x_i^*, or it is punished (properly)
  - Requiring users to take minimum security measures
  - Incentivizing ISP's to do more
    - Filtering traffic, blocking infected users

Thank you~

Questions & Comments?

Weighted POA (optional)

- POA is 1-dimensional information
- "Weighted POA" Q_w: an upper bound of

$$\rho_{\mathbf{w}} := \frac{\bar{G}_{\mathbf{w}}}{G^*_{\mathbf{w}}} = \frac{\sum_i w_i \cdot g_i(\bar{\mathbf{x}})}{\sum_i w_i \cdot g_i(\mathbf{x}^*_{\mathbf{w}})} \leq Q_{\mathbf{w}}$$

  - where $\mathbf{w} \in \mathcal{R}^n_{++}$. $\bar{\mathbf{x}}$ is a NE of the original game; and $\mathbf{x}^*_{\mathbf{w}}$ minimizes $G_{\mathbf{w}}(\mathbf{x}) := \sum_i w_i \cdot g_i(\mathbf{x})$. We have

$$\sum_i w_i \cdot g_i(\mathbf{x}) \geq \sum_i w_i \cdot g_i(\mathbf{x}^*_{\mathbf{w}}) \geq \sum_i w_i \cdot g_i(\bar{\mathbf{x}})/Q_{\mathbf{w}}, \forall \mathbf{x}$$

- Proposition 3
  - Given any NE, then all feasible cost vectors are in

$$\mathcal{B} := \{\mathbf{g} | \mathbf{w}^T \mathbf{g} \geq \mathbf{w}^T \bar{\mathbf{g}}/Q_{\mathbf{w}}, \forall \mathbf{w} \in \mathcal{R}^n_{++}\}$$

  - Given a feasible cost vector, all NE cost vectors are in

$$\bar{\mathcal{B}} := \{\bar{\mathbf{g}} | \mathbf{w}^T \bar{\mathbf{g}} \leq \mathbf{w}^T \mathbf{g} \cdot Q_{\mathbf{w}}, \forall \mathbf{w} \in \mathcal{R}^n_{++}\}$$