EcoBroker: An Economic Incentive-Based Brokerage Model for Efficiently Handling Multiple-Item Queries to Improve Data Availability via Replication in Mobile-P2P Networks

Anirban Mondal[1], Kuldeep Yadav[1], and Sanjay Kumar Madria[2]

[1] Indraprastha Institute of Information Technology
New Delhi, India
{anirban,kuldeep}@iiitd.ac.in
[2] University of Missouri-Rolla
Rolla, USA
madrias@mst.edu

Abstract. This work proposes EcoBroker, a novel economic incentive-based brokerage model for improving data availability via replication for multiple-item queries in Mobile-P2P networks. In Ecobroker, data requestors need to pay the price (in virtual currency) of their requested data items to data-providers. The main contributions of EcoBroker are two-fold. First, its economic incentive model effectively combats free-riding by incentivizing MPs to become brokers and to host replicated data, thereby improving data availability. Second, its brokerage model facilitates efficient processing of queries involving multiple data items. Our performance evaluation indicates that EcoBroker indeed improves data availability and querying-related communication overhead in Mobile-P2P networks.

1 Introduction

In a Mobile Ad-hoc Peer-to-Peer (M-P2P) network, mobile peers (MPs) interact with each other in a peer-to-peer (P2P) fashion. Proliferation of mobile devices (e.g., laptops, PDAs, mobile phones) coupled with the ever-increasing popularity of the P2P paradigm (e.g., Kazaa [14], Gnutella [7]) strongly motivate M-P2P network applications. Mobile devices with support for wireless device-to-device P2P communication are beginning to be deployed such as Microsoft's Zune [12].

M-P2P applications facilitate mobile users in sharing information with each other *on-the-fly* in a P2P manner. Notably, M-P2P users are often interested in issuing queries involving multiple (possibly related) data items. For example, a music fan Jim, who is moving in a shopping mall, could issue a request for obtaining several songs and video-clips of the Beatles music band. In the same vein, John, who has recently been intrigued by the paintings of Salvadore Dali (at a fine arts exhibition), may want to issue a query for multiple paintings of Dali. Jane, who likes nature photography, may want to request multiple photos of sunsets. Incidentally, the issuance of multiple-item queries is also common in static P2P systems such as Kazaa, where peers often issue queries for downloading entire albums of a specific artist. Notably, P2P interactions among mobile users are generally not freely supported by existing wireless communication infrastructures.

S. Kikuchi, S. Sachdeva, and S. Bhalla (Eds.): DNIS 2010, LNCS 5999, pp. 274–283, 2010.
© Springer-Verlag Berlin Heidelberg 2010

Our target applications mainly concern slow-moving objects e.g., cars on busy streets, people moving in a market-place or students in a campus. This work does not consider updates to data items. Items shared in our M-P2P application scenarios mostly involve MP3 songs, video-clips and photos, which are typically not updated in practice.

Data availability in M-P2P networks is typically lower than in fixed networks due to frequent network partitioning arising from user movement and users switching 'on'/'off' their mobile devices. The problem of data availability is further exacerbated by free-riding. Free-riding has been known to be rampant in static P2P environments i.e., a large percentage of the peers do not provide any data/services to the network [13]. The adverse effect of free-riding on data availability becomes even more pronounced in M-P2P environments due to the generally limited resources (e.g., energy, memory space, bandwidth) of MPs. Thus, an economic incentive model for enticing MPs to provide data and services becomes a necessity for improving M-P2P data availability.

Recall that our application scenarios call for efficient handling of multiple-item queries. Given a query involving k items, a naive way would be to issue k separate queries, thereby incurring considerable communication overhead as well as significantly taxing the limited resources of the M-P2P network. On the other hand, issuing a single query for k items would provide opportunities for optimizing communication overhead and preserving limited M-P2P resources such as MP energy and bandwidth. Observe that such optimizations would be possible if certain designated MPs acted as **brokers** and hosted replicas of 'hot' items at themselves. This provides a strong motivation for a brokerage model for improving data availability in M-P2P networks, especially in case of multiple-item queries. Hence, we propose **EcoBroker**, a novel economic incentive-based brokerage model for improving data availability via replication for multiple-item queries in M-P2P networks.

