Accessing Probabilistic Quorums in Dynamic Networks

Oksana Denysyuk
INESC-ID/IST
oksana.denysyuk@ist.utl.pt

Luís Rodrigues
INESC-ID/IST
ler@ist.utl.pt

ABSTRACT

Quorums are a fundamental building block for solving various fundamental problems such as consensus, distributed dictionaries, distributed storage, among others. In particular, probabilistic quorums have shown to be scalable, efficient, and suitable for dynamic environments [12]. Unfortunately, most existent analytic results for accessing probabilistic quorums are tailored to static networks [8]. However, we believe that the correct functioning of such systems must be assured in a much wider range of scenarios.

In this paper, we discuss the random walk based scheme for accessing probabilistic quorums in dynamic networks where an oblivious adversary changes the communication links arbitrarily in each round. We show that $\mathcal{O}(n^3 \log n)$ communication steps are needed to access a probabilistic quorum under these conditions.

Keywords

Quorums, Dynamic Networks, Random Walks

1. INTRODUCTION

Quorums are a basic construction mechanism in many distributed systems. They can be used as building blocks in solving various fundamental problems such as consensus [9], distributed dictionaries and location services [7], distributed storage [4], etc.

In many application areas, the networks in which these problems need to be solved are inherently dynamic, such as peer-to-peer overlays or mobile ad hoc networks. These networks are typically large and completely decentralized. As a result, each node can have an accurate view only of its local vicinity. Furthermore, the communication links may change unpredictably.

Given that deterministic algorithms may be very costly to implement in such dynamic environments, research has also addressed randomized approaches that give probabilistic guarantees. Probabilistic quorum system is not fixed a-priori, but is rather picked in a probabilistic manner for each interaction. By weakening the intersection guarantee to being probabilistic, these quorum systems are shown to be more scalable and efficient [12] and, therefore, more appropriate for dynamic environments.

Unfortunately, most existing probabilistic quorum systems are only guaranteed to work in static topologies. The protocols are usually formally analyzed in static graph models (e.g., random geometric graphs for wireless networks, random graphs for unstructured overlays, etc.). On the other hand, the simulations usually consider limited topology changes. For instance, Random Waypoint and Manhattan mobility models are commonly used to simulate wireless networks. These models are generally unable to generate worst-case network behavior. However, we believe that the correctness of the fundamental network protocols must be guaranteed for any possible scenario.

The problem of fault tolerance with regard to node failures has been widely studied and the bounds on the number of possible node failures tolerated by probabilistic quorum systems are well known [12]. However, if the nodes are not able to access the quorums correctly due to topology changes, the quorum systems can become unavailable even if no node crashes occur.

This paper addresses the need for providing results for networks with high and unpredictable dynamism. We model the network as a set of nodes that operate in synchronous rounds and the communication links change arbitrarily in each round according to the strategy chosen by an oblivious adversary. Therefore, in these conditions we cannot rely on routing.

In [8], it was shown that random walk is an efficient scheme to access probabilistic quorums for static networks. The random walk makes a necessary number of steps to reach a uniform distribution (i.e., there is an equal probability of being at any particular vertex), and then returns a sample to a client wishing to access a quorum. Furthermore, random walks do not rely on underlying routing protocols and therefore are suitable for our model.

Recent works have shown that asymptotic results for statics graph models differ significantly from those for dynamic

WRAS'10, July 29, 2010, Zurich, Switzerland.

graphs [5, 10]. For instance, it is well known that in undirected graphs the worst-case cover time of a simple random walk is $\mathcal{O}(n^3)$[3], where n is the size of the network. On the other hand, [5] has shown that simple random walks on undirected dynamic graphs may have exponential cover times. Therefore, the existing results for sampling in static networks do not apply to dynamic scenarios. In [5], it was shown that the cover time of maximum-degree random walk on (non-bipartite) dynamic graph is $\mathcal{O}(d_{max}^2 n^3 \log n)$, where d_{max} is the maximum degree of the graph.

