Controlled Epidemic-style Dissemination Middleware for Mobile Ad Hoc Networks

Mirco Musolesi and Cecilia Mascolo
Department of Computer Science, University College London
Gower Street, London WC1E 6BT, United Kingdom
{m.musolesi|c.mascolo}@cs.ucl.ac.uk

Abstract

Traditional middleware primitives offer very elementary information dissemination mechanisms, which, in the case of a decentralized and dynamic network such as a mobile ad hoc network, do not offer the ability to control the information spreading. Control over information dissemination could instead be very critical especially in terms of lifetime of the network. Gossip-based communication and epidemic-style algorithms, which are based on a store and forward approach, have been proposed to obtain message dissemination with probabilistic guarantees and lower overheads. However, epidemic algorithms have never been used to allow designers to control the spreading of the information depending on the desired reliability and the network structure.

In this paper, we present a middleware for ad hoc networking, which uses epidemic-style information dissemination techniques to tune the reliability of the communication in mobile ad hoc networks. The approach is based on recent results of complex networks theory; the novelty of our idea resides in the evaluation and the exploitation of the structure of the underlying network for the automatic tuning of the dissemination process and its use in the design of the API offered by the middleware. We present a detailed analytical model supported by several simulation results.

1 Introduction

Traditional middleware primitives for information dissemination fail to offer the right abstractions to the programmer of mobile applications, especially if these are targeted to very decentralized systems such as mobile ad hoc networks [16]. One of the main capabilities which is missed is the ability to control the information spreading from the application program. Examples of applications in which this feature is essential are emergency and rescue operations in possibly crowded public areas (such as inside stations, airports or shopping centers) or during events that involve a large number of people gathered together (such as in occasion of major sport events in stadiums or arenas). If the network infrastructure has failed, firefighters and other helpers might want to relying on device to device connectivity of the people in the area: the spreading of messages might need to be controlled so to preserve the lifetime of the network for future messages. For example, it may be sufficient to send the messages only to a percentage of the rescue team members (e.g., 50% of the doctors). In other situations, there might be a need to reach all the deployed emergency personnel. Up to our knowledge, no solutions exploiting the minimal necessary and sufficient number of replicated messages given the emergent network structure to guarantee a desired level of reliability exist. This lack of ability to control message dissemination at the application level is partly due to the poor APIs offered by the middleware but also to the lack of algorithms that can implement this tuning over the network.

Mobile ad hoc networks can be frequently and temporarily partitioned and the traditional routing protocols, including the basic flooding, fail to offer any sort of reliability when this happens. Epidemic-style protocols, instead, being store and forward approaches, allow for communication in dynamic and mobile networks, also in presence of temporary disconnections or network partitions. The analogy between information dissemination in mobile systems and epidemics transmission in communities is evident and a host can be referred to as *infected* when it receives a piece of information and stores it, and *susceptible* (i.e, it could be infected) otherwise.

Epidemics-inspired techniques have received huge attention in recent years from the distributed systems community [10]. These algorithms and protocols rely on probabilistic message replication and redundancy to ensure reliable communication. Epidemic techniques were firstly applied to guarantee consistency in distributed databases [9]. A few attemps have been made to employ epidemic based techniques for information dissemination in mobile ad hoc networks [18, 8, 3]. However, existing epidemic algorithms do not permit to control the spreading of the information

depending on the wanted reliability and the network conditions such as in the needed scenario depicted above. In fact, these approaches are fundamentally based on empirical experiments and not on analytical models: the input parameters that control the dissemination process are selected by using experimental results and are not based on any mathematical model. This implies that the message replication process cannot be tuned with accuracy in a dynamic way: for instance, it is not possible to set the parameters of the dissemination process in order to reach only a certain desired percentage of the hosts. Furthermore, these algorithms do not exploit the information on the underlying network topology. This is due to the fact that many interesting works on epidemic modelling in complex networked systems are very recent [1].

The use of these recent complex network theories allows us to devise a more precise model of the dissemination and to control the reliability level that can be imposed on message delivery, by evaluating the distribution of the degree of connectivity of nodes. In other words, the number of the replicas in the network and their persistence can be controlled to support a delivery process that is characterised by the reliability specified by developers. Moreover, by using these results we designed algorithms that are able to adapt *dynamically* to possibly variable degrees of connectivity of the hosts.

The contribution of this paper can be summarised as follows:

- We design a dissemination algorithm for mobile ad hoc networks that relies on epidemic models taking into account the structure of the underlying network and using recent results in complex networks theory concerning the modelling of epidemics spreading;
- we define a middleware interface for probabilistic communication and information dissemination in mobile systems that allows the programmers to set the reliability for unicasting and anycasting based on these theoretical results with a high degree of accuracy, also in presence of failures.

Complex networks are usually classified in two main groups depending on the distribution of the degree of connectivity of the nodes (i.e., the number of the links of the hosts): *exponential networks* and *scale free networks*. The formers are characterised by a connectivity distribution $P(k)$ peaked at an average value $\langle k \rangle$. Typical examples are random graph model [6] and the small-world model proposed by Watts and Strogatz [1]. Scale free networks are characterised by fluctuations of the degree k that any given node may have. Exponential networks are characterised by very small fluctuations (i.e., the degree of every vertex can be approximated as $k \approx \langle k \rangle$); for this reason, they are also identified as *homogeneous* networks. On the other hand, for the inherent fluctuations of the degree of connectivity, scale-free networks are classified as *heterogeneous* networks. We will first assume a mobile system with homogeneous network structure. This structure is realistic for a number of typical of scenarios characterized by a high density of hosts and where the movements of the hosts can be approximately modeled as random, such as in large outdoor spaces (i.e., squares, stations, airports or around sport venues) [12]. Then, we will discuss a generalization of the model to heterogeneous networks. This is the case of scenarios with the presence of groups of hosts and solitary individuals. In other words, we prove that the middleware has a general applicability, since it can be used in presence of *both* homogeneous and heterogeneous networks.

