This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2007 proceedings.

Predicting the performance of mobile ad hoc networks using scaled-down replicas

Fragkiskos Papadopoulos, Konstantinos Psounis
University of Southern California, Los Angeles, CA 90089
E-mail: fpapadop, kpsounis@usc.edu.

***Abstract*—Experimentation with mobile ad hoc network testbeds is preferred to simulations for performing high fidelity testing. But, at the same time, realistic experimentation with large-scale network testbeds is more difficult, time-consuming, and expensive.**

To side-step some of these problems several researchers, *e.g.* [1], [2], have recently suggested experimentation on scaled-down replicas and have managed to downscale some *specific indoor* network realizations with *no* mobility. However, the question of whether a scaled-down replica can reproduce the behavior of an arbitrary large-scale *mobile* ad hoc network, in a timely manner, and under realistic *outdoor* conditions, remains an interesting open problem.

In this work we investigate ways of constructing suitably scaled-down replicas, both in *space* and *time*, that can predict the performance of large-scale mobile ad hoc networks with high accuracy. We consider both large- and small-scale fading effects. Further, we present necessary and sufficient conditions for the scaling to be possible, and identify some of the issues that may arise in practice. Finally, we argue that it is not possible to build arbitrarily smaller replicas, and that the factor by which one scales down the original network (in space) depends on the carrier frequency.

I. Introduction

The popularity and scale of mobile ad hoc networks have grown rapidly in recent years. The community has developed many protocols tailored for such networks and deployments with several hundreds of nodes over significantly sized regions already exist.

Measuring the performance of such large-scale networks and predicting their behavior under new protocols, architectures and load conditions are important research problems. These problems are made difficult by the complex nature of wireless channels, the increasing number of nodes, and the growing size of the deployment area.

A commonly accepted practice in the research community is to use simulations for testing and evaluating the performance of such networks. However, it is very expensive and inefficient to accurately run large-scale simulations (*e.g.* with several hundreds of nodes), which incorporate realistic models for the wireless channel that faithfully capture radio propagation and error characteristics.

For some of these reasons, experimentation with wireless network testbeds seems to be preferred by many researchers for performing high fidelity testing and performance evaluation, *e.g.* [3], [4], [5]. But, at the same time, realistic experimentation with testbeds deployed on large-sized physical areas introduces another set of limitations. First, it is quite difficult to configure, manage and troubleshoot such testbeds. Second, such testbeds can have high maintenance costs. And third, it is very time-consuming to gather network statistics and hence to evaluate the performance of new protocols and architectures in a timely manner.

To side-step some of these problems, researchers, *e.g.* [1], [2], [6], have recently suggested experimentation on smaller-scale testbed miniatures. They have accomplished to shrink a wireless network testbed into a smaller space, by reducing the communication range of nodes and the distance between them, while maintaining link characteristics. However, these studies have only considered downscaling some *specific indoor* network realizations with *no* mobility and remained mostly focused on the implementation and deployment aspects of the testbed miniatures. Hence, the question of whether a scaled-down replica can reproduce the behavior of an *arbitrary* large-scale *mobile* network under realistic *outdoor* propagation mechanisms, remains an interesting open problem. And, the question of whether one can expedite experimentation while preserving network behavior has not been studied either.

In this paper we attempt to answer the following fundamental questions: Consider an arbitrary mobile ad hoc network deployed in an area of size A, with n nodes, arbitrary mobility, traffic and routing protocol, and a realistic outdoor propagation model that incorporates both large- and small-scale fading effects. (i) Can we deploy the same network in an area of size αA ($0 < \alpha \leq 1$) and yield the same performance (i.e. can we perform space-downscaling)? (ii) What are the necessary and sufficient conditions? (iii) In what sense is performance preserved? (iv) How small can we go, i.e. how small can α get? (v) Can we speed-up the scaled replica by some factor $\delta > 1$ (or scale down the experimentation time by $1/\delta$), and yield the same performance (i.e. can we perform time-downscaling)?

Our theoretical analysis coupled with realistic simulations give concrete answers to all of these questions. Interestingly enough we find that both space and time downscaling is possible, under fairly general assumptions, and that the factor α by which one can scale down the original network in space, depends on the carrier frequency f_c.

