

Cheap Talk and Selection in Dynamic Networks: A Performance-base Mechanism

Abstract

The performance of large-scale distributed systems, such as peer-to-peer networks, critically depends on the efficient collaboration of multiple peers in dynamic networks. The concept of server selection has been proposed to collect information about peers' performance and select the best node. It is costly to generate information from network monitoring or computational markets; but the information voluntarily reported by each node, called *cheap talk*, is generally available. However, since self-interested peers may misreport their local information, selection directly based on cheap talk will lead to an inefficient allocation. In this paper we propose a performance-based mechanism to control and reduce information misrepresentation in cheap talk. In repeated games of incomplete information, we illustrate how to construct punishment functions based on the inconsistency between reported information and realization, and therefore extract reasonably accurate information from distributed peers. Compared to monitoring and market-based pricing, this performance-based mechanism provides a computationally effective, inexpensive and scalable approach to utilize both self-reported information and historical transactions.

Keywords: truthful information revelation, cheap talk, mechanism design, repeated games

1 Introduction

The wide deployment of global communication networks, especially the Internet, has led to the creation and subsequent usage of many distributed electronic-commerce applications such as the World Wide Web and peer-to-peer (P2P) knowledge sharing systems. Unlike a centralized computing architecture, the effectiveness of distributed applications critically depends on *server selection*, which chooses the best resource/peer from multiple alternatives in networks according to their locations, quality, and performance. For example, in grid computing, the ideal scenario is to assign computational tasks with fixed awards to hosts which have the highest performance and the least backlogs; in selfish routing, the efficient strategy is to forward data packets at flat unit-prices to nodes which have the highest throughput rate and the lowest network latency.

There are three important characteristics of server selection in distributed scenarios: private information, imperfect monitoring, and repeated transactions. Some information on local resources, such as server workloads, service costs, and server performance, is private and can not be directly observed by remote parties. In large-scale distributed environments, many measurements of networks and resources, such as computational capacity, bandwidth, and delay, are often constantly changing because of spontaneous usage, network traffic, and exogenous noise. Therefore, acquiring complete information on all transactions is costly. However, historical transactions do provide information about the performance of network nodes. How to retrieve information from repeated transactions becomes a major issue in Dynamic Mechanism Design (DMD) [Feigenbaum and Shenker 2002].

There are two approaches to collecting information for server selection: network monitoring and market mechanisms. Through synthetic or real transactions, network-aware applications monitor network performance and adjust their resource demands in response to performance variations [Bolliger and Gross 1998]. On the other hand, to avoid the overhead of monitoring, computational markets and auctions are proposed to reveal true information about peers in

distributed computing [Nisan and Ronen 2001, Feigenbaum et al. 2002]. Most of these solutions belong to the Vickrey-Clarke-Groves (VCG) class of mechanisms [Vickrey 1961, Clarke 1971, Groves 1973], which use information rents to stimulate agents to report their true types.

Both approaches have their shortcomings. Network monitoring requires full control of all entities or at least all obedient entities. Most distributed applications consist of multiple entities with conflicting interests. They will not cooperate in reporting their true information without compensation. Since the members of an overlay network are dynamic, getting complete measurements is costly. For VCG mechanisms, a central authority is needed to gather global information and calculate payments. In most cases the calculation is Non-deterministic Polynomial-time hard. When the number of members in the network is large, gathering complete information on all transactions is difficult. Even for polynomial approximations, the need for global information will overload computational capabilities. Moreover, the VCG class of mechanisms is a one-shot game which applies to independent stages like multiple rounds of auctions. Due to the dynamic nature of distributed applications, VCGs are unable to ensure truth telling in repeated transactions.

In this paper, we model information exchange in distributed networks as *cheap talk*, which sends messages without *direct* payoffs and affects outcomes through people's responses [Farrell and Rabin 1996]. For example, for network applications, the responses of performance probing or querying are cheap talk because the contingent contracts binding reported values with payments are hard to implement. Although message senders may lie according to their preferences, with proper incentive mechanisms, cheap talk can be informative and efficient in server selection.