This paper is structured as follows. Section 2 introduces the middleware interfaces for controlled information dissemination. Section 3 provides a brief introduction to epidemic spreading models proposed in the recent complex networks studies and discusses the design of possible information dissemination strategies based on them. The implementation of our algorithm supporting he middleware primitives is described in Section 4. Section 5 shows an analytical study of our approach and presents several simulation results that confirm the validity of the theoretical model. In Section 6 we compare our approach to existing work, underlining its novelty and possible extensions of the model to heterogeneous networks scenarios. Section 7 concludes the paper, summarizing its contribution.

2 Middleware Primitives for Controlled Epidemic Dissemination

Our goal is to provide a set of primitives that allows developers to tune information dissemination in mobile ad hoc networks according to their specific application requirements. This problem can be evaluated from two different perspectives. In fact, the spreading of information from a source A to a certain percentage Ψ of the mobile hosts of the system can be seen as the problem of sending a message from host A to another randomly chosen host B with a certain probability Ψ. This probability can be interpreted as the reliability of the delivery mechanism.

We designed two primitives to support controlled communication in mobile systems that capture these two complementary perspectives. First of all, we design a primitive for *probabilistic unicast communication*:

```
epsend(message,recipient,reliability,time)
```

where `message` is the message that has to be sent to the `recipient` with a certain probability measured by the value `reliability` (that has to be chosen in the range $[0, 1]$) in a bounded time interval defined by the `time` field. The field `reliability` is used to set the value of Ψ. The validity of the message corresponding to the interval of time during which the infection will spread is specified by the field `time`.

Similarly, we introduce a primitive for *probabilistic anycast communication* as follows:

```
epcast(message,percentageOfHosts,time)
```

where `message` is the message that has to be sent to a certain percentage of hosts equal to the value defined in `percentageOfHosts` in a bounded time interval equal to `time`. In this case the field `percentageOfHosts` is used to set the value of Ψ.

The infectivity of the epidemics (i.e., the probability of being infected by a host that is in the same radio range, like in human diseases spreading) can be used to control the reliability of the unicast probabilistic communication mechanism. In other words, given an expected reliability (or percentage of hosts that has to be infected) equal to Ψ, we are able to accurately calculate the value of the infectivity in order to obtain an infection rate equal to a proportion of the total number of the hosts in the network. It is also worth noting that, as we will discuss in the next section, these primitives rely on a probabilistic algorithm based on the transmission of a *minimal*, and, at the same time, sufficient, number of messages. In other words, the energy consumption due to transmissions is minimized.

The applications of these middleware primitives are many. For example, these can be used to reach only a percentage of hosts in a network. Using these primitives, a member of an emergency squad finding a person in relatively critical conditions while exploring a disaster area, can alert a fraction of his/her colleagues, so to let others attend other patients.

3 Dissemination Techniques based on Epidemic Models

In this section, we discuss our application of mathematical models of epidemic spreading to the problem of probabilistic communication and information dissemination in mobile ad hoc networks. We consider a system composed of nodes characterized by a finite buffer size, which is a realistic assumption. The communication in the system is message passing based. Messages are composed of a header, containing information that is used to perform the shipment and a body, containing the data that has to be sent to a specific host. Every message is characterized by a unique identifier. An expiration time field is used to specify its validity. Given the limited buffer size, every node can store a finite number of messages. These are inserted in the buffer only if not already present.

We now briefly introduce the mathematical models that we exploit to design the dissemination algorithms. These are at the basis of the design of the middleware API presented in Section 2. In order to model the replication mechanisms for the messages, we exploit mathematical models that have been devised to describe the dynamics of infections in human populations [2]. The study of mathematical models of biological phenomena has been pioneered by Kermack and McKendrick in the first half of the last century. In the following decades, their work has been considerably extended and, nowadays, the study of epidemiology from a mathematical point of view is a mature scientific discipline. In particular, mathematical models of infection spreading of human diseases have been developed and successfully exploited to predict the evolution of the epidemics with the aim of finding effective countermeasures. Very recently, researchers in the area of complex networks theory have focused their attention on the problem of modeling epidemics spreading in networks characterized by well-defined structures [4].

According to the Kermack and McKendrick model, an individual can be in three states: *infected*, (i.e., an individual is infected with the disease) *susceptible* (i.e., an individual is prone to be infected) and *removed* (i.e., an individual is immune, as it recovered from the disease). This kind of model is usually referred to as the Susceptible-Infective-Removed (SIR) model [2]. In this paper we use a simplified version of the model, according to which individuals can exist in only two possible states, *infected* and *susceptible*. In the literature, this model is usually referred to as Susceptible-Infective-Susceptible (SIS) model [2]. We now map this model onto a mobile network of communicating hosts, where messages are disseminated. In the remainder of this paper we will substitute the term *individual*, used by epidemiologists, with the term *host*. A host is considered infected, if it holds the message and susceptible if it does not. If the message is deleted from the host, the host becomes susceptible again.

