Gossiping on MANETs: the Beauty and the Beast

Roy Friedman
Dept. of Computer Science
Technion - Israel Institute of Technology
roy@cs.technion.ac.il

Daniela Gavidia
Dept. of Computer Science
Vrije Universiteit Amsterdam
daniela@few.vu.nl

Luis Rodrigues
Faculdade de Ciências
Universidade de Lisboa
ler@di.fc.ul.pt

Aline Carneiro Viana
ASAP research team
INRIA, Saclay, France
aline.viana@inria.fr

Spyros Voulgaris
Dept. of Computer Science
ETH Zurich, Switzerland
spyros@inf.ethz.ch

ABSTRACT

Gossip protocols have emerged as a powerful technique for implementing highly scalable and robust services, such as information dissemination and aggregation. The fact that gossip protocols require very little or no structure to operate makes them particularly appealing to apply in dynamic systems, where topology changes are common (for instance, due to frequent faults or high churn rates). Therefore, gossip protocols seem particularly well fit to operate in wireless self-organizing networks. Unfortunately, these networks have a number of characteristics that impede the deployment of gossip protocols designed for wired networks. In this work we identify the inherent differences in communication between wired and wireless networks and their impact on the design and implementation of gossip protocols. In particular, our comparison includes drawing a distinction between the gossiping primitives suitable for each of these environments. In the context of this analysis, we conclude by presenting a list of open research questions.

1. INTRODUCTION

Gossip is a known design pattern for developing scalable and robust distributed systems. Gossip protocols are based on the idea of iterative information exchange between the participants. In each protocol step, each node only communicates with a small subset of the participants, and only exchanges a small amount of information in each such interaction. Moreover, the subsets of nodes that exchange information are drawn using some randomized mechanism, and the protocol is expected to converge into satisfying some global property in a probabilistic manner.

The bounded communication requirements of each step coupled with the probabilistic convergence of gossip protocols makes them highly scalable and robust. Hence, these protocols have become a preferred design pattern in large scale applications and middleware. Gossip has been often applied to Internet-based systems[1], and more recently, to wireless ad hoc networks. In particular, the communication paradigm associated with gossip protocols makes gossip an interesting approach to tackle many MANETs' design issues, e.g., scalability, adaptability, unreliable links, multi-hop routing, and dynamically changing topologies.

Interestingly, there are several inherent differences between MANET and Internet settings. These include, for example, the fact that MANETs operate on a radio broadcast medium, where links are unreliable and may suffer from a high percentage of message losses. Additionally, in wireless networks, remote nodes can only communicate with each other with the help of intermediate nodes, while in the Internet, all the routing is handled by dedicated routers, which are typically separate entities from the devices on which applications (and middleware) run. Moreover, in mobile adhoc networks, the network topology is constantly changing, which makes routing extremely expensive. Hence, it is not clear a priori whether a technique that performs well in one setting would fare well in the other.

In this paper, we investigate the suitability of various gossip techniques in the context of wireless networks. Our aim is to identify the specific characteristics of wireless networks that should be considered when designing gossip protocols for them, and to explore how they affect the respective design choices. In particular, our work explores these characteristics through the prism of a set of gossiping deployment scenarios, while discussing the pros and cons of the various alternatives.

Based on the above analysis, we come up with specific recommendations regarding the usage of gossip in different types of wireless networks and identify open research problems. Interestingly, as discussed later in this paper, the basic gossip scheme for highly dynamic wireless networks should be slightly different than the one commonly used in wired networks. In summary, to the best of our knowledge, this is the first work that (i) examines if existing gossip approaches

[1] The terminology "Internet-based" will be used in the rest of the paper, to make reference to the traditional wired Internet.

from wired networks can be deployed, exactly as they are conceived, in mobile and wireless settings, and (*ii*) tries to define the characteristics of gossip-based wireless applications compared to Internet-based applications.

The rest of this paper is organized as follows. Section 2 discusses the advantages of using gossip-based protocols and identifies design properties for wired settings. Section 3 discusses the characteristics of MANETs and investigates potential scenarios for gossip deployment in MANETs. Section 4 enumerates some gossip primitives that are particularly well suited for MANET environments. Section 5 then, provides discussion about the way in which gossip-based applications differ between Internet-based and MANET systems. Finally, Section 6 discusses open issues relevant to the design of gossip protocols in MANETs and Section 7 concludes this paper.

