Exact Probability Distributions for Peer-to-Peer Epidemic Information Diffusion

Emine Şule Yazıcı* Selda Küçükçifçi* Öznur Özkasap† Mine Çağlar*

1. INTRODUCTION

An efficient approach for information diffusion in distributed systems is to utilize epidemic algorithms that involve pair-wise propagation of updates. Epidemic algorithms are fully distributed and randomized approaches such that every peer in an information diffusion session picks a (subset of the other) peer(s) randomly for efficient propagation of updates, through periodic rounds. The underlying epidemics theory for the biological systems studies the spreading of infectious diseases through a population [1,2]. When applied to an information diffusion application, such protocols have beneficial features such as scalability, robustness against failures and provision of eventual consistency. Exact as well as asymptotical distributions have been studied for different epidemic models in [3,4]. In contrast to such previous studies, we investigate variations of the epidemic algorithms used in the context of distributed information diffusion and derive exact diffusion probabilities for them.

One of the first studies that applies epidemic methods to computer systems used the idea for spreading updates in a replicated database [5]. Later, a class of epidemic models known as anti-entropy has been proposed as a mechanism that runs in the background for recovering errors, in particular for loss recovery in Bimodal Multicast protocol [6]. In [7], an overview of epidemic information dissemination is given. A modeling approach for epidemic dissemination is the use of fluid models based on differential equations. See for example [8,9,10] in the context of spreading of worms in the Internet. In [11], a binomial probability distribution model for the information flow dynamics is used for comparing pull and push anti-entropy approaches for update exchange. Fluid or binomial probability models might provide good approximations especially for large networks.

In this study, we develop exact diffusion probabilities for pull and push information diffusion models of anti-entropy, as well as for the hybrid approach. The number of peers lacking the information forms a Markov chain advancing in rounds. The exact transition probabilities on the chain are derived through elaborated counting techniques on a digraph exactly, with no resort to approximate probability distributions that rely on several independence assumptions. Our results form a basis for computing message latencies exactly, which can be compared with the approximate results of [11]. Such results would be beneficial for integrating in several distributed scenarios such as replicated servers, loss recovery, failure detection and group membership management.

*Department of Mathematics, Koç University, Istanbul, Turkey.
†Department of Computer Eng., Koç University, Istanbul, Turkey.

2. MODEL DESCRIPTIONS

A popular distribution model based on the theory of epidemics is the anti-entropy [1]. According to the terminology of epidemiology, a site or peer holding information or an update it is willing to share is called *infectious*. A peer is called *susceptible* if it has not yet received an update. In the anti-entropy process, non-faulty peers are always either susceptible or infectious. Data diffusion progresses periodically via rounds of epidemics. In each round, every peer picks another site at random, and exchanges its state information with the selected one. We study the following approaches for update-exchange

Pull Approach: When an infectious peer (holding data to be shared) picks a susceptible peer (lacking the specific data) randomly, this triggers data dissemination from infectious peer to the susceptible. Steps involved in the dissemination between two such peers is depicted in Fig.1(a) where infectious peer (on the left) has data labelled A. The infectious peer sends a digest (also referred to as gossip) message including its state information. On receiving digest and comparing it with its local data, the susceptible peer finds out it lacks A and sends a request for A back to the infectious. Upon getting request, infectious peer sends a retransmission of data A which causes the other peer to be infectious for A. In fact, each peer in the system performs state information exchange periodically and concurrently with the others. Moreover, each peer may have a set of data in its local buffer. Therefore, a digest message generated by a peer would consist of state information on the current contents of its message buffer. In that respect, the figure simplifies the scenario and illustrates the communication between two sample peers for one piece of data A. Spreading updates is triggered by susceptible peers when they are picked as gossip destinations by infectious peers.

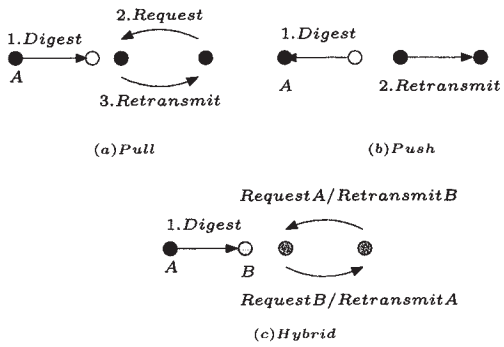

Figure 1: Model Illustrations

Push Approach: If a susceptible peer picks an infectious peer randomly, and sends its state information, this triggers information dissemination from infectious peer to the susceptible. Steps involved in the dissemination between two such peers is depicted in Fig.1(b) where infectious peer (on the left) has data labelled A. The infectious peer, on receiving digest and comparing it with its local data, finds out that the digest owner lacks A and directly retransmits, or pushes data A which causes the other peer to become infectious for A. As illustrated in the figure, in the push approach, no request messages are used. Spreading updates is triggered by infectious peers when they are selected as gossip targets by susceptible peers.