The main assumptions of our model are the following:

- all susceptibles in the population are equally at risk of infection from any infected host (this hypothesis is usually defined by epidemiologists as *homogeneous mixing*);
- the infectivity of a single host, per message, is constant[1];
- every host collaborates to the delivery process and no malicious nodes are present;
- each node has a buffer of the same size;
- the initial number of hosts and the host failure rate are known *a priori* by each host[2];

[1] Note that the infectivity per single message (i.e., a disease) is constant, but not per single host. In other words, a host usually stores messages characterized by different infectivities in its buffer.

[2] The initial number of hosts can be usually estimated in occasion of sport events, rallies, etc. for example by evaluating the seating capacity of the venues or the size of the area when the event takes place. Statistical data are also usually available for many application scenarios, such as number of passengers that uses a station or an airport in a certain time of the day, etc. Alternatively, this number can be estimated using distributed algorithms for the calculation of the approximated network size such as [13].

- the host failure rate can be approximated as a stationary process within the time interval of infection spreading (i.e., the number of hosts is considered constant during the spreading of the infection)[3];
- the failures of the nodes are uniformly random distributed and permanent.

Under the assumptions above, the dynamics of the infectives and susceptibles in the case of a scenario composed of $N(t)$ active hosts (i.e., not failed) can be approximately[4] described by means of a system of differential equations which we refined from [2] as follows:

$$\begin{cases} \frac{dS(t)}{dt} = -\beta S(t)I(t) + \gamma(t)I(t) \\ \frac{dI(t)}{dt} = \beta S(t)I(t) - \gamma(t)I(t) \\ \frac{dN(t)}{dt} = -\phi N(t) \\ S(t) + I(t) = N(t) \end{cases} \quad (1)$$

where $I(t)$ is the number of infected hosts at time t, $S(t)$ is the number of susceptible hosts at time t, β is the average number of contacts with susceptible hosts that leads to a new infected host per unit of time per infective in the population, γ is the average rate of removal of infectives from circulation per unit of time per infectives in the population and ϕ is the failure rate (i.e., the probability that one host fails per unit of time).

The equations of the system state that the decaying rate of susceptibles and the growth rate of infectives are calculated by considering two competing effects: the former, proportional to the infectivity β, the number of susceptibles $S(t)$ and the number of infectives $I(t)$; the latter, proportional to the removal rate γ and the number of infectives $I(t)$. The third equation is a consequence of the hypothesis of closed system (i.e., the nodes are the same and the number of hosts is constant over the interval of time taken into consideration). If we solve the system by using the initial condition $I(t) = I_0$ (where I_0 is the number of initial hosts infected), we obtain that the number of infectives at time t is described by the following equation:

$$I(t) = \frac{I_0 e^{\alpha\beta t}}{1 + \frac{I_0}{\alpha}(e^{\alpha\beta t} - 1)} \quad (2)$$

with $\alpha = N(t) - \frac{\gamma}{\beta}$. $N(t)$ is considered approximately constant during the entire epidemic process described by the system 1, since we assume that the failure process is stationary considering the interval of time during which the epidemics spreading happens (i.e., we assume $N(t) \approx N^*$ with N^* equal to the number of hosts present in the system at the beginning of the epidemics). In our case the initial condition is $I_0 = 1$: this represents the first copy of the message that is inserted in its buffer by the sender. This result can be used to calculate the number of infectives at instant t with a given infectivity β and a given removal rate γ, or, more interestingly for our purposes, β and γ can be tuned in order to obtain a certain epidemics spreading, after a specific length of time has passed. The infectivity β is the fundamental parameter of the message replication algorithm. In fact, a certain infectivity β can be selected in order to obtain, at time t^*, a number of infectives (i.e., hosts that have received the message) equal to $I(t^*)$ or, in other words, a percentage of infectives[5] equal to $I(t^*)/N(t^*)$. The parameter γ can be interpreted as the deletion rate of the messages from the buffer of the hosts. In fact, since the message buffers have limited size, it may be necessary to delete some messages according to a certain policy. Thus, from the average removal rate of messages from buffer, it is possible to derive the infectivity that is necessary and sufficient to spread the infection. In case the absence of overflow phenomena (i.e., in the case of sufficiently large buffers) can be assumed, the model can be simplified with $\gamma = 0$.

In order to effectively exploit the model just described, the actual connectivity of each host should be kept into account. As discussed in Section 1, the node degree k for each node can be approximated quite precisely with the average degree of connectivity $\langle k \rangle$ of the network. Therefore, in case of homogeneous networks, in order to take into account the effect of the connectivity, it is possible to rewrite the system (1), substituting β with $\lambda\frac{\langle k \rangle}{N}$ as follows, as discussed in [4]:

$$\begin{cases} \frac{dS(t)}{dt} = -\lambda\frac{\langle k \rangle}{N} S(t)I(t) + \gamma(t)I(t) \\ \frac{dI(t)}{dt} = \lambda\frac{\langle k \rangle}{N} S(t)I(t) - \gamma(t)I(t) \\ \frac{dN(t)}{dt} = -\phi N(t) \\ S(t) + I(t) = N(t) \end{cases} \quad (3)$$

λ represents the probability of infecting a neighboring host. $\frac{\langle k \rangle}{N}$ gives the probability of being in contact with a certain host. In other words, in this model, by substituting β with $\lambda\frac{\langle k \rangle}{N}$, we have separated, in a sense, the event of being connected to a certain host and the infective process [4].

