

Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks

Mirco Musolesi, Stephen Hailes, Cecilia Mascolo
Department of Computer Science, University College London
{m.musolesi|s.hailes|c.mascolo}@cs.ucl.ac.uk

Abstract

The vast majority of mobile ad hoc networking research makes a very large assumption: that communication can only take place between nodes that are simultaneously accessible within in the same connected cloud (i.e., that communication is synchronous). In reality, this assumption is likely to be a poor one, particularly for sparsely or irregularly populated environments.

In this paper, we present the Context-Aware Routing (CAR) algorithm. CAR is a novel approach to the provision of asynchronous communication in partially-connected mobile ad hoc networks, based on the intelligent placement of messages. We discuss the details of the algorithm, and then present simulation results demonstrating that it is possible for nodes to exploit context information in making local decisions that lead to good delivery ratios and latencies with small overheads.

1 Introduction

Mobile ad hoc networks represent the purest form of decentralised systems and, therefore, they impose many challenges to cooperative communication. As a consequence, much ad hoc network research has focused on the investigation of fundamental algorithms for routing on which almost everything else relies. However, in order to make the problem tractable, almost all research on routing algorithms makes the oversimplistic assumption that it is only meaningful to attempt to exchange messages within connected clouds of nodes, in other words, that all communication is synchronous in nature. This assumption is overly constrained if one considers that there is a strong requirement for communication that is asynchronous in nature, as argued above. In such a case, the delay tolerant character of the traffic allows useful communication to still occur by using nodes moving between disconnected groups of nodes (clouds) to transport messages from one cloud to another. Thus, it is perfectly possible that two nodes may *never* be part of the same connected cloud and yet may still be able to exchange delay tolerant information by making use of

predicted mobility patterns as an indicator of which other nodes might make good carriers for this information.

In the absence of special information, the problem of predicting which nodes might make good carriers in ad hoc networks is a very challenging one. Likely future mobility patterns must be inferred from past mobility patterns, but this alone is inadequate; parameters such as remaining battery lifetime are also key in determining which potential carriers are most likely to result in successful delivery. In this paper, we consider what types of information are available to nodes in deciding on a carrier. We use this analysis in the design of a Context-aware Adaptive Routing algorithm (CAR), a general framework for the evaluation and prediction of context information, aimed at achieving efficient and timely delivery of messages. Using simulations, we explore the performance of the CAR algorithm with respect to epidemic routing [13] and flooding. Whilst, in the developed world, synchronous communication (in the form of phone and Internet) is generally cheap and easy to come by, there are several real scenarios in less developed parts of the world in which different portions of a logical network are physically disconnected [4, 12].

This paper is organized as follows. Section 2 presents our approach. The details related to the evaluation of context information are discussed in Section 3. The description of the simulations carried out to evaluate CAR is provided in Section 4, together with an analysis of the results. Finally, in Section 5 we compare CAR with previous work in this area.

2 Context-aware Routing for Mobile Ad Hoc Networks

The Context-Aware Routing algorithm is built on the assumption that the only information a host has about its position is logical connectivity information. In particular, we assume that a host is not aware of its absolute geographical location nor of the location of those to whom it might deliver the message. Although this information could potentially be useful, and, indeed, we plan to examine its utility in the near future, it is currently unreasonable to assume the existence of GPS for all potential application domains for

this technology. Another basic assumption is that the hosts present in the system cooperate to deliver the message. In other words, we do not consider the case of hosts that may refuse to deliver a message.

The delivery process depends on whether or not the recipient is present in the same cloud as the sender. If both are currently in the same connected portion of the network, the message is delivered using the underlying synchronous routing protocol to determine a forwarding path. In the remainder of this paper we assume that a proactive routing protocol is used (in our simulations we employed DSDV [10]). Reactive protocols require different approaches to optimisation that would simply confuse the presentation and so are deemed to be outside the scope of this particular work.

