A Framework for Opportunistic Forwarding in Disconnected Networks

Iacopo Carreras, Daniele Miorandi and Imrich Chlamtac

CREATE–NET

via Solteri 38, 38100 Trento, Italy

email: {iacopo.carreras,daniele.miorandi,imrich.chlamtac}@create-net.org

Abstract— In this paper, we analyze the performance of a family of opportunistic forwarding schemes (the K-copy relaying strategies) over disconnected wireless networks. We introduce a classification of mobility models based on their dynamic properties, and characterize the M^2 (Marks-Memoryless) class. Statistical tools are combined with numerical simulations to show that some of the most used mobility models in the literature fall within the M^2 class. A mathematical framework is provided for evaluating the performance of opportunistic forwarding schemes in the presence of M^2 mobility, and it is shown that the finiteness of the mean time necessary to deliver a message depends only on the mobility characteristics and not on the relaying protocol specification.

Index Terms— wireless networks, opportunistic forwarding, disconnected networks, mobility models.

I. INTRODUCTION

The emerging of a novel pervasive computing environment imposes big challenges to the current ICT technologies, calling for novel paradigms of communication, computing and service provisioning [1].

Pervasive computing environments will be composed of millions of devices, making scalability and complexity key issues to address. It is expected that these devices will communicate via wireless links, with all the problems related to the broadcast and random nature of the wireless channels. More important, it is a widely acknowledged fact, after the seminal work by Gupta and Kumar [2], that *fully connected* wireless networks scale poorly. This fact is rooted in the fact that, in order to maintain connectivity, the transmission power should be scaled in such a way that interference becomes the limiting factor in determining network performance. This problem has been shown to be solvable in the presence of mobile hosts, by trading off delay for capacity. The landmark work in the area is that of Grossglauser and Tse [3], which showed that wireless networks could be made scalable (in the sense that the communication rate per connection scales as $\Theta(1)$ with the number of nodes) by exploiting relaying schemes based on opportunistic communications. The basic idea is to exploit the mobility of the nodes for "carrying around" the information messages, therefore providing a macroscopic-level diversity scheme. In this way, we can give up the full connectivity constraint, and reduce the transmission range in such a way that interference stops limiting the network performance. The price to pay for such a scale-free system is in terms of delay;

it has indeed been shown that a fundamental delay/capacity tradeoff arises in large-scale wireless networks [4, 5]. The idea of exploiting mobility for providing connectivity at the expense of delay is also at the basis of the research in the area of Delay Tolerant Networks (DTNs) [6]. DTNs, also referred to as "challenged Internets" are based on scenarios in which intermittent connectivity is imposed from some system constraints; the aim is to maintain the conventional end-to-end semantics typical of the Internet protocols under such constraints.

All these considerations led to an increased interest in relaying protocols. In this work, we will focus on the class of K-copy relaying strategies. According to such protocols, any node that receives a packet that is not destined to it can make K copies of such message and forward it to the nodes it encounters. In this way, redundant packets are injected into the network, in order to possibly speed up message delivery. This family of relaying strategies is well understood in the case of connected topologies, where it is usually referred to as "limited/controlled flooding" [7] and studied in the context of routing protocols. Here the focus is on the case of mobile, highly disconnected networks. In particular, we consider a scenario composed by extremely sparse mobile nodes, where communication occurs mostly between pairs of isolated nodes. For this situation, a few works has appeared in the literature dealing with the performance of such schemes and the impact of the considered mobility models [8, 9].

In this paper, we introduce a class of mobility models, the Marks-Memoryless (M^2) mobility models, characterized by the following two properties:

(i) The IDs of the nodes meeting at the n-th meeting instant is independent of the IDs of the nodes meeting at meeting instants $n-1, n-2, \ldots$;

(i) The sequence of meeting times of *all nodes* form a renewal process.

The models falling within the M^2 class lend themselves to be analytically treated, by means of rather simple stochastic process models. On the other hand, as it will be shown later on by using statistical analysis tools, a wide range of models proposed for mobile wireless networks fall within this class (including, e.g., random waypoint and random direction mobility models). In particular, we found that, for mobility models falling within the M^2 class, *the finiteness of the mean time to deliver a message to all nodes (respectively: any single node) does not depend on the particular relaying strategy*

employed, but just on the mobility features.

The contribution of this paper is threefold. First, it introduces a new approach to the classification of mobility models, based on their dynamic behavior. Second, it introduces a new class of mobility models, the M^2 class, and shows, by means of a detailed statistical analysis on the outcomes of numerical simulations, that some of the most used mobility models in the literature on ad hoc networks fall within such class. Last, it proposes a general framework for analyzing the performance of opportunistic forwarding schemes based on K-copy relaying strategies. The framework covers both the case of "epidemic" dissemination of data (in which a message originated by a node is directed to all the other devices in the network) and the classical unicast communications, in which intermediate nodes perform copy-and-forward operations. In particular, necessary and sufficient conditions for the finiteness of the mean time necessary to broadcast/unicast a message will be provided.

The remainder of this work is organized as follows. Sec. II introduces the M^2 mobility models class. Sec. III reports an in-depth statistical characterization of various classical mobility models and discusses their belonging to the M^2 class. Sec. IV describes the general framework for analyzing the performance of opportunistic forwarding schemes in the presence of M^2 mobility models. Sec. V concludes the paper pointing out some promising directions for future research in the field.