In this paper we study the communication complexity of accessing a probabilistic quorum by random walks in a dynamic graph with oblivious adversary. In order to guarantee uniform sampling, the random walk must converge to the uniform distribution before picking a sample. We use a random walk strategy proposed in [6] and show its convergence to the uniform stationary distribution even in such a dynamic setting. The time the random walk takes to converge to the stationary distribution is called *mixing time*. We show that in the present model $\mathcal{O}(d_{max} n^2 \log n)$ communication steps are needed to access a probabilistic quorum. It is known that the worst-case mixing time for a simple walk on a static graph is $\mathcal{O}(n^3)$. We show that by increasing to $\mathcal{O}(d_{max} n^2 \log n)$ the time the random walk is required to run before picking a sample, the availability of the quorum system is guaranteed in a much wider range of scenarios covered by a dynamic graph model.

The rest of the paper is organized as follows. Section 2 introduces some basic concepts and describes the model we use for our results. In Section 3 we make a brief overview of the related work. Section 4 discusses the problem of accessing probabilistic quorums by a random walk over dynamic graphs and establishes an upper bound on access time. Finally, Section 5 presents the conclusions and outlines the directions for future work.

2. MODEL

Probabilistic Quorum System

Let $\mathcal{Q}$ be a set system defined on all nodes in the network, i.e., $Q \in \mathcal{Q}$ is a subset of all the nodes in the system.

Let ω be an access strategy for $\mathcal{Q}$. An access strategy defines the rule used to pick a set Q from $\mathcal{Q}$.

Finally, let $0 < \epsilon < 1$ be given. The tuple $\langle \mathcal{Q}, \omega \rangle$ is an ϵ-intersecting quorum system if $P(Q \cap Q' = \emptyset) = 1 - \epsilon$, where the probability is taken with respect to the strategy ω.

Communication Model

We model a network as a fixed set of nodes that operate in synchronous rounds. The edges represent bidirectional communication links from one node to another. Let $V = \{v_1, v_2, \cdots, v_n\}$ be a set of vertices. We assume that the nodes only know their neighbors and have an estimate of a network size. This estimate can be computed in distributed manner without relying on routing [10].

$\mathcal{G} = G_1, G_2, \cdots$ denotes a sequence of undirected graphs where G_t is a static undirected graph on V in round t. Further, we consider oblivious adversary that can make arbitrary changes in every round as long as the graph remains connected in each round. We say $\mathcal{G}$ is connected if each G_t is connected. This model captures the dynamic nature of connected wireless and overlay networks. It can be also used to model static networks where packet losses may occur.

3. RELATED WORK

The use of probabilistic quorums for implementing fault-tolerant distributed systems has been discussed in [12]. Subsequently, a variety of probabilistic quorum system implementations has been proposed for different environments, such as ad-hoc networks [11, 1], peer-to-peer overlays [13].

The main schemes of accessing probabilistic quorums are through direct sampling, view-based or by flooding [8].

In [6], the authors propose a maximum-degree random walk as a method of constructing uniform partial views. Periodically, nodes send maximum-degree random walks with TTL equal to network mixing time (the authors use TTL of $\mathcal{O}(n)$). However, the given analytic results are only valid for static graphs. If we assume the existence of topology changes, the TTL of $\mathcal{O}(n)$ is not sufficient to provide uniform network samples.

PAN [11] is a well known probabilistic quorum system designed for ad hoc networks. PAN uses access strategies based on a membership service and assumes the availability of an underlying routing protocol. However, in our model, since communication links change unpredictably, no routing is available.

In [1], the authors assume that the underlying communication graph is complete and construct an incident overlay routing graph. This approach serves as a dynamic membership management scheme. The routing graph is constructed based on node identifiers. This approach is not comparable to ours since in our model the network is multihop; therefore, nodes cannot create links to arbitrarily chosen nodes but only to their neighbors (the set of available communication links is defined by the oblivious adversary).

In [8], the authors studied combinations of the above access schemes for probabilistic quorums. The network topology is modeled by a random geometric graph. Again, the given asymptotic results on communication complexity only apply to static topologies. The authors also complete their results with extensive simulations by using Random Waypoint mobility model. However, in our work we require correctness in the presence of unpredictable topology changes.