Hybrid Approach: This is a hybrid of two approaches described above. As illustrated in Fig.1(c), when a peer sends its digest to a randomly selected peer in the population, this may trigger data dissemination at both peers. Consider the case where a peer has data A and the other has data B. When the former selects the latter as the digest target in a given round, information A and B would be disseminated to the peer that lack it using pull-based or push-based approaches together. This approach is useful for delay sensitive applications since it decreases overall delay during data dissemination at the cost of possible duplicate data transmissions.

3. EXACT DIFFUSION PROBABILITIES

In this section, we will restrict our attention to the process of distributing a single data message. Therefore, a peer with a copy of the data message is referred to as infectious; otherwise, it would be susceptible. Each step of this diffusion process can be represented by a digraph D where a peer in the population of size n corresponds to a node of the digraph. If the node u chooses to communicate with the node v then there will be an arc with the tail u and the head v in D. Since each node chooses exactly one node at each step the out degree of each node will be 1 in D. The number of all possible such digraphs with n nodes is $(n-1)^n$. All of these digraphs are equally likely for each step of this process. Therefore, we will count the number of digraphs that infect i more nodes and take the ratio of this number with the number of all possible digraphs to find the probability of infecting i more nodes. Note that if there are k infectious nodes, after one step there will be $k+i$ infectious nodes, where $k+i = 1, 2, \ldots, n$.

Let S be the set of all susceptible nodes and I be the set of all infectious nodes with $|I| = k$ and $|S| = n-k$. For simplicity, we will denote arcs with susceptible heads and infectious tails by IS-arcs, similarly arcs with infectious heads and susceptible tails, infectious heads and infectious tails, and susceptible heads and susceptible tails will be represented by SI-arcs, II-arcs, and SS-arcs, respectively. Note that D is the disjoint union of four subgraphs formed by IS-arcs, SI-arcs, II-arcs, and SS-arcs.

In the following subsections, we will need the definition of Stirling number of the second kind. The number of ways of partitioning a set of n elements into k nonempty subsets is called Stirling number of the second kind and denoted by $S(n,k)$. It follows that $S(n,k) = kS(n-1,k) + S(n-1,k-1)$ and $S(0,0) = 1$. For further information on these numbers see [12].

Pull Case: We form the digraph D as above. In the pull case, a susceptible node s will be infected if and only if there exists a IS-arc in D with the head s. So SI-arcs, II-arcs, and SS-arcs will not contribute to the number of new infectious nodes. Fig.2 (a) illustrates the pull case. We will determine the number digraphs representing a step that results in i more infectious nodes. The number of different possible subgraphs formed by the union of SI-arcs and SS-arcs is $(n-1)^{n-k}$.

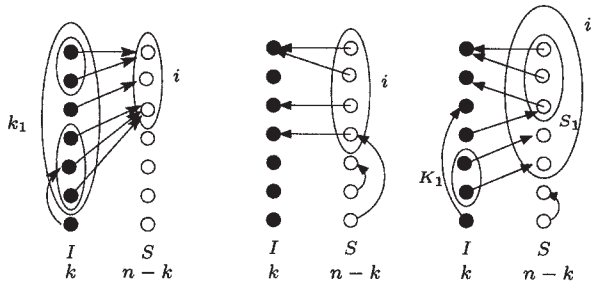

Figure 2: (a) Pull Case (b) Push Case (c) Hybrid Case

Now we need to count the number of different possible subgraphs that can be formed by IS-arcs and II-arcs. Let k_1 be the number of IS-arcs. Note that k_1 has to be at least i since each IS-arc infects at most one new node in S also there are $\binom{k}{k_1}$ such k_1-subsets of I. We have $k-k_1$ II-arcs. The number of different possible subgraphs formed by these arcs is $(k-1)^{k-k_1}$. Finally we will count the number of different subgraphs that can be formed by IS-arcs. Among $n-k$ susceptible nodes there are $\binom{n-k}{i}$ different i-subsets of S that may be infected. There are $S(k_1,i)i!$ different ways for k_1 nodes to infect exactly i new nodes. So the number of different subgraphs that can be formed by IS-arcs and II-arcs is

