Cumulative Reputation Systems for Peer-to-Peer Content Distribution

B. Mortazavi* and G. Kesidis

Pennsylvania State University
*also with Verizon Wireless

CISS Princeton
March 23, 2006

Outline

- P2P CDNs
- Literature review of incentive mechanisms
- Cumulative Reputation Mechanism
  - Transaction process
  - Convergence of reputations
  - Hierarchical trust groups and security issues
- Model of a cumulative reputation system
  - Non-hierarchical structure
  - Hierarchical structure
  - Misrepresentations
- A simple sequential transaction game
- A game involving up-link bandwidth resources

P2P CDNs

- P2P networks can be used for a variety of applications: Routing, QoS mgmt, VOIP (e.g., Skype), distributed computation (e.g., seti@home), content distribution networks (CDNs), etc.

- P2P CDNs
  - Structured vs unstructured
  - Decentralized, centralized, partially centralized, and hybrid centralized
  - Chord, CAN, Pastry and Tapestry: decentralized, structured
  - Gnutella: decentralized, unstructured
  - Kazaa: partially centralized, unstructured
  - Napster: Hybrid centralized, unstructured

Incentive Systems for P2P CDNs

- Rationally selfish behavior, e.g., free riding, degrades performance of P2P networks
- P2P resources similar to "public good" in economics
- Incentive Mechanisms:
  - Rule based,
  - Monetary based or
  - Reputation based.

Review of Incentive Mechanisms

- Rule based: Bit-torrent
- Monetary based, e.g., using micro-payments
- Cumulative reputation based:
  - Centralized (reputations stored and communicated through a server), e.g., eBay
  - Decentralized (reputation values stored at peers), e.g., EigenTrust, Kazaa
  - Partially decentralized (some peers responsible for holding and advertising the reputation rankings of certain others)
- Incentive mechanisms for rationally selfish peers have been evaluated using game models
- Reputation systems can isolate misbehaving peers
- But attacks are possible on reputation systems themselves

A Cumulative Reputation System: Definitions

- $\pi_j > 0$ is the propensity to cooperate of peer j.
- R_{ij} is the reputation of j from i's point of view.
- All reputations are normalized at each node.
- $G_j(\pi_j, \bar{R}_i)$ is the probability that j responds positively to i's query.
- The response function has the following properties:
  - G is nondecreasing in both arguments,
  - $G(\pi, \bar{R}) = 0$ and $\pi > 0$ imply $\bar{R} = 0$, and
  - $G(\pi, \bar{R}) \leq \pi$ for all $\bar{R} \in [0, 1]$.

A Cumulative Reputation System: Definitions (cont)

- Reputations are modified as a result of successful transactions:

$$R_{ik}(n) = \begin{cases} \frac{R_{ij}(n-1)+C}{1+C}, & k = j \neq i \\ \frac{R_{ik}(n-1)}{1+C}, & k \neq j, i \end{cases}$$

  with probability $G_j(\pi_j, \bar{R}_i(n-1))$.

- otherwise reputations are left unmodified if the transaction is unsuccessful.

- Note the trade-off in the choice of fixed $C > 0$.

Transaction Process

- Query resolution is not covered in this work.
- Successive transactions assumed independent.
- ρ_{ij} is the probability that peer j is "on" the system and that peer i is querying j.
- Peer arrivals and departures and the effect on query resolution are not considered.
- Ignoring a query and not being present in the P2P system are modeled in combination.

Convergence of Reputations

- Theorem: if $\varepsilon\pi \leq G(\pi, \bar{R}) \leq \pi$ and $G(\pi, \bar{R}) = \pi g(\bar{R})$ then for completely disseminated and honest polling,

$$\lim_{n\to\infty} \mathsf{E} R_{ij}(n) \;=\; \frac{\rho_{ij}\pi_j}{\sum_{k,k\neq i} \;\rho_{ik}\pi_k} \quad \text{for all } i \neq j.$$

- So, if all types of queries are equally likely, then the reputation of each peer i is a consistent estimator of its propensity to cooperate π_j.

Sketch of Proof of Convergence

- Define
$$X_{ij}(n) \equiv \frac{\rho_{ij}\pi_j}{\sum_{k,k\neq i} \rho_{ik}\pi_k} - R_{ij}(n)$$
for all $i \neq j$ and $n \geq 0$.