4. QUORUM CONSTRUCTION

We use the original quorum construction strategy presented in [12]. Given a universe of n servers, the quorums are all sets of size $|Q| = l\sqrt{n}$. The constant l is chosen to make ϵ sufficiently small.

The access strategy ω is defined as follows. Q is selected with uniform probability. As was shown in [12], the probability that some element appears in two randomly chosen quorums is least $1 - e^{-l^2}$.

4.1 Direct Sampling

As discussed above, in a dynamic graph with an oblivious adversary no routing is available as the graph changes arbitrarily within each round. For that reason only the mechanisms that do not rely on routing can be used in this model.

We define the quorum access strategy ω as follows. To choose a Q uniformly at random we use direct sampling approach. When the node wishes to access a quorum, it chooses $|Q|$ nodes uniformly at random from the set of all nodes V. Each node $q \in Q$ is sampled by using a different random walk, as described below. Therefore, to access a quorum, a node initiates at least $|Q|$ random walks.

To provide the uniform sampling we use the following random walk strategy, similar to [6]. Let d_{max} be the maximum degree of graph G_t. We define a *maximum-degree self-biased random walk* as a random walk with an added self-loop where the weight of the self-loop is $k = 1 - \frac{d}{d_{max}}$, where d is node degree, and the transition probability of $\frac{1}{d_{max}}$ is assigned to each outgoing edge. If d_{max} is not known, we can take $d_{max} = n$.

Adding self-loops to a random walk can be seen as augmenting the graph with a self-loop edge at each vertex. In this way, we assure that the resulting graph in each round is non-bipartite. Furthermore, the maximum-degree strategy is used to provide a uniform probability distribution over the network nodes.

As mentioned above, the sampling must be done to guarantee the uniform probability distribution over the nodes. Every uniform sample is obtained by a single random walk whose length equals the network mixing time, i.e. the number of steps the random walk takes to reach the uniform distribution. However, different random walk can pick the same node as a sample. One way to solve this problem is by starting more random walks; it was shown in [12], $\mathcal{O}(n)$ random walks are sufficient to get $|Q|$ uniform sample w.h.p. Alternatively, the random walk arriving at an already sampled node can restart the sampling process.

4.2 Access Times

In the following, we give the asymptotic results on access times of probabilistic quorum through direct sampling. We first calculate the estimated time the random walk takes to reach a uniform distribution. After reaching the uniform distribution, the random walk picks a sample and continues execution until reaching the source node.

Let, in each round t, A_{G_t} be transition probability matrix of a random walk on G_t and $P_t = (p_1, p_2, \cdots, p_n)$ be a probability distribution on vertices. The transition at each step can be expressed in the following way.

$$P_{t+1} = P_t A_{G_t} = \begin{pmatrix} p_1, & \cdots, & p_n \end{pmatrix} \begin{pmatrix} a_{1,1} & \cdots & a_{1,n} \\ \vdots & \ddots & \vdots \\ a_{n,1} & \cdots & a_{n,n} \end{pmatrix} \quad (1)$$

From (1), it is easy to see that a random walk on a dynamic graph can be seen as a stochastic process that holds Markov property, i.e., each transition of the random walk only depends on its current position and the transition probabilities at a given step.

In the proof of our results, we use the spectral theory techniques (similar to [5]). An important observation is that due to the maximum-degree self-biased strategy, the transition probability matrices A_{G_t} in each round are symmetric and doubly stochastic (every row sums to one and every column sums to one). Therefore, each A_{G_t} will have a left eigenvector $\frac{\mathbb{I}}{n} = (\frac{1}{n}, \frac{1}{n}, \cdots)$ corresponding to an eigenvalue $\lambda = 1$ and all the remainder eigenvalues are real. As a result, $\pi = \frac{\mathbb{I}}{n}$ is a stationary distribution of the given random walk.

In the proof of our result, we will use the following lemmas. Lemma 4.1 bounds the absolute values of eigenvalues of A_{G_t}. Lemma 4.2 bounds the convergence rate of the probability distribution to the uniform distribution.

