

Energy Efficient Hierarchical Epidemics in Peer-to-Peer Systems

Emrah Cem, Tugba Koc, and Ozgur Ozkasap
Department of Computer Engineering
Koc University, Istanbul, Turkey
ecem, tugbakoc, oozkasap@ku.edu.tr

Abstract—Epidemic or gossip-based mechanisms are preferred in several distributed protocols for their ease of deployment, simplicity, robustness against failures, load-balancing and limited resource usage. In flat neighborhood epidemics, peers have similar responsibilities and all participate in gossiping via neighboring peers. We have proposed an energy cost model for a generic peer using flat neighborhood epidemics, and examined the effect of protocol parameters to characterize energy consumption. Although it has been shown that a peers power consumption amount is independent of population size, peers always need to be active to process incoming gossip messages. In this study, we consider power awareness features of flat and hierarchical epidemics in peer-to-peer (P2P) systems, and propose a power-aware hierarchical epidemic approach with its energy cost model and analysis. In this adaptive approach, only a subset of peer population is active in gossiping by forming an overlay, so that the other peers can switch to idle state. It also allows data aggregation that can be utilized to reduce gossip message size. As a case study for epidemic protocol, we use our approach and simulation model for frequent item set discovery in unstructured P2P networks.

Index Terms—energy cost model, energy efficiency, peer-to-peer, gossip-based, epidemic, dominating set.

I. INTRODUCTION

Epidemic mechanisms have been used for various applications such as reliable data dissemination [1] [2], membership maintenance, overlay topology construction, failure detection, P2P streaming and data aggregation. Epidemic algorithms consist of rounds, which are the time intervals that nodes periodically communicate among each other. In each round, each peer contacts one or a few nodes, called neighbors, to exchange state. Algorithm finishes in multiple rounds, and data is disseminated to the network like an epidemic disease. An advantage of this approach is its robustness against peer failures. Removal, or a failure of a peer does not affect the dissemination speed significantly. Epidemic algorithms are preferred for their ease of deployment, simplicity and robustness. On the other hand, their disadvantage is communication overhead since any two neighbors may communicate multiple times during the algorithm, hence resulting in redundant information exchange when compared to hierarchical approaches.

There exist two main groups of epidemic algorithms, namely flat and hierarchical. Flat algorithms involves basic and neighborhood epidemics. The basic epidemic requires global knowledge of peer population and performs uniform gossiping, therefore it is not practical. On the other hand, neighborhood epidemic uses local knowledge which is more practical and performs gossiping with neighbors. Although neighborhood epidemic is better when compared to basic, it still has the problem of redundant communication. However, hierarchical epidemic makes use of structure among peers and aims to reduce communication overhead. In addition, it provides the possibility of active/passive peers to save energy.

Designing energy efficient epidemic protocols and services has become significant due to their wide usage in large scale distributed

systems. There is a lack of studies in terms of power usage and energy efficiency of epidemic approaches. In terms of their power usage, the efficiency of three models of epidemic protocols, namely basic, neighborhood and hierarchical epidemics, has been examined in [3]. However, it evaluates these through simulations only and provides results on latency and power.

In this study, we propose a dominating-set based and power-aware hierarchical epidemic approach that eliminates significant number of peers from gossiping. In contrast to the prior works on hierarchical epidemics, we use dominating-set to construct a hierarchy, and to choose peers performing gossip operation for energy efficiency. In this adaptive approach, only a subset of peer population is active in gossiping by forming an overlay consisting of dominating set peers, so that the other peers can switch to idle state. There exist prior studies using dominating set in P2P networks. For instance, a dominating-set-based P2P protocol is proposed in [4], [5] that aim to use minimum number of forwarding nodes for data delivery. Another study [6] uses dominating set to solve searching problem in P2P networks. To the best of our knowledge, using dominating set to build a hierarchy for epidemic communication and provide energy efficiency in P2P networks is a novel approach.