The main contributions of EcoBroker are two-fold:

1. Its economic incentive model effectively combats free-riding by incentivizing MPs to become brokers and to host replicated data, thereby improving data availability.
2. Its brokerage model facilitates efficient processing of queries involving multiple data items.

In the incentive-based model of EcoBroker, each data item has a *price* (in *virtual currency*). Data item price depends on access frequency and data quality [16] (e.g., image resolution, audio quality). A query issuing MP pays the *price* of the queried data item to the query-serving MP, a commission to the successful broker and a relay commission to each MP in the successful query path. Thus, EcoBroker provides an incentive for free-riding MPs to provide data as well as brokerage and relay services so that they can earn currency (i.e., broker commissions) for issuing their own requests. **Revenue** of an MP is defined as the difference between the amount of virtual currency that it earns (by providing data, brokerage and relay services) and the amount that it spends (by requesting data). Notably, virtual currency is suitable for M-P2P environments due to high transaction costs of micro-payments in real currency [22]. Secure virtual currency payments have been discussed in [4].

In the brokerage model of EcoBroker, brokers select and obtain the items to replicate at themselves to facilitate them in earning currency. Notably, brokers replicate items at themselves based on the currency-earning potential of the items. Querying in EcoBroker proceeds via broadcast until the queries are intercepted by brokers. When a broker intercepts a multiple-item query, it can quickly serve the query partially from the replicated items existing at itself. For the remaining queried items, whose replicas it does not contain, it uses its index to locate and obtain those items. In this paper, we do not describe the index used by brokers due to space constraints. Instead, we use an existing index for M-P2P environments, namely the CR*-tree index, which we have proposed in our previous work [17]. As such, our focus in this paper is on performing effective replication at brokers for facilitating multiple-item queries.

Notably, each item has only one original owner and brokers are allowed to replicate for brokerage purposes only by obtaining the items from their respective original owners. Hence, brokers require to provide an incentive to the original owners of those items. In the absence of any incentive, the original owners would not want to replicate their items at the broker MPs because they would lose revenues on future accesses to their own items. Thus, broker MPs make a one-time payment to the original owners for the right to replicate their items. Incidentally, replicating their items at the broker MPs does not preclude the original owners from hosting their own items.

In EcoBroker, brokers could be pre-designated in accordance with the application scenario under consideration, and there could be multiple pre-designated brokers. For example, recall our application scenario concerning an M-P2P user (in a shopping mall) searching for multiple songs of the Beatles music band. In this case, the shop-owners or the shopping mall administrators can act as brokers. If songs or movies are shared among M-P2P users in a University campus setting, some of the students (e.g., student organization leaders) can act as brokers.

Our performance evaluation demonstrates that EcoBroker indeed improves data availability and querying-related communication overhead in Mobile-P2P networks. The remainder of this paper is organized as follows. Section 2 provides an overview of relevant existing works, while Section 3 discusses the economic incentive-based model of EcoBroker. Section 4 presents the brokerage model of EcoBroker, while Section 5 reports our performance evaluation. Finally, we conclude in Section 6.

2 Related Work

Economic models for distributed systems [5,15] primarily focus on resource allocation. These models do not address unique M-P2P issues such as node mobility, mobile resource constraints, frequent network partitioning, free-riding and incentives for peer participation. Economic schemes for resource allocation in wireless ad hoc networks [25] do not consider free-riding. The works in [2,3,21,26] provide incentives to peers for forwarding messages, but they do not incentivize free-riders to host data.

Schemes for combating free-riding in static P2P networks [8,10,13] are too static to be deployed in M-P2P networks as they assume peers' availability and fixed topology. Schemes for improving data availability in mobile ad hoc networks (MANETs) (e.g., the **E-DCG+** approach [11]) primarily focus on replica allocation, but do not consider economic incentive schemes, M-P2P architecture and brokerage model for improving data availability. Interestingly, the proposals in [23,24] consider economic schemes in M-P2P networks. However, they focus on data dissemination with the aim of reaching as many peers as possible, while we consider on-demand services. Furthermore, they do not consider free-riders and M-P2P brokerage models.