$$\sum_{k_1=i}^{k} \binom{k}{k_1} (k-1)^{k-k_1} \binom{n-k}{i} S(k_1,i)i! \,.$$

Hence the probability of infecting i more nodes in the next step given $|I| = k$ is

$$p(i|k) = \frac{\binom{n-k}{i} i! \sum_{k_1=i}^{k} \binom{k}{k_1} (k-1)^{k-k_1} S(k_1,i)}{(n-1)^k}. \quad (1)$$

where $k = 2, 3 \ldots, n-1$ and $i = 0, 1, \ldots, n-k$.

For the case $k = 1$ we can easily see that $p(0|1) = 0$ and $p(1|1) = 1$.

Push Case: In the push case, a susceptible node s will be infected if and only if there exists a SI-arc with the tail s. So IS-arcs, II-arcs, and SS-arcs will not contribute to the number of new infectious nodes. Fig.2 (b) illustrates the push case. The number of different possible subgraphs formed by IS-arcs and II-arcs is $(n-1)^k$. There are i SI-arcs and $\binom{n-k}{i}$ different i-subsets of S. For each SI-arc there are k different choices for the head of the arc, therefore there are $\binom{n-k}{i}k^i$ different possible subgraphs formed by these arcs. Finally, the number of different possible subgraphs formed by SS-arcs is $(n-k-1)^{n-k-i}$.

So the probability of infecting i more nodes after the step given $|I| = k$ is

$$p(i|k) = \frac{\binom{n-k}{i} k^i (n-k-1)^{n-k-i}}{(n-1)^{n-k}} \quad (2)$$

where $k = 1, 2, \ldots, n-2$ and $i = 0, 1, \ldots, n-k$. Clearly when $k = n-1$, we get $p(0|k) = 0$ and $p(1|k) = 1$.

This probability distribution can be rewritten as $\binom{n-k}{i} \left(\frac{k}{n-1}\right)^i \left(\frac{n-k-1}{n-1}\right)^{n-k-i}$ which can now be recognized as the binomial distribution with parameters $n-k$ and success probability $k/(n-1)$. This coincides with the distribution modeled in [11] through probabilistic arguments. The only difference is that the number of possible nodes among which an infectious node chooses to communi-

cate, is rounded to n in [11]. In fact, it is $n-1$ as given in the present analysis.

Hybrid Case: In the hybrid case, a susceptible node s will be infected if and only if there exists either a SI-arc with the tail s or an IS-arc with the head s. So II-arcs and SS-arcs will not contribute to the number of new infectious nodes. Fig.2 (c) illustrates the hybrid case. There are i new infectious nodes and $\binom{n-k}{i}$ different i-subsets of S. Let S_1 be the set of the tails of SI-arcs where $|S_1| = i_1$. The number of different possible subgraphs formed by SI-arcs and SS-arcs is $\binom{n-k}{i}\binom{i}{i_1}k^{i_1}(n-k-1)^{n-k-i_1}$.

Let K_1 be the set of nodes that are the tails of the IS-arcs whose heads are in $S \setminus S_1$, where $|K_1| = k_1$. There are $\binom{k}{k_1}$ different ways to choose K_1. These k_1 arcs will infect $i-i_1$ new nodes and there are $S(k_1, i-i_1)(i-i_1)!$ different ways to do this. Finally the remaining $k-k_1$ arcs can be chosen in $(k-1+i_1)^{k-k_1}$ different ways. So the number of different possible subgraphs formed by IS and II-arcs can be calculated as

$$\Theta_{k,i}(k_1, i_1) = \sum_{k_1=i-i_1}^{k} \binom{k}{k_1}(k-1+i_1)^{k-k_1} S(k_1, i-i_1)(i-i_1)!.$$

Hence the probability of infecting i more nodes in the next step given $|I| = k$ is

$$p(i|k) = \frac{\binom{n-k}{i}}{(n-1)^n} \sum_{i_1=0}^{i} \binom{i}{i_1} k^{i_1}(n-k-1)^{n-k-i_1}\Theta_{k,i}(k_1, i_1)$$

where $k = 2, 3 \ldots, n-2$ and $i = 0, 1, \ldots, n-k$.