The solution of this system is similar to (2) (i.e., it is sufficient to substitute β with $\lambda\frac{\langle k \rangle}{N}$). Thus, it is possible to calculate λ as function of $I(t^*)$ and $\langle k \rangle$. Finally, it is interesting to note that in homogeneous networks, every

[3] This is a realistic assumption, since users usually require that the information will be disseminated in a limited time.

[4] This is rigorously justifiable only for complete graphs in large population limit. However, the model provides a good approximation also in scenarios composed of a limited number of hosts.

[5] Note that $\beta = f(I(t))$ is not defined for $I(t) = N(t)$. Therefore, from a practical point of view, in the case of a message sent to all the hosts of the system, we will use the approximation $I(t) = N(t) - \epsilon$, with $\epsilon > 0$, in the expression used to calculate β.

host knows its value of k and, consequently, of $\langle k \rangle$. We will exploit this property to tune the spreading of message replicas in the system.

4 Implementing the Middleware Interface

Every time one of the two middleware primitives defined in Section 2 is invoked, the middleware calculates the value of the infectivity λ that is necessary and sufficient to spread the information with the desired reliability in the specified time interval, by evaluating the current average degree of connectivity and the current removal rate of messages from the buffer. The message identifiers, the value of the calculated infectivity, the timestamp containing the value specified in `time` expressing its temporal validity are inserted in the corresponding headers of the message in the *infectivity* field. Then, the message is inserted in the local buffer.

```
avDegreeOfConnectivity=System.getAvDegreeOfConn();
deletionRate=System.getDeletionRate();
infectivity=
 calculateInfectivity(reliability,deletionRate,
 avDegreeOfConnectivity, time);
basicReproductiveNumber=System.getBasicReprNumber();
if (basicReproductiveNumber>1) {
   m=new Message();
   m.setMessageId(System.generateMessageId());
   m.setRecipient(recipient);
   m.setContent(messageContent);
   m.setInfectivity(infectivity);
   m.setTimeStamp(time);
   System.addToBuffer(m);
} else throw new deliveryException();
```

Program 1: Calculation of the parameters of the message.

A fundamental parameter in epidemiology is the basic reproductive number R_0 [2]. This can be interpreted as the number of secondary infected hosts generated by one primary infective. In epidemiology, this is generally used to evaluate the conditions which give rise an epidemic outbreak in a population. Under the given assumptions, the basic reproductive number is intuitively defined as:

$$R_0 = \frac{\lambda \langle k \rangle}{\gamma} \tag{4}$$

It can be deduced that the epidemics will spread only if $R_0 > 1$ [4]. In this case, in fact, the epidemics will be able to generate a number of infected hosts (represented by the numerator) larger than those which have become susceptibles again (represented by the denominator) per unit of time, leading to a monotonic increase of the number of infectives $I(t)$. By evaluating the basic reproductive number, if it is not possible to ensure the specified reliability (i.e., the basic reproductive number is less than 1), an exception is thrown. The conditions under which R_0 is greater than 1 are discussed in Section 5.1.1 A possible implementation using an object-oriented programming style is presented in Program 1.

Program 2 contains the epidemic spreading algorithm. This procedure is executed periodically with a period equal to τ. With respect to the calculation of the message infectivity, it is possible to assume τ as time unit in the formulae presented in Section 3. In other words, assuming, for example, $\tau = 10$, a timestamp equal to one minute corresponds to six time units. The value of τ can be set by the application developer during the deployment of the platform. Clearly, the choice of the values of τ influences the accuracy of the model, since it relies on a probabilistic process. For this reason, given a minimum value of timestamp equal to t_{MIN}, developers should ensure $\tau << t_{MIN}$. The number of rounds will be equal to t^*/τ. For the Law of the Large Numbers, we obtain a better accuracy of the estimation of the evolution of the epidemics as the number of rounds (i.e., from a probabilistic point of view, the number of trials) increases. We implemented the epidemic algorithm and the middleware interface using Java SDK and we tested the functionalities of the framework with laptops connected by a wireless ad hoc network. However, in order to validate the epidemic algorithm, we studied its properties from an analytical point of view and we tested it in more realistic large-scale scenarios by means of simulations as described in the next section.

```
for (int i=0;i<numberOfMessagesStored;i++) {
  m=System.getMessageAtPosition(i);
  infectivity=m.getInfectivity();
  for (int k=0;k<numberOfHostsInReach;k++) {
    rValue=random(0,1);
    if (rValue<=infectivity)
      System.sendMessage(m,k);
  }
}
```

Program 2: Epidemic Spreading Algorithm.

5 Evaluation

We now present the evaluation of the proposed approach based on the analytical derivation of some characterizing properties of the system and on simulation results that confirm the theoretical model. We do not show and compare the results obtained with existing epidemic protocols and spanning-tree based multicast algorithms, since the goal of our work is different. The main aim of those protocols is to achieve 100% reliability communication with all of the participating hosts, whereas we are interested in achieving an accurate tuning of the dissemination process given the network structure in order to be able to reach only a percentage of the nodes.

