P2P Cache-and-Forward Mechanisms for Mobile Ad Hoc Networks

Claudio Casetti, Carla-Fabiana Chiasserini, Marco Fiore
Politecnico di Torino, Italy

Chi-Anh La, Pietro Michiardi
EURECOM, France

***Abstract*—We investigate the problem of spreading information contents in a wireless ad hoc network. In our vision, information dissemination should satisfy the following requirements: (i) it should result in a desirable distribution of information replicas in the network and (ii) the information should be evenly and fairly carried by all nodes in their turn. In this paper, we show that these goals can be achieved by simple cache-and-forward mechanisms inspired by well-known node mobility models, provided that a sufficient number of information replicas are injected into the network. The proposed approach works under different network scenarios, is fully distributed and comes at a very low cost in terms of protocol overhead.**

I. Introduction

It is commonly acknowledged that portable user devices are rapidly becoming tantamount to a communication hub, sporting arrays of GPS navigators, multiple wireless interfaces and web-based applications. Such devices are capable of carrying latest news, traffic alerts or local sightseeing information. In this context, most pieces of information are likely to be of general use, and therefore a sensible dissemination and caching policy would be desirable.

In this work, we focus on such an environment: few and far between access points, or gateway nodes, in a highly-populated network area where user devices are equipped with a data cache and communicate according to an ad hoc networking paradigm. We assume that users create a cooperative environment where information is exchanged among nodes in a peer-to-peer fashion. The nodes storing an information copy are supposed to act as *providers* for this content to nearby nodes. Nodes act as providers for a limited time, before handing over the information to other nodes.

We then try to answer the following questions:

- Regardless of how the information is distributed at the outset, can simple cache-and-forward mechanisms achieve a target information distribution?
- Given such cache-and-forward mechanisms, do nodes evenly share the role of provider? And, are they equally burdened when they take on the role of provider?

Traditional approaches to information caching in communication networks [1]–[3] are based on the solution of linear programming problems, which often require global knowledge of the network condition, or lead to quite complex solutions that involve significant communication overhead. Unlike previous approaches, we propose and analyze a solution addressing the above issues that is fully distributed and is information-oblivious, i.e., not requiring knowledge of the content stored by users.

In particular, motivated by the need of a balanced load distribution among the provider nodes and of an equal quality of service provisioning to the users, we target a uniform distribution of contents, either over the network spatial area or over the network nodes. With this aim in mind, in Section III we investigate the applicability of two cache-and-forward mechanisms inspired by well-known node mobility models to disseminate information across the network. Both strategies, using the simulation setup in Section IV, are proven to yield a distribution of the information copies that is close to the target distribution, regardless of the considered network scenario (Section V). Also, the obtained results show that the level of fairness in distributing the burden among provider nodes depends on the number of information copies stored in the network.

II. Related work

Our study is related to the problem of optimal cache placement in wireless networks. Several works have addressed this issue by exploiting its similarity to the facility location and the k-median problems. Both these problems are NP-hard and a number of constant-factor approximation algorithms have been proposed for each of them [1], [4]; these algorithms however are not amenable to an efficient distributed implementation.

Distributed algorithms for allocation of information replicas are proposed, among others, in [3], [5]–[7]. These solutions typically involve significant communication overhead, especially when applied to mobile environments, and focus on minimizing the information access cost or the query delay. In our work, instead, we consider a cooperative environment and aim at a uniform distribution of the information copies, while evenly distributing the load among the nodes acting as providers.

Relevant to our study is also the work in [8], which computes the (near) optimal number of replicas of video clips in wireless networks, based on the bandwidth required for clip display and their access statistics. However, the strategy proposed in [8] requires a centralized implementation and applies only to strip or grid topologies.

In the context of sensor networks, approaches based on active queries following a trajectory through the network, or agents propagating information on local events have been proposed, respectively, in [9] and [10]. Note that both these works focus on the forwarding of these messages through the network, while our scope is to make the desired information available by letting it move through nodes caches.

III. Achieving the desired information distribution

We start by addressing the problem of where the information copies should be cached in the network so as to obtain the desired content distribution.