Now we will consider the end points. If $k = n-1$, $p(0|k) = 0$ and $p(1|k) = 1$. If $k = 1$, then $p(0|1) = 0$ and $p(i|1) = \frac{\binom{n-1}{i} i (n-2)^{n-i-1}(n-1)}{(n-1)^n}$ for all $i \geq 1$. There can be $\binom{n-1}{i}$ different i-subsets of S and i different possibilities for the head of the SI-arc, call this node u. There are $n-1$ possibilities for the arc with the tail u. The arcs coming out of the rest of the $i-1$ nodes will have heads in I and there is a unique way to do this. Finally the remaining $n-i-1$ arcs can be chosen in $(n-2)^{n-i-1}$ different ways.

4. CONCLUSIONS AND FUTURE WORK

We have derived the exact probability distributions for pull, push and hybrid information diffusion models of anti-entropy. To the best of our knowledge, this study is the first one deriving exact distributions which would be helpful in performance analysis of these epidemic diffusion models. In contrast, previous results rely on simplified models of epidemics usually requiring estimation of several parameters. Our findings show that the binomial model used previously for pull case is not accurate whereas the model for push case is exact. There exists no previous probability model for hybrid case, the exact distribution of which is derived in this paper. An extension of this work for computing message latency appears in [13].

The duplicate messages associated with the hybrid case due to both pull and push deliveries are worth counting in order to determine any trade off. Dissemination of only one message has been considered. Initialization with a bigger volume of content such as in file sharing applications, the possibility of gossiping to more than one peer and partial membership knowledge among peers could also be incorporated as future work. The effect of network topology is another aspect to be considered [14]. Although evaluating the probabilities $p(i|k)$ are computationally intensive, large group sizes are feasible with current computing abilities. Analyzing the complexity of the algorithm for numerical evaluation and associated numerical accuracy is left as future work. On the other hand, asymptotical analytical expressions for large n would be also useful.

Acknowledgment. This work is supported in part by TUBITAK (The Scientific and Technical Research Council of Turkey) under CAREER Award Grant 104E064 and by TUBA (Turkish Academy of Sciences).

References

[1] N. T. J. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Charles Griffin and Compan, London, 1975.

[2] D.J. Daley, J. Gani, *Epidemic Modeling: An Introduction*. Cambridge University Press, 1999.

[3] Jerzy Jaworski. Epidemic Processes on Digraphs of Random Mappings. *J. Appl. Prob.*, 36: 780–798, 1999.

[4] Boris Pittel. On Spreading a Rumor. *SIAM Journal on Applied Mathematics*, 47(1): 213–223, 1987.

[5] A. Demers, D.Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. of the Sixth ACM Symp. on Principles of Distributed Computing*, 1–12, 1987.

[6] K. Birman, M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu, and Y. Minsky. Bimodal multicast. *ACM Transactions on Computer Systems*, 17(2): 41–88, 1999.

[7] P. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulie. Epidemic information dissemination in distributed systems. *IEEE Computer*, pages 60–67, May 2004.

[8] C.C. Zou, L. Gao, W. Gong, D. Towsley, Monitoring and early warning for Internet worms. Proceedings of 10th ACM Conference on Computer and Communications Security (CCS03), October 2003.

[9] C.C. Zou, D. Towsley, W. Gong, On the performance of Internet worm scanning strategies. *Performance Evaluation* 63: 700-723, 2006.

[10] L. Massoulie, A. Ganesh, D. Gunawardena, P. Key and J. Scott, Efficient quarantining of scanning worms: optimal detection and coordination. *Proc. of IEEE INFOCOM*, 2006.

[11] M. Çağlar, Ö. Özkasap. A chain-binomial model for pull and push-based information diffusion. *Proc. of IEEE ICC*, 2006.

[12] J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, 1992.

[13] Ö. Özkasap, E. Ş. Yazıcı, S. Küçükçifçi, M. Çağlar, Exact Performance Measures for Peer-to-Peer Epidemic Information Diffusion. LNCS 4263, *Proc. of ISCIS'06*, 2006.

[14] L. Massoulie, A. Ganesh and D. Towsley. The effect of Network Topology on the Spread of Epidemics. *Proc. of IEEE INFOCOM*, 2005.