We propose a model based on cheap talk in which a regulator selects one of M servers according to their self-reported performance. In each period, the regulator assigns one transaction to one of the servers based on their self-reported information about local performance and pays a fixed award. By allocating transactions and payments to the servers with the best performance, the regulator maximizes long-term social welfare. If one server cheats by misreporting information, the regulator will punish it by reducing its future transactions.

We illustrate how to construct punishment functions based on the inconsistency between reported information and historical transactions; therefore promoting truth telling and extracting reasonably accurate information on these servers. Compared to monitoring and market-based pricing, this performance-based mechanism provides a computationally effective, inexpensive and scalable approach to utilize both self-reported information and historical transactions.

The paper will be organized as follows: In Section 2, we model the selection problem as an infinitely repeated game of incomplete information. Based on the pre-play communication – cheap talk – and the realized payoff in each period, we construct a performance-based mechanism with punishment in finite periods to the truth telling outcome; then in Section 3, we evaluate the effectiveness of this mechanism through simulations; finally, in the conclusion, we explain the intuition of this performance-based mechanism and describe future extensions.

2 The Model

In an infinitely repeated game of incomplete information, a regulator, which could be a user, a collection of users, or a moderator/social planner, performs a sequence of tasks on M servers by assigning one task to one of the servers in each period with a fixed payment. Each server observes its intertemporal performance, which is unknown to the regulator, and the server sends

a message to the regulator through pre-play cheap talk. The regulator will choose his strategies according to servers' self-reporting and their credibility.

When a server deviates by misreporting its performance, the regulator will punish it by stopping or reducing further transactions. Because servers do not clearly perceive a well-defined final period, i.e. the length of tasks is unknown, after each period they believe that the game will continue for an additional period. In this infinite horizon, the misinformed regulator can implement many punishment schemes to prevent cheating; however, some punishments may be inefficient and non-creditable.

In this section, we characterize regulator's probabilistic selection strategies, which extract information concealed in cheap talk and enforce the truth telling equilibrium.

2.1 Model Setting

The timing of the stage game in period t is a cheap-talk game:

1. Nature draws a type q_i^t for server i from a set of feasible types $Q = \{0,1\}$ according to a probability distribution $p(q_i)$, where $p(0) = 1 - p_0$ and $p(1) = p_0$.
2. Server i observes q_i^t and then chooses a message m_i^t from a set of feasible messages $M = \{0,1\}$.
3. The regulator observes m_i^t (but not q_i^t) for $i = 1 \dots M$ and then chooses an action a^t , where $a^t \in \Delta^{M-1} = \{(a_1, a_2, \dots, a_M) \in \mathbb{R}^M \mid \sum_i a_i = 1 \text{ and } a_i \geq 0 \text{ for all } i\}$ and a_i is the probability to assign a task to server i .
4. Payoffs are given by $u_i(q^t, a^t)$ and $u_R(\tilde{q}^t, a^t)$, where $q^t = (q_1^t, \dots, q_M^t)$ and \tilde{q}^t is the realization of q^t .

In period t , the regulator assigns a task to server i with a fixed payment c . Therefore, the expected utility of server i is $E(u_i(q^t, a^t)) = ca_i^t$; the expected utility of the regulator is $E(u_R(q^t, a^t)) = \tilde{q}^t(a^t)^T$. Both the regulator and servers have *overtaking* preferences, i.e. their utilities in T periods are $U_i^T = \sum_{t=1}^T u_i(q^t, a^t)$ and $U_R^T = \sum_{t=1}^T u_R(\tilde{q}^t, a^t)$. Servers' payoffs do not depend on their types and outcomes; therefore, they have incentive to cheat and earn more transactions. Regulator's utility depends on the realized quality of servers and maximizing regulator's utility is equivalent to maximizing social welfare.

Servers' strategy is the probability to cheat, i.e. $p_i^c = P(m = 1 \mid q = 0)$; Regulator's strategy is the selection action $a^t(m^0, \dots, m^t, \tilde{q}^0, \dots, \tilde{q}^t)$, which is a function of all messages and historical transactions.