5.1 Analytical Study of Properties of the System

5.1.1 Spreading and Persistence of Messages

It is interesting to derive under which conditions the reproductive number R_0 (defined in Section 4) is greater than 1. In this case, we will be sure that the epidemics will propagate until the expiration time[6]. With $P_{replacement}$ we indicate the probability that a message will have to be deleted from the buffer in order to free space when it is full. This will happen when a message is received, which is not already in the (full) buffer. With $P_{hit}(t)$ we indicate the probability of receiving from a neighbor a message that is already in the buffer (with a size equal to *BufferSize*) at time t. Therefore, the probability that a message in a full buffer will be deleted and replaced is equal to:

$$P_{replacement} = 1 - P_{hit} \tag{5}$$

γ represents the deletion rate from the buffer that is proportional to the arrival rate of a new message and the probability that this message is already in the buffer (represented by $P_{replacement}$). Since the arrival rate of new messages from any link is proportional to the spreading rate λ and the average number of connections $\langle k \rangle$, we re-write (4) as follows:

$$R_0 = \frac{\lambda\langle k \rangle}{\gamma} = \frac{\lambda\langle k \rangle}{\lambda\langle k \rangle P_{replacement}} = \frac{1}{P_{replacement}} \tag{6}$$

Thus, R_0 will be greater than 1 if and only if $P_{replacement} < 1$. In other words, if the buffer is large enough to ensure that the average removal rate is less than 1, the messages will remain in the system until their expiration time.

If the removal rate is higher than this threshold, the system will not be able to guarantee the persistence of the messages. It is possible to use this result to design a mechanism for determining when a notification that the message cannot be disseminated needs to be issued to the application. In general, the value of $P_{replacement}$ is dependent on the number of types of messages, their infectivities and the different stages of the dissemination processes (i.e., infections) that are present in the system. However, if the traffic behavior in terms of quantity and types of messages is homogeneous, the replacement rate observed at local level can be taken as a reasonable indicator of the average global replacement rate.

5.1.2 Number of Messages in the Network

Another interesting quantitative parameter is the total number of messages needed to disseminate messages to a certain percentage of hosts. In particular, we now estimate the number of replicas sent, per message, in the case of infinite buffers (i.e., $\gamma = 0$). This is the case of systems which are characterized by well-dimensioned buffers or where the traffic is low so the buffers are able to store all the incoming messages without the necessity of freeing space for them.

Considering an infection process repeated for a number of times equal to r number of rounds, indicating with t_r the time length of the r^{th} round, the total number of replicas per single type of message can be estimated as follows:

$$\textit{NumberOfReplicas} = \int_{t=0}^{t=t_r} \lambda\langle k \rangle I(t)dt \tag{7}$$

By substituting the value of $I(t)$ (obtained by solving the system (3)), we solve the integral obtaining the following estimation for the number of replicas:

$$\textit{NumberOfReplicas} = Nln(1 + \frac{1}{N}(e^{\lambda\langle k \rangle t_r} - 1)) \tag{8}$$

This can be approximated as follows:

$$\textit{NumberOfReplicas} = O(N\langle k \rangle) \tag{9}$$

It is interesting to note that in the case of a fully meshed network (i.e., all the hosts are in the transmission range), we obtain the worst case approximation:

$$\textit{NumberOfReplicas} = O(N^2) \tag{10}$$

Another interesting case is when $\langle k \rangle \approx lnN$. In this case the number of replicas is approximately linear:

$$\textit{NumberOfReplicas} = O(N) \tag{11}$$

Finally, if $\langle k \rangle$ is not dependent from N[7], the number of messages remains approximately *constant* as N increases[8].

5.2 Simulation Results

We evaluated the proposed system and model by considering the case of unicast communication with a given reliability specified by the user. We do not consider the case of anycast communication, since, as discussed, it relies on the same delivery process.

5.2.1 Description of the Simulation

In order to test the performance of these techniques we considered a mobile scenarios composed of a realistic number of hosts and we implemented and ran a series of simulations by using the popular open source discrete-event simulator OMNeT++ [19]. We defined a square simulation area with a side of 1 km and a transmission range equal

[6]It is interesting to note that, in theory, the message dissemination would continue also after the expiration time. However, since the replicas are deleted from the buffer after the expiration time, the epidemic process terminates.

[7]This is the case of scenarios where the hosts occupy a larger area as the population increases, so that the density of population and, consequently, $\langle k \rangle$ remain approximately constant.

[8]This result can be directly derived by applying L'Hospital's rule to calculate the limit.

to 200 m. The simulation was set to run several replicates for each mobile scenario in order to obtain a statistically meaningful set of results (with a maximum 5% error). The intervals between each message are modeled as a Poisson process. We studied scenarios characterized by different number of hosts (more precisely 32, 64, 96, 128). These input parameters model typical deployment settings of mobile ad hoc networked systems. We do not model explicitly the failures in the system, since we assume that during the infection process, the number of hosts remains constant.

All the messages are sent in the first 20 seconds in order to create enough traffic to saturate the buffer. The sender and receiver of each message are chosen randomly. We tested the algorithm with both finite and infinite (i.e. equal to 100) buffer size. The buffer for each node is set to 100 messages (i.e, infinite buffer), unless otherwise specified. The execution interval of the epidemic spreading procedure (presented in the box Program 2) is 10 seconds. The expiration time (i.e., the value of `time`) is equal to 10 minutes. Therefore, the number of rounds is 60. We simulated only the cases with the basic reproductive ratio R_0 greater than 1, since the middleware primitives simply return an exception if this value does not reach the threshold.