- Condition on $\mathbf{R}(n-1)$ to show
$$\mathsf{E}|X_{ij}(n)| \leq (1-\alpha)\mathsf{E}|X_{ij}(n-1)|.$$
for a scalar $0 < \alpha < 1$.

- Thus, $\mathsf{E}|X_{ij}(n)| \to 0$ exponentially quickly.

Reputation Systems: Vulnerabilities and Trust Groups

- Potential attacks on cumulative distributed reputation systems include but, are not limited to, badmouthing and ballot box stuffing which are variations of Byzantine attacks (false reputation referrals and associated collusions).

- Similar to [Buchegger et al. '04, Kamva et al. '03, Marti et al. '04, Yu et al. '04], our model can be extended to account for misrepresentation, weighted sampling, and sub-sampling of reputations.

- As a special case, we can model federations that are used by peers for feasible and reliable reputation polling (in the presence of both lying and spoofing of reputation referrals); light weight authentication could be used intra-group.

Sybil Attacks

- Sybil attacks [Castro et al. '02] occur when one end-user employs many different identities.

- A typical solution to Sybil attacks involves a centralized registration server that authenticates a unique identifier upon registration for each peer.

- Also,
  - if authenticated reputations based only on positive feedback and reputation referrals only among trusted peers,
  - then multiple identities will dilute the reputation of the end-user,
  - thereby providing a natural disincentive for Sybil attacks.

Simulations: Reputation System Models

- Response function $G(\pi, \bar{R}) = \pi * \text{min}(1, (N/2)\bar{R})$.

- A first-order autoregressive estimator, with forgetting factor $0 < \beta < 1$, is adjusted on a transaction-by-transaction basis:

$$\tilde{\mathbf{R}}(n) \;\; = \;\; \beta\tilde{\mathbf{R}}(n-1) + (1-\beta)\mathbf{R}(n)$$

- Note the trade-off in the choice of β is similar to that of C.

- We simulated a non-hierarchical structure where all nodes were polled by all for reputation referrals.

- We also simulated a hierarchical structure wherein:
  - nodes were arranged in trust groups;
  - intra-group transactions involved only intra-group polling for reputation referrals; and
  - inter-group transactions involved the group reputations instead of the individual ones.

Simulations: Non-Hierarchical Structure

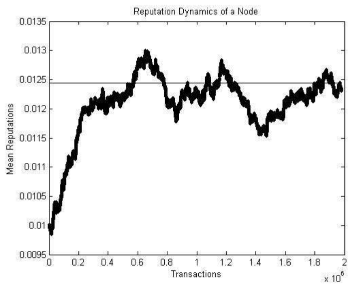

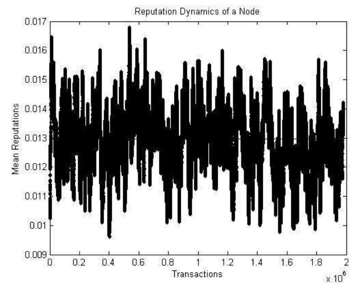

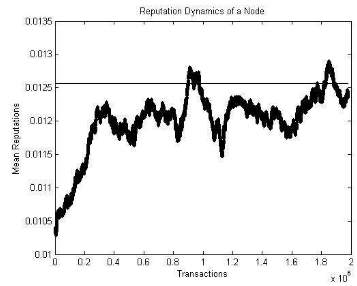

- $N = 100$ nodes.
- π_j/Π_{-i} is the horizontal line in the figures.
- First two figures depict sample path of $\bar{R}_j(\cdot)$, the mean reputation of a specific node j, with forgetting factors $\beta = 0.95$ and 0.15.
- Last figure depicts sample path of R_{ji}, the reputations of a specific node j from the point of view of a node i without referrals, with $\beta = 0.95$.
- Individual increases (directly related to C chosen to be $3/N$) indicate successful transactions for which j was a provider.

Non-Hierarchical Structure (cont)

- Reductions in sample paths occur upon successful transactions for which node j was not involved, and the general effect of forgetting factor β.

- As expected, reputations converge to the nodes' propensity to cooperate.

Non-Hierarchical Weighted Voting

- Alternatively, the reputations received from other notes (during voting) were weighted using the reputations of the voters themselves.