II. FLAT NEIGHBORHOOD EPIDEMICS

In flat neighborhood epidemics, all peers periodically exchange local state with neighbors until convergence. A peer i is the neighbor of peer j if they are directly connected in the overlay network. Main advantage of this approach over basic epidemics is that scalability problem is solved since no peer needs global knowledge. On the other hand, all peers may communicate more than once with their neighbors, which is necessary for convergence. Thus, high communication overhead is a drawback of this approach.

A. Cost Model for Flat Neighborhood Epidemics

Inspired by studies [7], [8], we have proposed an energy cost model for a generic peer using epidemic communication in ProFID protocol [9]. Different than prior works, we developed energy cost model for a peer using gossip-based communication and consider the effects of gossip parameters on the cost representation. ProFID protocol depends on three main components of operations performed by each peer: energy consumed while (1) computing new state (E_{comp}), (2) sending messages (E_{send}), and (3) receiving messages (E_{recv}).

Energy consumption of a peer P_i during an atomic pairwise averaging with P_j can be written as:

$$E_{i,j} = E_{send,j} + E_{receive,j} + E_{comp,i} + C \quad (1)$$

where $E_{send,j}$ is the energy consumed while sending a gossip message to P_j , $E_{receive,j}$ is the energy consumed while receiving a gossip message from P_j , and $E_{comp,i}$ is the local computation cost of P_i . A constant C is also added to the equation to include extra

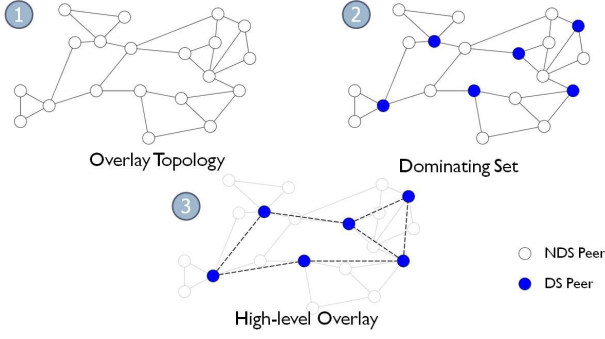


Fig. 1. Construction of high level overlay using dominating set

factors such as CPU's or other peripherals' energy consumption. The formula was given with respect to basic conditions (gossip to one neighbor, one round, one item). If we take all parameters such as $fanout$ and mms into consideration, we obtain the following equation

$$E_{P_i}(\text{single round}) \leq mms \cdot \sum_{j \in V \cup W} E_{i,j} \quad (2)$$

where V is the set of neighbors chosen by P_i as gossip targets, and W is the set of neighbors that initiates an atomic pairwise averaging with P_i . Note that the number of elements in V corresponds to the $fanout$ value. Details of the energy cost model are given in [9]. Since a peer repeats those operations in every round, number of rounds R would increase the energy cost of a peer proportionally. Hence, the overall energy cost of P_i can be written as:

$$E_{P_i} \leq R \cdot mms \cdot \sum_{j \in V \cup W} E_{i,j} \quad (3)$$

III. DOMINATING SET BASED HIERARCHICAL EPIDEMICS

In our hierarchical model, we use dominating set idea to build a high-level overlay as illustrated in Fig. 1. Dominating Set (DS) can be defined as a subset S of a graph $G=(V,E)$ such that every vertex in G is either in S or adjacent to a vertex in S . In our approach, a DS corresponds to a subset of peers such that a peer in the system is either in DS or a neighbor of a DS peer. There are two types of peers in the system, namely dominating set (DS) peer and non-dominating set (NDS) peer that can be chosen through distributed algorithms for DS construction [5], [10]. In this model, we aim to save energy by reducing the number of peers performing gossip operation. NDS peers are excluded from gossip operation. DS peers collect local states of NDS neighbors. Consequently, DS peers form a high level overlay topology on which gossiping is performed. There are two main advantages of this approach. (1) NDS peers send their local state to one or multiple DS neighbors and then they switch to passive mode in which they just wait for the result to be announced by a DS neighbor. During this time period, they contribute reducing the energy consumption. (2) Only DS peers participate in gossiping and hence the convergence time and message complexity of the epidemics reduces. Alg. 1 shows hierarchical epidemic algorithm used for frequent item discovery. Next, we develop energy cost model for two types of peers in the hierarchy: DS and NDS peers.