The movements of the hosts are generated by using a Random Way-Point mobility model [7]; every host moves at a speed that is randomly generated by using a uniform distribution. The range of the possible speeds is $[1, 6]m/s$. We selected this mobility model, since as discussed in [12], its emergent topology has an exponential structures, with Poisson-like distributions. Therefore, in this scenario, the properties of the network can be studied with a good approximation by assuming a homogeneous networks model. The accuracy of the approximation increases as the density of population increases, since, considering the finite and limited simulated time, we obtain a scenario characterized by a time series of degree of connectivity values characterized by lower variance. Moreover, the so-called border effects, due to the host that moves at the boundaries of the simulated scenarios, have less influence as the density of population increases. This also means that as the number of failures in the system increases, the accuracy of the model decreases. In fact, considering uniformly randomly distributed failures, a scenario composed of 32 nodes can be used to model the case of a scenario with an initial number of 64 nodes, where half of them have failed. Figure 1 shows the distribution of the degree of connectivity in the simulated scenarios composed of different numbers of hosts.

5.2.2 Analysis of Simulation Results

In this subsection we will analyze the results of our simulations, discussing the performance of the proposed techniques. We will study the variations of some performance indicators, such as the delivery ratio and the number of messages sent as functions of the density of hosts (i.e., the number of the hosts in the simulation area), considering different buffer size (and consequently different removal rates).

Figure 2 shows a comparison with the estimated epidemic spreading (i.e., the number of infectives $I(t^*)$) and the data obtained from the simulation of a mobile scenario composed of 128 nodes, with $t^* = 10min$ and $\gamma = 0$. It is interesting to note that the values of the theoretical curve are higher than the experimental ones. This is due to the fact that the degree of connectivity is not perfectly homogeneous in the simulated scenarios. For example, if a message is sent by a host that has a degree of connectivity $\overline{k} > \langle k \rangle$, the value of β will be lower than the infectivity associated to the average degree of connectivity $\langle k \rangle$[9].

Figure 3 and 4 show the delivery ratio in terms of population density, for the case of a desired reliability equal to 100 and 50, respectively, with $t^* = 10min$ and $\gamma = 0$. The obtained delivery ratios are really close to the values expected from our model analysis. Also in this case, the better approximation of the assumption of homogeneous network, obtained when the density of population increases, leads to better results (i.e., a more accurate estimation) for the case of 128 nodes. Figure 5 and 6 show the number of messages as function of population density. This confirms the analytical results presented in Section 5.1.2. In fact, the curve is approximately linear, as justified by the fact that, in our simulations $\langle k \rangle \ll N$. The number of replicas per host per message are plotted in Figure 7 and 8. These diagrams illustrate the scalability of our approach, since the number of replicas per host per message can be approximated as $O(\langle k \rangle)$. The influence of the buffer size is presented in Figure 9 and Figure 10. The first shows the comparison between the cases of infinite and limited (with a size equal to 20) buffers. The effect of the non perfect network homogeneity is present also here and is more evident for the scenarios composed of a lower number of hosts. In fact, if the actual degree of connectivity is higher than the assumed $\langle k \rangle$ the probability of deletion of messages from the buffer increases. In this case, the assumptions at the basis of the model in (3) are not valid. Figure 10 shows that the number of messages is greater than in the case of infinite buffers. In fact, an increased infectivity is needed in order to spread the messages also in presence of the removal phenomena, due to the limited buffer size.

6 Related Work and Discussion

In this section, we compare our solution with existing work and applications of the proposed model and outlining our current research directions.

[9] From a practical point of view, in order to cope with this issue, it is sufficient to increase β, for example by adding a correction equal to a percentage of the value calculated by using the theoretical model. However, for illustration purposes, in the simulations presented in the remainder of this paper, we used values of β derived directly from the model presented in Section 3 without corrections.

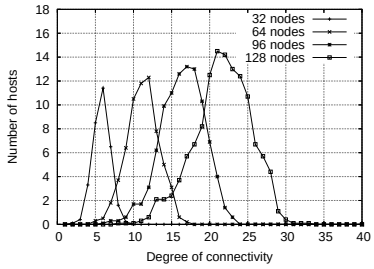

Figure 1. Distribution of the degree of connectivity in the simulated mobile scenarios.

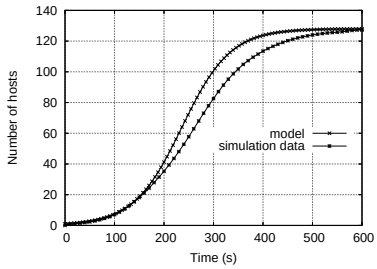

Figure 2. Comparison between epidemics model curve and simulation data of infection spreading in the 128 hosts scenario with desired reliability equal to 100, $t^* = 10min$ and $\gamma = 0$.

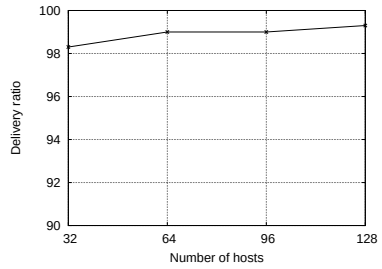

Figure 3. Delivery ratio Vs population density with desired reliability equal to 100 and $\gamma = 0$.