2. THE BEAUTY (GOSSIP)

Gossip protocols were pioneered at Xerox PARC, as a part of the Clearinghouse project [4], where gossiping was used to remove inconsistencies in tables in wide-area database systems. Since then, the popularity of gossip protocols in the distributed systems domain has flourished. In addition to their elegant simplicity, the appeal of gossip protocols lies in the fact that they can be easily implemented in a fully-decentralized way and exhibit desirable properties, namely reliability, robustness and scalability.

In this section, we discuss why gossip protocols are appealing to build highly scalable, robust, flexible, adaptive and easily managed architectures. We further proceed to identify a number of network properties that are typically assumed when designing gossip protocols for wired settings.

2.1 Motivation

We motivate the use of gossip-based protocols using a concrete application. Consider the case of information dissemination among a large set of processes. One way to implement this service consists in building a minimal cost spanning tree (MST) connecting all the processes. Using the tree, we can then deterministically disseminate the data without redundancy. In static networks, a minimal spanning tree can be built in polynomial time. The real problem of the MST approach is its cost in dynamic settings. A failure of a single process disconnects the overlay network, and a new tree must be built. If failures are frequent, especially in networks with many nodes, the cost of maintaining and repairing the tree may dominate the cost of the protocols. Naturally, one can design trees with multiple redundant paths: unfortunately, such trees are even more expensive to build and maintain, and are advantageous only when faults are rare.

In contrast, gossip-based dissemination operates as follows [3, 6]. When a node wishes to broadcast a message, it selects t nodes from the system at random (this is a configuration parameter called *fanout*) and sends the message to them; upon receiving a message for the first time, each node repeats this procedure. This proactive approach results in message redundancy which in turn provides probabilistic guarantees for reliable message dissemination. Moreover, message delivery degrades gracefully in the face of node failures and message loss. It should be noted that this sort of protocols require very little structure: just a probabilistic approximation of the entire system membership.[2] Furthermore, these protocols have an inherent redundancy that makes them very robust to the effects of system dynamics, such as faults, churn, transient network failures, etc.

2.2 Advantages of Gossip

We can enumerate some of the main advantages of gossip-based protocols:

Scalability: Gossip protocols are inherently scalable. Each node performs a fixed set of operations at a fixed rate, regardless of the network size. Consequently, gossip protocols are virtually infinitely scalable with respect to node load.

Network load balancing: In a gossip protocol, each and every node communicates with a limited number of peers for a limited number of times, in order to provide a service. In particular, there is no special node in charge of distributing and collecting information on behalf of other nodes, which avoids routing hot-spots. Additionally, as gossip peers are selected at random, the load is equally distributed among all links.

Resilience to node failures: Gossip protocols are inherently redundant: a node is likely to be contacted by more than one peer in a given execution. Therefore, any failure of a single node has little impact on the reliability of the protocol. In fact, it has been shown that some gossip-based protocols can even tolerate failure rates as high as 80% [25].

Resilience to transient network failures: Due to the random selection of peers, multiple paths are explored in a gossip-based protocol. This makes the protocol highly resilient to link failures, including non-transitive network partitioning (a scenario where node A is partitioned from node B, B is partitioned from C, but A is not partitioned from C) .

Symmetric nature: Gossip protocols are symmetric by nature. That is, nodes are roughly equivalent, with no node playing a specialized role. As a consequence, gossip protocols are inherently simple.

Simple low cost management: In order to operate, each node only needs to maintain an approximation of the system membership, in order to select its gossip peers at random [9, 25]. This membership is much simpler to maintain than structures with much stronger constraints (such as a spanning tree).

2.3 Network Assumptions

Many gossip based protocols have been designed for wired environments. Most of these protocols make the following assumptions about the underlying network.

- *Clique connectivity* – Gossip protocols for wired networks typically assume IP routing as the underlying layer. Hence, any process can communicate with any other process in the system.

[2]Coincidentally, membership management has also been addressed using a gossip-based solution [25].

- *The cost of reaching a node is (almost) the same for any node* – Although the network latency changes from one peer to another, the difference does not reach the point of hampering the protocol operation. Therefore, the "locality" of peer processes becomes an optimization issue and not a fundamental design constraint.

- *Communication is not constrained by physical connectivity* – As a consequence of the full network connectivity and roughly equal routing costs assumptions, as mentioned above, the overlays are typically built without considering the underlying structure of the network and node proximity.