If a message cannot be delivered synchronously, the best carriers for a message are those that have the highest chance of successful delivery, i.e., the highest *delivery probabilities*. The message is sent to one or more of these hosts using the underlying synchronous mechanism. Delivery probabilities are synthesized locally from context information such as the rate of change of connectivity of a host (i.e., the likelihood of it meeting other hosts) and its current energy level (i.e., the likelihood of it remaining alive to deliver the message). We define *context* as the set of attributes that describe the aspects of the system that can be used to optimize the process of message delivery. Since we assume a proactive routing protocol, every host periodically sends both the information related to the underlying synchronous routing (in DSDV this is the routing tables with distances, next hop host identifier, etc.), and a list containing its delivery probabilities for the other hosts. When a host receives this information, it updates its routing tables. With respect to the table for asynchronous routing, each host maintains a list of entries, each of which is a tuple that includes the fields (*destination*, *bestHost*, *deliveryProbability*). In this paper, we choose to explore the scenario in which each message is placed with only a single carrier rather than with a set, with the consequence that there is only a single list entry for each destination. When a host is selected as a carrier and receives the message, it inserts it into a buffer. The size of this buffer is important, and represents a trade-off between storage overhead and likely performance.

3 Prediction and Evaluation of Context

The general problem from the point of view of the sender of a message is to find the host with the best delivery probability, as calculated using the predicted values of a range of context attributes. Instead of using the available context information as it is, CAR is optimized by using *predicted* future values for the context, so to have more realistic values. The process of prediction and evaluation of the context information can be summarized as follows.

- Each host calculates its delivery probabilities. This process is based on the *prediction* of the future values of the attributes describing the context (see Section 3.2) and on the *composition* of these estimated values using multi-attribute utility theory [7] (see Section 3.1). The calculated delivery probabilities are periodically sent to the other hosts in the connected cloud as part of the update of routing information.
- Each host maintains a logical forwarding table of tuples describing the next logical hop, and its associated delivery probability, for all known destinations.
- Each host uses local prediction of delivery probabilities between updates of information. The prediction process is used during temporary disconnections and it is carried out until it is possible to guarantee a certain accuracy. Moreover, in the case of hosts within reach, the interval between update shipment is based on the possibility or not to make accurate predictions. In other words, hosts send updates only when the evolution of the mobile scenario follows a certain trend. This is done by evaluating the sampled values of context information.
- If a host does not store any information about the message recipient, it sends the message to the host in the cloud that has the highest mobility (or it keeps the message if it has the highest mobility). In other words, the message is stored by the host that in average meets the highest number of hosts and so that has the highest probability to get in reach with the recipient.
- If the carrier, while moving, meets a host with a higher delivery probability, the message is transferred to the host with the higher delivery probability.

3.1 Local evaluation of context information

There are several techniques that can be used to combine and evaluate the multiple dimensions of context in order to decide which nodes are the best candidates for carrying a particular message. The simplest is to allow application developers to define a static hierarchy among the predicted context attributes.

A possible alternative to this method is to use goal programming, exploiting the so-called *preemptive methodology*. With respect to a single attribute, our goal is to maximize its value. The optimization process is based on the evaluation of one goal at a time such that the optimum value of a higher priority goal is never degraded by a lower priority goal [11]. However, in general, the definition of static priorities is inflexible. For more realistic situations, we expect to need to attempt simultaneous maximization of a range of different attributes, as opposed to using a predefined hierarchy of priorities.

Significance-based evaluation of context-aware information The priority based technique just mentioned seems too simplistic because, in general, our decision problem involves multiple conflicting objectives [7]. For example, if we wish to determine which host has the best delivery probability, considering both the battery energy level and the rate of change of connectivity, it may happen that the host char-

acterized by the highest mobility has scarce residual battery energy and vice versa. In general, maximization across all parameters will not be possible and, instead, we must trade off the achievement of one objective (i.e., the maximization of a single attribute) against others.

The context information related to a certain host can be defined using a set of attributes (X_1, X_2, \dots, X_n) . Those attributes denoted with a capital letter (e.g., X_1) refer to the set of all possible values for the attribute, whereas those denoted with a lower case letter (e.g., x_1) refer to a particular value within this set. In the remainder of this section we will use the classical notation of utility theory. Our goal is to allow each host locally to associate a utility function $U(x_1, x_2, \dots, x_n)$, representing the delivery probability, with every other host.

Our aim is to maximize each attribute, in other words, to choose the host that presents the best trade-off between the attributes representing the relevant aspects of the system for the message delivery. The combined goal function used in the Weights method can be defined as

$$\text{Maximise}\{f(U(x_i)) = \sum_{i=1}^n w_i U_i(x_i)\}$$

where w_1, w_2, \dots, w_n are *significance weights* reflecting the relative importance of each goal.