We consider a tagged[1] information content and we target two desirable distributions of information: the first uniform over the spatial area covered by the network (*spatial uniformity*), the second following the layout of the network topology (*nodal uniformity*). Spatial uniformity is motivated by the need to guarantee equal access to the information over the whole service area, while nodal uniformity allows the information density to match the node density and, therefore, to cluster the information where the demand is higher.

To achieve the target distribution, we let the information move across nodes according to two well-known mobility models, namely the random walk and the random direction [11] models, which are often used to represent the movement of user nodes in ad hoc networks. In our context, a mobile entity is not a network node but, rather, a copy of the tagged information which "hops" from a user node that just stopped being a provider for that information onto another node which will become the new content provider. We apply the two mobility models and develop the dissemination strategies detailed below.

The random walk dissemination (RWD) strategy. We consider the simplest random walk possible, in which each mobile entity, i.e., each copy of the information content, roams the network by moving from a node to a one-hop neighbor selected with equal probability. Each node caches the information for a fixed amount of time, and then hands it over to the next selected hop in the information copy visit pattern. This approach requires trivial node operations and introduces low overhead.

The random direction dissemination (RDD) strategy. Each mobile entity alternates periods of movement (move phase) to periods during which it pauses (pause phase). In our context, the pause phase corresponds to the time period during which the information copy is cached at a provider node. The move phase starts at the time instant when the current information provider hands over the content to one of its one-hop neighbors, and it ends when the new provider is reached by the information copy. The new provider is identified by first selecting a target location: the closest node to that location becomes the new provider. To this end, at the beginning of a move phase, the current provider independently selects the direction and the distance[2] for the movement of the information content, thus identifying a target location whose position is included in the content messages. We introduce a simple broadcast-based application-level routing scheme that allows information to be moved towards the target location, with each forwarder selecting as a next hop the neighbor that best fits the ideal trajectory designed by the original provider. The neighbor selection process is performed in a reactive manner, as it involves an exchange of advertisement (by the forwarder) and reply (by candidate next hop neighbors) messages at each movement hop. When a node has no neighbors closer than itself to the target position, it elects itself as the new provider, and the pause phase starts again. Some remarks are in order. First, this scheme requires nodes to be capable to estimate their position (i.e., through GPS), a fair assumption in most practical scenarios. Second, the information moves across user nodes, thus it may be transmitted along a direction that just approximates the planned trajectory, or it may be stored at a node that is nearby (but not exactly at) the selected geographical destination. Third, geographical areas devoid of nodes that can support the information movement may be encountered during move phases: in that case, the current forwarder assumes a *boundary* has been hit, and applies a reflection to the movement angle.

The characterization of the spatial distribution of randomized algorithms applied to node mobility has been investigated in the literature from an analytic perspective, in an ideal setting. If the network topology can be represented as an undirected, connected, non-bipartite graph, then the distribution of nodes moving according to the random walk model converges to a unique stationary distribution regardless of the initial distribution, and this stationary distribution is uniform over nodes in the case of regular graphs[3] [12]. As for the random direction model, in [11] it has been shown that, if at time $t = 0$ the position and the orientation of mobile nodes are independent and uniform over a finite square area, they remain uniformly distributed over the area for all time instants $t > 0$, provided that the entities move independently of each other.

In the context of this work, we cannot trivially use techniques similar to [11], [12] and show that the same randomized algorithms applied to information instead of mobile nodes achieve a uniform distribution. The dissemination mechanisms we apply to information operate on network deployments that do not have a regular structure, hence the results for the random walk model do not directly apply. Furthermore, the combined effects of node mobility and information mobility make the analysis harder, especially for the RDD strategy where the information only approximately reaches its geographical destination. Lastly, in this work we are interested in both spatial and nodal uniformity, and for the latter we are not aware of any previous studies that prove convergence to a uniform distribution in a general scenario. Therefore, in the following, we carry out a thorough simulation campaign to investigate the actual distribution of the information that is obtained through our approach and its distance from the target uniform distribution.

[1] I.e., we assume information content to be uniquely identifiable.

[2] Note that randomly selecting a travel distance is equivalent to randomly selecting speed and travel time.

[3] A graph is regular if each of its vertices has the same number of neighbors.