- *Links are relatively reliable* – If a correct process knows the ID of another correct process, it succeeds in establishing communication with that process with high probability.

3. THE BEAST (MANET)

Mobile ad hoc networks (MANETs) are dynamic systems. The nature of wireless communication makes them prone to different sources of interference, which may cause instability in the topology. Furthermore, MANETs include mobile nodes which form another important cause of frequent topology changes. In addition, the lack of a management infrastructure and the dynamics of the network impose two fundamental requirements on the design/operation of MANETs: (a) all nodes in a MANET should assume the same management capabilities, and (b) any network operation (*e.g.*, routing, location service, data dissemination, etc) should be decentralized [24].

At a first glance, the attractive properties of gossip-based protocols listed in the previous section seem to address the aforementioned issues of MANETs. Nevertheless, a number of MANET characteristics differentiate the porting of Internet-based gossip protocols to the wireless world. This section discusses the characteristics of MANETs that hinder or facilitate gossip protocols, and investigates potential deployment scenarios for gossip applications.

3.1 Limitations of MANETs

Unfortunately, many of the network assumptions that are made when building a gossip protocol for a wired network no longer hold in a MANET setting. This motivates the use of different gossip primitives tailored to these networks.

- *Expensive routing* – Routing is an expensive operation in MANETs. The lack of dedicated equipment to perform routing functions, and the dynamics of the network that cause routes to be unstable, makes routing a resource hungry operation. Furthermore, the more distant the corresponding node, the more expensive the communication.[3] This invalidates several assumptions made for wired networks, such as: communication is carried out over an arbitrary overlay, which is not constrained by physical connectivity; clique connectivity; and the cost of reaching a node is (almost) the same for all nodes. On the contrary, on MANETs, communication with physical neighbors (ideally within the communication range) is favored.

- *Links are unreliable* – Variables such as obstructions, interference, environmental factors, and mobility make determining connectivity a priori difficult in MANETs. Also, contrary to wired networks, where collision can be detected, in MANETs the wireless medium does not permit collision detection at the sender side, making collision avoidance hard to achieve. Thus, links may suffer a much higher percentage of message losses.

- *Constrained resources* – Due to the absence of wiring and their small physical size, wireless devices often have limitations in memory, processing, and, aboveall, power. Unfortunately, the inherent redundancy that makes gossip robust may be a severe limitation in an environment where the optimization of resources is strongly required in order to minimize energy consumption and network traffic.

3.2 Opportunities in MANETs

Fortunately, MANETs also offer opportunities that can be exploited to make gossip more efficient. Namely:

- *Broadcast communication* – Unlike the Internet, where the point-to-point model of communication is dominant, communication in wireless networks is inherently broadcast based. That is, when a node transmits some information, typically all nodes within its transmission range can receive it. Therefore, every time a process communicates with a (physical) neighbor, the gossip automatically reaches all its other neighbors too at no additional cost. This clearly favors gossip-style dissemination, particularly in densely populated wireless networks.

- *Mobility* – In MANETs, nodes may carry information while moving. Therefore, node mobility can assist with the spreading, mixing, and aggregation of information. By gossiping with neighbors at different physical locations, a mobile node may contribute in overcoming limitations of the underlying wireless topology, for instance, by connecting otherwise disconnected network partitions.

- *Position information* – In some MANETs, nodes are aware of their position in space, for instance by being equipped with integrated GPS receivers. In such cases, position knowledge can be utilized by a number of location-aware algorithms, such as location-oriented dissemination or geographic routing.

- *Data-centric routing* – Like certain peer-to-peer systems, some MANETs have adopted a data-centric communication model as opposed to the traditional point-to-point address-centric abstraction. In a data-centric abstraction, the nature of the data is more important than the identity of the source. In particular, routing decisions are taken primarily based on the type of data. This model of communication clearly favors in-network data aggregation, besides imposing changes on the way routing and storage are performed in the network.

[3]See, for example, the performance measurements in [1].

4. GOSSIP PRIMITIVES FOR MANETS

When considering the application of gossip-based protocols in mobile ad hoc networks, it is useful to distinguish between three gossip deployment scenarios.

Sparse gossip: In this deployment scenario, gossiping involves only a small fraction of the nodes that constitute the MANET. Therefore, it is unlikely that other participants are reached when a broadcast is executed at the network level. In this case, gossip needs to be performed on top of a routing layer.