The organization of the paper is as follows: Section II briefly discusses prior work on scaling down networks. Section III performs space-downscaling, under the assumption that the communication range is an ideal circle, and that the received signal power is a deterministic function of distance. Section IV extends the space-downscaling methodology for the case of realistic large- and small-scale propagation models. Under these models the received signal power is a random variable that can exhibit significant fluctuations and time-correlations due to mobility. Section V performs time-downscaling on top of space-downscaling. Section VI presents realistic simulations that verify our theoretical arguments, and Section VII concludes the paper giving directions for future work.

II. Related Work

The design of *performance-preserving network downscaling* techniques (that is, techniques where performance metrics of a large-scale network, *e.g.* throughput, packet delays, etc., are preserved by a suitably scaled-down replica), has been extensively studied for the case of *wireline* networks that resemble the Internet. For example, Psounis et al. [7] have introduced a method called SHRiNK that creates a slower version of the original network and can predict significant performance measures by observing the slower replica. Further, Papadopoulos et al. [8] have introduced two methods called DSCALEd and DSCALEs that perform topological downscaling by retaining only the congested links of the original network, and extrapolate from the performance of the downscaled network to that of the larger Internet. These techniques can be used to reduce the amount of traffic and the size of the network that one works with, thus enabling scalable performance prediction and simulation.

There have been very few studies of the question of whether one can design a performance-preserving network downscaling technique for the case of *wireless* networks. The most relevant to our work is the one by Naik et al. [2]. In this line of work the authors show via simple theoretical arguments and experiments that it is possible to deploy the Kansei testbed [9] into a smaller area and yield the same performance, as if the same testbed was deployed in a larger area. This is achieved by appropriately reducing the separation distances between nodes and their transmission range. However, the Kansei testbed comprises a *static indoor* wireless network with a *specific* symmetric topology [9]. And because the network is static, the authors have completely ignored small-scale fading effects that are due to mobility and can cause significant fluctuations and time-correlations in the received signal power. In our work, we want to investigate if a similar downscaling is possible for *arbitrary mobile* ad hoc networks operating in *outdoor* environments in the presence of both large- and small-scale fading effects.

Other lines of work have mainly focused on the implementation and deployment aspects of their indoor miniaturized testbeds, *e.g.* [1], [6], [10]. In [10] the authors have also presented some preliminary thoughts of how they could build their testbed in order to emulate small-scale propagation effects. Our work is very different. We are not interested in building a particular indoor testbed, but instead we want to investigate whether the full behavior of an arbitrary mobile ad hoc network deployed in an outdoor environment at one spatial scale, can be preserved by a suitably scaled replica deployed in an outdoor environment at another spatial scale.

And, finally, in-contrast to any existing prior work, we also perform time-downscaling of the wireless network, which can be used to significantly expedite the time required to perform experiments with testbeds.

III. Space-Downscaling of Mobile ad hoc Networks

In this section we present the space-downscaling methodology. For ease of exposition we assume for now that: (a) the communication range is an ideal circle around the transmitter, and (b) the received signal power is a deterministic function of distance. These assumptions imply that if a receiver is within the communication circle, it receives all packets. Otherwise, it loses all packets.

A. General methodology

Consider an *arbitrary static* ad hoc network, which is deployed in an area of size $A = xy$ square units and where every node i ($i = 1...n$) has a transmission radius of TR_i units, as shown in Figure 1(i). Let $0 < \alpha \leq 1$ be a scaling factor and perform the following operations to this system: (i) Multiply the x and y dimensions of the area by $\sqrt{\alpha}$, (ii) multiply the x_i and y_i coordinates of every node i by $\sqrt{\alpha}$, and (iii) multiply the transmission radius of every node i by $\sqrt{\alpha}$. The result of this operation is the scaled-in-space by a factor α system, depicted in Figure 1(ii).