IV. Simulation set-up and methodology

In this work we use the *ns*-2 network simulator, where all nodes are equipped with standard 802.11b interfaces, with 11 Mbps fixed data transmission rate. To evaluate the behavior of the cache-and-forward strategies discussed in Section III, we implemented a simple application that allows nodes to query providers through limited-scope flooding. Queries can traverse a maximum number of hops, $h_{max} = 5$, before being discarded [4]. We improve the query propagation process by adopting the PGB technique [13] to select forwarding nodes that relay queries to their destinations. Sequence numbers are used to detect and discard duplicate queries and avoid the *broadcast storm* [14] phenomenon. Upon reception of a query, a provider replies with a probability that is inversely proportional to the number of hops traversed by the query message. This is done to further mitigate the overhead of any duplicate query that would reach multiple providers.

In the following, we define the simulation settings we consider in this work. Note that all results presented in the remainder of this paper are averaged out over 10 simulation runs, each with a randomized selection of initial information providers. Simulation time is set to 10,000 seconds, unless specified otherwise. Moreover, we assume a network composed of $N = 2000$ nodes that are spatially distributed on a square area $\mathcal{A}$ of 500 m side. Each node has a transmission range of 20 m. When employing the RDD scheme, providers characterize the information move phase by randomly choosing angles that are uniformly distributed in $[0, 2\pi]$, and exponentially distributed distances, with mean equal to 100 m. We study both *static* and *mobile* cases, as will be detailed below.

A. Static node spatial distribution

We define the following static node deployments:

- *Uniform distribution*: nodes are uniformly placed on $\mathcal{A}$;
- *Clustered distribution*: we assume nodes to be deployed in four equally sized clusters. Each cluster corresponds to a "point of interest" around which nodes are located. Nodes are also placed in-between clusters so as to ensure network connectivity. In practice, we implement the random trip model as defined in [15] and take a snapshot of the network topology as our initial node distribution.

B. Node mobility

The impact of node mobility on the dissemination mechanisms we designed is analyzed for the following mobility models:

- *Stationary Random Waypoint*: nodes are initially deployed according to the stationary distribution of the node mobility model [16] (resulting in nodes being more often located towards the center of the network area; then, each node selects a random destination in $\mathcal{A}$ and moves towards it at a constant speed selected uniformly at random from a distribution with a mean of 3 m/s. The pause time is set to 10 s.
- *Random trip*: following the definition in [15], nodes revolve around four "points of interest". The initial node deployment conforms to the clustered distribution defined for the static case. The stationary random waypoint model defined above guides node movements inside a cluster. Inter-cluster mobility is allowed with probability 0.3.

C. Parameter space

We now define the parameters used in our evaluation, accounting for the initial distribution and number of information providers in the network, as well as the query behavior of mobile nodes.

- *Number of information providers*: at the beginning of each simulation run, a predefined number of providers $\mathcal{C}$ is randomly chosen among all nodes in the network; we choose $\mathcal{C} \in \{20, 50, 100, 200, 400\}$.
- *Information caching time*: when taking up the role of information provider, a node i keeps a local copy for a time τ_i. In this work, we assume $\tau_i = \tau \ \forall i$ with $i = 1, \ldots, N$. In the following we present results for $\tau \in \{10, 100\}$ seconds.
- *Information demand*: we assume nodes to issue queries to information providers using the simple application defined above. Without loss of generality, we focus on one information content (of size equal to 1KB) that is made available in the network. Users' demand for the available information is modeled through a query rate which we assume to be common to all users, $\lambda_i = \lambda = 0.0025$ req/s $\forall i$ with $i = 1, \ldots, N$. The aggregate query rate Λ over all nodes depends on the number of information providers active in the network[5], i.e., $\Lambda = (N - \mathcal{C})\lambda$.

D. Evaluation metrics

To understand to which extent the information distribution achieved by our dissemination techniques resembles the desired content diffusion, we look at the distributions of inter-distance between information copies. As discussed before, we consider the following two reference distributions:

- *Spatial uniformity:* since we consider a square area where nodes are deployed and we seek a uniform dissemination of content over the network area, the target distribution is the solution to the bidimensional case of the hypercube line picking problem [17], which is known to be:

$$q(x) = \begin{cases} 2x\left(x^2 - 4x + \pi\right) & \text{if } 0 \leq x < 1, \\ 2x\left[4\gamma - \left(x^2 + 2 - \pi\right) - 4\tan^{-1}\gamma\right] & \text{if } 1 \leq x < \sqrt{2}, \end{cases}$$

with $\gamma = \sqrt{x^2 - 1}$.