Dense gossip: In this deployment scenario, the majority of nodes—or even all of them—participate in gossiping. Therefore, it is possible to exploit the broadcast nature of the wireless communication to improve gossip.

Delay tolerant gossip: In this deployment scenario, gossip participants are assumed to be mobile and their mobility patterns instrumental in preserving network connectivity. Moreover, the network may be disconnected for long durations. In this case, gossip via mobile nodes is the only way to ensure the dissemination of information in the system.

In this section, we describe some gossip primitives that are particularly well suited for MANET environments, through the prism of the previously described gossip deployment scenarios.

4.1 Broadcast Gossip

In a wired setting, gossip is typically performed by letting the source of a message select at random the nodes with which it is going to exchange information. Then, the source sends a point-to-point message to each of these targets.

This procedure is rather cumbersome in dense gossip MANET scenarios. In such scenarios, gossip targets are within the communication range: to send separate point-to-point messages to each target is an overkill, given that a single transmission may reach all the targets. Even if a message is sent to all target nodes in a single broadcast message, additional information should be included in this message to selectively specify which nodes are targeted. For instance, the message may carry an array with the IDs of all targets. This additional information is an additional cause of resource consumption. Moreover, due to mobility, the indicated nodes may already be out of range when the message is transmitted, which reduces the reliability of the gossip mechanism.

A much more natural primitive for dense gossip MANET deployment scenarios is what we call *broadcast gossip*. In broadcast gossip, the source of a message does not select target nodes. Instead, it simply broadcasts the message to all nodes in its communication range. Then, each and every recipient uses some algorithm to decide if it is willing to act as a forwarder of the message or not. Different decision algorithms may be used:

- *Probabilistic* – Each recipient of the broadcast tosses a coin to decide if it should become a forwarder or not [5, 11, 23].

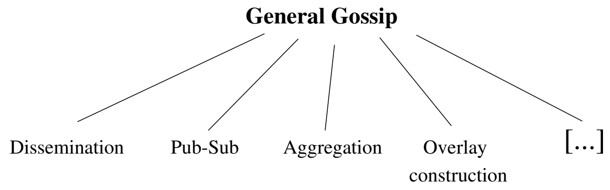

Figure 1: Gossip-based applications.

- *Counter based* – Each recipient monitors how many neighbors have previously decided to gossip a given message, and decides to forward it if it hears fewer than M such gossips [5, 11].
- *Distance based* – Each recipient estimates its distance to the source of the gossip. More distant nodes are more likely to become gossip peers than closer nodes[20].

Some known protocols combine several of these techniques [5, 11].

4.2 Opportunistic Gossip

Opportunistic gossip is unique to mobile ad hoc networks. In this scheme, mobility, and in particular the knowledge of the mobility pattern of nodes, is utilized to propagate messages in the system. That is, messages are forwarded to nodes that are likely to be moving towards a desired region of the network. During their mobility, these nodes offload the message opportunistically to other nodes they encounter. The decision regarding which nodes to offload a message to can be taken either by means of some heuristic or probabilistically. Hence, with opportunistic gossip, mobility becomes a substitute for dissemination along connected paths. Here, mobility becomes an asset of data dissemination rather than an obstacle.

Notice that in extremely fast moving networks, as well as when the network does not exhibit continuous connectivity for long durations (and sometimes is never fully connected), opportunistic gossip is in fact the only way to disseminate messages. In the literature, these types of networks have been referred to as *delay-tolerant*, *intermittently-connected*, or *highly-partitioned*. Let us note that most existing wireless routing techniques assume that there is (almost) always a connected path from a source to a destination. This assumption is clearly not valid in very fast moving networks and sparse mobile networks.

5. APPLICATION SCENARIOS

In this section, we explore ways in which gossip-based applications differ between Internet-based and MANET systems through the prism of a number of fundamental applications, depicted in Figure 1.

5.1 Overlay construction

Gossip protocols can be used to exchange membership information. This information can be used to build and maintain overlays with different classes of constraints [12, 21, 25, 26].

As explained in Section 2, Internet-based systems are by default routing-enabled and follow the point-to-point communication model. A node can initiate communication to another node only if the former knows the network address of the latter. Consequently, in order to gossip, each node needs to maintain a list of *links* to other nodes, that is, a list of network addresses of other nodes that participate in the system. These links collectively form an *overlay*, which defines the "who knows whom" relationship, and determines the potential communication paths for gossiping.