II. THE M^2 MOBILITY MODELS CLASS

We introduce a probability space $\{\Omega, \mathcal{F}, \mathbb{P}\}$, on which all the random processes of interest will be defined. We assume that there are N nodes in the network, their IDs being uniquely mapped on $\mathcal{S} = \{1, 2, \ldots, N\}$. A *mobility pattern* is defined as a set of N rules (measurable with respect to $\mathbb{P}$), each one corresponding to a node. At time t, rule k defines, given the position and speed vector $\underline{v}$ of node k over the interval $(-\infty, t^-)$, the new velocity $\underline{v}(t^+)$. The proposed definition is very general, in that it accounts for mobility models with a possibly infinite memory. In particular, it comprises the most common mobility models used in ad hoc networking, such as Random Waypoint Mobility (RWM), Brownian Motion (BM), Random Direction Model (RDM) together with their generalizations [10]. Different nodes may obey to different rules, which enables us to include the cases of mobile nodes moving at different speeds, following different mobility models etc.

Given a mobility pattern, we associate to it a *marked point process* $\{Z_n\}_{n \in \mathbb{Z}} = \{T_n, \sigma_n\}_{n \in \mathbb{Z}}$ [11], where:

- $\{T_n\}_{n \in \mathbb{Z}}$ denotes the sequence of *meeting times*, i.e., the time instants at which any two nodes of the network come within mutual communication range;
- the marks of the sequence $\{\sigma_n\}_{n \in \mathbb{Z}}$ take the form $\sigma_n = (i, j)$, where $i, j \in \mathcal{S}$ denotes the IDs of the devices getting into communication range.

We assume that no simultaneous meetings can take place, so that all meeting times are different. Further, we assume, without any loss of generality, $\ldots < T_{-1} < T_0 \leq 0 < T_1 < \ldots$.

We associate with T_n a counting process $\mathcal{N}$, defined as:

$$\mathcal{N}(A) = \sum_{n \in \mathbb{Z}} \delta_{T_n}, \tag{1}$$

for each $A \subset \mathbb{R}$, δ_x being the Dirac measure at $x \in \mathbb{R}$, and we define as *intrameeting times*[1] the sequence $\{Y_n\}_{n \in \mathbb{Z}}$, where $Y_n = T_n - T_{n-1}$.

We define as Marks-Memoryless (M^2) the class of all mobility models for which:

(i) The marks σ_n constitute an i.i.d. sequence;

(ii) The sequence of *intrameeting times* $\{Y_n\}_{n \in \mathbb{Z}}$ is i.i.d..

Such characterization defines a class of mobility models which is "memoryless" in the sense that (i) there is no correlation between the IDs of the nodes that meet in successive meeting times (ii) the process $\{T_n\}_{n \in \mathbb{Z}}$ is a renewal process. This does not imply that the *mobility pattern* is memoryless. On the contrary, it somewhat implies some form of regularity in the mobility patterns, such that the sequences $\{\sigma_n\}_{n \in \mathbb{Z}}$ and $\{Y_n\}_{n \in \mathbb{Z}}$ are i.i.d.

The intensity of the process $\{T_n\}_{n \in \mathbb{Z}}$ is defined as:

$$\lambda = \mathbb{E}\left[\mathcal{N}\left([0, 1)\right)\right], \tag{2}$$

where $\mathbb{E}[\cdot]$ denotes the expectation taken with respect to the measure induced by $\mathbb{P}$. For the M^2 class, we have that $\mathbb{E}[Y_n] = \lambda^{-1}$ [12]; further, in the following, we will assume $0 < \lambda < +\infty$.

We denote by π the stationary probability distribution of the marks σ_n, so that we have:

$$\pi(i, j) = \mathbb{P}[\sigma_0 = (i, j)], \quad i, j \in \mathcal{S}. \tag{3}$$

III. CHARACTERIZATION OF THE M^2 MOBILITY MODELS: STATISTICAL ANALYSIS

In order to verify the soundness of the proposed M^2 class, we compared its assumptions with what happens according to some of the most popular mobility models adopted in MANETs research, namely the Random Waypoint [13], Random Direction [10] and Random Walk [10] mobility models. This has been accomplished running extensive simulations (using the freely available simulation tool Omnet++ [14]), and performing a statistical analysis of the simulation outcomes. In all the simulated scenarios, mobile nodes (MNs) are moving over a square playground of size $L = 500$ m. Distinct nodes are considered to be in communication range if the mutual distance is less than $R = 25$ m. We varied the number and speed of mobile nodes moving according to a predefined mobility model. The simulated scenario is uniform, in that all the nodes of the network move according to the same mobility rules. In the Random Waypoint (RWP) mobility model, MNs can be either active or pausing. Passing from the pausing to the active state, they select a destination at random (usually according to a uniform distribution) and move, on a straight line, till they reach it. Then, they move to the pausing state, and repeat the process. In our implementation of RWP, MNs move at a constant speed and without pausing. We reproduced a

---

[1]The notion of intrameeting times is introduced to differentiate from the conventional notion of intermeeting times, which refer to the successive meetings of a *specific* pair of nodes.

Mobility model	Marks i.i.d.	Intrameetings i.i.d.
RWP	PASS	PASS
RD	PASS	PASS (with high number of nodes)
RW	NOT PASS	NOT PASS

TABLE II

SUMMARY OF THE OUTCOMES OF THE STATISTICAL ANALYSIS OF THE RWP, RD AND RW MOBILITY MODELS AGAINST THE M^2 ASSUMPTIONS.

perfect simulation [15], sampling the initial location of nodes according to the corresponding stationary distributions (which is not uniform), following the approach in [16]. Subsequent destinations are then sampled from the uniform distribution. This approach eliminates the time needed for the simulation to reach the stationary regime.

In the Random Direction (RD) mobility model each MN, when reaching the boundary of the simulation area, pauses for a randomly chose pausing time, and, afterwards, chooses uniformly an angular direction in the range $[0, \pi]$ radians and starts to move. In our RD implementation, we considered MNs to be moving at a constant speed with no pausing.