It is worth noting that, in our case, the solution is very simple, since it consists in the evaluation of the function $f(U_1, \dots, U_n)$ using the values predicted for each host and in the selection of the host i with the maximum such value.

Autonomic adaptation of the utility function As it stands, the utility function weights are fixed in advance, reflecting the relative importance of the different context attributes. However, such a formulation is still too static, since it fails to take into account the values of the attributes. Thus, for example, a small drop in battery voltage may be indicative of the imminent exhaustion of the battery; consequently, it would be useful to reduce the weight of this attribute non-linearly to reflect this.

In general, we wish to adapt the weights of each parameter *dynamically* and in ways that are dependent on the values of those parameters. In other words, we need a runtime self-adaptation of the weightings used for this evaluation process that could be categorized as a typical autonomic mechanism [1]. A simple solution to this problem is the introduction of adaptive weights a_i into the previous formula, in order to modify the utility function according to the variation of the context.

$$\text{Maximise}\{f(U(x_i)) = \sum_{i=1}^n a_i(x_i) w_i U_i(x_i)\}$$

$a_i(x_i)$ is a parameter that may itself be composite. For our purposes, we define it to have three important aspects that

help to determine its value, though the model could easily be expanded to incorporate other aspects deemed to be of importance:

- a) Criticality of certain ranges of values, $a_{range_i}(x_i)$
- b) Predictability of the context information, $a_{predictability_i}(x_i)$
- c) Availability of the context information, $a_{availability_i}(x_i)$

We now compose the a_i weights as factors in the following formula:

$$a_i(x_i) = a_{range_i}(x_i) \cdot a_{predictability_i}(x_i) \cdot a_{availability_i}(x_i)$$

Adaptive weights related to the possible ranges of values assumed by the attributes We can model the adaptive weights $a_{range}(x_i)$ as a function whose domain is $[0, 1]$. For example, with respect to the battery energy level (modeled using the percentage of residual battery energy), we would use a monotonically decreasing (though not necessarily linear) function to assign a decreasing adaptive weight that is, in turn, used to ensure that the corresponding utility function decreases as the residual energy tends towards zero.

Adaptive weights related to the predictability of the context information In general, it is possible to exploit different statistical attributes for the analysis of time series. One could, for example, use the *autocorrelation function* to describe the degree of association between values of the time series at different lags. In short, this gives a measure of the predictability of the context information. Furthermore, there are clear guidelines for adapting the use of the autocorrelation function for non-stationary data with both trends and seasonal variations.

In building the autocorrelation function, we first consider the auto-covariance of the time series at lag k . The *lag* represents the time difference (in terms of the number of samples) between the two instants being considered. Therefore, we use the *autocorrelation coefficient* ρ_k , at lag k defined as follows

$$\rho_k \equiv Cov(X_t, X_{t+k}) / Var(X_t)$$

It is worth noting that it is possible to prove that $0 \leq |\rho_k| \leq 1$. The absolute value of ρ_k is exactly 1 for a perfect autocorrelation, whereas an autocorrelation coefficient close to zero (either positive or negative) indicates little or no correlation between two samples X_t and X_{t+k} (i.e., the impossibility of making predictions). In the case of a so-called random series, for a large number n of samples, the value of ρ_k is approximately equal to 0. We therefore determine parameter $a_{predictability_i}$, thus $a_{predictability_i} = |\rho_k|$. An interesting issue is the choice of the value of the lag k . It is possible that autocorrelation signals will drift slowly over time and, consequently, the value of k will also need to change to reflect this. However, we expect the underlying processes that

determine the nature of the original signal to change slowly if at all.