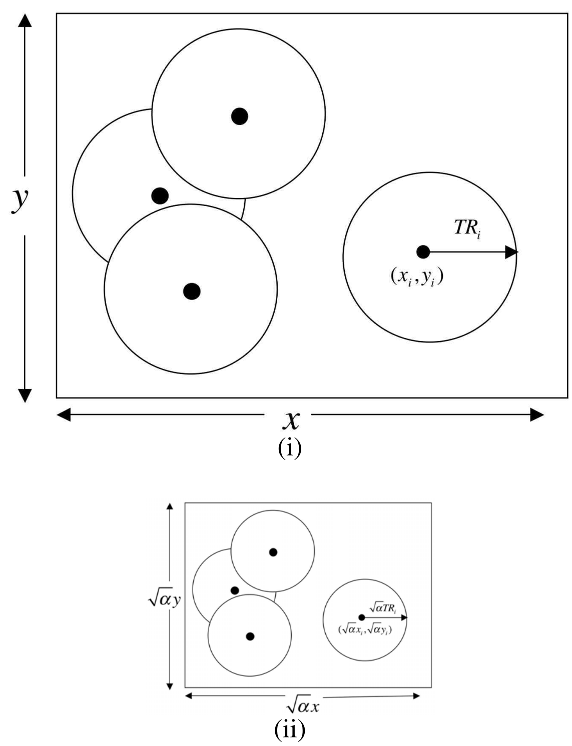

Fig. 1. (i) Original system (area size $= A$ square units), and (ii) scaled-in-space system (area size $= \alpha A$ square units). (In this example $\alpha = 0.25$.)

It is easy to see that because all node inter-distances and the transmission radius are scaled by our operations by the same factor, the structure of the configuration of the original system is not affected by the downscaling we perform.

Now, lets consider mobility. In this case we want to ensure that the structure of the two configurations is the same at *every* time t. Assume that at time $t = 0$ the original system is the one shown in Figure 1(i) and that we have performed operations (i)...(iii) and built the scaled replica shown in Figure 1(ii). Now, let $d_i \geq 0$ denote the distance that node i will travel in the original system and let $v_i(t)$ be its traveling speed at time $t \geq 0$. Because all distances in the original system were scaled by $\sqrt{\alpha}$ we should also scale d_i by the same factor, $\forall i$. And, since we scale the travel distance by $\sqrt{\alpha}$, to ensure that the topologies of the two systems change at the same speed we should also scale $v_i(t)$ by $\sqrt{\alpha}$, $\forall i, t$. We are now ready to state our first theorem:

Theorem 1: Perform the following operations to an *arbitrary mobile* ad hoc network, in order to build a scaled-in-space by a factor α replica: (i) Multiply the x and y dimensions of the area by $\sqrt{\alpha}$, (ii) multiply the x_i and y_i coordinates of every node i by $\sqrt{\alpha}$, (iii) multiply all distances travelled by nodes by $\sqrt{\alpha}$, (iv) multiply the speed of nodes by $\sqrt{\alpha}$, and (v) multiply the transmission radius of every node by $\sqrt{\alpha}$. If the propagation delays are insignificant and the same external arbitrary traffic is applied in the original and scaled systems, then (under assumptions (a) and (b)) the two systems will yield the same performance as a function of time, $\forall\ 0 < \alpha \leq 1$.

Proof: From the previous discussion we can see that the configuration of the original system is just a zoomed version of the configuration of the scaled system at every time $t \geq 0$. Further, if the propagation delays are insignificant in the original system, they will remain insignificant in the scaled system, since in the latter all distances are decreased. Under the same *external* traffic it is easy to see that packets are successfully received at each node at exactly the same times between the two systems. Hence, the packet arrivals (and transmissions) at each node also occur at exactly the same times, and thus the two systems will yield the same performance as a function of time. ∎

Remark: While the insignificant propagation delay assumption holds for the majority of ad hoc networks, there are some classes of such networks, *e.g.* underwater networks, where this assumption is no longer true [11]. Hence, in such cases the propagation delays of the scaled network will be smaller than those of the original network. While this may not be a significant issue for the case of arbitrary traffic without feedback, it is unclear whether one can preserve performance in the case of arbitrary traffic with feedback, as the round-trip times in the scaled system will be different.

B. Scaling down the transmission range

So far, we have described the general space-downscaling methodology. However, we have not described yet how to achieve the downscaling of the transmission range of nodes, which might seem trivial, but in practice, it requires some thought as it depends on the propagation model under consideration.

For illustration purposes lets consider two simple models, namely the free-space propagation model and the two-ray ground reflection model [12]. Under these models, assumptions (a) and (b) are both satisfied. (We consider more realistic models in the next section.)