2.2 Basic Model – M Symmetric Servers with Accurate Observation

In this symmetric model, all servers have the same p_0 , which is public information. Observed types will not change within each period, i.e. $\tilde{q}^t = q^t$. Without loss of generality, we assume that $q = 1$ represents high quality or good performance of transactions, $m = 1$ denotes a message of $q = 1$, and $c = 1$.

In each period, with complete information or truth telling, the regulator can maximize social welfare and get the first best payoff $E(u_R) = 1 - (1 - p_0)^M$.

Although information is incomplete in this game and servers have incentive to cheat, it is possible to reach the first best with certain punishment. One example is extreme punishment: When detect cheating, the regulator will exclude the server for future transactions forever. Extreme punishment is a mechanism to facilitate truth telling because, with overtaking utility, the gain in a single period is always less than the expected payoff of infinite periods in the future; however two problems make the punishment inefficient:

1. The extreme punishment is not robust and cannot recover from off-equilibrium paths. When an error occurs in observation, extreme punishment will exclude a server permanently and result in a loss of social welfare.
2. The extreme punishment is not credible because it is not re-negotiate proof. The misinformed regulator would improve the outcome by putting positive selection probability on the cheating server.

When the regulator can verify types through transactions, his optimal strategy has to generate finite periods of punishments to implement incentive constraints and to minimize the loss of efficiency. A random selection strategy with T-period blackout can be constructed as following: the regulator randomly choose a server from $H^t = \{i \mid m_i^t = 1, i = 1, \dots, M\}$ if $H^t \neq \emptyset$, otherwise, the regulator choose one of M servers randomly; when server i cheats, e.g. $\tilde{q}_i^t \neq q_i^t$, the regulator will ignore server i 's messages in coming T periods by setting $m_i^{t+1}, \dots, m_i^{t+T}$ to zero.

Proposition 1. In the game of accurate observation, where the regulator can detect cheating, excluding a cheater for $T^m = \left\lceil \frac{1-p_0}{p_0} \right\rceil$ periods is the minimal punishment to enforce the truth telling equilibrium, which is a sub-game perfect Nash equilibrium.

The random selection strategy with a T^m -period blackout suggests that a small punishment would be sufficient if the average performance of servers is high. For example, when $p_0 > 0.5$, one-period punishment is more then enough; when p_0 is small, more periods are needed.

Although a T^m -period blackout generates adequate punishment, it is not an optimal strategy when servers have errors in observing their types. A small chance of inconsistency between \tilde{q}^t and q^t may induce punishment, which eliminates the messages in blackout periods and reduces regulator's expected payoff.

2.3 Extended Models – M Symmetric Servers with Imperfect Observation

To extend the basic model, we consider the case that errors would occur in observation: even if a server observes high quality $q = 1$, the realized outcome, when the regulator performs a task on this server, may be low. We assume that, for any i and t , $P(\tilde{q}_i^t = 0 \mid q_i^t = 1) = \varepsilon$. This can be interpreted as a system error, imperfect observation, dynamic nature of quality, or the time lag between observation and realization.

Because of idiosyncratic errors, the regulator cannot detect cheating explicitly. Therefore, excluding a server completely based on the inconsistency in one transaction is not an optimal strategy. In this extension, we use random selection based on credibility to separate cheating from errors in observation.

At the beginning of this game, the regulator has a prior belief $\beta^0 = (\beta_1^0, \dots, \beta_M^0) = (1, \dots, 1)$ on the credibility of all servers. If in period t , server i 's realized type is not consistent with its message, the regulator will downgrade this server's credibility and upgrade the credibility of others, i.e. $\beta^{t+1} = (\beta_1^t + \delta/(M-1), \dots, \beta_i^t - \delta, \dots, \beta_M^t + \delta/(M-1))$. In each period, the belief of credibility β is used as weights for probabilistic server selection. For example, in period t , if only servers 1, 2, and 3 report $m_1^t = m_2^t = m_3^t = 1$, regulator's strategy is

$$a^t = \left(\frac{\beta_1^t}{\beta_1^t + \beta_2^t + \beta_3^t}, \frac{\beta_2^t}{\beta_1^t + \beta_2^t + \beta_3^t}, \frac{\beta_3^t}{\beta_1^t + \beta_2^t + \beta_3^t}, 0, \dots, 0 \right).$$

In the process of updating beliefs, δ carries out a small punishment and spreads positive selection weights over the rest of server. If the deduction is caused by an idiosyncratic error, it will be compensated when other servers make mistakes. However, cheating can only reduce credibility. Therefore, random selection based on credibility can punish cheating concealed in errors of observation. The following simulation shows this strategy enforces the truth telling equilibrium.