Thus, in order to adapt the lag value to retain a strongly correlated signal, we adopt a very simple adaptive technique. At the initial instant t_0 , k is set to 1. This is increased, up to a value of k_{MAX} , if the autocorrelation coefficient is below a given lower bound threshold $\rho_{strongCorrLB}$. The process wraps on reaching k_{MAX} , setting the value of k back to unity in order to ensure that the entire space is searched. If, on the other hand, the autocorrelation coefficient exceeds an upper bound threshold $\rho_{strongCorrUB}$, k is decreased until it reaches the value 1. We can summarize these concepts using the following update equation for the lag k :

$$k(t+1) = \begin{cases} 1 & \text{if } t = t_0 \text{ or } k(t) = k_{MAX} \\ k(t) + 1 & \text{if } \rho(t) \leq \rho_{strongCorrLB}, k(t) < k_{MAX}, t \neq t_0 \\ k(t) - 1 & \text{if } \rho(t) > \rho_{strongCorrUB}, k(t) \geq 1, t \neq t_0 \\ k(t) & \text{otherwise} \end{cases}$$

Adaptive weights related to the availability of the context information It is unreasonable to assume that all context attributes have the same degree of availability. Thus, we expect to have a time-varying set of attributes available whose values are known or predictable. Attributes may drop out of this set if meaningful values can no longer be predicted for them, since the information on which the prediction would have been based is too old. The simplest approach to this problem is to ensure that missing context information carries an adaptive weight a_i equal to 0:

$$a_{availability_i} = \begin{cases} 1 & \text{if the context information is currently available} \\ 0 & \text{if the context information is not currently available} \end{cases}$$

Formally, to date, we have implicitly assumed that a static set of attributes will be defined. However, it is worth noting that, using this approach, we can dynamically incorporate new attribute values, simply by assuming that they were always there, but had zero weight for $a_{availability_i}$.

Automatic adaptation of the refresh period of routing tables and context information In wired networks, routing table state update is often done on an unvarying regular basis as well as on a by-need basis. However, this approach is wasteful in mobile ad hoc environments. Thus, we consider how to adapt the rate of context information dissemination by noting that we already know that such information is predicted by recipients and that such predictions are likely to be most accurate when the signal on which they are based is most predictable. Thus, a possible function for the determination of refresh time is given by:

$$t(x_1, x_2, \dots, x_n) = c \sum_{i=1}^n |\rho_{k_i}|$$

where c is a constant of proportionality.

There are several possible extensions of this model. For example, one might wish to take account of the absolute value of a parameter in determining update rates. Thus, for example, as battery energy levels decline, one might wish to update information increasingly less frequently despite the consequent unpredictability at the other end, in order to conserve remaining energy. If information at the recipients becomes totally outdated, then $a_{availability_i}$ will be set to zero for all our attributes and the result is that we will not be likely recipients of messages to transfer, which is in line with the behavior we would expect. Thus, we could replace the simple constant in the above equation with a generic function of values of individual attributes. Likewise, we could obtain a more refined model associating different weights with the autocorrelation coefficient for each attribute in a way similar to that applied previously for composing the utility functions for evaluating which host has the best message delivery probability.

3.2 Prediction of the context information attributes using Kalman filters

Kalman filter prediction techniques [6] were originally developed in automatic control systems theory. These are essentially a method of discrete signal processing that provides optimal estimates of the current state of a dynamic system described by a *state vector*. The state is updated using periodic observations of the system, if available, using a set of *prediction recursive equations*.

Kalman filter theory is used in CAR both to achieve more realistic prediction of the evolution of the context of a host and to optimize the bandwidth use. As discussed above, the exchange of the routing tables that store the delivery probabilities is a potentially expensive process, and unnecessarily so where such information is relatively easily predictable.

If it is possible to predict future values of the attributes describing the context, it is possible to update the delivery probabilities stored in the routing tables, even if fresh information is unavailable. Fortunately, it is possible to express this prediction problem in the form of a state space model. We have a time series of observed values that represent context information. From this it is possible to derive a prediction model based on an inner state that is represented by a set of vectors, and to add to this both trend and seasonal components [2]. It is worth noting that one of the main advantages of the Kalman filter is that it does not require the storage of the entire past history of the system, making it suitable for a mobile setting in which memory resources may potentially be very limited.

4 Simulation and Results

We evaluated the CAR algorithm by using the OmNet++ discrete event simulator [14]. In order to obtain credible results and to test the peculiar characteristics of our protocol, it was also necessary for us to develop a new group mobility

model, that will be presented in Section 4.