The free-space propagation model assumes the ideal propagation condition that there is only one clear line-of-sight path between the transmitter and the receiver. The received signal power in free space at distance d from the transmitter is given by:

$$P_r(d) = \frac{P_t G_t G_r \lambda_c^2}{(4\pi)^2 d^2 L}, \tag{1}$$

where P_t is the transmitted signal power, G_t and G_r are the antenna gains of the transmitter and the receiver respectively, $L(L \geq 1)$ is the system loss factor not related to propagation, and λ_c is the carrier wavelength. The two-ray ground reflection model considers both the line-of-sight path between the transmitter and the receiver and a ground reflection path. The received power at distance d from the transmitter is predicted by:

$$P_r(d) = \frac{P_t G_t G_r h_t^2 h_r^2}{d^4 L}. \tag{2}$$

In Equation (2), h_t and h_r are the heights of the transmit and receive antennas respectively.

One way to scale down the transmission radius by $\sqrt{\alpha}$ under the previous two models, is to ensure that the received signal power at distance $\sqrt{\alpha}d$ from the transmitter in the scaled system equals the received signal power at distance d from the transmitter in the original system, $\forall d$. To accomplish this, we choose to scale down the transmission power P_t in the scaled system appropriately, keeping the rest of the antenna characteristics unaltered. Hence, if the free-space model is assumed we see from Equation (1) that we should multiply P_t in the scaled system by α since $P_r(d)$ is inversely proportional to d^2. If the two-ray ground reflection model is assumed, P_t in the scaled system should be multiplied by α^2 since $P_r(d)$ is now inversely proportional to d^4.

IV. Space-Downscaling of Mobile ad hoc Networks under Realistic Propagation Models

In this section we investigate in what sense is performance preserved by the scaled-in-space replica under more realistic propagation models. In general, these models can be classified into two categories: (i) large-scale, and (ii) small-scale propagation models.

Large-scale propagation models usually characterize the received signal power over large transmitter-receiver separation distances (*e.g.* over several hundreds or thousands of meters). The free-space model and the two-ray ground reflection model we saw earlier can be classified as large-scale models. However, they are unrealistic since they predict the received signal power as a deterministic function of distance and represent the communication range as an ideal circle around the transmitter. In reality however, the received signal power at a certain distance from a transmitter is a random variable, *e.g.* due to shadowing effects caused by obstructions in the environment such as buildings, hills, etc. A realistic large-scale propagation model that incorporates randomness in the received signal power is the log-normal shadowing model [12], which we will shortly study.

In many cases, the received signal power over very short travel distances (a few wavelengths) or short time durations (on the order of seconds), may also exhibit rapid fluctuations and time-correlations. These phenomena are mainly due to mobility and to multipath propagation effects caused by reflecting objects and scatterers in the environment. These phenomena can have a significant impact on system's performance. While large-scale models can be used in this case for predicting the local average signal strength over certain transmitter-receiver separation distances, the fluctuations around the local average as well as the time-correlations, are characterized via small-scale propagation models [12]. A realistic small-scale model that we study in this paper is the Ricean/Rayleigh fading model.

A. Space-downscaling under log-normal shadowing

The shadowing propagation model consists of two parts. The first part is known as the path loss model and predicts the *average* received signal power at distance d from the transmitter, denoted by $\overline{P_r(d)}$. It uses a close-in reference distance d_0. $\overline{Pr(d)}$ is computed relative to $Pr(d_0)$ as follows:

$$\overline{P_r(d)} = Pr(d_0)\left(\frac{d_0}{d}\right)^n, \tag{3}$$

where n is called the path loss exponent and $Pr(d_0)$ is computed from Equation (1).

The second part of the shadowing model reflects the variation of the received signal power around the average. It is a log-normal random variable, that is, it is of Gaussian distribution if measured in dB. The overall shadowing model is represented (in dB) by:

$$P_r(d)_{dB} = \overline{P_r(d)}_{dB} + X_{dB}, \tag{4}$$

where $\overline{P_r(d)}$ is defined by Equation (3) and X_{dB} is a Gaussian random variable with zero mean and standard deviation σ_{dB}, called the shadowing deviation. Both the path loss exponent and the shadowing deviation, depend on the environment and are obtained by field measurement [12]. Hence, its natural to make the following assumption:

Assumption 1: The scaled system is deployed in an environment that yields the same path loss exponent n and shadowing deviation σ_{dB}, as the original environment.