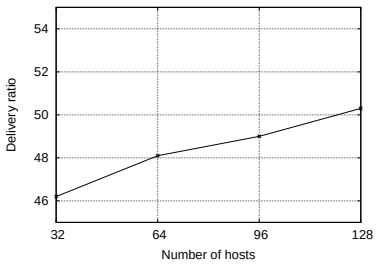

Figure 4. Delivery ratio Vs population density with desired reliability equal to 50 and $\gamma = 0$.

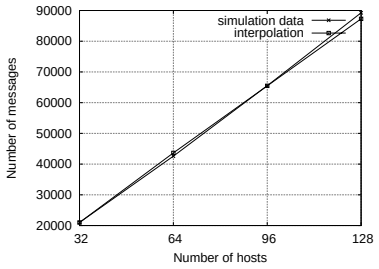

Figure 5. Number of messages Vs population densFity with desired reliability equal to 100 and $\gamma = 0$.

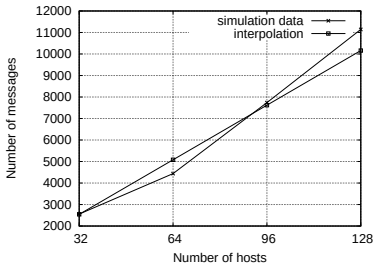

Figure 6. Number of messages Vs population density with desired reliability equal to 50, $t^* = 10min$ and $\gamma = 0$.

6.1 Comparison with the State of the Art

As far as mobile systems are concerned, a first study of the possible application of epidemic techniques in MANETs is presented in [18] by Vahdat and Becker. Many refinements of this approach have been proposed. A study of the information dissemination based on epidemic models in mobile ad hoc networks is presented in [14]. However, the authors discuss only a theoretical framework, without proposing concrete implementation of the model. Moreover, they do not take into account the influence of the structure of the network in the dissemination process.

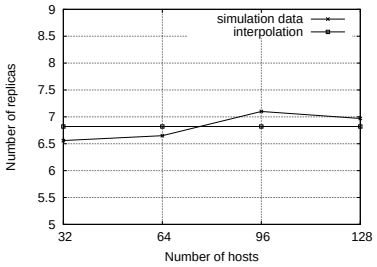

Figure 7. Number of replicas per host per message Vs population density with desired reliability equal to 100, $t^* = 10min$ and $\gamma = 0$.

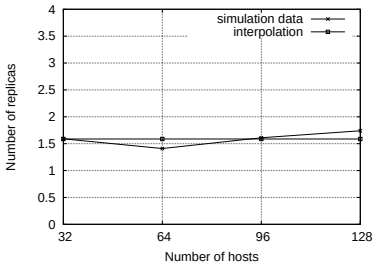

Figure 8. Number of replicas per host per message Vs population density with desired reliability equal to 50 and $\gamma = 0$.

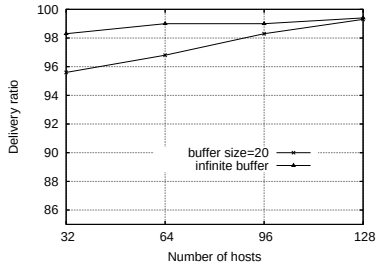

Figure 9. Influence of the buffer size in the 128 hosts scenario with desired reliability equal to 100 and $t^* = 10min$: delivery ratio Vs population density with buffer size equal to 20.

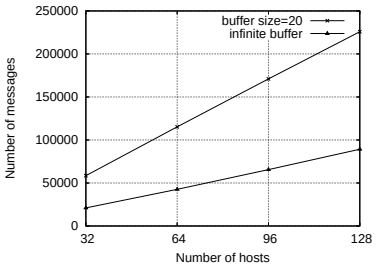

Figure 10. Influence of the buffer size in the 128 hosts scenario with desired reliability equal to 100: number of messages Vs population density with buffer size equal to 20.

Epidemic-style techniques have been applied to the design of publish-subscribe systems for highly dynamic environments; two recent interesting examples of such systems are presented in [8] and [3]. Our approach can be used to improve the performance of this class of systems in terms of resource consumption, since it allows for a precise tuning of the dissemination process.

In terms of more general distributed systems, the seminal paper on the application of epidemic techniques is [9], where these algorithms are used to maintain consistency in replicated databases. A general introduction to epidemic algorithms for information dissemination in distributed systems can be found in [10]. Much work addressing different faces of the problem have been proposed, including the remarkable contributions presented in [5, 11]. In general, in these works, the authors consider the structure of the underlying network topology only marginally, or from empirical and experimental perspectives. A notable exception is [15], where the authors discuss the application of the Harari graphs to the design of protocols for broadcasting in fixed networks.

With respect to these works, the novelty of this paper resides in the evaluation of the structure of the network by using accurate models to control and tune the dissemination process according to a desired reliability. We also underline that the design of our system is based on theoretical results confirmed by experimental evidence, whereas in some the existing works, mathematical models are only used to understand the emergent behavior of the system *a posteriori*. Moreover, up to our knowledge, this work can be considered the first concrete application of the recent results on epidemics spreading in complex networks [4].

We believe that these epidemic techniques should be applied only in the cases where useful context information cannot be inferred. In another work [17], we have applied prediction techniques to adapt the communication mechanisms by evaluating the evolution of the mobile scenarios.

