Mathematical Modeling of Incentive Policies in P2P Systems

John C.S. Lui
cslui@cse.cuhk.edu.hk

Department of Computer Science & Engineering
The Chinese University of Hong Kong

Outline

Motivation

- Cooperation plays an essential role in many developing large-scale network systems and application.
  - Wireless mesh networks (e.g., forward packets).
  - P2P file sharing systems (e.g., BitTorrent [Performance 2007]).
  - P2P streaming, VoD (e.g., PPLive, P2P-VoD [Sigcomm 2008]).
- Individuals are selfish.
- Important to consider incentive protocols to encourage cooperation.

Background

- Micro-payment in Napster. Weakness: central authority.
- Tit-for-tat in Bit-torrent. Free-riding is still possible.
- Reputation-based policies. Concern: collusion.

Background: continue

- Natural for nodes to *learn* from the environment.
- Shared history based incentive mechanisms can overcome scalability problem of private history based mechanisms.
- Designing/testing a "good" incentive is difficult.
- Design and evaluation of incentive protocols: ad-hoc

Contribution

- A general (and simple) mathematical framework to analyze and evaluate incentive protocols for P2P systems.
- Analysis of several incentive policies using this framework.
- Performance evaluation for these incentive policies.
- Connection with *evolutionary game theory*.

Assumptions

- **Finite strategies:** Given an incentive policy $\mathcal{P}$ which has a finite strategy set

$$\mathcal{P} = \{s_1, s_2, \ldots, s_n\},$$

where s_i is the i^{th} strategy. All users in a P2P system can use any $s_i \in \mathcal{P}$. A user chooses s_i is of type i.

- **Service model:** The system runs in discrete time slots. At the beginning of each time slot, each peer randomly selects another peer in the system and requests for a service.
- Denote $g_i(j)$ as the probability that a peer of type s_i will provide a service to a peer of type s_j.

Assumptions cont.

- **Gain and loss model:** at each time slot, a peer gains $\alpha > 0$ points when it receives a service from another peer, while loses β points when it provides a service to another.Without loss of generality, one can normalize β by setting $\beta = 1$.
- **Learning model:**
  - At the end of a time slot, a peer can choose to switch (or adapt) to the current best strategy s_h.
  - Let $\mathcal{G}_i(t)$ be the expected gain of using strategy s_i at time slot t, then a peer using strategy s_i will switch to strategy s_h at time slot $t + 1$ with probability

$$\gamma(\mathcal{G}_h(t) - \mathcal{G}_i(t)),$$

where $\gamma > 0$ is the learning rate.

General Model

- Let $x_i(t)$ be the fraction of type s_i peers at time t.
- If a peer is of type s_i, the expected services it receives, denoted by $E[R_i(t)]$, can be simply expressed as:

$$E[R_i(t)] = \sum_{j=1}^{n} x_j(t) g_j(i) \quad \text{for } i = 1, \ldots, n. \tag{1}$$

- The expected number of services provided by type s_i peer at time t is $E[S_i(t)]$, which is:

$$E[S_i(t)] \approx \sum_{j=1}^{n} x_j(t) g_i(j) \quad \text{for } i = 1, 2, \ldots, n. \tag{2}$$

General Model

- Since a peer receives α points for each service it receives and loses $\beta = 1$ point for each service it provides, the expected gain per slot at time t is $\mathcal{G}_i(t)$:

$$\mathcal{G}_i(t) = \alpha \sum_{j=1}^{n} x_j(t) g_j(i) - \sum_{j=1}^{n} x_j(t) g_i(j) \quad i = 1, 2, \ldots, n. \tag{3}$$

- We can put the above expression in matrix form and derive $\mathcal{G}(t)$, the expected gain per slot for the whole P2P system at time t as

$$\mathcal{G}(t) = \sum_{i=1}^{n} x_i(t) \mathcal{G}_i(t) = (\alpha - 1) \boldsymbol{x}^T(t) G \boldsymbol{x}(t), \tag{4}$$

where $\boldsymbol{x}(t)$ is a column vector of $(x_1(t), \ldots, x_n(t))$ and G is an $n \times n$ matrix with $G_{ij} = g_i(j)$.