In the Random Walk (RW) mobility model each MN performs a "movement step" (consisting in moving in a given direction for either a constant time or for the time necessary to cover a given distance), after which it samples a new direction (again, uniformly distributed in $[0, 2\pi]$) and performs a movement step. In our implementation, each movement occurs with a constant traveled distance of 10 m and at a constant speed.

Table I summarizes the parameters we used for performing the simulations. The position of every mobile node in the simulation is updated every 0.001 s, which corresponds to the time granularity of the simulated scenarios. Simulations are run over the CREATE-NET grid infrastructure so to minimize the execution time, and for a duration sufficient for collecting 10^6 meetings. For each simulation, the IDs (i, j) of the mobile nodes meeting, together with the corresponding meeting timestamp, are traced and used for the statistical analysis of the M^2 mobility model assumptions.

Simulation area	500×500 m
Mobile nodes communication range	25 m
Mobile nodes speed	$5, 7.5, 10$ m/s
Number of mobile nodes	$15, 30, 45, 60$
Mobility Models	Random Waypoint, Random Direction and Random Walk

TABLE I

SIMULATION PARAMETERS.

In order to reflect the symmetry of the marks (i.e., $\sigma(i, j) = \sigma(j, i)$), the following transformation has been applied to trace the IDs of the meeting pair, mapping $\mathcal{S} \times \mathcal{S}$ onto the set $\{1, \ldots, \binom{N}{2}\}$:

$$\sigma(i, j) = \begin{cases} (N-1)i + j - \frac{i(i+1)}{2} & i < j \\ (N-1)j + i - \frac{j(j+1)}{2} & i > j \end{cases}$$

where i and j are the IDs of the 2 nodes meeting, with $i, j \in \mathcal{S} = \{1, 2, \ldots., N\}$.

The aim of the statistical analysis is twofold. First, we are interested in verifying the independence of the marks. In particular, the autocorrelation has been adopted as a first measure of the marks independence, and the marks uniform distribution as a further assessment.[2]

As a second step, we have considered the intrameeting times series, and conducted a statistical analysis for verifying:

- the intrameeting times independence;
- the fit of the intrameeting times distribution with some well-known ones.

Before going into the details, we reported in Tab. II a summary of the results of the statistical analysis tests. As we can see, not all the mobility models considered satisfy the M^2 hypothesis, and some of them only with an adequate number of nodes.

In the following the details of the statistical methods adopted are presented. Complete results are presented for the RWP mobility model, while for the RD and RW mobility models only the general conclusions are reported due to lack of space (the complete results can be found in [17]).

A. Marks Independence and Distribution

We have first analyzed the autocorrelation function of the marks process. This, while not guaranteeing independence, measures the strength of the linear dependence of the values of a time series. With respect to the autocorrelation, we have consider the *autocorrelation plot* [18], where the normalized autocorrelation at lag k for a time series x_i composed of L symbols is defined as:

$$r_k = \frac{\sum_{i=1}^{L-k} (x_i - \mu)(x_{i+k} - \mu)}{\sigma}, \qquad (4)$$

where σ and μ are, respectively, the variance and mean of the considered time series. We considered, as test parameter, the *autocorrelation peak*, defined as :

$$r_{peak} = \max_k \{|r_k| : k \neq 0\}. \qquad (5)$$

We have then verified by means of statistical inference methods that the marks are uniformly distributed. The adopted tool is the *Chi Square Goodness of fit test* [19] with a significance of the test of 0.95. In Tab. III, the autocorrelation peak is reported in the case of a variable number of mobile nodes (15, 30 45 and 60, respectively), moving according to the RWP mobility model and a speed of 5, 7.5 and 10 m/s, respectively. For all the considered settings, the autocorrelation peak is below 0.0104, which allows us to reasonably exclude any linear dependence among the samples of the marks process.

---

[2]In a uniform scenario, by symmetry, the marks are uniformly distributed. If after a long simulation time the marks sample distribution is not uniform, this can be interpreted as a signal of the presence of some form of memory in the process (in that the sample distribution is not converging over time). It turns out that some models, which were not memoryless, e.g., random walk, passed the autocorrelation test but failed the uniformity one.

Number of Nodes	Autocorrelation peak		
	v=5 m/s	v=7.5 m/s	v=10 m/s
15	0.0100	0.0066	0.0077
30	0.0104	0.0073	0.0075
45	0.0072	0.0069	0.0070
60	0.0056	0.0051	0.0058

TABLE III

MARKS AUTOCORRELATION PEAK IN THE CASE OF A VARIABLE NUMBER OF MOBILE USERS MOVING ACCORDING TO THE RANDOM WAYPOINT MOBILITY MODEL AND A SPEED OF 5, 7.5 AND 10 M/S, RESPECTIVELY.

Number of Nodes	Chi square goodness of fit test			C-S test threshold
	v=5 m/s	v=7.5 m/s	v=10 m/s	
15	103.365	94.1686	65.041	128.803
30	356.270	356.841	348.833	483.570
45	878.750	877.670	876.231	1063.273
60	1583.634	1643.063	1513.304	1867.961

TABLE IV

CHI SQUARE GOODNESS OF FIT TEST BETWEEN THE MARKS AND UNIFORM DISTRIBUTIONS, IN THE CASE OF A VARIABLE NUMBER OF NODES MOVING ACCORDING TO THE RANDOM WAYPOINT MOBILITY MODEL WITH A SPEED OF 5,7.5 AND 10 M/S, RESPECTIVELY.