In the previous section we saw that in a scaled-in-space by a factor α replica, the received signal power at distance $\sqrt{\alpha}d$ from the transmitter should be equal to the received signal power at distance d from the transmitter in the original system, $\forall d$. However, there, the received signal power was a deterministic function of distance whereas now it is a random variable (given by Equation (4)). Therefore, we now require that the *distribution* of the received signal power at distance $\sqrt{\alpha}d$ from the transmitter in the scaled system be the same with the distribution of the received signal power at distance d from the transmitter in the original system, $\forall d$. As X_{dB} is the same between the two systems, to accomplish this, we only need to scale down the average received power $\overline{P_r(d)}$. Like before, we choose to do this by scaling down the transmission power P_t. Let d_0 be equal in the two systems. Then, by looking at Equations (1) and (3), it is easy to deduce that P_t in the scaled system should be multiplied by $(\sqrt{\alpha})^n$.

We can now state the following proposition, whose proof follows immediately from the above discussion.

Proposition 1: To construct a scaled-in-space by a factor α replica under the log-normal shadowing propagation model, perform steps (i)...(iv) of Theorem 1, and (step (v)) multiply the transmission power of every node by $(\sqrt{\alpha})^n$. If the propagation delays are insignificant and the same arbitrary traffic is applied in the original and scaled systems, then under Assumption 1, the two systems will yield the same performance in *distribution*, $\forall\ 0 < \alpha \leq 1$.

B. Space-downscaling under Ricean/Rayleigh fading

Under Ricean fading the received signal power at a certain distance d from the transmitter, is also dependent on time t [12] and is given as follows [13]:

$$P_r(t,d) = \frac{P_r^{ls}(d)}{2(K+1)}\left[(x_1(t)+\sqrt{2K})^2 + x_2^2(t)\right], \tag{5}$$

where $P_r^{ls}(d)$ is the received signal power as predicted by a large-scale model at distance d from the transmitter, K is the Ricean K-factor that depends on the structure of the environment and determined by measurement, and $x_1(t)$, $x_2(t)$ are zero-mean Gaussian random variables with unit variance. (For $K = 0$ the model is the well-known Rayleigh fading model.) The important difference here comparing to the previous model we studied, is that the random variable $x_i(t)$ $(i = 1, 2)$ can exhibit *time-correlations* due to Doppler spreading that is caused by the mobility of the nodes [12]. [1] Hence, for a fixed d, $\{P_r(t,d), t \geq 0\}$ is now a stochastic process.

Before proceeding, we first state the following assumption, which implies that the scaled system is deployed in an environment with similar structural properties (*e.g.* with respect to reflecting objects, scatterers, etc.) as the original environment.

Assumption 2: The scaled system is deployed in an environment that yields the same Ricean K-factor as the original environment.

Now, as before, we require that the *distribution* of the received signal power at distance $\sqrt{\alpha}d$ from the transmitter in the scaled system be the same with the distribution of the received signal power at distance d from the transmitter in the original system, $\forall d$. However, now we also need to ensure that the time-correlations in the received signal power, due

[1] For simplicity, we assume that the motion of other objects in the environment is negligible compared to the motion of the mobile nodes.

to Doppler spreading, remain the same between the two systems. (This is important, as these time-correlations determine significant performance measures, such as the *level crossing rate*, the *average fade duration*, and so on.) Mathematically speaking, we have to ensure that for any fixed d the finite-dimensional distributions of $\{P_r(t,d), t \geq 0\}$ are the same between the two systems, i.e. the stochastic process remains the same. To accomplish this for the model under study, it is enough to ensure that the maximum Doppler shift remains the same [12].

The maximum Doppler shift is defined as $f_m = \frac{v f_c}{c}$, where v is the perceived relative velocity between the transmitter and the receiver, f_c the carrier frequency, and c the speed of light. (Recall that $f_c = \frac{c}{\lambda_c}$.) Since all node speeds in the scaled system are multiplied by $\sqrt{\alpha}$, for f_m to remain the same, f_c must be divided by the same factor.