Description of the simulation For reasons of space and in order to allow for fair comparison with existing research, we report results based on simulations that use only part of the full generality of the CAR algorithm. Thus, we simulated the CAR model using a utility function based on the evaluation of two attributes: (i) the change rate of connectivity and (ii) the probability of being located in the same cloud as the destination. We made the assumption that these factors have the same relevance, so assigned them the same weights in the evaluation of the overall utility (i.e., $w_i = 0.5$). Moreover, we also assumed that all the possible values in the range had the same importance (i.e., $a_{range_i}(x_i) = 1$) and that the values of attributes are always available during the simulation (i.e., $a_{availability_i}(x_i) = 1$). The change rate of connectivity attribute is locally calculated by examining the percentage of a node's neighbors that have changed their connectivity status (connected to disconnected, or vice versa) between two instants. The collocation attribute measures the percentage of time that two hosts have been in reach. To calculate it, we periodically run a Kalman filtering process, assuming that the value is 1 if the host is currently in reach or 0 if not. Clearly, the resultant predicted values will be in the range $[0, 1]$ and they will directly express an estimation of the probability of being in reach of the host in the future.

We implemented a simplified version of the DSDV protocol [10] in order to simulate and test the synchronous delivery in connected portions of the network, as described in Section 2. Each host maintains a *routing and context information table* used for asynchronous and synchronous (DSDV) routing. Each entry of this table has the following structure:

```
(targetHostId, nextHopId, dist, bestHostId, delProb)
```

The first field is the recipient of the message, the second and the third are the typical values calculated in accordance with the DSDV specification, whereas the fourth is the identifier of the host with the best delivery probability, the value of which is stored in the last field. The overhead due to the addition of this field to the standard format of routing tables of DSDV is negligible. It is worth noting that all the automatic mechanisms, such as the variable refresh period of routing tables, described previously, were implemented.

We also simulated flooding and the epidemic protocols in order to provide comparators for the performance of the CAR solution. The implementation of the epidemic protocol follows the description presented in [13].

We evaluated the performance of each protocol sending 100 messages with a simulation time equal to 300 seconds. The messages were sent after 40 seconds, in order to allow for the settling of initial routing table exchanges, and the intervals between each message were modeled as a Poisson

process, with $\lambda = 5s^{-1}$, and the consequence that all messages are sent in about 20 seconds. The sender and receiver of each message are chosen randomly.

In the CAR simulation, each message has a field that is similar to a *time to live* value that is decreased each time that the message is transferred to another host (the initial value being 15). Moreover, in this case, we also introduced a *split horizon* mechanism to prevent messages from being retransmitted unnecessarily. The buffer for each node was set to 20 messages, unless otherwise specified.

The one key aspect of the simulation not yet addressed is that of the mobility model. Clearly, the random way-point mobility model, which is used extensively in such studies, largely for reasons of simplicity, does not accurately reflect human behaviour and annihilate the effect of the prediction since movement is entirely random. Consequently, we devised a new group-based mobility model [9]. This is presented briefly in the following section.

Mobility model Mobility models that assume that individuals move independently of one another in random ways are unrealistic in terms of the deployment scenarios for ad hoc networks that are most commonly expounded. For example, on a battlefield, it would be indicative of a very troubled army if each soldier were to move randomly with respect to all others. Thus, we have developed a new model with a form of hierarchical clustering that better reflects the ways in which collections of people are structured at an organizational level and, consequently, the ways in which they move [9]. Thus, we introduce the concept of a collection of nodes, which has its own motion overlaid on a form of random motion within the cloud. By parameterizing this model differently, we can represent different archetypes: for example, one would expect to use different parameters for an academic who spends her life traveling between home and the university, interacting with a very closed set of people, as opposed to a salesman who travels much more extensively and interacts less discriminately. A host that belongs to a cloud moves inside it towards a goal (i.e., a point randomly chosen in the cloud space) using the standard random way-point model. When a host reaches a goal, it also implicitly reaches a decision point about whether to remain within the cloud, and, if leaving, to where it should go. Each of these decisions is taken by generating a random number and comparing it to a threshold (which is a parameter of the model). It is worth noting that clouds also move towards randomly chosen goals in the simulation space.

