Near Rationality and Competitive Equilibria in Networked Systems

Nicolas Christin, Jens Grossklags and John Chuang
University of California, Berkeley
School of Information Management and Systems
{christin,jensg,chuang}@sims.berkeley.edu

Background

- Game: Model of interactions between
  - participants (*players*) in a network, who
  - (based on available info) select *strategies*, that
  - yield individual *payoffs* (utilities)
- Nash equilibrium: set of strategies from which no utility-maximizing user has any incentive to deviate

$$\forall \sigma_i' \in \Sigma_i \ , \ u_i(\sigma_i', \sigma_{-i}^*) - u_i(\sigma_i^*, \sigma_{-i}^*) \leq 0$$

Traditional assumptions

- Each player has perfect knowledge of structure of the game
  - Strategies available to all other users
  - Payoffs associated with each strategy
- Each player is perfectly rational
- Players perfectly execute their strategies

Very stringent assumptions…
… especially in the context of large networks!

Thesis statement

Slightly expanding the (Nash) solution concept allows to:

1) evaluate the robustness of a model to small perturbations

2) help reconcile empirical data with predictions

For instance, by considering:

$$\exists \varepsilon > 0 \;\; : \;\; \forall \sigma_i' \in \Sigma_i \;\; , \;\; u_i(\sigma_i', \sigma_{-i}^*) - u_i(\sigma_i^*, \sigma_{-i}^*) \leq \varepsilon$$

(sometimes called ε equilibrium)

Case studies

- Three case studies to illustrate our point
  - Network formation
  - Response to a security threat
  - TCP congestion control
- Methodology
  - Describe the game
  - Compare Nash and ε-equilibrium outcomes
  - Discuss findings

Case study 1: Network formation

- Players
  - Network nodes interested in creating peering connections to other nodes (e.g., ad-hoc network)
- Utility of a player
  - Parameterized function of
    - its distance to other nodes (hop count),
    - its connectivity (out degree), and
    - its path loading (number of routes passing through the node)
- Strategies available to a player
  - Choose the set of connections maintained with other nodes

Case study 1: Network formation

- Nash outcome

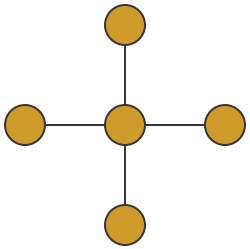

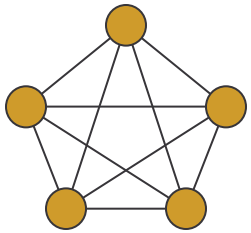

Nodes either organize in a star-shaped network, or in a fully connected network (classical result)

- ε-equilibrium outcome

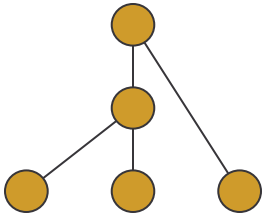

For a range of parameters (depending on ε), *any* topology forms an ε-equilibrium

- Findings
  - Even with a small uncertainty ε, the outcome of the model may be significantly different
  - Model needs to be refined for complex networks (e.g., the Internet)

Case study 2: Response to security threat

- Simplified model of first stage of a (manual or) semi-automated DDoS attack
- Players
  - Network users, all subject to a security threat of unknown severity
  - The least protected user(s) is (are) compromised
- Utility of a player
  - If not compromised, utility is characterized by a cost function *C(s)* increasing with their security level *s*
  - If compromised, utility is characterized by a large penalty *P >> C(s)* for any feasible *s*
- Strategies available to a player
  - Select security level *s*

Case study 2: Response to security threat

- Nash outcome

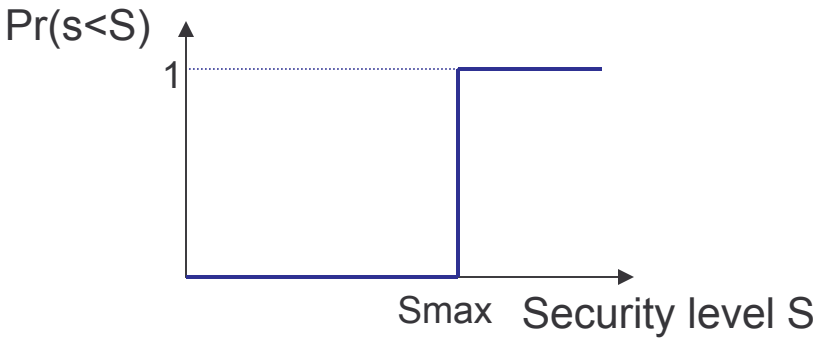

All users choose the highest security level available

- ε-equilibrium outcome

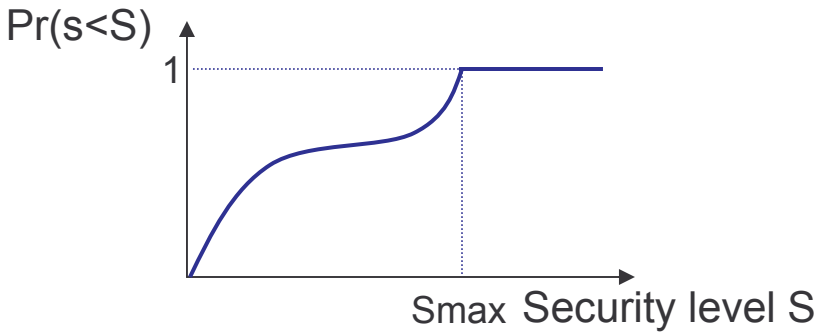

Dispersion of security levels between no security and highest security

- Findings
  - Empirical data doesn't corroborate predicted Nash outcome
  - Simple refinement captures that users are not *perfectly* rational but instead "near rational"

Case study 3: TCP congestion control

- Model initially proposed by [Akella et al. ’02]

- Players
  - TCP (SACK) flows sharing bottleneck link(s) in the network
- Utility of a player
  - Throughput (successful transmission rate) obtained by the flow
- Strategies available to a player
  - Change the TCP congestion control parameters
    - additive increase/multiplicative decrease coefficients

Case study 3: TCP congestion control

- Nash outcome

Congestion control is *completely disabled*
(no matter what)

- ε-equilibrium outcome

No change to the default congestion control parameters for large number of flows
(assuming changes to congestion control parameters are bounded over finite time intervals)

- Findings
  - Empirical data doesn't corroborate predicted Nash outcome [Akella et al. '02]
    - …otherwise the Internet would collapse due to congestion!
  - Combining near rationality and model refinement enables us to match observed behavior
  - Tweaking TCP stack not worth the hassle

Summary

- Model designers should test (or even allow) for slight deviations from Nash equilibrium (near rationality)
  - Can help assess the robustness of a model
    - Network formation
    - Response to security threat
  - Can help determine if a model is too simplistic
    - Network formation
  - Can help reconcile empirical data with predicted outcome
    - TCP congestion control
    - Response to security threat

Open questions and research agenda

- This is not a plea for using ε-equilibria!
  - Only a convenient (i.e., very simple) tool for our argument
  - Other types of equilibrium (e.g., QRE) potentially useful
- Model design methodology
  - Accuracy vs. complexity of a model
  - Refine the game definition or the equilibrium concept
- Mechanism design methodology
  - Impact of near rationality on mechanism performance
- Application: Security model
  - Refine our model (distinguish between different threats)
  - Gather empirical data (user behavior)

Questions?