However, a scaling factor α that will cause a significant alteration of the carrier frequency may not be desirable or possible in practice. For example, a factor α that will shift the carrier frequency to another frequency band may alter the propagation characteristics of the signal significantly. [2] As a rule of thumb, we believe that the scaled carrier frequency should, at least, belong to the same frequency band as the original carrier frequency. This is based on the well-known fact that in practice, propagation characteristics of frequencies belonging to the same band are similar, *e.g.* see [14] and references therein. Given the original frequency band of operation and carrier frequency, this requirement provides a lower bound on the scaling factor α that we can have, and hence an answer to the question of *"How small can we go?"*.

However, since a change of the carrier frequency may cause problems, one might also think that it may be preferable not to scale down the speed of nodes in the scaled system, so that we could maintain the same maximum Doppler shift, without the need of altering the carrier frequency. However, this will cause the topology of the scaled system to change at a faster speed comparing to the topology of the original system. This in turn, will cause proactive routing protocols, *e.g.* DSDV [15], to perform poorer in the scaled system, since they won't be updating the routing tables fast enough to compensate for the more dynamic topology. This, for example, may cause more packet drops in the scaled system, as a larger number of stale routing table entries will direct them to be forwarded over more broken links. And, reactive routing protocols, *e.g.* DSR [16], will impose a larger routing overhead to the scaled system in order to keep up with the more dynamic topology. This implies a larger number of routing packets, which in turn, could increase the probability of packet collisions and delay data packets in network interface transmission queues, making again the scaled system performing worse.

Hence, to avoid the above issues, we prefer to scale down the speed of nodes (by $\sqrt{\alpha}$), and hence scale the carrier frequency (by $1/\sqrt{\alpha}$), under the following assumption:

[2]Other reasons may include the need for significant hardware changes, the unavailability of the requested frequencies, etc.

Assumption 3: If α is the desired scaling factor, then scaling the original carrier frequency f_c by $1/\sqrt{\alpha}$ is feasible, i.e. $\frac{f_c}{\sqrt{\alpha}}$ does not belong to another frequency band, and we can operate the scaled system on this new frequency. (In this case the corresponding scaling factor α is also called feasible.)

We now scale down the transmission power in the scaled system so that the value of $P_r^{ls}(.)$ (see Equation (5)) at distance $\sqrt{\alpha}d$ from the transmitter, is equal to the value of $P_r^{ls}(.)$ at distance d from the transmitter in the original system. (We have described earlier how this can be done depending on the large-scale model under consideration.) This ensures that for an arbitrary time t, the distribution of the received signal power at distance $\sqrt{\alpha}d$ from the transmitter in the scaled system is the same to the distribution of the received signal power at distance d from the transmitter in the original system, $\forall d$. And since we have scaled the carrier frequency by $1/\sqrt{\alpha}$, we have ensured similar time-correlations of the received signal power in the two systems. We now state the following proposition whose proof immediately follows from the above arguments.

Proposition 2: To construct a scaled-in-space by a factor α replica under Ricean/Rayleigh fading, perform steps (i)...(iv) of Theorem 1, and (step (v)) multiply the transmission power of every node as described earlier according to the large-scale model assumed, and (step (vi)) multiply the carrier frequency by $1/\sqrt{\alpha}$. If the propagation delays are insignificant and the same arbitrary traffic is applied in the original and scaled systems, then under Assumptions 2 and 3, the two systems will yield the same performance in *distribution*, for every feasible $0 < \alpha \leq 1$.

V. Time-Downscaling of Mobile ad hoc Networks

We now present the time-downscaling methodology. The method can be used in conjunction with space-downscaling in order to expedite experimentation.

Consider an original system under either a deterministic propagation model, or the log-normal shadowing model, and suppose that we have built a scaled-in-space by a factor α replica, according to the procedures we have described earlier. Also, suppose that some arbitrary routing protocol is employed, and let $\delta > 1$ be a scaling factor.

In order to perform time-downscaling, we perform the following operations to the scaled-in-space replica: (i) Multiply all times $t \geq 0$ at which nodes will *initiate* a trip by $\frac{1}{\delta}$, (ii) multiply all node speeds by δ (and hence the duration of each trip by $\frac{1}{\delta}$), (iii) multiply the inter-arrival times of the *external* network arrivals by $\frac{1}{\delta}$, (iv) multiply the transmission bandwidth by δ (and hence the packet transmission time by $\frac{1}{\delta}$), (v) speed-up all the operations of the routing protocol by δ, [3] and (vi) speed-up all the MAC layer operations by δ. [4]

The result of the above operations is a scaled-in-time replica, by a factor of $\frac{1}{\delta}$. It is easy to see that since both the

[3]For example if the DSDV routing protocol is used, multiply all protocol timeouts, *e.g.* the periodic update time-interval, the minimum triggered update period, etc., by $\frac{1}{\delta}$.