In Fig. 1, the autocorrelation plot is presented in the case of 60 mobile nodes moving according to the RWP mobility model, and with a speed of 10 m/s. The presented plot confirms what is briefly summarized in the autocorrelation peak, that is that the marks do not present any linear dependence among them.
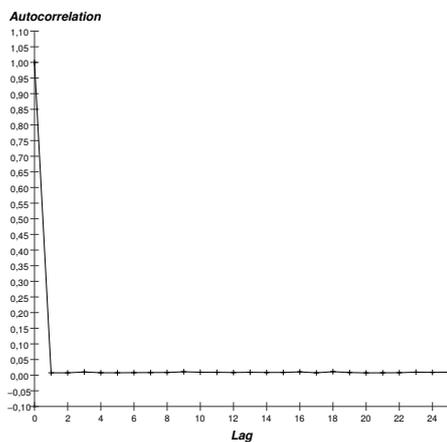

Fig. 1. Marks autocorrelation plot in the case of 60 mobile nodes moving according to the Random Waypoint mobility model with a speed of 10 m/s.

Similar autocorrelation results can be observed for the RD mobility model, while for RW the mobility model the autocorrelation of the marks process results to be low only in the case of a high number (> 45) of mobile nodes. For all the simulated scenarios, the autocorrelation peak decreases while increasing the number of mobile nodes. This can be easily explained considering that to a higher number of mobile nodes corresponds a higher number of meetings over a larger alphabet (the defined marks process is considering every meeting occurred in the entire playground) and, thus, a reduced autocorrelation.

In Tab. IV the results of the chi-square goodness of fit test between the marks distribution and the uniform distribution are presented in the case of a variable number of nodes moving according to the RWP mobility model with a speed of 5, 7.5 and 10 m/s, respectively.

As it is intuitively clear from the table, for all the considered settings the distribution of the marks verifies the hypothesis that the marks are distributed uniformly with a rejection level of 0.95. Similar conclusions hold for the RD mobility model, but not for the RW mobility model. This is due to the characteristics of the Random Walk mobility model, where, if the specified distance (or time) traveled by a mobile node

before choosing the next destination is short, then the resulting movement pattern is a random roaming within a limited region. This results in an extremely low drift of the nodes position and, consequently, in a high correlation and not-uniform distribution of the marks. This effect is clearly visible from the marks histogram. In Fig. 2 and Fig. 3, the marks histogram is presented in the case of 60 mobile nodes moving according to the RWP and RW mobility modes, respectively. The latter results to be characterized by a high number of peaks. Each peak corresponds to a couple of nodes that were in the same region when ending the simulation.
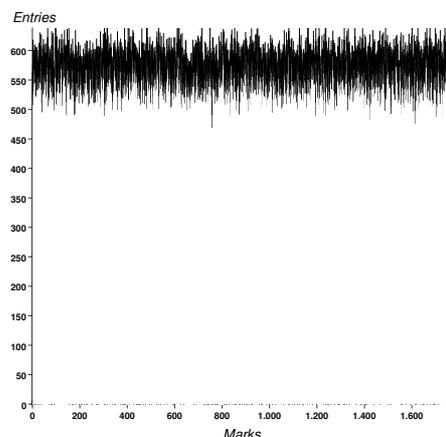

Fig. 2. Marks histogram in the case of 60 mobile nodes moving according to the Random Waypoint mobility model with a speed of 10 m/s.

¿From the autocorrelation and the marks distribution statistical analysis we can conclude that the marks i.i.d. assumption holds for both the RWP and RD mobility models. The same conclusion cannot be applied to the RW mobility model, where the "memoryless" mobility rule turns up to provide a considerable amount of memory to the marks process.

B. Intrameeting Times Independence and Distribution

We have then considered the statistical analysis of the intrameeting times process. An *intrameeting*, as defined in Sec. II, corresponds to the time elapsed between two distinct meetings occurred in the simulated area.
As a first step, we have adopted the autocorrelation analysis to verify the serial dependency of the samples of the intrameetings process. Then, we have followed non parametric statistical tests for verifying the Empirical Cumulative Distribution
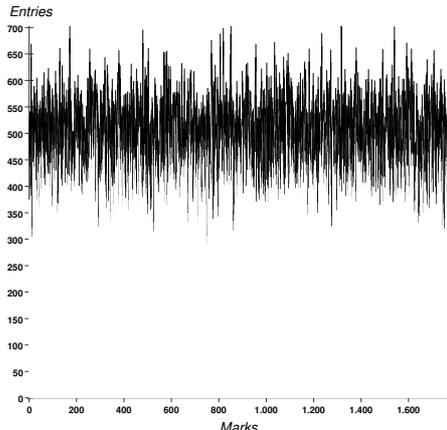
Fig. 3. Marks histogram in the case of 60 mobile nodes moving according to the Random Walk mobility model with a speed of 10 m/s.

Function (ECDF) of the intrameeting times series against well known cumulative distribution functions, e.g., the exponential distribution and the Pareto distribution. The methodology we have followed is graphically summarized in Fig. 4.
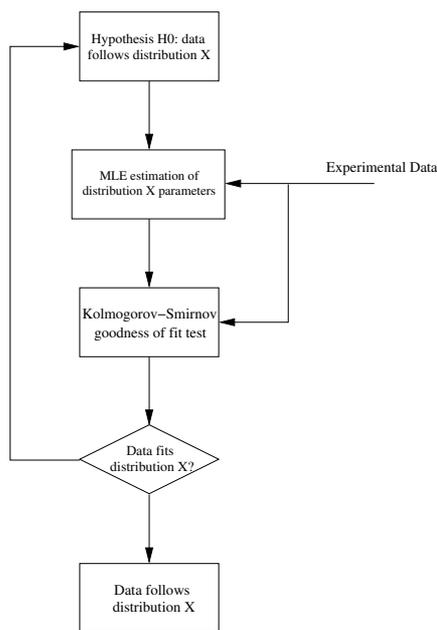

Fig. 4. Graphical representation of the methodologies adopted for evaluating the distribution of the intrameetings time series.