LEMMA 4.1. *Let A_{G_t} be the transition probability matrix of the maximum-degree self-biased random walk on an undirected graph with n vertices. If $|\lambda_1| \geq \cdots \geq |\lambda_n|$ are the left eigenvalues of A_{G_t}, then*

$$\lambda_1 = 1 \textit{ and } |\lambda_i| \leq \left(1 - \frac{1}{n^2 d_{max}}\right), \textit{ for } i \geq 2$$

PROOF. Let λ_* denote the spectral gap (i.e. the difference between the largest and the second largest eigenvalues in absolute value) of A_{G_t}. As shown in [14],

$$\lambda_* \geq \frac{1}{n^2 d_{max}}$$

Let $\alpha_1, \cdots, \alpha_n$ be left eigenvectors with corresponding eigenvalues $|\lambda_1| > \cdots > |\lambda_n|$ of A_{G_t}. As A_{G_t} is a doubly stochastic matrix, and $\alpha_1 = \frac{\mathbb{I}}{n} = \left(\frac{1}{n}, \cdots, \frac{1}{n}\right)$ corresponding to the eigenvalue $\lambda_1 = 1$. Moreover, as the random walk is self-biased, A_{G_t} is aperiodic. Therefore, we have

$$\lambda_1 > |\lambda_i|, \text{ for } i \geq 2$$

Thus,

$$\lambda_* = 1 - |\lambda_i| \geq \frac{1}{n^2 d_{max}}, \text{ for } i \geq 2$$

□

LEMMA 4.2. *Let G_t be a connected graph on V and $P_t = (p_1, \cdots, p_n)$ be a probability distribution on its vertices in round t, for $t \geq 1$. Let A_{G_t} be a transition probability matrix of a maximum-degree self-biased random walk on G_t. Then:*

$$\left\| P_{t+1} - \frac{\mathbb{I}}{n} \right\|_2^2 \leq \left(1 - \frac{1}{n^2 d_{max}}\right)^t$$

PROOF. Let $A = \{\alpha_1, \cdots, \alpha_n\}$ be an orthonormal set of left eigenvectors of A_{G_t} with corresponding eigenvalues $\lambda_1, \cdots, \lambda_n \in \mathbb{R}$ ordered by their absolute values so that $|\lambda_1| \geq \cdots \geq |\lambda_n|$. As A_{G_t} is a doubly stochastic matrix,

$$\begin{aligned} \left\| P_{t+1} - \frac{\mathbb{I}}{n} \right\|_2^2 &= \left\| P_t A_{G_t} - \frac{\mathbb{I}}{n} \right\|_2^2 \\ &= \left\| P_t A_{G_t} - \frac{\mathbb{I}}{n} A_{G_t} \right\|_2^2 \\ &= \left\| \left(P_t - \frac{\mathbb{I}}{n}\right) A_{G_t} \right\|_2^2 \end{aligned} \tag{2}$$

Since $A = \{\alpha_1, \cdots, \alpha_n\}$ is an orthonormal system in $\mathbb{R}^n$ and $\left(P_t - \frac{\mathbb{I}}{n}\right)$ is orthogonal to $\alpha_1 = \frac{\mathbb{I}}{\sqrt{n}}$, there exists some $B = \{\beta_1, \cdots, \beta_n\} \in \mathbb{R}^n$ such that

$$P_t - \frac{\mathbb{I}}{n} = \sum_{i=2}^n \beta_i \alpha_i \tag{3}$$

By standard calculation we have,

$$\begin{aligned} \left\| P_t - \frac{\mathbb{I}}{n} \right\|_2^2 &= \langle P_t - \frac{\mathbb{I}}{n}, P_t - \frac{\mathbb{I}}{n} \rangle \\ &= \langle \sum_{i=2}^n \beta_i \alpha_i, \sum_{i=2}^n \beta_i \alpha_i \rangle \\ &= \sum_{i=2}^n |\beta_i|^2 \end{aligned} \tag{4}$$

On the other hand,

$$\begin{aligned} \left\| \left(P_t - \frac{\mathbb{I}}{n}\right) A_{G_t} \right\|_2^2 &= \left\| \sum_{i=2}^n \beta_i \alpha_i A_{G_t} \right\|_2^2 \\ &= \left\| \sum_{i=2}^n \lambda_i \beta_i \alpha_i \right\|_2^2 \\ &\leq \sum_{i=2}^n |\lambda_i|^2 |\beta_i|^2 \end{aligned} \tag{5}$$

The last inequality follows from Cauchy-Schwartz bound.