- I.e., aggregated reputations can be defined as:

$$\bar{R}_{ji}(n) = \frac{\sum_{k,k \neq i} \; h(R_{jk}(n))R_{ki}(n)}{\sum_{k,k \neq i} \; h(R_{jk}(n))}$$

- For all nodes, we took $h(R) \equiv R$ in one set of trials and $h(R) = \mathbf{1}\{R > \theta\}$ with $\theta = 0.01$ in another.

- For both cases, the reputation dynamics were observed to be similar to those when no weighting was used.

- This was expected since all transactions were equally likely and the π_j values were selected independently using the same distribution.

Simulations: Hierarchical Structure

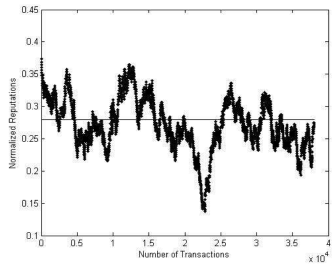

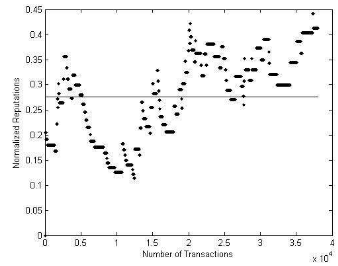

- We simulated 20 nodes partitioned into 5 trust-groups of 4.
- Inter-group transactions involved normalized individual reputations, and intra-group transactions involved normalized group reputations, maintained at the group supernodes.
- Figure 1 is the sample path of one group's reputation from the point of view of another.
- Figure 2 depicts an individual node's mean reputation (within a group).

Hierarchical Structure (cont)

- Intra-group transactions were more frequent than inter-groups, hence group reputation sample path appears smoother and has a shorter transient phase than the individual reputation sample path.

- Decreases in the sample path occur less frequently than in non-hierarchical system counterpart because inter-group transactions had no effect (i.e., a lower transaction rate).

- Reputations fluctuate about their expected mean cooperation value as in the non-hierarchical experiments.

Simulations: Considering Misrepresentations

- Parameter λ is introduced into aggregated reputations as: $\bar{R}_{ji}(n) = \sum_{k,k\neq i} \lambda_{jki} h(R_{jk}(n)) R_{ki}(n) / \sum_{k,k\neq i} h(R_{jk}(n$

- Assumptions:
  - 25% of the nodes lied unfavorably about all nodes to all nodes ($\lambda = 0.75$)
  - 10% lied in favor ($\lambda = 1.25$)
  - 65% were honest ($\lambda = 1$)

- Expect that the limiting sample paths will slightly deviate from the propensity to cooperate because:

  $$(0.25 \cdot 0.75 + 0.10 \cdot 1.25 + 0.65 \cdot 1)\frac{\pi_j}{\Pi_{-i}} = 0.9625\frac{\pi_j}{\Pi_{-i}}.$$

- Indeed, we observed nodes received about 4% fewer successful transactions compared to completely honest reporting.

- Note that when $h(R) = R$, the highly reputable nodes can lie more effectively.

A Simple Sequential Transaction Game

- Nodes modify their own cooperation level π as they file-swap.
- In each round of the game, peer nodes request files from each other.
- Again, whether the requestee of a transaction grants or denies a request is based on the requestee's cooperation level and the requester's reputation ranking in the system.
- For a successful transaction, the requester increases the requestee's reputation level as described above.
- Peers evaluate their success rate S (number of successfully received files versus the number requested) for every round.
- If $S > \sigma$ (the peer's "satisfaction" level), π may be reduced to conserve resources (uplink bandwidth in particular [Feldman et al. '04],[Ma et al. '04]).
- Else, π is increased in order to obtain higher satisfaction (as a result of subsequently higher reputation).

Game Assumptions and Dynamics

- Satisfaction levels were chosen in $\{15\%, 50\%, 75\%\}$

- Cooperation levels were initialized and randomly chosen from the interval $[0.5, 1]$.

- Satisfactory success rates at the end of each round resulted in a small decrease (to conserve resources) in propensity to cooperate and in a small increase otherwise ($\varepsilon = 0.01$).

- A user with desired success rate of 50% maximizes their cooperation level ($\pi \rightarrow 1$).

- The cooperation parameters for easily satisfied (15%) users oscillated about a mean of approximately 0.3.