We have first assumed the intrameetings to follow a given distribution (e.g., exponential), and we have adopted a maximum likelihood technique for estimating the corresponding parameters (e.g., the mean in the case of an exponential distribution). We have then applied the Kolmogorov-Smirnov test [19] for verifying the matching of the ECDF with the estimated CDF. Refer to the [17] for a more detailed description of the statistical analysis followed.

As for the analysis of the marks, in the following we present the results of the RWP mobility model, and we refer to [17] for the remaining simulated scenarios.

In Tab. V, the intrameeting times autocorrelation peak is reported in the case of a variable number of mobile nodes

moving according to the RWP mobility model with a speed of 5, 7.5, 10 m/s, respectively. The autocorrelation peak is below 0.03 for all the considered settings. This can be interpreted as a very low linear dependence among subsequent samples of the intrameeting process.

Number of Nodes	Autocorrelation Peak		
	v=5 m/s	v=7.5 m/s	v=10 m/s
15	0.0297	0.0280	0.0259
30	0.0203	0.0027	0.0029
45	0.0138	0.0135	0.0127
60	0.0107	0.0106	0.0111

TABLE V

INTRAMEETING TIMES AUTOCORRELATION PEAK IN THE CASE OF A VARIABLE NUMBER OF MOBILE NODES MOVING ACCORDING TO THE RANDOM WAYPOINT MOBILITY MODEL AND A SPEED OF 5, 7.5 AND 10 M/S, RESPECTIVELY.

Similar conclusions apply to the autocorrelation analysis of the RD mobility model, and to the RW mobility model for a sufficiently high number of nodes. For the RW mobility modes, only when there are several pairs of nodes meeting in the playground we can achieve a diversification in the meetings pattern, with a corresponding uncorrelation of the intrameeting times.

A first graphical analysis of the probability density function of the RWP intrameeting time series suggested us that the exponential distribution is a good candidate to be taken into account. We have then considered the empirical cumulative distribution (ECDF) of the observed intrameeting times, and compared it with the ideal exponential cumulative distribution function estimated from the simulation data. In Fig. 5, the two distributions are plotted in the case of 60 mobile nodes moving at a speed of 10 m/s. As it is clear, the two distributions present a good matching.
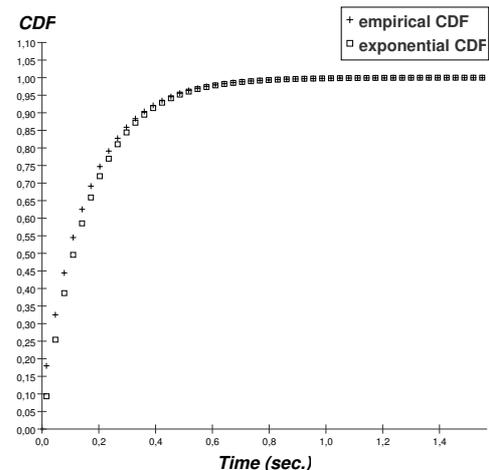

Fig. 5. Empirical Cumulative Distribution Function and estimated exponential cumulative distribution function in the case of 60 mobile nodes moving according to the RWP mobility model with a speed of 10 m/s.

As a further assessment, we have applied the Kolmogorov-Smirnov goodness of fit test between the RWP intrameeting times ECDF and the estimated exponential cumulative distribution function. The results of the test are summarized in Tab.VI.

For most of the considered settings, the two distributions present a good fit, passing the test with a significance level of 0.95.

Number of Nodes	K-S goodness of fit test			K-S test threshold
	v=5 m/s	v=7.5 m/s	v=10 m/s	
15	0.0035	0.0037	0.0040	0.0045
30	0.0041	0.0047	0.0059	0.0045
45	0.0031	0.0042	0.0043	0.0045
60	0.0035	0.0041	0.0046	0.0045

TABLE VI

KOLMOGOROV-SMIRNOV GOODNESS OF FIT TEST BETWEEN THE INTER-MEETING TIMES DISTRIBUTION AND EXPONENTIAL DISTRIBUTION, IN THE CASE OF A VARIABLE NUMBER OF NODES MOVING ACCORDING TO THE RANDOM WAYPOINT MOBILITY MODEL WITH A SPEED OF 5, 7.5 AND 10 M/S., RESPECTIVELY.

Th RD mobility model also confirms the intrameetings exponential ECDF, but only with a number of nodes greater than 30. Conversely, as expected, the RW mobility model never succeeds in the Kolmogorov-Smirnov test with exponential distribution.

IV. K-RELAYING STRATEGIES

We assume that communications between nodes take place by means of an opportunistic forwarding mechanism [8]. In particular, we consider the family of K-copy relaying strategies. Under the K-copy relaying strategy, each node receiving a message not destined to it can forward such packet to K other nodes encountered on the way. The parameter K, which defines the number of copies of a message a node is allowed to make and disseminate, controls the delay-storage tradeoff of the system. The larger K, the faster the diffusion of a message in the network and the larger the amount of network resources (i.e., bandwidth, storage on devices etc.) consumed. The case $K = 0$ corresponds to the case in which direct source-destination communication only is possible. The case $K = 1$ corresponds to the multi-hop case, while $K = N - 1$ corresponds to the case of flooding. This family of protocols is of particular interest in that it naturally lends itself to a distributed implementation. Also, assuming that the value of K can be stamped in the packet header, there is the possibility of using different values of K for different applications, thus achieving application-dependent delay-storage tradeoffs.