[4]The choice of $h_{max} = 5$ is arbitrary: queries can roughly propagate over half of the network diameter, given our settings.

[5]Indeed, providers do not issue requests to access the content

- *Nodal uniformity:* in order to test the uniformity of providers over the network nodes, we take as a reference distribution the empiric distribution of node inter-distances measured in simulation.

Using inter-distances instead of actual coordinates allows us to handle a much larger number of samples (e.g., $C(C-1)$ instead of just C samples), making results more accurate. To compare information and reference distributions, we will resort to a visual comparison of the PDFs, as well as to the χ^2 goodness-of-fit test.

Then, we provide a basic performance evaluation of the information query process achieved by our application, and focus on the following metrics.

- *Cumulative provider time:* we evaluate the load balancing properties of the different information dissemination strategies by computing the cumulative time $\hat{\tau}_i$ each node i spends as an information provider. Given that the cache time τ is deterministic, we can compute $\hat{\tau}_i = \tau \times \mathcal{I}_i$, where $\mathcal{I}_i$ accounts for the number of times node i takes up the role of information provider during the simulation time.
- *Served queries at each information provider*: we measure the cumulative number of served queries for each information provider. Note that this metric is also useful to understand the impact of the hop-based reply policy implemented by provider nodes (i.e., the likelihood of replies decreases with the increase of hops traversed by the query).
- *Euclidean distance to access information*: we measure the cumulative Euclidean distance from a node to its *closest* information provider, every τ seconds. The distance to access information is the result of the spatial distribution of information in the network and can be used to measure how "fair" our mechanisms are toward each querying node.

V. The Logistics of Information

In this section, we look at how the RWD and RDD strategies can achieve the two objectives outlined in the introduction of the paper: a desirable distribution of the information and a fair distribution of information burden across the provider nodes. In the set of results we present, no information drop is allowed; indeed, for both the RWD and RDD strategies, a provider that hands the information over to another node considers the transfer as successful only if it receives an acknowledgment message, otherwise it repeats the procedure by selecting a different neighbor. The duplication probability we obtained by implementing such an application was negligible (order of 10^{-5}). Thus, we can consider that the overall number of providers $\mathcal{C}$ does not change during the simulation time.

Information distribution. We now focus on the properties of the information distribution of our cache-and-forward strategies, that is, we study where information replicas are cached in the network. The following results are obtained for different static node deployments (uniform and clustered) and node mobility models (stationary random waypoint and random trip models) when $\mathcal{C} = 200$ and the caching time $\tau = 10$ s. Note that the probability density functions (PDF) we show hereafter are computed from samples collected over all the simulation time.

Both Fig. 1 and Fig. 2 indicate a target PDF corresponding to the two information distributions we take as reference, *spatial uniformity* and *nodal uniformity*, as defined in Sec. IV-D.

Fig. 1 (a) shows that both the RWD and the RDD strategies yield information distributions that closely overlap both the nodal and spatial uniformity targets. Indeed, if nodes are static and uniformly distributed on the network area, then information mobility can be thought of as equivalent to node mobility, with the constraint that information can only move in well-defined positions that are given by the original network deployment. Our simulation results, backed up by prior analytical studies [11], [12] on node mobility, indicate that information replicas achieve a uniform distribution, both in space and over the nodes.

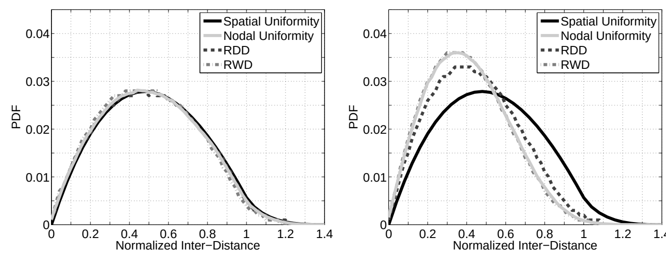

(a) Static scenario (b) Mobile scenario

Fig. 1. PDF of the inter-distance between information copies normalized to $\mathcal{A}$, for the RWD and RDD dissemination policies in static uniform and mobile scenarios (random waypoint) when $\mathcal{C} = 200$ and $\tau = 10$ s.

