A Dynamic Pricing Mechanism for P2P Referral Systems

Bin Yu, Cuihong Li
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

Munindar P. Singh
Department of Computer Science
North Carolina State University
Raleigh, NC 27695, USA

Katia Sycara
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

Abstract

Most existing research on peer-to-peer systems focuses on protocol design. In this paper, we consider the issue of free riding in peer-to-peer referral systems. Free riders are agents that refuse either to answer a query or to give referrals. Free riding is detrimental to the system, since it may prevent requesters from finding high quality providers efficiently. To mitigate the issue of free riding, we propose a dynamic pricing mechanism to motivate the agents to behave rationally. Service providers learn appropriate prices of referrals and answers in order to maximize their payoffs through stochastic iterative learning algorithms.

1. Introduction

This paper presents a P2P referral system, in which peers cooperate to search the relevant information in the system through referrals. Referral systems have been studied by [4, 5, 7], but the problem of *free riding* remains to be addressed [1]. Peers often cannot find suitable services since many peers choose to decline requests from others. A natural approach to control free riding is to introduce a micropayment protocol into referral systems, in which each peer has to pay for the services it receives from others, e.g., [3]. The prices in a micropayment protocol could be static (fixed) or dynamic. Fixed pricings in micropayment protocols can reduce free riding in P2P systems, but they don't consider quality of services provided by different peers and the supply-demand relationship in the system. In this paper we propose a dynamic pricing mechanism to motivate each peer to behave rationally in peer-to-peer search processes [6]. The requesting peer can decide which referral or answer it wants to pay in the micropayment protocol, but the prices from the same peer are dynamic for different search processes. The prices of referrals and answers are chosen from a discrete space of prices and determined by qualities of services and the supply-demand relationship in the system [2].

AAMAS'04, July 19-23, 2004, New York, New York, USA.

2. Dynamic Pricing Mechanism

We consider two research challenges in the design of the dynamic pricing mechanism: (1) how to choose the intermediate peers P_k so the requesting peer knows the service P_k provides deserve the price it declares; (2) what are the best pricing strategies for the intermediate peers and service providers to maximize their profits.

2.1. Micropayment Protocol

Suppose P_r is the requesting peer, and $\{P_1, P_2, \ldots, P_n\}$ are a set of acquaintances of P_r. P_r has two *selling prices* in its profile: Y'_r for a referral and X'_r for an answer. X'_r and Y'_r are chosen from a discrete space of prices, respectively. For any peer P_i, $1 \leq i \leq n$, P_r has two *reserve prices* he is willing to pay: Y_i for a referral and X_i for an answer. Given a peer P_r and any of its acquaintances P_i, the reserve prices for P_i are initialized as: $X_i = X_0$, $Y_i = Y_0$, where X_0 and Y_0 are constants in the system.

The reserve prices are used to estimate if the selling prices from other peers are reasonable. For example, at time t, peer P_r receives a set of prices for answers from sellers $\{P_l, P_{l+1}, \ldots, P_m\}$. For any seller P_j, the estimated margin of the reward for choosing peer P_j is $X_j - X'_j$, where X_j is the reserve price of an answer in the acquaintance model for P_j at time t ($X_j = X_0$ if there is no acquaintance model for P_j) and X'_j is the selling price of an answer in the profile of peer P_j. Similarly, for referring services, the requesting peer P_r computes the estimated margin for a referral from P_j as $Y_j - Y'_j$.

Given a set of referring services, the requesting peer will choose referrals from the highest to the lowest estimated margins for a single referral. For a given referral graph G, the requesting peer will choose the referral with the highest estimated margin to expand the graph, where the requesting peer pays all referrals P_j has. After the requesting peer receives an answer, it or its user will evaluate the quality of the answer and revise the reserve price for the service from the answering peer. On the other hand service providers adjust the estimated values for prices of referrals and answers.

2.2. Pricing Strategies

In this section we present a stochastic iterative algorithm as an effective dynamic pricing strategy for a peer. In this algorithm a peer estimates the value of each price, and most of the time chooses the price with the higher value estimation. The algorithm is composed of two parts: the first is *value estimation* and the second is *price selection*.