Algorithm 1: Hierarchical Epidemics Algorithm

Require: $fanout, ui, convLimit, \epsilon, T, S$
Ensure: F (Set of frequent items)

DS Peer
 x =the number of NDS neighbors;
for $i=1$ to x **do**
 send(req(LSI,DS));
end for
do for NDS neighbor state collection
 msg = accept();
 DS.update(msg);
 S.update(msg);
if completed then
 do periodically
 if !converged then
 targets = getNeighbors(fanout);
 for $i=1$ to $fanout$ **do**
 send(push, S, targets(i));
 end for
 end if
 end if
do until convergence
 msg=accept();
 if msg.Type == push **then**
 avg = AVG(S, msg); S.update(avg); send(pull, avg, sender);
 else if msg.Type == pull **then**
 S.update(msg);
 end if
 currSizeEstim=msg.getVal(ui);
 if ISCONVERGED(convLimit, ϵ , currSizeEstim) **then**
 converged=true;
 end if
 NDS Peer
 msg = accept();
 if msg.Type == request(LSI,DS) **then**
 send(LSI,DS, sender);
 else if msg.Type == request(DS) **then**
 send(DS, sender);
 else if msg.Type == result **then**
 S.update(msg);
 end if

A. Energy Cost Model for DS Peer

There are three main operations for DS peers causing energy consumption. Firstly, DS peers send request messages to their NDS neighbors to get their local state information and the id set of their DS neighbors. During this operation energy cost of a DS peer P_i during a request operation from a NDS peer P_j is:

$$E_{i,j,1} = E_{send,j} + (s + d)(E_{receive,j} + E_{comp,i}) + C \quad (4)$$

where $E_{send,j}$ is the energy consumed while sending request message to a NDS peer P_j , s is the number of item tuples in local state information, d is the number of DS neighbors, $E_{receive,j}$ is the energy consumed while receiving the local state information and the id set of DS neighbors of a NDS peer P_j , and $E_{comp,i}$ is the energy consumed to aggregate local state information of NDS neighbor. A DS peer send request messages to all its NDS neighbors and if we extend the formula:

$$E_{P_i,1} = \sum_{j \in X} E_{i,j,1} \quad (5)$$

where X is the set of NDS neighbors. Secondly, after getting local state information and set of other DS neighbors, DS peers start gossiping with other DS peers. This operation can be thought of as a flat gossip among DS peers in high level of the hierarchy. Therefore, the energy consumption of a DS peer P_i during an atomic pairwise averaging with P_j can be written as:

$$E_{i,j,2} = E_{send,j} + E_{receive,j} + E_{comp,i} + C \quad (6)$$

where $E_{send,j}$ is the energy consumed while sending a gossip message to P_j , $E_{receive,j}$ is the energy consumed while receiving

a gossip message from P_j , and $E_{comp,i}$ is the local computation cost of P_i while choosing and preparing the tuple, computing the average and updating the state. To represent the energy cost of a DS peer during an atomic pairwise averaging operation, the formula was given with respect to the basic conditions (gossip to one neighbor, one round, one item). So, if we extend this cost model, the formula can be written as:

$$E_{P_i,2} = R_H \cdot mms \cdot \sum_{j \in Y \cup Z} E_{i,j,2} \quad (7)$$

where Y is the set of DS neighbors chosen by P_i as gossip targets and Z is the set of DS neighbors that initiates an atomic pairwise averaging with P_i . In general, a gossip message comprises multiple item tuples whose number is upper-bounded by maximum message size (mms) parameter. Since a peer repeats those operations in every round, number of rounds R_H would increase the energy cost of a peer proportionally. Finally, after finishing gossip and convergence in high level, DS peers should send the result to their DS neighbors. During this send operation, energy cost is:

$$E_{i,j,3} = s(E_{send,j} + E_{comp,i}) + C \quad (8)$$

where $E_{send,j}$ is the energy consumed while sending final state information to a NDS peer P_j , s is the number of item tuples in local state information, $E_{comp,i}$ is the energy consumed while preparing the tuple to send. A DS peer send result message to all its NDS neighbors; hence the extended version of the formula is represented by:

$$E_{P_i,3} = \sum_{j \in X} E_{i,j,3} \quad (9)$$

The overall energy cost of P_i can be written as:

$$E_{P_i} = \sum_{j \in X} E_{i,j,1} + R_H \cdot mms \cdot \sum_{j \in Y \cup Z} E_{i,j,2} + \sum_{j \in X} E_{i,j,3} \quad (10)$$

In this formula, $E_{P_i,1}$ and $E_{P_i,3}$ do not have a significant effect on the overall energy consumption. The important effect is due to $E_{P_i,2}$ which is related to atomic pairwise averaging operation. However, in hierarchical model the number of peers performing atomic pairwise averaging operation even less than those in the flat model and due to less number of peers gossiping each other R_H (the number of rounds in hierarchical gossip) is smaller than R_F (the number of rounds in flat gossip), because the number of rounds and the number of gossiping peers, N_H , are proportional based on the following formula:

$$R \approx (1/\log \epsilon) \cdot \log N_H \cdot convLimit \cdot (1/fanout) \cdot (1/mms) \quad (11)$$

B. Energy Cost Model for NDS Peer

NDS peers are not fully active in the system, and their energy consumption is far less when compared to DS peers, because they do not take part in the gossip operations. Their responsibility in the system is to receive request messages, send their local state information and the id set of their DS neighbors. At the end of the protocol, they receive final state information from their DS neighbors and update themselves. Energy consumption of a NDS peer can be

analyzed in two parts. The first part is to send response messages to DS neighbors.

$$E_{i,j,1} = E_{receive,j} + (s + d)(E_{send,j} + E_{comp,i}) + C \quad (12)$$

where $E_{receive,j}$ is the energy consumed while receiving request message from a DS neighbor P_j , $E_{send,j}$ is the energy consumed while sending local state information and the id set of the other DS neighbors, s is the number of item tuples, d is the number of DS neighbors, and $E_{comp,i}$ is the energy consumed while preparing the tuple and DS neighbor set. A NDS peer gives response to all its DS neighbors, so the formula can be extended as:

$$E_{P_i,1} = \sum_{j \in K} E_{i,j,1} \quad (13)$$

where K is the set of DS neighbors. In the second part, NDS peers consume energy when receiving result messages from their DS neighbors. So, the energy cost can be written as:

$$E_{i,j,2} = s(E_{receive,j} + E_{comp,i}) + C \quad (14)$$

where $E_{receive,j}$ is the energy consumed while receiving final state information from a DS neighbor, s is the number of item tuples, and $E_{comp,i}$ is the energy consumed while updating local state information. The extended formula can be represented as:

$$E_{P_i,2} = \sum_{j \in K} E_{i,j,2} \quad (15)$$

Therefore, the overall energy consumption of a NDS peer P_i can be written as:

$$E_{P_i} = \sum_{j \in K} (E_{i,j,1} + E_{i,j,2}) \quad (16)$$

IV. EXPERIMENTAL RESULTS

We have performed various experiments in PeerSim, a scalable P2P simulation environment, to compare the energy efficiencies of flat neighborhood and hierarchical epidemic approaches in terms of message overhead and convergence time. Moreover, we conducted experiments to observe the characteristics of high-level overlay consisting of only DS peers. As a case study, we used data aggregation problem in which all peers keep multiple $\langle item, frequency \rangle$ tuples and aim to compute the sum of each individual item. We used Barabasi-Albert (BA) model [11] with average degree 10 while constructing the overlay topologies. Each peer contacts with a single neighbor at each round and sends its local state in a single gossip message. Note that sampling algorithms can be used to reduce the content to be sent in a single message.