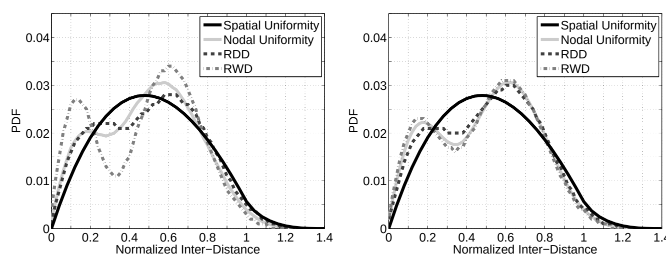

(a) Static scenario (b) Mobile scenario

Fig. 2. PDF of the inter-distance between information copies normalized to $\mathcal{A}$, for the RWD and RDD dissemination policies in static (*clustered*) and mobile (*random trip*) scenarios when $\mathcal{C} = 200$ and $\tau = 10$ s.

Fig. 1 (b) illustrates the implications of the combined effects of node and information mobility. In this case, the RWD cache-and-forward strategy approximates very well the nodal uniformity target, whereas spatial uniformity, represented now by a different distribution, is not achieved. The reason lies in the fact that, by moving the information of a single random hop at a time, the RWD scheme is strictly bound to the nodal distribution. The RWD distribution and the reference distribution for nodal uniformity exhibit a lower mean than the spatially uniform distribution: indeed, node mobility reduces the inter-distance between nodes, which are more likely to be moving around the center of the network area. When considering the RDD strategy, we observe that the information distribution it achieves falls in between the two reference distributions. As a matter of fact, the RDD strategy tends to a uniform distribution over space; however, the movement of carriers biases such distribution towards that shaped on the nodes layout.

The results for the clustered scenario are shown in Fig. 2. When nodes are static and distributed in clusters, there are some parts of the network area that are not home to any provider. Hence, it is reasonable that spatial uniformity cannot be achieved. In this case, having a cache storing information where there are no nodes to access it would serve little purpose. For this reason, in a clustered scenario, nodal uniformity appears to be a more sensible target and our results confirm the effectiveness of our cache-and-forward strategies to approximate a desirable distribution of information[6]. While Fig. 2(a) shows that the RDD policy is more accurate than the RWD strategy in approximating nodal uniformity, in Fig. 2(b) the two cache-and-forward schemes achieve similar results. Indeed, when nodes revolve around several points of interest and are free to move from one cluster to another, provider nodes can also be found in parts of the network area with a low node density.

In order to assess the impact of the number of information replicas, i.e., the number of providers, in the network, we focus on a single scenario, the static uniform one, where we already observed that the two reference distributions match (see Fig. 1 (a)). We use the time series of the χ^2 index computed considering the observed and the target PDF, with smaller values indicating a better fit. Fig. 3 also reports the average and standard deviation of the χ^2 index.

We note that an increase of 1 order of magnitude (from 20 to 200) in the number of providers, which implies more nodes bearing the cost of serving information, differently affects our mechanisms: Fig. 3 (a) indicates that for the RWD policy the mean χ^2 index improves by almost 4 times while Fig. 3 (b) shows a tenfold improvement for the RDD mechanism, although both schemes exhibit similar average values when the number of providers is low. It should be noted that an increased number of providers greatly helps in stabilizing the information distribution, as testified by the standard deviation of the χ^2 index in both schemes. Fig. 3 (b) also pinpoints an important property of the RDD strategy: a small standard deviation of the χ^2 index implies that at all times, the achieved distribution does not diverge too much from the target, which ensures a fair access to information by client nodes.