General Model

- According to the learning mechanism, we can describe the dynamics as this fluid model:

$$\dot{x}_h = \gamma \sum_{i \neq h} x_i(t) \left(\mathcal{G}_h(t) - \mathcal{G}_i(t)\right)$$
$$= \gamma \left(\mathcal{G}_h(t) - \sum_{i=1}^{n} x_i(t)\mathcal{G}_i(t)\right) = \gamma \left(\mathcal{G}_h(t) - \mathcal{G}(t)\right) \qquad (5)$$
$$\dot{x}_i = -\gamma x_i(t) \left(\mathcal{G}_h(t) - \mathcal{G}_i(t)\right), \quad i \neq h. \qquad (6)$$

Key ideas

- Given an incentive policy $\mathcal{P}$, we have to first find out all $g_i(j)$, or *all* entries in $\boldsymbol{G}$.
- Once we found G, we can derive:

$$\begin{aligned}
\boldsymbol{x}(t) &= [x_1(t), x_2(t), x_3(t), \ldots,], \\
\mathcal{G}_i(t) &= \text{Performance measure of each strategy} \\
\mathcal{G}(t) &= \text{Performance measure of the incentive policy}
\end{aligned}$$

Three types of peers

In a typical P2P system, one can classify peers according to their *behavior* upon receiving a request:

- **cooperator:** a peer has a cooperative behavior when it serves other peers unconditionally.
- **defector:** a peer has a defective behavior when it refuses to serve any request from other peers.
- **reciprocator:** a peer has a reciprocative behavior when it serves according to the requester's contribution level. In short, it tries to make the system fair.

Image Policy $\mathcal{P}_{image}$

- Image incentive policy $\mathcal{P}_{image}$ has three pure strategies:
  1. s_1, or pure cooperation,
  2. s_2, or image reciprocation,
  3. s_3, or pure defection.
- Under this policy, when a reciprocative peer receives a request for service:
  - this peer checks (or infers) the requester's reputation, and
  - it will only provide service with the *same probability* as this requester serves other peers.

Image Policy $\mathcal{P}_{image}$: continue

- To model this incentive policy, we have to derive $g_i(j)$.
- For s_1 (pure cooperation), we have:

$$g_1(j) = 1 \quad j = 1, 2, 3.$$

- For s_3 (pure defection), we have:

$$g_3(j) = 0 \quad j = 1, 2, 3.$$

- For s_2 (image reciprocation):
  - $g_2(1) = 1$.
  - $g_2(3) = 0$.
  - $g_2(2) = ?$

Image Policy cont.

- To derive $g_2(2)$:

$$\begin{aligned} g_2(2) &= \text{Prob[a reciprocator will grant a request]} \\ &= \sum_{i=1}^{3} \text{Prob[the requester is of type } s_i] \times \\ &\quad \text{Prob[granting the request} | \text{type } s_i \text{ requests]} \\ &= x_1(t)g_2(1) + x_2(t)g_2(2) + x_3(t)g_2(3) \\ &= x_1(t) + x_2(t)g_2(2). \end{aligned}$$

- Solving the above equation, we have

$$g_2(2) = \frac{x_1(t)}{1 - x_2(t)}. \qquad (7)$$

Proportional Policy $\mathcal{P}_{prop}$

- Three types of peers:
  1. s_1 (cooperator);
  2. s_2 (reciprocator);
  3. s_3 (defector);
- Reciprocative peers serve the requester with the probability equal to the requester's consumption to contribution ratio, or $E[S_j]/E[R_j]$.
- In case the ratio is larger than one, the probability to serve the request is set to one.

Proportional Policy $\mathcal{P}_{prop}$: continue

- For s_1 (pure cooperation), we have:

$$g_1(j) = 1 \quad j = 1, 2, 3.$$

- For s_3 (pure defection), we have:

$$g_3(j) = 0 \quad j = 1, 2, 3.$$

- For s_2 (reciprocator)
  - If the requester is a cooperator, its ratio is ≥ 1, thus $g_2(1) = 1$.
  - If the requester is a defector, its ratio is zero, hence $g_2(3) = 0$.
  - $g_2(2) = ?$

Proportional Policy ($\mathcal{P}_{prop}$) cont.

- For $g_2(2)$, we have:

$$
\begin{aligned}
E[R_2(t)] &= x_1(t)g_1(2) + x_2(t)g_2(2) + x_3(t)g_3(2) \\
&= x_1(t) + x_2(t)g_2(2), \\
E[S_2(t)] &= x_1(t)g_2(1) + x_2(t)g_2(2) + x_3(t)g_2(3) \\
&= x_1(t) + x_2(t)g_2(2).
\end{aligned}
$$

- Since $E[R_2(t)] = E[S_2(t)]$, $g_2(2) = 1$.