A. High-level overlay characteristics

In this section, we analyze the characteristics of high-level overlay in terms of peer average degree and the proportion of peers that are in DS for different network size. These are important because they directly affect the convergence time and the scalability of the algorithm. If the average degree at the high-level is too low then algorithm will converge slowly. If the average degree at the high-level is too high, then algorithm will not be scalable since it will behave like basic epidemics at high-level. Proportion of peers in DS is also important because, the less the number of peers in DS, the

faster the convergence. On the other hand, the more the number of peers in DS, the more robust the algorithm is. Therefore, there is a trade off between the speed and the robustness of the algorithm.

While constructing the topology, we use Barabasi-Albert model [11] which produces power-law distribution with exponent $\gamma = 3$. Based on this model, the network starts with m_0 peers, where m_0 is much less than the network size. Then, at each step a new peer is added to the existing network and the probability that newly added peer is connected to peer i is defined as:

$$\Pi(k_i) = \frac{k_i}{\sum_j k_j} \quad (17)$$

where k_i is the degree of peer i . The above equation states that peer with higher degree has more probability to connect to the newly added peer. This phenomena is also known as *preferential attachment*. In order to modify the average degree of the topology, we simply change the number of peers to which a newly added peer connect.

As depicted in Fig. 2, percentage of peers that are in DS decreases when the average degree of peers increases, which is expected because less peers can cover all the peers in the network if peers have more neighbors. Another deduction is that more than half of peers are eliminated from gossiping, which means more than half of the peers will not consume energy during gossiping. Moreover, convergence time of the algorithm is expected to reduce since the number of peers participating in gossiping decreases.




Fig. 2. Percentage of peers in DS (Barabasi-Albert)

Fig. 3 shows that average degree of peers in high-level overlay logarithmically increases. This property would have a positive effect for both scalability and robustness of the algorithm. The more the number of peers in DS, the more robust the algorithm is because computation involves in more peers and the failure of one of peers causes less information loss. On the other hand, the more the number of peers in DS, the more peer will share the load; hence, no peer will be a bottleneck. In order to see the effect of topology on the




Fig. 3. Average degree of DS peers

percentage of peers in DS, we have also tested Erdos-Renyi (ER) model which is one of the well known models for generating random topology. In this model, an edge is added between each pair of peers with a probability p independently of other edges. We selected p such that average degree is 10 for fair comparison with power-law topology. We observed that percentage of peers in DS is around 20% of the total number of peers on the overlay. This is less in comparison to power-law BA topology (which is around 30% as shown in Fig. 2). The reason is likely that most peers in power-law topology has low degree, while most peers have degree around 10 in case of random topology. As the definition of dominating set states, a peer is either in dominating set or the neighbor of other peer in the dominating set. When a peer is added to dominating set, its neighbors also satisfy the definition. Therefore, the more the number of peers with high degree, the less the size of dominating set.

B. Scalability and energy consumption comparison

In this section, we compare the scalability of flat and hierarchical epidemic algorithms. These results are also directly related with energy consumption because the less the number of messages sent and convergence time, the less energy consumption. Fig. 4 and 5 depict the energy consumption characteristics of flat and hierarchical epidemics on two different topology models (BA and ER) in terms of convergence time and message cost.

In Fig. 4, when the effect of algorithm is considered, hierarchical epidemics outperforms flat case in both topology types. Reasons of faster convergence are pre-aggregation before gossiping and less number of peers participating in gossiping. Faster convergence also means less number of messages sent per peer since gossip rate is the same. Fig. 5 shows that number of messages sent per peer is much less in hierarchical epidemics when compared to flat epidemics. Since the number of gossiping peers is affected from the topology, the number of messages sent per peer is also directly affected. In ER model, peer degrees are uniformly distributed as opposed to BA model in which most of peers have a low degree. Depending on the algorithm that chooses gossiping peers, number of gossiping peers change. Results show that slightly more peers participate in gossiping in BA in comparison to ER model.