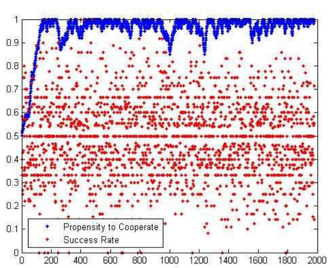

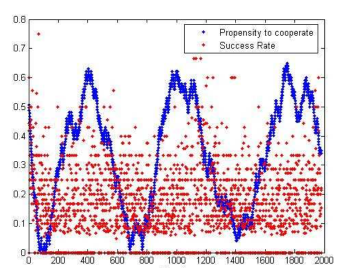

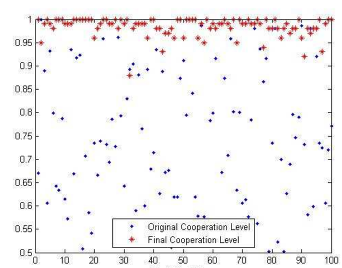

A Game Involving Uplink Resources

- As before, R_{ij} is the reputation of j from the point of view of i

- and response function G is a function of the reputation of the requester and the propensity to cooperate of the requestee.

- Upon each successful transaction, the requester increases requestee's reputation level as described above

- At the end of each round, peers evaluate their success rate and adjust their uplink rate accordingly to maximize their net utility.

Uplink Game Assumptions

- Each peer i has a fixed aggregate upload capacity U_i.
- Peer i's selected upload capacity for the P2P network $u_i < U_i$ changes dynamically.
- Peers play a game in hopes of maximizing a net utility $V(\xi) - au$ where ξ is an estimation of the user's success (quasi-stationarity assumption).
- Requested rates r are scaled according to uplinks u to cause congestion.

Uplink Game Dynamics

- Step 1: Peer i queries peer j with probability ρ_{ij}:
  - Requester i is informed of the current uplink rate u_i of requestee j.
  - The file size requested by i is r_i.
  - File size is implicitly assumed to be divided by a common unit of discrete time and, so, is interpreted as a rate.

Uplink Game Dynamics (cont)

- Step 2: Each peer j is in receipt of a set M_j of queries:
  - Note M_j may be empty.
  - Each requester $i \in M_j$ *deserves* an allocation:

$$\delta_i \equiv \frac{R_{ji}}{\sum_{k \in M_j} R_{jk}} u_j$$

  - The *excess demand* is

$$\varepsilon \equiv \left[\sum_{k \in M_j} r_k - u_j \right]^+$$

  - The penalty for excess demand is distributed among those nodes who have requested transfer rates, r_i, more than what they deserve: $r_i > \delta_i$, i.e., $i \in M_j^* \subset M_j$.
  - The allocation for such nodes i is:

$$x_{ji} = r_i - \frac{g(R_{ji})}{\sum_{k \in M_j^*} g(R_{jk})} \varepsilon,$$

    where g is positive non-increasing,
  - otherwise $x_{ji} = r_i$.

Uplink Game Dynamics (cont)

- Step 3: Each requester i adjusts their reputation of requestee j by adding $R_{ij} + cx_{ij}$ for some constant c and then normalizing all reputations stored at i.
- Step 4: At this point, a requestee can assess the value obtained from the CDN, from the result x of the current transaction (or by accumulation of past and present transactions) via a utility function: $V(x/r)$ or $V(x)$.
  - The utility function is assumed to be non-decreasing but concave (law of diminishing returns) with $V(0) = 0$.
  - A game is formulated here where peers iteratively adjusts their uplink rate, u, to maximize their unimodal net utility assessed over an interval.
  - Net utility is $V(X_i(n)) - au_i(n)$
  - where $X_i(n) = \sum_{k=nT}^{(n+1)T-1} x_i(k)$.
  - Each user i sequentially modifies their constrained uplink rate $u_i \in [1, u_i^{\max}]$ to maximize their net utility using, e.g., an annealing strategy.

Uplink game designed so that...

- A peer does not receive at a rate larger than requested.
- Under overloaded conditions, nodes with:
  - higher reputation and lower requested rate receive at the rates requested and
  - low reputation and high requested rate receive at a rate much lower than requested.
- Choose g so that peer with no reputation (initialization) will be granted a positive uplink capacity when $R = 0$.
- If a user reduces their uplink rate u, their reputation R will eventually be reduced and their allocation x will be reduced during periods of excessive demand.
- Dynamics are continuous so can apply Brouwer's theorem to obtain existence of a fixed point.
- Basic assumption that excess demand ε is frequently > 0.