Linear Incentive Policy Class $\mathcal{C}_{LIP}$

- $\mathcal{P}_{prop}$ belongs to the *linear incentive policy class*.
- Any policy in $\mathcal{C}_{LIP}$ has a constant generosity matrix $\boldsymbol{G} = [G_{ij}]$.
- Any incentive policy of $\mathcal{C}_{LIP}$, we have

$$\boldsymbol{G} = \begin{bmatrix} 1 & 1 & 1 \\ p_c & p_r & p_d \\ 0 & 0 & 0 \end{bmatrix}$$

- This gives us a larger design space for incentive protocol.

Dynamics and Robustness of Image Policy $\mathcal{P}_{image}$

- Consider the performance gap of different strategies:

$$
\begin{aligned}
\mathcal{G}_3(t) - \mathcal{G}_1(t) &= 1 - \alpha x_2(t), \\
\mathcal{G}_3(t) - \mathcal{G}_2(t) &= \left[x_1(t)(1 - \alpha x_2(t))\right]\left[1 - x_2(t)\right]^{-1}, \\
\mathcal{G}_2(t) - \mathcal{G}_1(t) &= \left[(1 - \alpha x_2(t))(1 - x_1(t) - x_2(t))\right]\left[1 - x_2(t)\right]^{-1}.
\end{aligned}
$$

- **Case A:** when $x_2(t) > 1/\alpha$, $\mathcal{G}_1(t) > \mathcal{G}_2(t) > \mathcal{G}_3(t)$. Defectors and reciprocative peers will continue to adapt to cooperative strategy until $x_2(t) = 1/\alpha$ which is case B.
- **Case B:** when $x_2(t) = 1/\alpha$, it is an unstable equilibrium. Either go to A or go to C.
- **Case C:** when $x_2(t) < 1/\alpha$, $\mathcal{G}_3(t) > \mathcal{G}_2(t) > \mathcal{G}_1(t)$, cooperators and reciprocative peers switch to defective strategy. System collapses.

Dynamics and Robustness of Image Policy $\mathcal{P}_{image}$

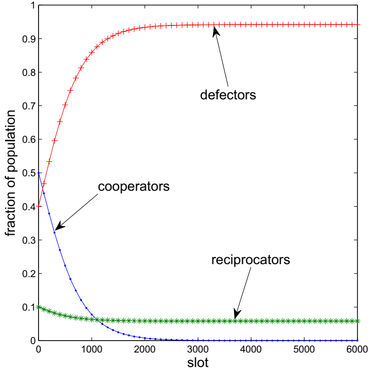

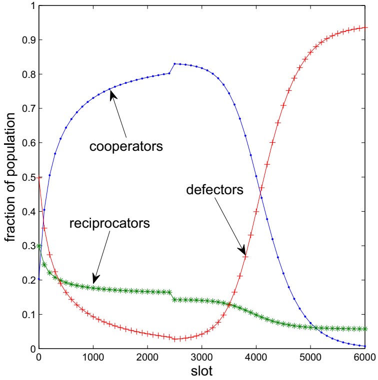

Dynamics and Robustness of Proportional Policy $\mathcal{P}_{prop}$

- Consider the performance gap of different strategies:

$$\begin{aligned} \mathcal{G}_3(t) - \mathcal{G}_2(t) &= x_1(t) - (\alpha - 1)x_2(t), \\ \mathcal{G}_2(t) - \mathcal{G}_1(t) &= 1 - x_1(t) - x_2(t) \geq 0, \\ \mathcal{G}_3(t) - \mathcal{G}_1(t) &= 1 - \alpha x_2(t). \end{aligned}$$

- **Case A:** when $x_2(t) > \frac{1}{\alpha-1}x_1(t)$, $\mathcal{G}_2(t) > \mathcal{G}_3(t)$, so the fraction of reciprocative peers $x_2(t)$ will keep increasing until they dominate the P2P system.
- **Case B:** when $x_2(t) = \frac{1}{\alpha-1}x_1(t)$, $\mathcal{G}_3(t) = \mathcal{G}_2(t) > \mathcal{G}_1(t)$, so cooperators peers adapt to s_2 and s_3. The system go to case A.
- **Case C:** when $x_2(t) < \frac{1}{\alpha-1}x_1(t)$, defectors win. Since s_2 has a higher performance than s_1, $x_1(t)$ will decrease at a faster rate than $x_2(t)$, and the system will go to case B.