By Lemma 4.1,

$$|\lambda_i|^2 \leq |\lambda_i| \leq \left(1 - \frac{1}{n^2 d_{max}}\right), \text{ for } i \geq 2$$

Thus,

$$\left\| \left(P_t - \frac{\mathbb{I}}{n}\right) A_{G_t} \right\|_2^2 \leq \left(1 - \frac{1}{n^2 d_{max}}\right) \sum_{i=2}^n |\beta_i|^2 \tag{6}$$

From (2) , (4) and (6) we have,

$$\begin{aligned} \left\| P_{t+1} - \frac{\mathbb{I}}{n} \right\|_2^2 &= \left\| \left(P_t - \frac{\mathbb{I}}{n}\right) A_{G_t} \right\|_2^2 \\ &\leq \left(1 - \frac{1}{n^2 d_{max}}\right) \left\| P_t - \frac{\mathbb{I}}{n} \right\|_2^2 \\ &\leq \left(1 - \frac{1}{n^2 d_{max}}\right)^t \left\| P_1 - \frac{\mathbb{I}}{n} \right\|_2^2 \end{aligned} \tag{7}$$

□

The following theorem gives an upper bound on mixing time of a random walk on $\mathcal{G}$. In our proof we use the definition of a mixing time given in [2]:

$$t_{mix}(v_i) = min\{t : \left\| P_t - \pi \right\|_{TV} < \frac{1}{4}\}$$

THEOREM 4.3. *Let $\mathcal{G} = G_1, G_2, \cdots$ be a connected undirected dynamic graph with maximum degree d_{max}.*

If $t_{mix}(v_i) = min\{t : \left\| P_t - \pi \right\|_{TV} < \frac{1}{4}\}$ then

$$t_{mix}(v_i) = \mathcal{O}(d_{max} n^2 \ln n)$$

PROOF. $\pi = \frac{\mathbb{I}}{n}$ is a uniform stationary distribution for a maximum-degree self-biased random walk on $\mathcal{G}$.

By Lemma 4.2,

$$\left\| P_{t+1} - \frac{\mathbb{I}}{n} \right\|_2^2 \leq \left(1 - \frac{1}{d_{max} n^2}\right)^t \tag{8}$$

Hence, for $t' = 2d_{max} n^2 \ln 2n + 1$ we have

$$\begin{aligned} \left\| P_{t'} - \frac{\mathbb{I}}{n} \right\|_2^2 &\leq \left(1 - \frac{1}{d_{max} n^2}\right)^{2d_{max} n^2 \ln 2n} \\ &< \frac{1}{4n^2} \end{aligned} \tag{9}$$

Let $P_{t'} - \frac{\mathbb{I}}{n} = (p_1, \cdots, p_n)$.

From (9), $|p_i| \leq \frac{1}{2n}, \forall_{1 \leq i \leq n}$. Thus,

$$\begin{aligned} \left\| P_{t'} - \frac{\mathbb{I}}{n} \right\|_{TV} &= \frac{1}{2} \sum_{i=1}^{n} |p_i| \\ &< \frac{1}{4} \end{aligned} \tag{10}$$

Therefore,

$$t_{mix}(v_i) \leq d_{max} n^2 \ln 2n + 1$$

□

The following theorem establishes an upper bound on a hitting time of any vertex of $\mathcal{G}$ by the maximum-degree self-biased random walk starting from a uniform distribution.

THEOREM 4.4. *Let $Y_1, Y_2, \cdots$ be a random walk on a connected undirected dynamic graph $\mathcal{G} = G_1, G_2, \cdots$ and $\pi = \frac{\mathbb{I}}{n}$ initial distribution on V.*

Let $t_{hit_\pi}(v_i) = min\{t : v_i \in \{Y_1, Y_2, \cdots, Y_t\}\}$, then

$$E\left[t_{hit_\pi}(v_i)\right] \leq \mathcal{O}(n^2 \ln n)$$

.