[4]This can be accomplished by multiplying all the time constants that are maintained by the MAC layer, *e.g.* the slot time duration, etc., by the factor $\frac{1}{\delta}$.

external arrival process and all network operations are scaled in time by the same factor, the only difference between the scaled-in-time replica and the original system, is that, whatever happens in the latter at time t, happens in the former at time $\frac{t}{\delta}$.

Now, suppose that the scaled-in-space replica was constructed assuming the Ricean/Rayleigh fading model, and that we have performed the operations (i)...(vi) in order to build the scaled-in-time replica. Also, recall that under small-scale fading, the received signal power exhibits time-correlations. Since the scaled-in-time replica runs on a different time-scale compared to the original system, we have to ensure that similar time-correlations exist in the *new* time-scale. In other words, if the received signal power is correlated for a time duration τ in the original system, it should be correlated for a time duration of $\frac{\tau}{\delta}$ in the scaled-in-time replica. To check this, we consider the *coherence time* T_c, which is the time domain dual to Doppler spread [12].

The coherence time represents the time duration over which two signals have strong potential for amplitude correlation and it is approximately given by $T_c \approx \frac{C}{f_m}$, where $C > 0$ a constant, and f_m the maximum Doppler shift [12]. (Recall that $f_m = \frac{v f_c}{c}$.) Hence, to have similar time-correlations, we have to check that T_c in the scaled-in-time replica is divided by δ.

Recall from the previous section that f_m in the original system is equal to that of the scaled-in-space system. Now, since we are multiplying all node speeds by δ, f_m is also multiplied by δ. Thus, T_c is divided by the same factor, as required. We can now state our final theorem. Its proof follows immediately from our discussion.

Theorem 2: Under a deterministic propagation model, the scaled-in-time replica preserves the performance of the original system as a function of time. Under log-normal shadowing or Ricean/Rayleigh fading, performance is preserved in distribution. These results hold $\forall \ \delta > 1$.

Hence, rather than performing experiments with the original system for some time duration T, one can perform experiments with the scaled-in-time replica for a much smaller time duration $\frac{T}{\delta} \ll T$ (for δ large), hence reaching to conclusions about the network behavior much faster.

VI. SIMULATIONS

In this section we use the ns-2 simulator [17] to verify our theoretical arguments and to demonstrate how accurately the scaled replica can predict the performance of the original network under realistic settings. We show results with and without time-downscaling. We compare the end-to-end packet delays, the system's throughput, and the drop ratio (defined as the ratio of the total number of packets received over the total number of packets sent), between the original and scaled systems, under the log-normal shadowing model, and the Ricean/Rayleigh fading model. Current ns-2 implementations only simulate the large-scale log-normal shadowing model. To simulate small-scale Ricean/Rayleigh fading we have applied the ns-2 extension described in [13].

The original network consists of 50 nodes and it is deployed in an area of some size A. Each node has an omnidirectional antenna with gain $G = 1$ and height $h = 1.5$m. The transmission power is $P_t = 0.2818$W, the carrier frequency $f_c = 914$MHz, and the system loss $L = 1$. Under these values the receiving threshold at the physical layer is set such that the transmission range as predicted by a deterministic large-scale model (*e.g.* see Equations (1) and (2)) is 250m. The interface queue length is set to 250 packets, the packet size is 84bytes, and the wireless bandwidth is 1Mbps.

Half of the nodes are sources generating traffic according to a Poisson process at a rate of 0.2packets/sec, destined to a randomly selected destination. The nodes are moving according to the Random Waypoint mobility model [18]. The time that each node pauses after a trip is uniformly distributed between 0sec and 10sec. And the speed of each node during a trip is uniformly distributed between 1m/s and 19m/s. Finally, the routing protocol used is the DSDV protocol [15]. [5]

The scaled-in-space replica is deployed in an area of size αA, according to the procedure of the propagation model under consideration, as described earlier. We set $\alpha = 0.01$. For performing time-downscaling to the scaled-in-space replica, we used $\delta = 10$, and followed the procedure described in the previous section.