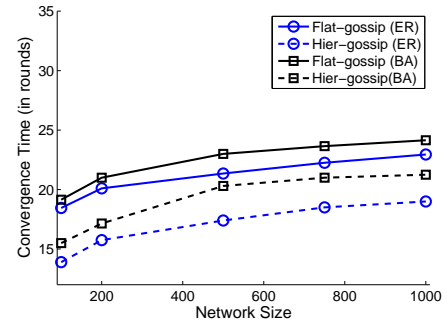


Fig. 4. Convergence time: Flat and hierarchical epidemics on BA and ER topologies

V. CONCLUSION

We studied the energy efficiency aspect of epidemic protocols, and proposed a novel hierarchical epidemic approach that uses dominating set while constructing the hierarchy. Proposed approach utilizes the benefits of both epidemic and hierarchical approaches. It uses only local knowledge and provides the possibility of active/passive peers




Fig. 5. Average number of gossip messages sent per peer: Flat and hierarchical epidemics on BA and ER topologies

to save energy. Since the constructed hierarchy is 2-level, it is more robust against node failures than pure hierarchical approaches. Moreover, message overhead is significantly reduced thanks to the hierarchy. As future work, we will analyze the effect of churn and test the hierarchical epidemic approach in PlanetLab, which is a real network test-bed. Moreover, we plan to develop an algorithm for dynamic DS update in order to deal with churn. We also plan to propose a smart distributed DS construction algorithm which considers not only neighborhood information but also stability and the capabilities of peers in terms of capacity and computational power.

ACKNOWLEDGMENT

This work was partially supported by the COST (European Cooperation in Science and Technology) framework, under Action IC0804: Energy efficiency in large scale distributed systems, and by TUBITAK (The Scientific and Technical Research Council of Turkey) under Grant 109M761. A preliminary version of this study appeared in [12].

REFERENCES

- [1] K. Birman, M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu, and Y. Minsky, "Bimodal multicast," *ACM Transactions on Computer Systems*, vol. 17, no. 2, pp. 41–88, 1999.
- [2] O. Ozkasap, Z. Genc, and E. Atsan, "Epidemic-based reliable and adaptive multicast for mobile ad hoc networks," *Computer Networks*, vol. 53, pp. 1409–1430, 2009.
- [3] R. van Renesse, "Power-aware epidemics," in *Proc. of IEEE Symposium on Reliable Distributed Systems*, 2002.
- [4] D. T. Ahmed, S. Shirmohammadi, and A. El-Saddik, "A dominating set based peer-to-peer protocol for real-time multi-source collaboration," in *WETICE*, 2007, pp. 119–124.
- [5] L. Jia, R. Rajaraman, and T. Suel, "An efficient distributed algorithm for constructing small dominating sets," *Distributed Computing*, vol. 15, no. 4, pp. 193–205, 2002.
- [6] C. Yang and X. Li, "'dominating-set-based' searching in peer-to-peer networks," *IJHPCN*, vol. 3, no. 4, pp. 205–210, 2005.
- [7] C. Seo, G. Edwards, D. Popescu, S. Malek, and N. Medvidovic, "A framework for estimating the energy consumption induced by a distributed system's architectural style," in *Proc. of SAVCBS*, 2009.
- [8] R. Basmadjian and H. de Meer, "An approach to reduce the energy cost of the arbitrary tree replication protocol," in *Proc. of e-Energy*, 2010, pp. 151–158.
- [9] E. Cem, E. Demirkaya, E. Esiner, B. Ozaydin, and O. Ozkasap, "Energy cost model for frequent item set discovery in unstructured p2p networks," in *Proc. of ISCS*, 2011.
- [10] F. Kuhn and R. Wattenhofer, "Constant-time distributed dominating set approximation," *Distributed Computing*, 17(4), pp. 303–310, 2005.
- [11] A. L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [12] E. Cem, T. Koc, and O. Ozkasap, "Dominating-set based and power-aware hierarchical epidemics in p2p systems," in *Proc. of e-Energy, poster, New York*, 2011.