3 Simulation

Using EZSIM, we design a discrete system simulation, in which ten servers are connected to the regulator through a network. Each server has a fixed time to process one task and the length of queues indicates the performance of servers. We denote zero tasks in server's queue as high performance, which has the same probability 0.6 for all servers. When intends to assign a task, the regulator routinely checks performance reports from servers and chooses one using random selection based on credibility. The network delay between the regulator and servers implements errors in observation. If the time to finish this task is longer than the fixed processing time, e.g. waiting in the queue, the regulator will adjust the credibility of related servers. We are interest in finding the optimal strategy of servers to report their performance.

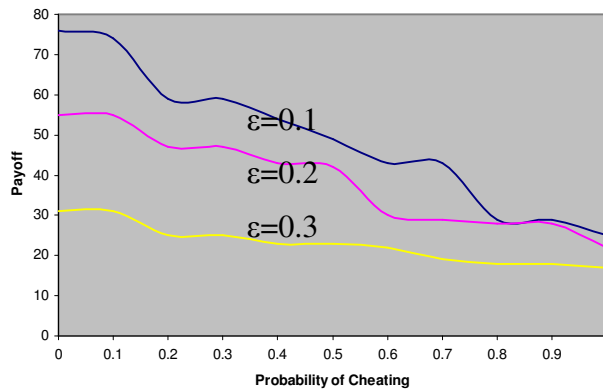


Figure 1. Server's strategies and payoffs in the simulation

Figure 1 shows that servers' payoffs decrease when they choose higher probability to cheat and zero probability of cheating is the dominant strategy – the truth telling equilibrium. Even with large network delay, i.e. bigger errors in observation, the random selection based on credibility can still prevent small cheating attempts.

There are several aspects of these simulation results that are worth noting. First, cheap talk with reputation incentives utilizes both historical and self reported information. Second, when the reputation (credibility) of a server is above a certain point, the cheap talk about the types of servers is credible. Finally, this performance-based mechanism improves the efficiency without the expenses of monitoring or pricing.

4 Conclusion and Future Work

This paper discusses the use of a performance-based mechanism to foster truthful information revelation from cheap talk in distributed applications. The basic analytical model we proposed illustrates how punishments motivate peers to report their information without commitment. The numerical simulation shows that the performance-based mechanism with cheap talk enforces the truth telling equilibrium and improves the efficiency of server selection in distributed collaboration.

Besides facilitating decentralized server selection, performance-based mechanisms can be extended to many eCommerce applications. For long-term interactions in a decentralized network with self-reported and non-committed message exchange, the performance-based mechanism acts as a reputation system to regulate the information exchange overtime. By revealing information and taking account of past behavior, the regulator constructs beliefs on nodes' credibility and uses punishment to exclude the ones taking advantage of society. This paper provides a new starting point for dynamic mechanism design on the Internet.

5 References

- Bolliger, J. and T. Gross. A Framework-Based Approach to the Development of Network-Aware Applications. *IEEE Transactions on Software Engineering* 24 (1998).
- Clarke, E.. Multipart pricing of public goods. *Public Choice* 11 (1971), 17-33.
- Farrell, J. and M. Rabin. Cheap Talk. *The Journal of Economic Perspectives* 10:3 (1996) 103-118.
- Feigenbaum, J., C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based Mechanism for Lowest-Cost Routing, *Proc. of ACM PODC* (2002).
- Feigenbaum, J. and S. Shenker. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. in *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications* (2002), ACM Press, New York, 1-13.
- Groves, T. Incentives in teams. *Econometrica* 41 (1973), 617-663.
- Nisan, N. and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior* 35 (2001), 166-196.
- Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1961), 8-37.