The results we presented in this section support the idea we advocate in our work: exploiting well-known mobility models to derive cache-and-forward mechanisms is indeed an efficient, light-weight alternative to complex (and centralized) techniques akin to facility location and k-median problems previously appeared in the literature. Provided that simple distributed schemes can achieve a desirable information distribution, we now move forward and examine the implications of our mechanisms from the perspective of provider and client nodes.

[6]It is not straightforward to conclude that our mechanisms are able to mimic nodal uniformity, and we are not aware of any theoretic studies in the literature that support our simulation results

Load balancing. We now turn our attention to the important question of load balancing across providers. For brevity, below we present just a subset of the results we derived. In particular, since the RDD manages to provide a better approximation to the target information distribution than RWD, we only show the performance of the RDD policy. Also, we present results only for the static uniform scenario and the mobile network with random waypoint mobility, since a similar performance is achieved under clustered network topologies.

In Fig. 4 we plot the complementary distribution function (CCDF) of the cumulative time a node is serving as an information provider (i.e., the provider time) over the whole duration of our simulations, that is, we normalize the provider time to the simulation time. The results are presented for the RDD policy in a static uniform scenario, for different values of the caching time τ and when the initial number of information providers sums to $\mathcal{C} = 20$ and to $\mathcal{C} = 200$ (which correspond, respectively, to 1% and 10% of the total number of nodes). Looking at the figure, we observe that when we increase $\mathcal{C}$ from 20 to 200, the load is spread more uniformly across the nodes since there is an increased opportunity for being (randomly) selected as information provider. The effect of an increased caching time τ from 10 s to 100 s, is, instead, a translation of the CCDF to higher values, without affecting the load distribution.

We now take a deeper look at the impact of different scenarios and simulation parameters on the effective load that an information provider supports in terms of number of served queries. Note that the number of served queries is not equal to the number of queries a provider receives because of the reply behavior described in Section IV. Figs. 5 a) and b) present the CCDF of the number of queries served by the provider nodes, respectively, when $\mathcal{C} = 20$ and $\mathcal{C} = 200$. Both the static uniform scenario and the mobile scenario with random waypoint mobility are considered. Looking at the plots, we note that an increased number of initial providers is effective in spreading the query load more evenly, especially in the static case. In the case of the static topology, when $\mathcal{C} = 20$, roughly 50% of providers never get a chance to satisfy a user request, whereas with $\mathcal{C} = 200$, about 60% of providers are serving a number of queries comprised in the interval $[70, 150]$. The combined effect of node mobility and an increased number of initial providers is striking: Fig. 5 b) indicates that approximately 95% of providers serve roughly the same amount of queries. Thus, node mobility, which at first could be considered harmful to information distribution mechanisms, turns out to be a good ally in terms of load balancing.

Clients' perspective. Lastly, we take the perspective of

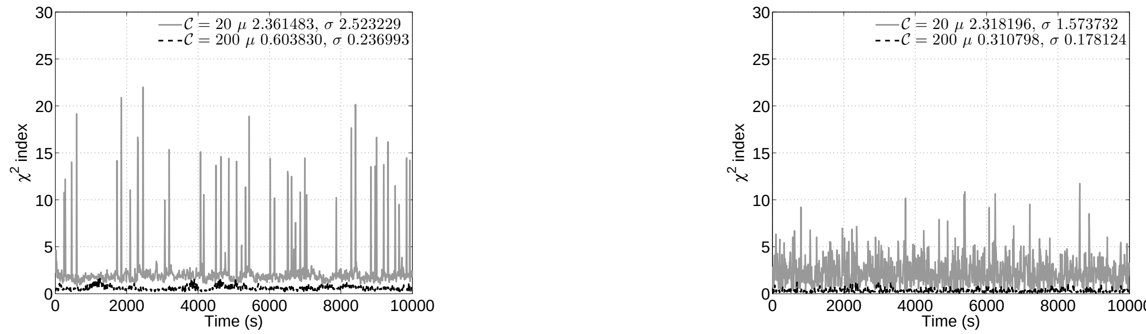

(a) RWD strategy

(b) RDD strategy

Fig. 3. χ^2 index for the static uniform scenario: mean (μ) and standard deviation (σ) with 10 s observation intervals, for $\mathcal{C} = \{20, 200\}$ and $\tau = 10$ s.