The main assumption needed to analyze the performance of K-copy relaying strategies under the M^2 mobility models is that data communications take place at meeting times only. This assumption models well situations in which all the nodes are isolated with high probability, such as very sparse networks and/or networks whose communication range is underdimensioned for achieving stringent power consumption constraints. Clearly, in most situations connected components will arise, facilitating the diffusion of messages in the network. In such cases, the results obtained in our framework will provide lower bounds on the achievable system performance. We further assume that the duration of each meeting is sufficient for transferring all the data to be exchanged.
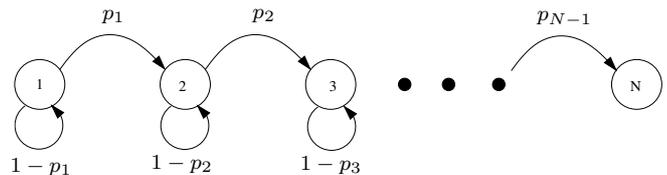

Fig. 6. Model for the "infected nodes" in the network, $Z(t)$, embedded at $\{T_n > U_0\}$.

We consider two possible models for data diffusion. The first one deals with the dissemination of data in the network, following an epidemic model [20], so that information is diffused throughout the whole network. The second one deals with the case of end-to-end transmissions, in which a packet generated at a node i needs to be delivered to node j, exploiting the mobility of the nodes to convey the packet.

A. Epidemic Dissemination

In this case, we introduce the stochastic process $Z(t)$, taking values in $\mathcal{S} \cup \{0\}$, and denoting the number of nodes, including the source one, which has correctly received the message at time t. We assume that the message is generated by node i at a randomly chosen time instant U_0, so that $Z(t) = 0 \ \forall t < U_0$. Let us consider the restriction of $Z(t)$ on $[U_0, +\infty)$, and let us sample the process at the meeting times $\{T_n > U_0\}$. The embedded process $\{Z_n\} = \{Z(T_n), T_n > U_0\}$ evolves as depicted in Fig. 6. At the first meeting time, the counter Z_n increases by one if the source node i takes part in the meeting, which happens with probability p_1, which equals $\sum_{j \neq i} \pi(i,j)$, and so on. Clearly, when $Z(t)$ becomes larger than K, we need to account for the possibility that the diffusion cannot take place since the "infected" device involved in the meeting has already distributed K copies of the message.

Considering the tree of all possible meeting pairs, we can compute the probability mass distribution of the number of meeting times necessary for broadcasting the message to all devices in the network. Given the statistical characterization of the random variable Y_0 and the measure π, we can compute, after some cumbersome algebra, the distribution of the quantity T_K^{epid}, the time necessary for the "epidemic" to spread by means of a K-copies relaying strategy. Such a quantity is indeed equal to the sum of the time elapsed from the generation of the message to the first meeting time, T^{first} plus the time necessary to "infect" all other nodes.[3] Let us denote by Ψ the number of meeting times necessary to flood the network, i.e., $Z_\Psi = N, Z_{\Psi-1} = N-1$, and by $\mathcal{G}_\Psi(\cdot)$ its generating function,

$$\mathcal{G}_\Psi(z) = \sum_{k=N-1}^{+\infty} z^k \mathbb{P}[\Psi = k], \qquad (6)$$

where the summation starts from $(N-1)$, which equals the minimum number of meeting instants needed to reach all the hosts. Further, we denote by $H_Y(\cdot)$ the Laplace-Stjelties

---

[3]Please note that the *first meeting time* does not necessarily involve the node that actually generated the packet. While this may look strange, it will become apparent that such decomposition provides an easy and elegant way to obtain some results on the finiteness of the mean time for the message to diffuse.

Transform (LST) [12] of the probability density function of Y_0,

$$\mathcal{H}_Y(s) = \mathbb{E}[e^{-sY}]. \quad (7)$$

Then, we can express the LST of T^{epid} as:

$$\mathcal{H}_{T_K^{epid}}(s) = \mathbb{E}[e^{-sT_K^{epid}}] = \mathcal{H}_{T^{first}}(s) \cdot \mathcal{G}_\Psi\left(\mathcal{H}_Y(s)\right). \quad (8)$$

Deriving, we obtain the following relationship for the mean epidemic dissemination time:

$$\mathbb{E}[T_K^{epid}] = \mathbb{E}[T^{first}] + \mathbb{E}[\Psi] \cdot \mathbb{E}[Y_0]. \quad (9)$$

One conventional performance metric of interest is the mean time to reach all the nodes. We make the additional assumption that *the packet generation process does not depend on the meeting process*. The assumption is reasonable and covers a wide range of situations, in which the creation of a message does not depend on the possible presence of other devices in the area and/or on the time elapsed since the last meeting time. In this case, we can apply Arrivals See Time Averages (ASTA) [21] property for evaluating the mean time elapsed between the generation of a message and the first meeting with any other node [21]. In particular, if ASTA holds, we have [12]:

$$\mathbb{E}[T^{first}] = \frac{\mathbb{E}[Y_0^2]}{2\mathbb{E}[Y_0]}. \quad (10)$$

We are interested in finding necessary and sufficient conditions for ensuring the finiteness of the mean broadcast time. We first need a preliminary lemma.

Lemma 1: Given that the marks process $\{\sigma_n\}_{n\in\mathbb{Z}}$ admits an invariant probability measure π such that $\pi(i,j) > 0$ for all $i,j \in \mathcal{S}$, the mean intermeeting time between node i and node j is $\mathbb{P}$-a.s. finite.

Proof: Let us denote by $\{T(i,j)_n\}_{n\in\mathbb{Z}}$ the sequence of meeting times between nodes i and j. Due to the properties of the M^2 mobility model, the inter-meeting times of i and j are sums of an i.i.d. number of i.i.d. quantities (the intrameeting times), so that they are also i.i.d. The 0-th intermeeting time of nodes i and j, $Y_0(i,j) = T_1(i,j) - T_0(i,j)$ is the sum of η intermeeting times. η is a geometric random variable of parameter $\pi(i,j)$. By Wald's lemma [12], we have $\mathbb{E}[Y_0(i,j)] = \mathbb{E}[\eta] \cdot \mathbb{E}[Y_0] = \frac{\mathbb{E}[Y_0]}{\pi(i,j)}$. Hence if $\pi(i,j) > 0$, then the inter-meeting time is finite $\mathbb{P}$-a.s. ∎

Thanks to the independence of the marks σ_n, we have that:

Corollary 1: Given that the marks process $\{\sigma_n\}_{n\in\mathbb{Z}}$ admits an invariant probability measure π such that $\pi(i,j) > 0$ for all $i,j \in \mathbb{S}$, the mean time elapsed between a meeting instant $T_{\hat{n}}$ and the first subsequent meeting between node i and j is $\mathbb{P}$-a.s. finite.