P2P replication suitable for mobile environments has been incorporated in systems such as ROAM [19], Clique [20] and Rumor [9]. However, these systems do not incorporate economic schemes and brokerage models.

3 The EcoBroker Economic Incentive-Based Model

The architecture of EcoBroker consists of query-issuing MPs, relay MPs, broker MPs and data-providing MPs. Relay MPs forward messages (e.g., queries, data) in lieu of a *relay commission*. Broker MPs facilitate query-issuing MPs in obtaining their queried data items in lieu of a *broker commission*.

In EcoBroker, the price μ_d of a data item d is computed as follows:

$$\mu_d = \eta_d \times DQ_d \tag{1}$$

where η_d represents the access frequency of d. Notably, data item price increases with increasing access frequency due to frequently accessed items being more important to the network as a whole. DQ_d reflects the quality of d (e.g., image resolution, audio quality). The value of DQ is determined as in our previous work in [16], where we considered three discrete levels of DQ i.e., *high*, *medium* and *low*, their values being 1, 0.5 and 0.25 respectively. Understandably, higher-quality data items command higher prices. As a single instance, an MP requesting for MP3 songs would typically be willing to pay a higher price for obtaining better audio quality. Notably, there can be alternative approaches to defining data item prices, and some of these approaches have been examined in our previous works [17,18].

Broker commission is 5% of the price of each item retrieved by the broker. In EcoBroker, a broker only earns its commission if it retrieves all the items in a given query, hence partially answered queries do not entail a broker commission. Relay commission is a constant r, which is application-dependent. In particular, the value of r is significantly less than that of the average data item prices, thereby implying that EcoBroker provides more incentives to MPs for hosting data than for relaying messages.

4 Brokerage Model of EcoBroker for Efficiently Handling Multiple-Item Queries

This section presents the brokerage model of EcoBroker. In particular, we discuss how brokers select and obtain the items to replicate at themselves for earning currency.

For selecting the data items to replicate at itself, a given broker MP M keeps track of the queries that pass through itself. This enables M to determine the currency-earning potential for various items in the network. M maintains a list L, each entry of which is of the form $\{d_{id}, \mu_d, size_d, LL\}$, where d_{id} refers to the unique identifier of a given data item d, μ_d is the price of d and $size_d$ is the size of d. LL is a pointer to a list, each entry of which is of the form $\{k, \eta_k\}$. Here, k represents the number of items in each query issued for d and η_k is the access frequency of each k-item query.

Note that the other items in the queries issued for k could differ since η_k only considers the number of k-item queries issued for d as opposed to the actual items in those queries.

Selection of Candidate Data Items for Replication

M uses the information in list L to assign a score λ to each item. As the value of λ for an item increases, its currency-earning potential for M also increases. Hence, M prefers to replicate items with higher values of λ. For each item, M computes λ as follows.

$$\lambda = \sum_{k=1}^{max_k} \ (\ \mu \times \eta_k \ ) \ / \ (\ size \times k \ ) \tag{2}$$

where μ represents the price of the data item d and η_k is the *estimated* access frequency of d for k-item queries. M estimates the value of η_k for d based on the number of queries that had recently passed through it for d. The term $size$ refers to the size of d. k represents the number of queried items in each query (for d) intercepted by M. As a single instance, if a query for 5 items (one of which is d) is intercepted by M, k = 5.

In Equation 2, max_k refers to the maximum number of items allowable in a single query. The value of max_k is application-dependent. We found max_k = 7 to be reasonable for our application scenarios. Intuitively, if the value of max_k is too high, queries would be likely to fail because it may not be practically feasible to successfully obtain all the queried items.