Simulation Configuration 50% of the hosts are initially placed randomly in a cloud, whereas the others are positioned randomly in the simulation area. Each cloud is defined using a squared area with a side length of 200 m. In other words, we randomly select the point $(minX, minY)$ that, together with the length of the side, defines the cloud area. For these simulations, there is only a single level

of cloud. Every host is characterized by two values, P_{escape} , indicating the probability of escaping from the current cloud, and $P_{escapeCloud}$ describing the probability of choosing a new goal in the space between clouds. Each cloud moves with a random speed (with a value in the range 1-2 m/s); moreover, each host moves with a randomly generated different speed (with a value in the range 1-3 m/s). It is worth noting that the movement of a host is the result of the composition of these speeds. In our simulation, the positions of all the hosts and clouds are updated every second. When a cloud reaches its goal, a new goal is chosen in the simulation space. When a host reaches its goal, a threshold probability $P_{escapeThreshold}$ is generated randomly (its range is clearly $[0, 1]$). If its P_{escape} is greater than $P_{escapeThreshold}$ the new goal is chosen outside the current cloud, else inside. If outside, we randomly generate $P_{escapeCloudThreshold}$ and compare it to $P_{escapeCloud}$ to determine whether or not the goal should be chosen in some other cloud or in the open space between clouds. For those hosts that are already outside a cloud, the choice of a new goal is done in an analogous way.

Analysis of results In this subsection we will analyze the results of our simulations, comparing the performance of CAR with the flooding and epidemic protocols. We will discuss the variation of some performance indicators as functions dependent on the density of hosts (i.e., the number of the hosts in the simulation area) and the size of the buffers used to store messages in both the epidemic and CAR.

In Figure 1.a, there is a comparison between the delivery ratios of the three protocols in each of three different scenarios (with 16, 24 and 32 hosts). In all cases, the number of messages that may coexist within a node's buffer is unconstrained.

CAR achieves a performance between that of flooding and epidemic routing, as expected. Flooding suffers from the inability to deliver messages to recipients that are in other clouds when the messages are sent but is here simply as a comparator to demonstrate the numbers of messages being delivered that cannot be delivered directly, because the recipient is in a cloud different from the cloud of the sender. The epidemic protocol can be considered optimal in terms of delivery ratio, simply because each message is propagated to all accessible hosts, all of which have buffers large enough to hold it. In CAR, we have chosen to operate under the most stringent conditions: there is only ever a single copy of each message, which represents the worst case for this protocol. Clearly, it would be possible to trade off a small amount of intelligent replication (to improve the delivery ratio) against an increase in overhead.

The dependency of the delivery ratios on the buffer size is similar for all the protocols (see in Figure 1.b the results for the 32 hosts scenario). Both of these demonstrate a substantial degradation of their performance as buffer size de-

creases; however, this phenomenon is more evident in the epidemic approach as a result of the degree of replication of messages.

Figure 1.c is interesting because there are two competing effects at work for the epidemic protocol. When the buffer size is small, there is a high probability that messages will be eliminated due to overflow, as discussed above. Consequently, the number of messages exchanged is also low. At the other end of the scale, as the buffer size increases to a point where it can accommodate all the messages in the system, there is no repeated exchange of messages, so the number is also low. In the middle of the range, however, the buffer size is insufficient to hold all messages and there is a cycle in which messages are eliminated by buffer overflow and then reinstated by other nodes, resulting in very high overhead. In the case of CAR, it is worth noting that the overhead in terms of the number of messages exchanged is more or less constant, regardless of buffer size, demonstrating its *scalability*. CAR will always be the limiting case for performance under this metric because it only creates a single copy of each message. Thus, even at the point where buffer size becomes effectively infinite, the epidemic protocol will necessarily exchange more messages than ours, simply as a result of the replication.

Figure 1.d shows the distribution of the number of messages with respect to their delivery latency in the 32 hosts scenario. It is possible to observe that a proportion of the messages are delivered more or less immediately, since the recipients are in the same cloud as the sender. This phenomenon is more evident for the epidemic protocol since it allows for a rapid dissemination of the information in connected clouds.