A significant amount of research has focused on devising algorithms for building and maintaining a multitude of overlay types, possessing various properties appropriate for one application or another. These include overlays that resemble random graphs [25] or small-worlds [12], super-peer structures [21], geographic-based topologies, semantic overlays[26], etc.

The MANET world is, however, fundamentally different. When routing is either not available, or too expensive, communication remains inherently limited to physically neighboring nodes. More specifically, a node can gossip with another node only if the latter is in the transmission range of the former. Let alone that transmission ranges are not always symmetric. Thus, in the *dense gossip* scenario (for static and slow moving networks), the overlay is mainly constrained by the physical topology.

As a result, overlay construction in static and slow moving networks is more relevant for the *sparse gossip* scenario. In this case, gossip can be used to disseminate membership and routing information on participating nodes. This information must then be complemented with information regarding the proximity of nodes, such that the resulting overlay has a strong correlation with the physical topology [18].

5.2 Information dissemination

Let us now look at information dissemination, which is the most widely studied application of gossip, e.g., [4, 3, 6, 10]. In Internet-based systems, gossip-based information dissemination typically relies on the periodic exchange of information between *random* pairs of nodes. It depends, in other words, either on full knowledge of the network, or on the creation of an overlay resembling a random graph, serving as a basis for information propagation. Random overlays typically have logarithmic diameter, which results in dissemination of a message to the whole network in a logarithmic number of steps.

In MANETs, on the contrary, topology constrained communication *prevents*, or at least *hinders*, communication between arbitrary pairs of nodes. Consequently, information flows along paths determined by the physical topology of the MANET. Thus, in the dense gossip scenario (for static and slow moving networks), the number of steps it takes information to spread across the whole network is directly proportional to the diameter of the physical topology. In such cases, cooperation among nodes is strongly required in order to limit the dissemination overhead. Broadcast gossip primitives have served as the technique of choice to achieve this goal.

On the other hand, in sparse gossip scenarios, gossip nodes can be strategically chosen in order to decrease dissemination steps. Nevertheless, for a "good" distribution of gossip nodes, knowledge of the network topology may be required.

Opportunistic-based gossip protocols may be more appropriate for MANETs. In such networks, gossip takes advantage of node mobility to expedite the dissemination of information to the entire network. In the case of information dissemination to a given, confined area, a gossip mechanism should take care of controlling the dissemination overhead, e.g., by having a good estimate for the destination's position. Yet, mobile destinations change and it is not always feasible to predict nodes' mobility patterns.

It is important to notice that the consequence of location constrained communication extends beyond the speed of dissemination, namely to the reliability of dissemination and its resilience to node failures. Consider, for example, the case in which nodes are placed such that some of them are central to the connectivity of the network. A failure of a strategically located node in such a network may prove disastrous to the successful dissemination of a message to the whole network. In contrast, in Internet-based networks the failure of a single node, or a small number of arbitrary nodes, typically does not prevent messages from reaching all alive nodes. Note however that mobility and opportunistic gossip may help to mitigate the effect of node failures in MANETs.

5.3 Topic-based publish/subscribe (Multicasting)

In topic-based publish/subscribe, publishers issue events, each one being associated with a certain topic. Subscribers register with one or more topics and expect to receive all events concerning topics of their interest.

This is essentially equivalent to multicasting messages (i.e., events) to all members (i.e., subscribers) of a multicast group (i.e., topic). We will follow the publish/subscribe terminology.

In Internet-based systems, subscribers of each topic typically form a separate overlay, enabling the dissemination of messages among themselves. That is, the dissemination of an event to the subscribers of its respective topic is essentially reduced to basic information dissemination among these subscribers only, without involving any non-interested subscribers.

The consequences of the infeasibility of overlay construction in MANETs become more obvious in this context. As direct communication between any arbitrary pair of nodes may generally be prohibitively high, or even impossible, the dissemination of an event to subscribers of a particular topic will inevitably involve unrelated subscribers. The brute force solution is to disseminate all events to the whole network. This way the event of a given topic will reach all subscribers interested in it. However, it will also reach all other subscribers, imposing unnecessary load on the MANET. An optimization would be to spread information concerning nodes' subscriptions, such that the flooding of an event can be pruned in the directions of the network where no node is interested in it. Such mechanisms, however, can prove particularly complex, especially when node mobility is involved.

The sparse gossip deployment scenario is a better candidate for the topic-based publish/subscribe, as it permits gossiping among an arbitrary subset of the nodes, namely the nodes subscribed to a given topic.