The value of λ increases with increase in μ and η_k because higher-priced items and items with higher access frequencies imply more currency-earning potential for M. λ decreases with increase in item size due to reduction in M's currency-earning potential per unit of its limited memory space. λ decreases with increase in k because queries with higher values of k entail an increased probability of failure in locating the other items in the query. Recall that partially answered queries entail no earnings for brokers.

M sorts the items in L in descending order of the values of λ and selects those items (from L), whose values of λ exceed λ_{Th}, where λ_{Th} is the average value of λ for all the data items in L. Thus, λ_{Th} = ($(1/n_d) \sum_{i=1}^{n_d} \lambda_d$), where n_d is the number of data items in L and λ_d is the value of λ for a given item d. Items, for which the value of λ falls below λ_{Th}, are then deleted from L. Thus, now list L contains the items, which M wants to replicate at itself. List L is refreshed periodically to ensure that replication is performed based on recent access information. Notably, M may not necessarily be able to obtain and host all the items in L due to M's memory space constraints as well as item price constraints. Now let us see how M determines the maximum price, which it is willing to pay, for obtaining each item in L.

Determination of the Maximum Price for Obtaining a Replica from the Original Data-Provider

Let $cost_{max_d}$ be the maximum price that an MP M is willing to pay for obtaining the replica of a given item d in list L. For each item d, M computes $cost_{max_d}$ as follows:

$$cost_{max_d} = 0.5 \times (\ \mu \times \eta \ ) \ / \ k_{avg} \tag{3}$$

where μ is the price of a given item d and η is the total estimated access frequency of d. The term $(\mu \times \eta)$ is the total estimated (future) currency-earning potential of d. The factor of 0.5 reflects that M is willing to pay 50% of the estimated currency-earning potential of d to d's owner. Thus, the estimated (future) currency from the replicated data item is shared *equally* between M and the owner of d to ensure fairness. This acts as an incentive for d's owner to sell d to M because it would earn currency without having to expend its limited energy and bandwidth for serving queries on d. Additionally, d's owner could use its energy and bandwidth to serve queries on other items that it owns, thereby enabling it to earn more currency.

In Equation 3, k_{avg} is the weighted average value of k for all the queries for d that passed through M. M computes k_{avg} as follows:

$$k_{avg} = (\sum_{k=1}^{max_k} (k \times \eta_k)) \; / \; (\sum_{k=1}^{max_k} \eta_k) \tag{4}$$

For example, if there were 20 three-item queries (in which d was one of the queried items) and 40 such five-item queries involving d, the value of $k = (20 \times 3 + 40 \times 5)/(20 + 40)$ i.e., 4.33. Notably, the value of $cost_{max_d}$ decreases with increasing value of k_{avg} because as the number of *other items* included in queries involving d increases, the probability of M earning currency by hosting d decreases. This is because of increased chances of query failure on any of the other items.

Algorithm for Obtaining the Selected Items for Replication

Now let us examine the algorithm used by M to obtain its desired items. M broadcasts a query request for the items in L. M's broadcast request is a list of the form $\{ d_{id}, cost_{max_d} \}$, where d_{id} is the unique identifier of the item and $cost_{max_d}$ is the maximum price that M is willing to pay for hosting the item.

Recall that each item has only one original owner. Hence, when the owner M_O of an item d intercepts M's broadcast query, it evaluates d's currency-earning potential α_d at itself. If $\alpha_d < cost_{max_d}$, it decides to provide d to M in lieu of a payment of $cost_{max_d}$. M_O computes the value of α_d as follows.

$$\alpha_d = (\mu_d \times \eta_d) \; / \; k_{avg} \tag{5}$$

where μ_d and η_d are d's price and estimated access frequency respectively. The value of k_{avg} is computed by Equation 4. The value of α_d decreases with increase in k_{avg} due to the reasons explained for Equation 3. Notably, the values of η_d and k_{avg} as computed by M_O and M are likely to be different because each MP estimates these values based on the queries that pass through them. Furthermore, M_O might also be willing to provide d to M if its energy is low or if it lacks adequate bandwidth to serve queries on d.