5 Discussion and Concluding Remarks

A number of approaches have been proposed to enable asynchronous communication in mobile ad hoc networks. The epidemic routing protocol [13], described earlier, that forms the basis for much of the work in this field, applied this early approach to the field of asynchronous message delivery, but in a rather naive fashion. Chen and Murphy refined the epidemic model, presenting the so-called Disconnected Transitive Communication paradigm [3]. Their approach is similar to ours, since it essentially argues for the use of utility functions, but it provides a general framework rather than a detailed instantiation, and so aspects related to the composition of calculated delivery probabilities are almost entirely missing. In [8], Lindgren et al. propose a probabilistic routing approach to enable asynchronous communication among intermittently connected clouds of hosts. Their approach is based on the fact that the exploited communication model is typically transitive and, for this reason, the probability of message delivery must be calculated accordingly. In [5], Fall proposes the Delay Tolerant Network




Figure 1. Publish-subscribe model (32 hosts scenario): (a) Delivery ratio vs population density. (b) Delivery ratio vs buffer size. (c) Number of messages vs buffer size. (d) Average delay vs population density.

architecture to solve the internetworking issues in scenarios where partitions are frequent and a connected path between message senders and receivers may be not present (such as satellite and interplanetary communication systems). Zhao et al. in [15] discuss the so-called Message Ferrying approach for message delivery in mobile ad hoc networks. The authors propose a proactive solution based on the exploitation of highly mobile nodes called ferries. These nodes move according to pre-defined routes, carrying messages between disconnected portions of the network.

With respect to the existing work in this research area, we have introduced a general framework for the evolution and the prediction of the mobile context to provide efficient and effective communication mechanisms in mobile ad hoc networks. Moreover, we believe that it is possible to integrate our techniques with these approaches, since they address orthogonal aspects of the problem. It is worth noting that we used lightweight mechanisms, because we believe that routing algorithms that are complex from a computational point of view are unsuitable for mobile devices, usually characterised by scarcity of resources. One example is the use of Kalman filter techniques that do not necessitate storing all the history of the evolution of the context information.

To conclude, in this paper, we presented a novel approach to the challenge of asynchronous ad hoc routing. Thus, we have designed a general and flexible framework for the evaluation of context information using probabilistic, statistical, autonomic and predictive techniques in order to optimize the consumption of the scarce resources of mobile devices whilst retaining good delivery performance. In order to assess our algorithm, a new mobility model [9], better reflecting the realities of human organization, was developed and used in simulations that give a feel for the relative performance of CAR relative to flooding and epidemic routing. The results demonstrated that, even without message replication, CAR performs respectably in terms of message delivery, with very much lower overheads than the alternatives.

References

- [1] D. F. Bantz, C. Bisdikian, D. Challener, J. P. Karidis, S. Masrianni, A. Mohindra, D. G. Shea, and M. Vanover. Autonomic personal computing. *IBM Systems Journal*, 42(1), 2003.
- [2] P. J. Brockwell and R. A. Davis. *Introduction to Time Series and Forecasting*. Springer, 1996.
- [3] X. Chen and A. L. Murphy. Enabling disconnected transitive communication in mobile ad hoc networks. In *Proceedings of the Workshop on Principles of Mobile Computing (POMC'01)*, pages 21–23, August 2001.
- [4] A. Doria, M. Uden, and D. P. Pandey. Providing connectivity to the Saami nomadic community. In *Proceedings of the Second International Conference on Open Collaborative Design for Sustainable Innovation*, December 2002.
- [5] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of SIGCOMM'03*, August 2004.
- [6] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME Journal of Basic Engineering*, March 1960.
- [7] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preference and Value Tradeoffs*. John Wiley and Sons, 1976.
- [8] A. Lindgren, A. Doria, and O. Schelen. Probabilistic routing in intermittently connected networks. *Mobile Computing and Communications Review*, 7(3), July 2003.
- [9] M. Musolesi, S. Hailes, and C. Mascolo. An ad hoc mobility model founded on social network theory. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*. ACM Press, October 2004.
- [10] C. E. Perkins and P. Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. In *Proceedings of the Conference on Communications Architecture, Protocols and Applications (SIGCOMM 94)*, August 1994.
- [11] H. A. Taha. *Operations Research: An Introduction*. Prentice Hall, 1996.
- [12] I. T. Union. Connecting remote communities. *Documents of the World Summit on Information Society*, 2003. <http://www.itu.int/osg/spu/wsis-themes>.
- [13] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-2000-06, Department of Computer Science, Duke University, 2000.
- [14] A. Varga. The OMNeT++ discrete event simulation system. In *Proceedings of the European Simulation Multiconference (ESM'2001)*, Prague, 2001.
- [15] Z. Venrui, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'04)*, May 2004.