5.4 Data aggregation

Data aggregation constitutes another important representative of gossip-based applications. *Averaging* constitutes a fundamental aggregation operation, in which each node is equipped with a numeric value, and the goal is to estimate the average of all nodes' values. Various previous works [16, 15, 13] have shown how averaging can be used as the basis for the computation of other aggregates, including generalized mean, variance, counting of nodes, sum, and product.

In averaging, a node updates its estimate to the average between its previous local estimate and the estimate received. That is, when nodes p and q with estimates s_p and s_q gossip, their estimates are updated as follows:

$$s_p \leftarrow \frac{s_p + s_q}{2} \quad , \quad s_q \leftarrow \frac{s_p + s_q}{2}$$

Note that the sum of the two nodes' estimates does not change, therefore neither does the global average. The variance, however, over the set of all nodes' estimates decreases, unless s_p and s_q were already equal, in which case it remains unaltered. Experiments and theoretical analysis in [13, 15, 17, 22] show that the variance converges to zero. Moreover, when applied on small-world or random graph topologies it converges at an exponential rate, whose exponent depends on the communication graph defining the nodes' neighbors.

In data aggregation, nodes are essentially exchanging their estimates, in an effort to influence the global estimate. The initial estimates of the nodes are, thus, diffused in the network, affecting, to some extent, the global average estimate. As a consequence, averaging speed is highly related to the speed of information propagation in a certain network. The rule of thumb is that the higher the link randomization in an overlay, the faster the aggregation convergence.

As such, it is clear that aggregation by data summarization—by using average, minimum/maximum, count, or sum operations [2, 8, 14, 19]—over a MANET is possible, despite the fixed and non-randomized communication paths. Some proposed solutions that aggregate data by the use of an average operator, for example, execute a depth-first tree traversal algorithm. In this algorithm, there is only one request in the network (at one time) and nodes which have already responded to the query have their identifiers stored in the request message. Other solutions count the number of different values in a network, or even compute the size of the network.

Nevertheless, averaging over a MANET will generally converge slower than over a randomized Internet-based overlay. When aggregation seeks for an estimation involving the value of all nodes, dense gossip is more appropriate for this application. However, in many MANET applications, one may be interested in obtaining local aggregation (for instance, the average temperature in a certain area of the network). The extent to which global and local aggregations techniques differ is still an open research issue.

6. OPEN ISSUES

As discussed before, gossip-based protocols have the potential to improve the scalability and fault tolerance of MANETs. Nevertheless, they also introduce new challenges for protocol designers, beyond what we have discussed in this paper. Some of these challenges are the following:

- The applicability of gossip-based broadcast algorithms in a practical setting is limited by their implicit assumption that each node has enough buffering resources. Indeed, in order to operate in a reliable manner, the nodes participating in the broadcast must be equipped with enough resources to ensure that messages are gossiped a sufficient number of times. If a node/link does not have enough resources, it may drop a large number of messages that are being forwarded. If several nodes/links do not have enough resources, reliability might end up being drastically affected. One might consider calibrating, a priori, the transmission rate of the senders according to the resources available at every node. The static flavor of this naive solution makes it unrealistic in fixed networks and unfeasible in MANETs. The study of flow-control techniques that can be applied to gossip protocols in MANETs is still an open topic.
- Probabilistic analysis often assumes certain "good" properties regarding the network, such as that nodes are distributed according to some easy to analyze distribution. Nevertheless, in practice, the network may not follow these assumptions. Additionally, probabilistic analysis only guarantees the common cases. Yet, for users it would be good if the gossip protocols were designed such that they could recover from "bad luck" occurrences, which are possible (even if unlikely) due to the use of randomness. A way to combat these phenomena is to combine probabilistic mechanisms with deterministic corrective measures, along the lines of [5]. Studying the right combination of probabilistic mechanisms with deterministic ones is therefore an interesting research direction.
- One measure of efficiency of gossip-based protocols is the resulting convergence time. Gossip protocols guarantee that in Internet-based systems, if each node forwards a message to $O(\log n)$ random peers (n being the size of the system), then with high probability, $O(\log n)$ rounds will be required to disseminate a message to the entire network [7]. The fan-out of $O(\log n)$ peers per node, not only represents a measure of the amount of state required to do routing, but it is also a measure of how much state needs to be adjusted when nodes join or leave the networks. As previously described, in order to avoid costly routing, in MANETs gossip peers are often chosen to be the physical neighbors rather than random nodes in the network. This simplifies the state adjustment caused by network dynamics, but may affect the convergence time as well. Convergence time also depends on the mobility pattern of nodes. In general, we can define the amount of *effort* a protocol invest as the number of gossip peers each node has and how many of these gossip peers are local vs. how many of them are only accessible through multiple hop routing. The formal analysis of the amount

of time required to converge under a given effort and mobility pattern is an interesting open issue.