Only the willing owners of the items reply to M. Upon receiving their replies, M sorts the items in descending order of λ. Then, M obtains the replicas from the corresponding owners one-by-one and makes the necessary payments to them. Thus, M keeps filling up its available memory space starting with the item with the highest value of λ, subject to memory space and item size constraints. M terminates the replication procedure when its available memory space for replication is exhausted.

5 Performance Evaluation

This section reports our performance evaluation by means of simulation.

For simulation purposes, we have used OMNeT++ 3.3p1 [6]. OMNeT++ is an object-oriented modular discrete event network simulation framework. It actively supports parallel distributed simulation and mobility, thereby making it a good choice for providing a realistic M-P2P simulation environment.

MPs move according to the *Random Waypoint Model* [1] within a region of area 1000 metres $\times$1000 metres. The *Random Waypoint Model* is appropriate for our application scenarios, which involve random movement of users. 10% of the peers were data-providers, while the others were free-riders. Thus, a total of 100 MPs comprised 10 data-providers and 90 free-riders (which provide no data). The number of brokers in our experiments was 10 i.e., out of the 90 free-riders, 10 peers decided to act as brokers. Each data-provider owns and hosts 10 items of different sizes. Each query is a request for k data items. 10 such k-item queries/second are issued in the network. We use a highly skewed Zipf distribution with zipf factor of 0.9 to determine the number of such k-item queries to be directed to each MP as well as to decide the frequency with which each individual item should occur in the queries. Communication range of all MPs is a circle of 100 metre radius. Table 1 summarizes our performance study parameters.

Table 1. Performance Study Parameters

Parameter	Default value	Variations
No. of MPs (N_{MP})	100	
No. of items in query (k)	5	2,3,4
Zipf factor (ZF)	0.9	
Queries/second	10	
Bandwidth between MPs	28 Kbps to 100 Kbps	
Probability of MP availability	60% to 80%	
Size of a data item	50 Kb to 350 Kb	
Memory space of each MP	1 MB to 2 MB	
Speed of an MP	1 metre/s to 10 metres/s	
Size of message headers	220 bytes	

Our performance metrics are **data availability** (**DA**) and **average querying hop-counts** (**HC**). DA equals $((N_S/N_Q) \times 100)$, where N_S is the number of successful queries and N_Q is the total number of queries. Thus, DA reflects the percentage of successful queries. Notably, a k-item query is considered to be successful if and only if all the queried items are retrieved. Hence, partially answered queries are considered to be unsuccessful. Queries can fail due to MPs being switched 'off' or due to network partitioning or due to failure in retrieval of at least one of the queried items. HC is the average number of hops per query.

As reference, we adapt a non-economic and non-broker-based model **NB (No_Broker)** since existing M-P2P proposals do not address economic broker-based models. In NB, brokerage is not performed and querying is broadcast-based. As NB does

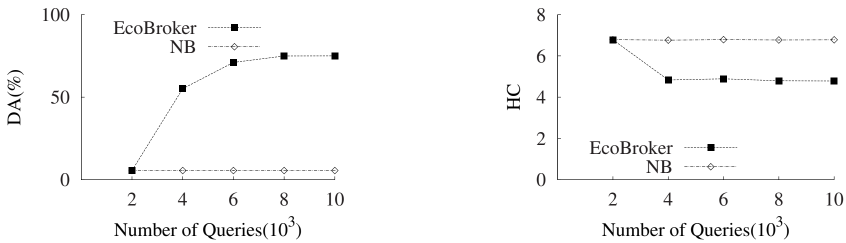

(a) Data Availability (b) Average querying hop-counts

Fig. 1. Performance of EcoBroker

not provide incentives for free-riders to host data, only a single copy of any given data item d exists at the owner of d.

Performance of EcoBroker

Figure 1 depicts the performance of EcoBroker using default values of the parameters in Table 1. Thus, all the queries in this experiment were 5-item queries. Broker-related replication procedures are initiated only after the first 2000 queries, hence both Eco-Broker and NB initially show comparable performance in terms of both DA and HC. Periodically, every 200 seconds (i.e., for every 2000 queries since there are 10 queries per second), brokers perform replication procedures.