Proof: The process $\{T_n\}_{n\in\mathbb{Z}}$ is a renewal process, and the marks $\{\sigma_n\}_{n\in\mathbb{Z}}$ are iid. It follows that the intermeeting time has the same distribution of the time elapsed between a meeting instant $T_{\hat{n}}$ and the first subsequent meeting between node i and j. The application of Lemma 1 concludes the proof. ∎

We hence come to the first main result of the paper:

Theorem 1: Given that the marks process $\{\sigma_n\}_{n\in\mathbb{Z}}$ admits an invariant probability measure π such that $\pi(i,j) > 0$ for all $i,j \in \mathbb{S}$, the mean time to reach all the nodes of the network is $\mathbb{P}$-a.s. finite *if and only if* $\mathbb{E}[Y_0^2]$ is finite.

Proof: We start proving the sufficiency of the condition. A simple sample-path argument leads to the consideration that $T_K^{epid} \geq T_{K+1}^{epid} \ \forall K = 0, 1, \ldots, N-2$. We focus therefore on $K = 0$. We can see that the message broadcast time can be upperbounded by the sum of the first meeting time and the maximum of the intermeeting times between the source node i and any other node in the network. Hence:

$$T_{N-1}^{epid} \leq T^{first} + \max_{j \neq i}\{T_0(i,j)\} \leq T^{first} + \sum_{j \neq i} T_0(i,j). \quad (11)$$

Passing to expectations, we then get:

$$\mathbb{E}[T_{N-1}^{epid}] \leq \mathbb{E}[T^{first}] + \mathbb{E}[\max_{j \neq i}\{T_0(i,j)\}]$$
$$\leq \mathbb{E}[T^{first}] + \mathbb{E}[\sum_{j \neq i} T_0(i,j)] \quad (12)$$

The second term in the right hand side of (12) is a finite sum of finite quantities (from Corollary 1) and is hence $\mathbb{P}$-a.s finite. If $\mathbb{E}[Y_0^2]$ is finite, then from (10) the first term is also finite. The proof of the necessity is much simpler, since it suffices to notice that $T_K^{epid} \geq T^{first} \ \forall K$. Under the assumption of independence between the message generation process and the meeting process, (10) holds, from which the claim follows. ∎

As a simple corollary, we get the following result, which complements the analysis in [8]:

Corollary 2: Given the independence assumption on the message generation process, and given that the intermeeting times are distributed according to the following law:

$$F_{Y_0}(a) = \mathbb{P}[Y_0 \leq a] = 1 - a^{-\beta}, \quad a \geq 0$$

with $\beta > 0$, the mean broadcast time is finite *if and only if* $\beta > 2$.

Proof: We have:

$$\mathbb{E}[Y_0^2] = \int_0^{+\infty} da \, a[1 - F_{Y_0}(a)] = \int_0^{+\infty} da \, a^{1-\beta}. \quad (13)$$

The integral converges if and only if $1 - \beta < -1$, i.e., $\beta > 2$. The application of Theorem 1 concludes the proof. ∎

B. End-to-end Communications

We now move to a different traffic pattern, in which a given node i wants to send a message to a specific node j. In order to deliver the message to its final destination, the sender exploits the presence of mobile devices, which can act as relays, in order to overcome the possible lack of a direct connection from the source to the destination. The problem is basically to perform an end-to-end communication over a possibly disconnected network.

In general, the case of end-to-end communications can be characterized according to the diagram in Fig. 7. The process $Z(t)$ described therein represents the number of copies of the message in the network at time t, and is sampled at the meeting instants $\{T_n\}_{n\in\mathbb{Z}}$. The sample space is $\{1, \ldots, N-2\} \bigcup \{D\}$, (absorbing) state D corresponding to the case of destination reached.

The characterization, in terms of delay, can be carried out using a reasoning similar to the one presented for the case of
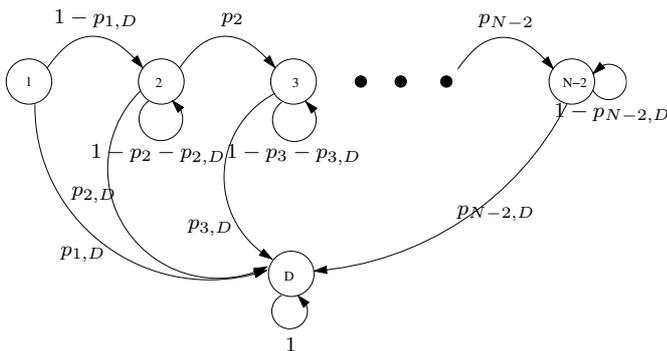
Fig. 7. Model for the "infected nodes" in the network, $Z(t)$, embedded at $\{T_n > U_0\}$, for the case of end-to-end communications.

epidemic dissemination. We denote by T_K^{e2e} the time elapsed between the generation of a message at a given node i and its delivery to the final destination j when using a K-copy relaying strategy. For the moment, let us focus on the mean expected delay, and in particular on the conditions ensuring its finiteness. We introduce the independence traffic generation assumption made for the case of epidemic dissemination, and get the following:

Proposition 1: Given the independence of the message generation process on the meeting process, the mean time taken by a message to reach its destination is finite *if and only if* $\mathbb{E}[Y_0^2]$ is finite.