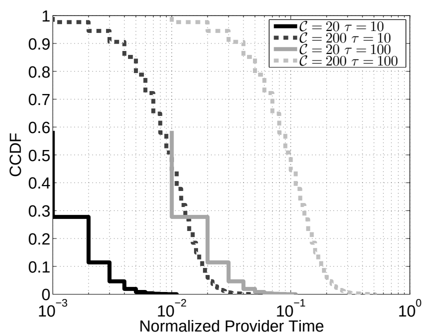

Fig. 4. CCDF of the time a node spends in provider mode, normalized to the total simulation time, for the RDD policy in the static uniform scenario.

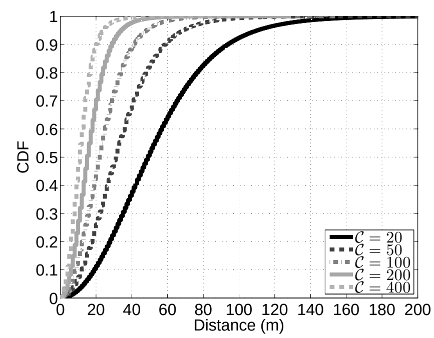

Fig. 6. CDF of the Euclidean distance to closest information replica, for the RDD policy in a mobile scenario with random waypoint mobility.

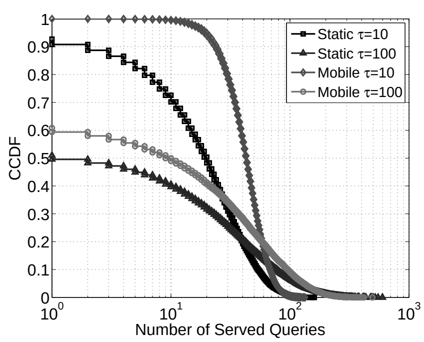

(a) $\mathcal{C} = 20$

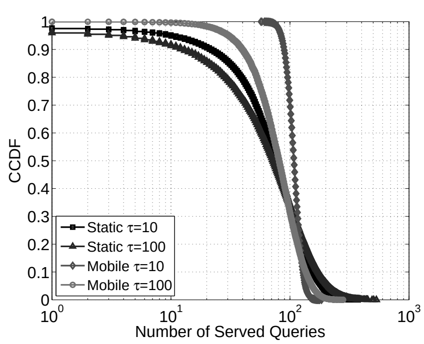

(b) $\mathcal{C} = 200$

Fig. 5. CCDF of the total number of queries served by information providers for the RDD policy.

users issuing queries to access information held by providers. In Fig. 6 we plot the cumulative distribution function (CDF) of the Euclidean distance from a querying node to the closest provider, for the mobile scenario with random waypoint mobility and $\tau = 10$ s. More specifically, we study the impact of an increasing number of initial providers $\mathcal{C}$, when we let this simulation parameter grow from 20 to 400 providers (i.e., from 1% to 20% of the total number of nodes). Both the mean distance to access information, ranging from 50 m to 15 m, and the variance of the CDF, shrink considerably when increasing the number of initial providers. Given that the node radio range is set to 20 m, the implications of this result are the following: when a sufficient number of initial providers is injected into the network (i.e., 200–400), nodes may access the information within one hop, whereas an insufficient number of initial information copies (i.e., 20–100) may constrain a node to propagate its query over multiple hops to retrieve the information.

Summary. The evaluation we carried out showed that, despite their simplicity and low overhead, the proposed cache-and-forward schemes achieve the two objectives defined in this work. Indeed, as long as enough providers are injected into the network (in our simulations, 10% of the total number of nodes), we have that:

(i) under static uniform scenarios, the information distri-

bution yielded by RWD and RDD well approximates the spatially uniform distribution; instead, both schemes closely match a uniform distribution on nodes when they are grouped around point of interests as simulated in the clustered scenario;

(ii) node mobility edges both cache-and-forward mechanisms toward achieving a good approximation of nodal uniformity in the clustered scenario; instead, when nodes move according to the random waypoint model, the quality of the approximation of the spatially uniform distribution deteriorates for both policies;

(iii) in terms of load balancing, both dissemination strategies evenly distribute the service load across the provider nodes and, again, mobility has a beneficial effect.