6.2 Relaxing the Assumption of Homogeneous Networks

The results and the solutions discussed in this paper rely on the assumption of homogeneous networks, that are emerging from the random movements of the nodes. We now show that the proposed approach can be generalized extended to the general case of heterogeneous networks. These structures are emerging in presence of small clusters of people or communities. The results that we are going to present for the case of heterogeneous networks are also valid for homogeneous network, since the latter can be treated as a particular case of the former.

For heterogeneous networks the approximation $k \approx \langle k \rangle$ is not valid. However, the same probabilistic communication primitives introduced in Section 4 could be used, with a different semantics. This relies on the following observations: given k fluctuating in the range $[k_{MIN}, k_{MAX}]$, we observe that for a value of the infectivity corresponding to $k = k_{MIN}$, the obtained spreading of the infection $I(t^*, k_{MIN})$ will satisfy the following property:

$$I(t^*, k) > I(t^*, k_{MIN}) \quad \forall k \in]k_{MIN}, k_{MAX}] \qquad (12)$$

In other words, if k_{MIN} is selected in the calculation of the value of the infectivity, the value of `Reliability` can be considered approximately as a guaranteed lower bound of the reliability level. The value of k_{MIN} can be dynamically retrieved and set by the middleware by monitoring the connectivity of the host in mobile systems. We plan to investigate these adaptive mechanisms further in the future.

7 Concluding Remarks

In this paper, we have introduced middleware primitives for controlled information dissemination in mobile ad hoc networks, which relies on optimized epidemic-style techniques. With respect to unicast communication, we have showed that protocols that statistically ensure the desired reliability level for the case of homogeneous networks can be designed. We have also showed that these results may be applied to the case of anycast and multicast communication to tune and optimize the replication process. We have evaluated our approach through simulation and have presented a possible generalization of the model discussing the relaxation of the assumption of homogeneous networks.

Acknowledgements: We would like to acknowledge the support of EPSRC CREAM and EU RUNES Projects.

References

[1] R. Albert and A.-L. Barabasi. Statistical Mechanics of Complex Networks. *Review of Modern Physics*, 74:47–97, 2002.

[2] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1992.

[3] S. Baehni, C. Chabra, and R. Guerraoui. Frugal Event Dissemination in a Mobile Environment. In *Proceedings of ACM Middleware'05*, 2005.

[4] M. Barthélemy, A. Barrat, R. Pastor-Satorras, and A. Vespignani. Dynamic Patterns of Epidemic Outbreaks in Complex Heterogeneous Networks. *Journal of Theoretical Biology*, 2005.

[5] K. P. Birman, M. Hayden, O. Ozkasp, Z. Xiao, M. Budiu, and I. Minsky. Bimodal Multicast. *ACM Trans. on Computer Systems*, 17(2):41–88, 1999.

[6] B. Bollobas. *Random Graphs*. Cambridge University Press, Second edition, 2001.

[7] T. Camp, J. Boleng, and V. Davies. A Survey of Mobility Models for Ad Hoc Network Research. *Wireless Communication and Mobile Computing*, 2(5):483–502, 2002.

[8] P. Costa and G. P. Picco. Semi-probabilistic Content-Based Publish-Subscribe. In *Proceedings of ICDCS'05*, pages 575–585, 2005.

[9] 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. *ACM SIGOPS Operating Systems Review*, 22(1), January 1988.

[10] P. T. Eugster, R. Guerraoui, A.-M. Kermarrec, and L. Massouli. Epidemic Information Dissemination in Distributed Systems. *IEEE Computer*, May 2004.

[11] P. T. Eugster, S. Handurukande, R. Guerraoui, A.-M. Kermarrec, and P. Kouznetsov. Lightweight Probabilistic Broadcast. *ACM Trans. on Computer Systems*, 21(4):341–374, 2003.

[12] I. Glauche, W. Krause, R. Sollacher, and M. Greiner. Continuum Percolation of Wireless Ad Hoc Communication Networks. *Physica A*, 325:577–600, 2003.

[13] M. Jelasity and A. Montresor. Epidemic-style proactive aggregation in large overlay networks. In *Proceedings of ICDCS'04*, pages 102–109, Tokyo, Japan, Mar. 2004. IEEE Computer Society.

[14] A. Khelil, C. Becker, J. Tian, and K. Rothermel. An Epidemic Model for Information Diffusion in MANETs. In *Proceedings of ACM MSWiM'02*, September 2002.

[15] M.-J. Lin, K. Marzullo, and S. Masini. Gossip Versus Deterministically Constrained Flooding on Small Networks. In *Proceedings of DISC 2000*, pages 253–267, October 2000.

[16] C. Mascolo, L. Capra, and W. Emmerich. Middleware for Mobile Computing (A Survey). In *Networking 2002 Tutorial Papers*, 2497, pages 20–58. E. Gregori and G. Anastasi and S. Basagni, 2002.

[17] M. Musolesi, S. Hailes, and C. Mascolo. Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks. In *Proceedings of the IEEE WoWMoM 2005. Taormina, Italy.* IEEE press, June 2005.

[18] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks. Technical Report CS-2000-06, Dept. of Computer Science, Duke University, 2000.

[19] A. Varga. The OMNeT++ Discrete Event Simulation System. In *Proceedings of ESM'01*, Prague, 2001.