Dynamics and Robustness of Proportional Policy

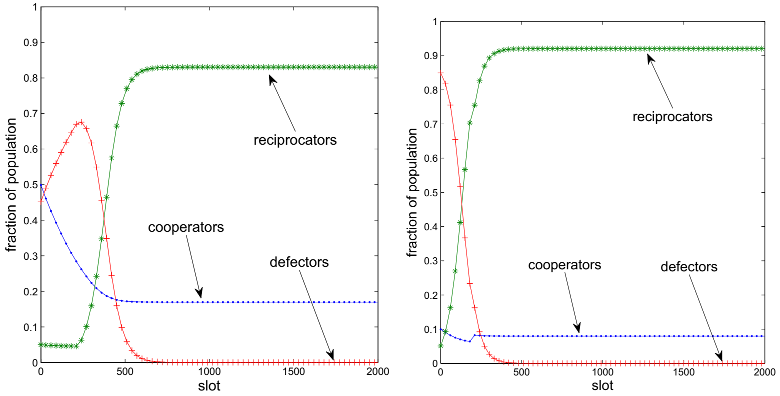

Dynamics and Robustness of $\mathcal{C}_{LIP}$

- Consider the performance gap of different strategies:

$$\begin{aligned} \mathcal{G}_1(t) &= \alpha(x_1(t) + p_c x_2(t)) - 1, \\ \mathcal{G}_2(t) &= \alpha(x_1(t)+p_r x_2(t))-(p_c x_1(t)+p_r x_2(t)+p_d x_3(t)), \\ \mathcal{G}_3(t) &= \alpha(x_1(t) + p_d x_2(t)) \end{aligned}$$

- The *sufficient condition* for robustness is:

$$p_d = 0; \quad p_r \geq p_c. \tag{8}$$

- When p_c is small, the system is more likely to be robust.
- Blind altruism of cooperator helps defectors to survive thus damages the system.

Dynamics and Robustness of $\mathcal{C}_{LIP}$

- Now we restrict our attention to linear strategies with $p_r, p_c > p_d > 0$.
- The robustness of these policies depends on the initial population, and this is especially true for the reciprocators.
- Let $c_{upper} = \frac{p_c}{(\alpha-1)(p_r-p_d)+p_c-p_d}$ and $c_{lower} = \frac{p_d}{(\alpha-1)(p_r-p_d)}$. It can be shown that for the given learning model,
  - when $x_2(0) > c_{upper}$, the system is robust.
  - when $x_2(0) < c_{lower}$, the system will collapse.
  - other initial conditions, the robustness depends on the learning mechanism and the fraction of other strategies.

Dynamics and Robustness of $\mathcal{C}_{LIP}$

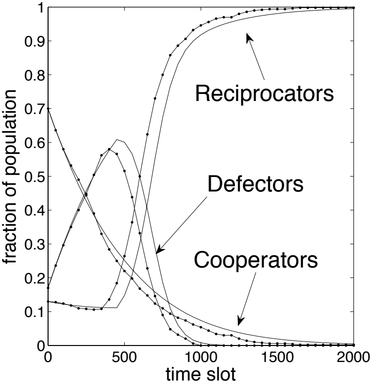

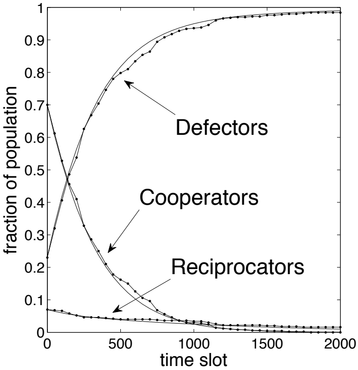

Connection to Evolutionary Game Theory

Theorem

A linear incentive policy can be mapped to a two-player symmetric game, and the Evolutionary stable strategy (ESS) of this game is an asymptotically stable fixed point (ASF).

Conclusion

- We present a *simple* mathematical framework to model the evolution and performance of incentive policies. Peers are assumed to be rational and are able to learn about the behavior of other peers.
- Image incentive policy usually leads to a complete system collapse.
- Proportional incentive policy, which takes into account of service consumption, can lead to a robust system.
- Performance and Dynamics of $\mathcal{C}_{LIP}$
- Connection with evolutionary game theory.
- Framework to design and analyze distributed incentive protocols.

Interesting Questions

- How do we model other *learning algorithms*?
- How about other incentive policies?
- How can we extend this framework to wireless mesh networks?
- How about incentive protocols for ISPs to cooperate?
- Once we know the dynamics and robustness of a given incentive policy, how can we enhance it?