PROOF. $Y_1, Y_2, Y_3, \cdots$ is a stochastic process, with possible states $V = \{v_1, \cdots, v_n\}$, holding the Markov property.

By Law of total probability,

$$P(Y_{t+1} = v_i) = \sum_{j=1}^{n} P(Y_{t+1} = v_i | Y_t = v_j) P(Y_t = v_j) \tag{11}$$

π is a stationary distribution of the given random walk. As in each round t, $P(Y_t = v_j) = \frac{1}{n}, \forall v_j \in V$ and v_i has at least one neighbor v_j, the probability of arriving at v_i in step $t+1$ is given as following

$$P(Y_{t+1} = v_i | Y_t = v_j) P(Y_t = v_j) \geq \frac{1}{n^2}, \forall t > 0 \tag{12}$$

Therefore,

$$\begin{aligned} P(v_i \notin \{Y_1, \cdots, Y_{n^2 \ln n}\}) &\leq \left(1 - \frac{1}{n^2}\right)^{n^2 \ln n} \\ &< \frac{1}{n} \end{aligned} \tag{13}$$

Therefore,

$$E\left[t_{hit_\pi}(v_i)\right] = \mathcal{O}(n^2 \ln n)$$

. □

The following theorem gives an upper bound on the time a random walk takes to get a random sample and return to a starting node.

THEOREM 4.5. *Let $\mathcal{G} = G_1, G_2, \cdots$ be a connected undirected dynamic graph with maximum degree d_{max}.*

Let $Y_t = Y_1, Y_2, \cdots$ be a maximum-degree self-biased random walk on $\mathcal{G}$. The time Y_t takes, starting at any vertex $v_i \in V$ to reach a uniform distribution and return to a starting node is $\mathcal{O}\left(d_{max} n^2 \ln n\right)$.

PROOF. Let $t_{access}(v_i)$ be the time the maximum-degree self-biased random walk starting at any vertex $v_i \in V$, takes to reach a uniform distribution and return to a starting node

$$t_{access}(v_i) = min\{t > t_{mix} : Y_t = v_i\}$$

$t_{access}(v_i)$ is a stopping time and

$$t_{access}(v_i) = t_{mix}(v_i) + t_{hit_\pi}(v_i)$$

From Theorem 4.3 and 4.4 we have the following,

$$\begin{aligned} E\left[t_{access}(v_i)\right] &= t_{mix}(v_i) + E\left[t_{hit_\pi}(v_i)\right] \\ &= \mathcal{O}\left(d_{max} n^2 \ln n + n^2 \ln n\right) \end{aligned} \tag{14}$$

□

From Theorem 4.5, we conclude that the random walk must be initiated with TTL $= \mathcal{O}\left(d_{max} n^2 \ln n\right)$ to pick a uniform sample. As $\mathcal{O}(n)$ random walks are needed to get $|Q|$ samples, the total message cost of accessing the quorum by a random walk strategy is $\mathcal{O}(d_{max} n^3 \ln n)$. However, it is important to observe that in a maximum-degree self-biased random walk a considerable probability weight is assigned to a self-loop; therefore, in practice, many transitions will not generate messages. As for memory usage, a node wishing to access a quorum only needs to store $\mathcal{O}(|Q|)$ node identifiers.

4.3 Other access schemes

In the following, we make a brief discussion on other alternative schemes of getting a uniform network sample of size $|Q|$ in a dynamic network where no routing is available. However, we defer for the future work a thorough analysis of all the possible techniques of accessing random quorums in dynamic graphs with an oblivious adversary.

View-based

The nodes can also construct random views. However, as no routing is available, each node in a view must be accessed either through network-wide flooding or by a random walk; therefore intuitively there is no advantage in constructing views if no routing is available.

Flooding

A network-wide flooding can be used to get the ids of all nodes in the network and choose a quorum uniformly at random. However, flooding itself can be seen as an alternative to quorum systems. The cost of flooding in dynamic networks is $\mathcal{O}(n)$ in terms of communication steps and $\mathcal{O}(d_{max}n^2)$ as for the number of messages, i.e. every node after receiving a flooded message must retransmit it for $\mathcal{O}(n)$ consecutive rounds in order to assure that all the nodes receive the message [10]. As a result, the total cost of receiving all the ids is $\mathcal{O}(2n)$ communication steps and $\mathcal{O}(d_{max}n^3)$ messages. Additionally, the nodes must have the capacity to store $\mathcal{O}(n)$ node identifiers.