- Gossip-based protocols rely on the repeated execution of an algorithm for the eventual convergence to a desired state, such as the calculation of an average or the dissemination of a message throughout the network, as studied in the pioneering work of Demers et al [4]. Without knowledge of the global properties of the network (such as network size), the point at which convergence is reached is not always clear. This raises the question: for how long should nodes gossip? Nodes in Internet-based systems can afford to gossip periodically for an indefinite amount of time. Doing so allows them to converge to the desired state and, if conditions change, adjust accordingly. Take as an example, a gossip-based protocol for building and maintaining an overlay. Thanks to continuous periodic gossiping, the overlay can be built and reconfigured in the event of nodes joining or leaving. In the case of MANETs, communication is an expensive operation due to energy and bandwidth constraints. For this reason, determining when and for how long to gossip is an interesting open issue. Ideally, gossip-based protocols should only run when necessary, possibly triggered by changes in the environment (such as a new sensor measurement or discovering a new neighbor). While some methods for estimating the size of a MANET exist, e.g., [1], the decision to stop gossiping may need to depend on additional local observations. Most likely, the dynamicity of the system (churn, mobility) will be an important factor in determining for how long to gossip.

- The overlay symmetric nature assures a good distribution of the load over peers. Consequently, routing hot-spots can be easily avoided. In sparse or dense gossip scenarios applied to static or low moving ad hoc networks, this can be also envisioned. Nevertheless, routing hot-spots are harder to deal with in mobile MANETs, since nodes do not have global information about the network load balancing. Thus, how nodes can deal with hot-spots and how load balancing can be taken into account during gossip, constitute interesting research directions.

- Certain MANETs are heterogeneous in terms of the capabilities of different nodes in the network (computing power, energy, radio transmission range, etc.). In such heterogeneous networks, it makes sense to bias the forwarding probability based on the capabilities of each node. A systematic approach for doing so and a formal analysis of its implications remains an important open research topic.

- The deployment characteristics of MANETs are not properly addressed or simply abstracted away in traditional models based on which algorithms for MANETs are currently analyzed. Some unrealistic considerations are that: nodes know their neighbors at the beginning of the algorithm, transmission range follows the Unit Disk Graph (UDG) modes, nodes are distributed uniformly at random in the plane, or still that nodes wake up or start the algorithm at the same time. These considerations may invalidate the use of analyzed algorithms in real-world deployment scenarios, regardless of how efficient they are in simulation scenarios. Thus, it is important to take this into account when designing gossip-based protocols for MANETs.

7. CONCLUSIONS

This paper examined the suitability of gossip-based protocols in the context of MANETs. We have identified the specific characteristics of these networks, and explored how they affect various potential gossiping strategies. Moreover, we have studied three major types of MANETs, sparse, dense, and delay tolerant, and examined the type of application and programming abstractions that suite each of them.

We believe that gossip should play a major role in both dense and delay tolerant MANETs. In particular, a distinguishing feature of gossip-based protocols is their use of randomization and redundant, ever-changing, implicit data paths. An interesting consequence of this is that gossip-based protocols often work better in mobile environments than in static ad hoc networks. In other words, they turn mobility from an obstacle into an asset.

Yet, some challenges still remain, as we described in Section 6. These include, for example, keeping the overhead of the randomized dissemination process at bay, using the right combination of probabilistic mechanisms with deterministic ones, evaluating gossip convergence complexity, the avoidance of routing hot-spots in mobile environments, determining stop gossip primitives in mobile environments, and considering real deployment characteristics.

8. REFERENCES

[1] Z. Bar-Yosef, R. Friedman, and G. Kliot. RaWMS: random walk based lightweight membership service for wireless ad hoc network. In *Proc. 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 238–249, May 2006.

[2] M. Bawa, A. Gionis, H. Garcia-Molina, and R. Motwani. The price of validity in dynamic networks. *Proceedings of ACM SIGMOD International Conference on Management of Data*, pages 515–526, June 2004.