In Figure 1a, DA remains relatively constant over time in case of NB primarily due to the absence of replication in NB, which implies that there is only copy of a given queried item in the network. On the other hand, DA increases significantly in case of EcoBroker due to its effective brokerage-related replication model, which boosts data availability. DA eventually plateaus for EcoBroker due to network partitioning and unavailability of some of the MPs.

As the results in Figure 1b indicate, HC was significantly close to the TTL of 7 hops in case of NB. This is because this experiment involved 5-item queries and in the absence of any replication, the queries needed more hops to retrieve the queried items. In contrast, HC decreased considerably in case of EcoBroker because of its effective economic incentive-based brokerage model, which incentivizes brokers to host replicas, thereby resulting in shorter query paths. Incidentally, after the initial decrease in HC for EcoBroker, HC hits a saturation point because of additional querying hops required for retrieving queried items, whose replicas do not exist at the broker. Since this experiment involved 5-item queries, it can be reasonably expected that brokers were not able to replicate all the queried items at themselves possibly due to memory space constraints as well as due to item price constraints.

Effect of Variations in *k*

Figure 2 depicts the results of varying k. Recall that k is the number of items per query. In Figure 2a, observe that as k increases, DA decreases for both EcoBroker and NB.

This is because at higher values of k, probability of query failure increases due to failure in the retrieval of at least one of the queried items. Observe that the decrease in DA with increasing k is significantly less pronounced for EcoBroker than for NB. This is because the effect of EcoBroker's economic incentive-based brokerage model becomes more prominent with increasing values of k. In particular, for 4-item queries and for 5-item queries, EcoBroker outperforms NB by upto 70% in terms of DA due to its economic brokerage model.

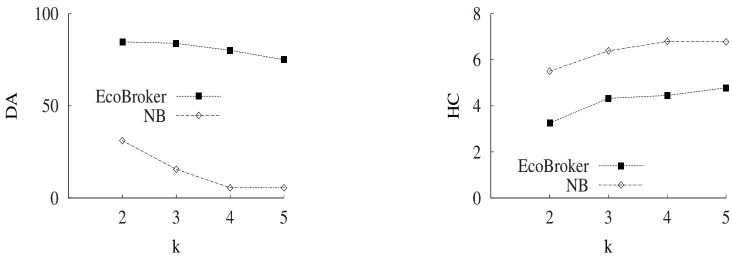

(a) Data Availability (b) Average querying hop-counts

Fig. 2. Effect of variations in k

As the results in Figure 2b indicate, HC increases with increasing values of k for both the approaches. Understandably, queries involving more items incur more hop-counts. Observe that the increase in HC is most significant as k increases from 2 to 3. This is because 2-items queries have a significantly higher probability of getting answered successfully (possibly at the same MP) as compared to the case for 3-item queries. EcoBroker outperforms NB in terms of HC due to its brokerage model, which which incentivizes brokers to host replicas, thereby facilitating queries in being answered within lower number of hops.

6 Conclusion

We have proposed EcoBroker, a novel economic incentive-based brokerage model for improving data availability via replication for multiple-item queries in M-P2P networks. In Ecobroker, data requestors pay the price (in virtual currency) of their requested data items to data-providers. The economic incentive model of EcoBroker effectively combats free-riding by incentivizing MPs to become brokers and to host replicated data, thereby improving data availability. Furthermore, EcoBroker's brokerage model facilitates efficient processing of queries involving multiple data items. Our performance evaluation demonstrates that EcoBroker indeed improves data availability and querying-related communication overhead in M-P2P networks.