Value estimation The estimated value of price a for a referral or an answer after t plays is denote by $Q_t(a)$ and r_{k_a} is the reward from choosing price a at the k_a-th time. A peer updates the estimated value of a price, if the price is chosen, based on the current reward and the previous estimated value. The update of the estimated value is based on the following rule:

$$Q_{k_a+1}(a) = Q_{k_a}(a) + \alpha(r_{k_a+1} - Q_{k_a}(a)) \qquad (1)$$

where step size α is a constant, $0 < \alpha \leq 1$. Based on this rule the recent rewards are weighted more heavily than long past ones. This is necessary in a non-stationary environment in which the mean reward of a price changes over time. The recent environment, e.g., acquaintances of each peer, is more similar to the environment today and therefore gives more information, than the environment of long past. With a constant step size α, the estimates never completely converge but continue to vary in response to the most recently received rewards [2].

Price selection Although a peer can exploit the estimated values of prices for price selection, exploring other prices is necessary to gain information on the prices that have been rarely chosen in the past. We suggest the ϵ-*greedy* method to choose an appropriate price. In this method a peer behaves greedily, i.e., chooses the price with the highest estimated value, most of the time, but every once a while, say with probability ϵ, select a price at random. A peer will not select among the other prices equally. The selection among other prices is directed by the uncertainty level of the value estimation to encourage exploration. One way to estimate the uncertainty of the price is to compare the number of times that the price has been chosen and the number of time that the price offered rewards. Suppose n_a is the number of price a being chosen and n'_a is the number of price a being offered reward, then the uncertainty for this price can be estimated as $(n_a - n'_a)/n_a$, where $n'_a \leq n_a$.[1] The potential rewards for price a is $Q(a) * n_a/(n_a - n'_a)$, where $Q(a)$ is the average rewards received at price a.

Another consideration in the design of pricing mechanisms is the initial value for the average reward $Q(a)$ of each price a. The algorithms we discussed so far are dependent on the initial price-value estimates, $Q_0(a)$. For example, for any price a, initially $Q_0(a) = 0$. The kind of design is problematic when we estimate the uncertainty of each price during exploration. In such cases, some prices could be never explored since the potential reward is zero. One solution is to use an optimistic estimate for any price a, e.g., $Q_0(a) = Y_0$ for any referring services. This optimistic estimate encourages the exploration of each price at least once.

1 $n_a = 1$ when price a is never chosen or is chosen once.

3. Conclusion

This paper proposes a dynamic pricing mechanism to motivate each peer to behave rationally in referral systems. Our paper only provides a preliminary study of pricing mechanism design in referral systems. For example, we simply assume that the qualities of services are consistent for each peer. Also, we don't consider how to effectively determine the upper bounds of selling prices for referrals and answers during exploration. In future work, we plan to focus on these problems and develop more efficient and incentive compatible mechanisms for referral systems and peer-to-peer systems in general.

Acknowledgements

This research was supported by the National Science Foundation under grant No. ITR-0081742, the AFOSR under grant No. F49620-01-1-0542 and by AFRL/MNK grant No. F08630-03-1-0005.

References

[1] Eytan Adar and Bernardo Huberman. Free riding on Gnutella. *First Monday*, 5(10), 2000.

[2] Dimitri P. Bertsekas and John N. Tsitsiklis. *Neural Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

[3] Philippe Golle, Kevin Leyton-Brown, Ilya Mironov, and Mark Lillibridge. Incentives for sharing in peer-to-peer networks. In *Proceedings of the Second International Workshop on Electronic Commerce*, pages 75–87, 2001.

[4] Michael N. Huhns, Uttam Mukhopadhyay, Larry M. Stephens, and Ronald D. Bonnell. DAI for document retrieval: The MINDS project. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, pages 249–283. Pitman/Morgan Kaufmann, London, 1987.

[5] Henry Kautz, Bart Selman, and Al Milewski. Agent amplified communication. In *Proceedings of the National Conference on Artificial Intelligence*, pages 3–9, 1996.

[6] Jeffrey Shneidman and David Parkes. Rationality and self-interest in peer-to-peer networks. In *Proceedings of Second International Workshop on Peer-to-Peer Systems*, 2003.

[7] Munindar P. Singh, Bin Yu, and Mahadevan Venkatraman. Community-based service location. *Communications of the ACM*, 44(4):49–54, 2001.