Proof: A simple sample-path argument leads to the following inequalities:

$$T^{first} \leq T_K^{e2e} \leq T_K^{epid}, \quad \forall K = 0, 1, \ldots, N-1. \quad (14)$$

If $\mathbb{E}[Y_0^2]$ is finite, T_K^{epid} is also finite by Theorem 1. Then, applying the expectation operator to (14), $\mathbb{E}[T_K^{e2e}] \leq \mathbb{E}[T_K^{epid}] < +\infty$.

If $\mathbb{E}[T_K^{e2e}]$ is finite, then from (14), $\mathbb{E}[T^{first}]$ is also finite. Under the independent message generation assumption, ASTA holds, so that $\mathbb{E}[T^{first}] = \frac{\mathbb{E}[Y_0^2]}{2\mathbb{E}[Y_0]}$, from which the finiteness of $\mathbb{E}[Y_0^2]$ follows. ∎

V. CONCLUSIONS

Networking paradigms able to support disconnected operations represent a promising directions for enabling the arising of pervasive computing/communication environments. In this paper, we have presented an analytical framework for analyzing the performance of a family of opportunistic forwarding schemes (the K-copy relaying strategies) over disconnected wireless networks. We have introduced a classification of mobility models based on their dynamic properties, namely the intrameeting times and the marks distribution. Further, we have presented a class of mobility models, called M^2 (Marks-Memoryless), whose feature enable a simple characterization in terms of stochastic processes. Statistical tools are combined with extensive numerical simulations (performed using a freely available simulation tool) to show that the M^2 class includes some of the most used mobility models in the literature on mobile ad hoc networks. A mathematical framework has been provided for evaluating the performance of opportunistic forwarding schemes in the presence of M^2 mobility models,

and it has been shown that the finiteness of the mean time necessary to deliver a message (to either all nodes or any node) depends only on the mobility characteristics and not on the relaying protocol specification.

Two research directions appear of main interest for enhancing the present work. The first one concerns the incorporation, in the simulation framework, of a more realistic application scenario, where the physical layer characterization incorporates the effect of random channel variations and interference from neighboring nodes' transmissions, and real chunks of data are transmitted over the wireless link. The second one deals with the study of the applicability of the M^2 framework to more realistic mobility models. While some examples are reported in [17] using the traces in [8], we are still far from getting a complete understanding on the scenarios that can be analyzed by means of the proposed framework.

REFERENCES

[1] M. Weiser, "The computer for the 21st century," *ACM Mob. Comput. Commun. Rev.*, vol. 3, no. 3, pp. 3–11, 1999.
[2] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. on Inf. Th.*, vol. 46, no. 2, pp. 388–404, Mar. 2000.
[3] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. on Netw.*, vol. 10, no. 4, pp. 477–486, Aug. 2002.
[4] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah, "Throughput-delay trade-off in wireless networks," in *Proc. of IEEE INFOCOM*, Hong Kong, 2004.
[5] M. Neely and E. Modiano, "Capacity and delay tradeoffs for ad hoc mobile networks," *IEEE Trans. on Inf. Th.*, vol. 51, no. 6, pp. 1917–1937, Jun. 2005.
[6] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proc. of ACM SIGCOMM*, Karlsruhe, DE, 2003.
[7] J. Luo, P. Eugster, and J.-P. Hubaux, "Probabilistic Reliable Multicast in Ad Hoc Networks," *Elsevier Ad Hoc Networks*, vol. 2, no. 4, pp. 369–386, 2004.
[8] A. Chaintreau, P. Jui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proc. of IEEE INFOCOM*, Barcelona, ES, 2006.
[9] R.Groenevelt, G. Koole, and P. Nain, "Message delay in mobile ad hoc networks," in *Proc. of Performance*, Juan-les-Pins, October 2005.
[10] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communications & Mobile Computing (WCMC)*, vol. 2, no. 5, pp. 483–502, 2002.
[11] F. Baccelli and P. Bremaud, *Elements of Queueing Theory*. Berlin: Springer–Verlag, 1994.
[12] S. M. Ross, *Stochastic Processes*. New York: J. Wiley & Sons, 1996.
[13] J. Yoon, M. Liu, and B. Noble, "Sound mobility models," in *Proc. of ACM MobiCom*, San Diego, CA, 2003.
[14] OMNeT++ discrete event simulation system. [Online]. Available: http://www.omnetpp.org
[15] J.-Y. L. Boudec and M. Vojnovic, "Perfect simulation and stationarity of a class of mobility models," in *Proc. of IEEE Infocom*, 2005.
[16] W. Navidi and T. Camp, "Stationary distributions for the random waypoint mobility model." *IEEE Trans. Mob. Comput.*, vol. 3, no. 1, pp. 99–108, 2004.
[17] I. Carreras, D. Miorandi and I. Chlamtac, "The effect of mobility on opportunistic forwarding," CREATE-NET, Tech. Rep., 2005. [Online]. Available: http://www.create-net.it/∼icarreras
[18] G. E. P. Box and G. Jenkins, *Time Series Analysis, Forecasting and Control*. Holden-Day, Incorporated, 1990.
[19] D. C. B. Alexander M. Mood, Franklin A. Graybill, *Introduction to the theory of statistics*. McGraw-Hill, 1974.
[20] A. Khelil, C. Becker, J. Tian, and K. Rothermel, "An epidemic model for information diffusion in manets," in *Proc. of MSWiM*. New York, NY, USA: ACM Press, 2002, pp. 54–60.
[21] B. Melamed and W. Whitt, "On arrivals that see time averages," *Operations Research*, vol. 38, pp. 156–172, Jan. 1990.