References

1. Broch, J., Maltz, D.A., Johnson, D.B., Hu, Y.C., Jetcheva, J.: A performance comparison of multi-hop wireless ad hoc network routing protocol. In: Proc. MOBICOM (1998)

2. Buttyan, L., Hubaux, J.P.: Stimulating cooperation in self-organizing mobile ad hoc networks. ACM/Kluwer Mobile Networks and Applications 8(5) (2003)
3. Crowcroft, J., Gibbens, R., Kelly, F., Ostring, S.: Modelling incentives for collaboration in mobile ad hoc networks. In: Proc. WiOpt (2003)
4. Elrufaie, E., Turner, D.: Bidding in P2P content distribution networks using the lightweight currency paradigm. In: Proc. ITCC (2004)
5. Ferguson, D.F., Yemini, Y., Nikolaou, C.: Microeconomic algorithms for load balancing in distributed computer systems. In: Proc. ICDCS, pp. 491–499 (1988)
6. OMNeT++ for simulation, `http://www.omnetpp.org/`
7. Gnutella, `http://www.gnutella.com/`
8. Golle, P., Brown, K.L., Mironov, I.: Incentives for sharing in peer-to-peer networks. In: Proc. Electronic Commerce (2001)
9. Guy, R., Reiher, P., Ratner, D., Gunter, M., Ma, W., Popek, G.: Rumor: Mobile data access through optimistic peer-to-peer replication. In: Proc. ER Workshops (1998)
10. Ham, M., Agha, G.: ARA: A robust audit to prevent free-riding in P2P networks. In: Proc. P2P, pp. 125–132 (2005)
11. Hara, T., Madria, S.K.: Data replication for improving data accessibility in ad hoc networks. IEEE Transactions on Mobile Computing 5(11) (2006)
12. `http://www.microsoft.com/presspass/presskits/zune/default.mspx`
13. Kamvar, S., Schlosser, M., Garcia-Molina, H.: Incentives for combatting free-riding on P2P networks. In: Kosch, H., Böszörményi, L., Hellwagner, H. (eds.) Euro-Par 2003. LNCS, vol. 2790, pp. 1273–1279. Springer, Heidelberg (2003)
14. Kazaa, `http://www.kazaa.com/`
15. Kurose, J.F., Simha, R.: A microeconomic approach to optimal resource allocation in distributed computer systems. IEEE Trans. Computers 38(5), 705–717 (1989)
16. Mondal, A., Madria, S.K., Kitsuregawa, M.: CADRE: A collaborative replica allocation and deallocation approach for Mobile-P2P networks. In: Proc. IDEAS (2006)
17. Mondal, A., Madria, S.K., Kitsuregawa, M.: ConQuer: A peer group-based incentive model for constraint querying in mobile-P2P networks. In: Proc. MDM (2006)
18. Mondal, A., Madria, S.K., Kitsuregawa, M.: EcoRep: An economic model for efficient dynamic replication in Mobile-P2P networks. In: Proc. COMAD (2006)
19. Ratner, D., Reiher, P.L., Popek, G.J., Kuenning, G.H.: Replication requirements in mobile environments. Mobile Networks and Applications 6(6) (2001)
20. Richard, B., Nioclais, D., Chalon, D.: Clique: A transparent, peer-to-peer replicated file system. In: Proc. MDM (2003)
21. Srinivasan, V., Nuggehalli, P., Chiasserini, C.F., Rao, R.R.: Cooperation in wireless ad hoc networks. In: Proc. INFOCOM (2003)
22. Turner, D.A., Ross, K.W.: A lightweight currency paradigm for the P2P resource market. In: Proc. Electronic Commerce Research (2004)
23. Wolfson, O., Xu, B., Sistla, A.P.: An economic model for resource exchange in mobile Peer-to-Peer networks. In: Proc. SSDBM (2004)
24. Xu, B., Wolfson, O., Rishe, N.: Benefit and pricing of spatio-temporal information in Mobile Peer-to-Peer networks. In: Proc. HICSS-39 (2006)
25. Xue, Y., Li, B., Nahrstedt, K.: Optimal resource allocation in wireless ad hoc networks: A price-based approach. IEEE Transactions on Mobile Computing (2005)
26. Zhong, S., Chen, J., Yang, Y.R.: Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. In: Proc. IEEE INFOCOM (2003)