5. CONCLUSIONS

In this paper we have addressed the problem of implementing access strategies for probabilistic quorums over dynamic graphs where the oblivious adversary may change the communication links arbitrarily in each round. Due to the absence of efficient routing over dynamic networks with arbitrary changes in every time step, the two most straightforward approaches are flooding and random walk. We studied the complexity of a self-biased maximum-degree random walk pick a uniform sample. We showed that in $\mathcal{O}\left(n^3 \ln n\right)$ communication steps a random walk gets a uniform sample and returns to a source node. We also showed that, in dynamic networks, the random walk is asymptotically comparable to flooding (an analogous phenomenon is observed in the static case).

As a future work we would like to make an in depth analysis of other access strategies of random quorums in dynamic graphs. Other interesting research directions are studying the complexity of solving consensus and mutual exclusion in dynamic graphs with communication links changing arbitrarily in each communication step.

Acknowledgments.

The authors are grateful to Paolo Romano for his comments on earlier drafts of this paper. This work was partially supported by FCT through PIDDAC Program funds via the INESC-ID multiannual funding and the project "Redico" (PTDC/EIA/71752/2006).

6. REFERENCES

[1] Ittai Abraham and Dahlia Malkhi. Probabilistic quorums for dynamic systems. *Distrib. Comput.*, 18(2):113–124, 2005.

[2] David Aldous and James Allen Fill. *Reversible Markov Chains and Random Walks on Graphs.* unpublished, 1995.

[3] Romas Aleliunas, Richard M. Karp, Richard J. Lipton, László Lovász, and Charles Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *FOCS '79: 20th Annual Symposium on Foundations of Computer Science*, pages 218–223, 1979.

[4] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. *J. ACM*, 42(1):124–142, 1995.

[5] Chen Avin, Michal Koucký, and Zvi Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 121–132, Berlin, Heidelberg, 2008. Springer-Verlag.

[6] Ziv Bar-Yossef, Roy Friedman, and Gabriel Kliot. Rawms - random walk based lightweight membership service for wireless ad hoc networks. *ACM Trans. Comput. Syst.*, 26(2):1–66, 2008.

[7] Roy Friedman and Gabriel Kliot. Location services in wireless ad hoc and hybrid networks: A survey. Technical report, Technion Computer Science, 2006.

[8] Roy Friedman, Gabriel Kliot, and Chen Avin. Probabilistic quorum systems in wireless ad hoc networks. In *In Proceedings of the 38th IEEE International Conference on Dependable Systems and Networks (DSN-DCCS)*, 2008.

[9] Rachid Guerraoui and Michel Raynal. The information structure of indulgent consensus. *IEEE Trans. Comput.*, 53(4):453–466, 2004.

[10] Fabian Kuhn, Nancy Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In *STOC '10: Proceedings of the 42nd ACM Symposium on Theory of Computing*, page to appear, Cambridge, Massachusetts, USA, 2010. ACM.

[11] Jun Luo, Jean-Pierre Hubaux, and Patrick Th. Eugster. Pan: providing reliable storage in mobile ad hoc networks with probabilistic quorum systems. In *MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing*, pages 1–12, New York, NY, USA, 2003. ACM.

[12] Dahlia Malkhi, Michael Reiter, and Rebecca Wright. Probabilistic quorum systems. In *PODC '97: Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing*, pages 267–273, New York, NY, USA, 1997. ACM.

[13] Ken Miura, Taro Tagawa, and Hirotsugu Kakugawa. A quorum-based protocol for searching objects in peer-to-peer networks. *IEEE Trans. Parallel Distrib. Syst.*, 17(1):25–37, 2006.

[14] Ravi Montenegro. The simple random walk and max-degree walk on a directed graph. *Random Struct. Algorithms*, 34(3):395–407, 2009.