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

[4] 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 ACM, editor, *Proceedings of the 6th Annual ACM Symposium on Principles of Distributed Computing (PODC '87)*, pages 1–12, Vancouver, Canada, August 1987.

[5] V. Drabkin, R. Friedman, G. Kliot, and M. Segal. RAPID: Reliable Probabilistic Dissemination in Wireless Ad-Hoc. In *Proceedings of the 26th IEEE International Symposium on Reliable Distributed Systems (SRDS)*, 2007.

[6] P. Eugster, R. Guerraoui, S. Handurukande, P. Kouznetsov, and A.-M. Kermarrec. Lightweight probabilistic broadcast. *ACM Transaction on Computer Systems*, 21(4):341–374, 2003.

[7] P. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massoulié. From Epidemics to Distributed Computing. In *IEEE Computer Magazine*, volume 37, pages 60–67. IEEE Computer Society, 2004.

[8] P. Flajolet and G. N. Martin. Probabilistic counting. *Symposium on Foundations of Computer Science*, pages 76–82, Nov. 1983.

[9] A. J. Ganesh, A.-M. Kermarrec, and L. Massoulié. Peer-to-peer membership management for gossip-based protocols. *IEEE Transactions on Computers*, 52(2):139–149, February 2003.

[10] D. Gavidia, S. Voulgaris, and M. van Steen. A gossip-based distributed news service for wireless mesh networks. In *Proceedings 3rd IEEE Conference on Wireless On demand Network Systems and Services (WONS)*, Les Menuires, France, January 2006.

[11] Z. Haas, J. Halpern, and L. Li. Gossip-based ad hoc routing. In *Proceedings of IEEE INFOCOM*, June 2002.

[12] M. Jelasity, W. Kowalczyk, and M. van Steen. Newscast Computing. Technical Report IR-CS-006, Vrije Universiteit Amsterdam, Department of Computer Science, Amsterdam, The Netherlands, November 2003.

[13] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transaction on Computer Systems*, 23(3):219–252, 2005.

[14] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems*, 23(3):219–252, 2005.

[15] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2003)*, pages 482–491. IEEE Computer Society, 2003.

[16] D. Kempe, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 163–172, New York, NY, USA, 2001. ACM Press.

[17] W. Kowalczyk and N. Vlassis. Newscast EM. In *Advances in Neural Information Processing Systems (NIPS) 17*, 2004.

[18] J. Luo, P. Eugster, and J.-P. Hubaux. Route driven gossip: Probabilistic reliable multicast in ad hoc networks. In *Proceedings of the INFOCOM*, 2003.

[19] L. Massoulié, E. L. Merrer, A.-M. Kermarrec, and A. Ganesh. Peer counting and sampling in overlay networks: random walk methods. *Proceedings of ACM Symposium on Principles of Distributed Computing*, pages 123–132, June 2006.

[20] H. Miranda, S. Leggio, L. Rodrigues, and K. Raatikainen. A power-aware broadcasting algorithm. In *Proceedings of the 17th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 1–5, Helsinki, Finland, Sept. 2006.

[21] A. Montresor. A robust protocol for building superpeer overlay topologies. In *P2P '04: Proceedings of the Fourth International Conference on Peer-to-Peer Computing (P2P'04)*, pages 202–209, Washington, DC, USA, 2004. IEEE Computer Society.

[22] A. Montresor, M. Jelasity, and O. Babaoglu. Robust Aggregation Protocols for Large-scale Overlay Networks. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 19–28. IEEE Computer Society, 2004.

[23] Y. Sasson, D. Cavin, and A. Schiper. Probabilistic broadcast for flooding in wireless mobile ad hoc networks. In *Wireless Communications and Networking Conference*, pages 1124– 1130, Mar. 2003.

[24] A. C. Viana, M. D. Amorim, Y. Viniotis, S. Fdida, and J. F. Rezende. Twins: a dual addressing space representation for self-organizing networks. *IEEE Transactions on Parallel and Distributed Systems*, 17(12):1468–1481, Dec. 2006.

[25] S. Voulgaris, D. Gavidia, and M. Steen. Cyclon: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 13(2):197–217, June 2005.

[26] S. Voulgaris and M. van Steen. Epidemic-style Management of Semantic Overlays for Content-Based Searching. In *EuroPar*, LNCS 3648, pages 1143–1152. Springer-Verlag, Aug./Sept. 2005.