VI. Conclusions

We considered a peer-to-peer wireless network, where nodes may act as both clients and providers to other network nodes. In such a cooperative environment, we addressed the problem of achieving a desired distribution of information and a fair load distribution among the provider nodes. We designed lightweight, content-transparent, distributed algorithms that regulate the information storage at the network nodes and allow a fair selection of the nodes acting as providers.

Our simulation results indicate that under a variety of scenarios including static, mobile, and clustered network topologies, our simple mechanisms are effective in approximating a desired information distribution. Mobility appears to be a useful ally, instead of a problematic phenomenon, since it helps to achieve an even distribution of the load on providers.

Additional issues, such as the dynamic adaptation of number of information replicas to time/space-varying content demand or information survival are left for future work.

References

[1] I. Baev, R. Rajaraman, "Approximation Algorithms for Data Placement in Arbitrary Networks," *SCM-SIAM SODA*, 2001.

[2] S. Prabh, T. Abdelzaher, "Energy-Conserving Data Cache Placement in Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 1, pp. 178–203.

[3] B. Tang, H. Gupta, S. Das, "Benefit-based Data Caching in Ad Hoc Networks," *Trans. on Mobile Computing,* vol. 7, no. 3, pp. 208–217, Mar. 2008.

[4] K. Jain, V. V. Vazirani, "Approximation Algorithms for Metric Facility Location and k-Median Problems Using the Primal-Dual Schema and Lagrangian Relaxation," *J. ACM,* vol. 48, no. 2, 2001.

[5] T. Hara, "Effective Replica Allocation in Ad Hoc Networks for Improving Data Accessibility," *IEEE INFOCOM,* 2001.

[6] L. Yin, G. Cao, "Balancing the Tradeoffs between Data Accessibility and Query Delay in Ad Hoc Networks," *IEEE SRDS,* 2004.

[7] J. Holliday, R. Steinke, D. Agrawal, A. E. Abbadi, "Epidemic Algorithms for Replicated Databases," *IEEE Transactions on Knowledge and Data Engineering,* vol. 15, no. 3, pp. 1–21, May-June 2003.

[8] A. Aazami, S. Ghandeharizadeh, T. Helmi, "Near Optimal Number of Replicas for Continuous Media in Ad-hoc Networks of Wireless Devices," *Proc. Intl. Workshop on Multimedia Information Systems,* 2004.

[9] N. Sadagopan, B. Krishnamachari, and A. Helmy, "Active Query Forwarding in Sensor Networks," *Ad Hoc Networks Journal, Elsevier,* vol. 3, no. 1, pp. 91-113, Jan. 2005.

[10] D. Braginsky and D. Estrin, "Rumor Routing Algorithm For Sensor Networks," *First Workshop on Sensor Networks and Applications (WSNA),* Atlanta, GA, Sept. 2002.

[11] P. Nain, D. Towsley, B. Liu, and Z. Liu, "Properties of Random Direction Model," *IEEE INFOCOM,* Mar. 2005.

[12] S. Bandyopadhyay, E. J. Coyle, T. Falck, "Stochastic Properties of Mobility Models in Mobile Ad Hoc Networks," *IEEE Transactions on Mobile Computing,* vol. 6, no. 11, pp. 1218–1229, Nov. 2007.

[13] V. Naumov, R. Baumann, T. Gross "An Evaluation of Inter-vehicle Ad Hoc networks Based on Realistic Vehicular Traces," *ACM MobiHoc,* 2006.

[14] S.Y. Ni, Y.C. Tseng, Y.S. Chen and J.P. Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network," *ACM Mobicom,* 1999.

[15] J.-Y. Le Boudec, M. Vojnovic, "The Random Trip Model: Stability, Stationary Regime, and Perfect Simulation," *ACM /IEEE Trans. on Networking,* vol. 14, no. 6, pp. 1153–1166, Dec. 2006.

[16] J.-Y. Le Boudec, M. Vojnovic, "Perfect Simulation and Stationarity of a Class of Mobility Models," *IEEE Infocom,* Mar. 2005.

[17] M. Trott, "The Mathematica Guidebooks Additional Material: Average Distance Distribution," 2004.