For the shadowing propagation model we have $d_0 = 10$m, $n = 4$ and $\sigma_{dB} = 6$dB, for both the original and scaled systems. This setting simulates a shadowed outdoor urban area. For the Ricean fading model we set $K = 6$ in both systems, and for Rayleigh fading $K = 0$. We first show results where we have performed space-downscaling *only*.

A. Simulation results with space-downscaling only

For the experiments presented here, the *simulation* time for both the original and scaled systems was 2000sec. Figure 2 shows the distribution of the end-to-end packet delays (up to a delay of 40ms), under the aforementioned propagation models, for an original area size $A = 1400$x1400m^2. (Similar results hold for other area sizes.)

It is evident from the plots that the scaled-in-space system can predict the distribution of the end-to-end delays of the original system with a high accuracy. For the cases of Ricean and Rayleigh fading, we also present the distributions of the end-to-end packet delays when the carrier frequency is not scaled according our procedure (i.e. not multiplied by $\frac{1}{\sqrt{\alpha}}$), but instead retained equal to that of the original system. As we can see, in these cases performance prediction is less accurate as expected, since the received signal power does not exhibit the same time-correlations between the two systems.

Figure 3 shows how accurately the scaled-in-space replica can predict the packet drop ratio of the original system for various area sizes. And Figure 4 does the same for the system's throughput. For the cases of Ricean and Rayleigh fading, we present again the results when the carrier frequency of the

[5] Similar results hold for non-Poisson arrival processes, other mobility models, and other routing protocols.

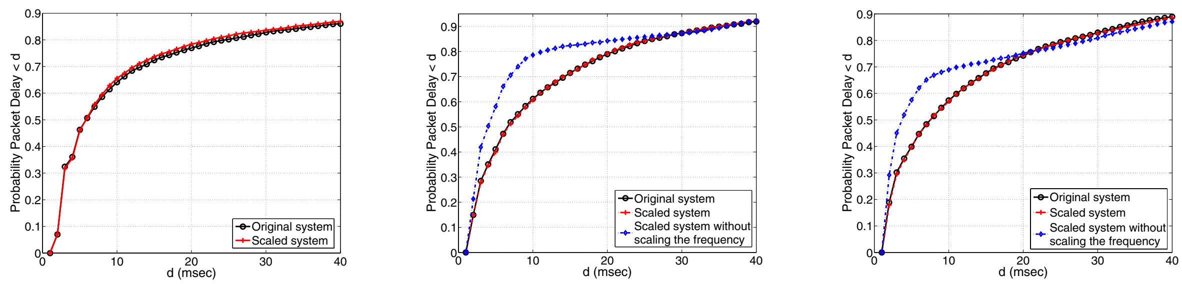

(i) (ii) (iii)

Fig. 2. Distribution of end-to-end packet delays under (i) log-normal shadowing, (ii) Ricean fading, and (iii) Rayleigh fading. (Space-downscaling only.)

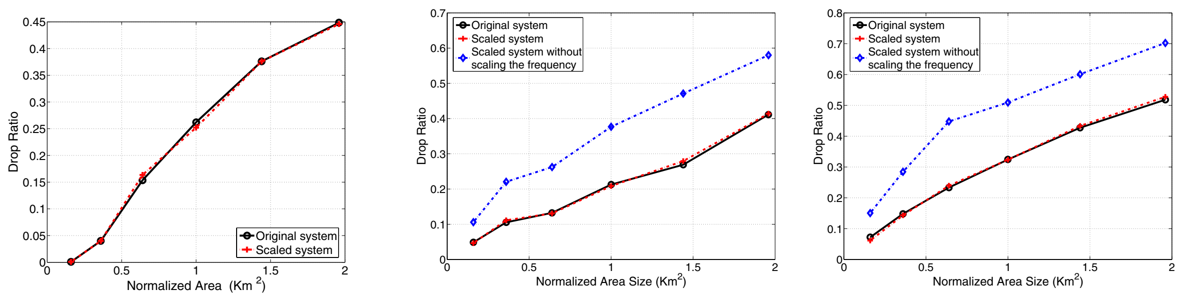

(i) (ii) (iii)

Fig. 3. Packet drop ratio under (i) log-normal shadowing, (ii) Ricean fading, and (iii) Rayleigh fading. (Space-downscaling only.)

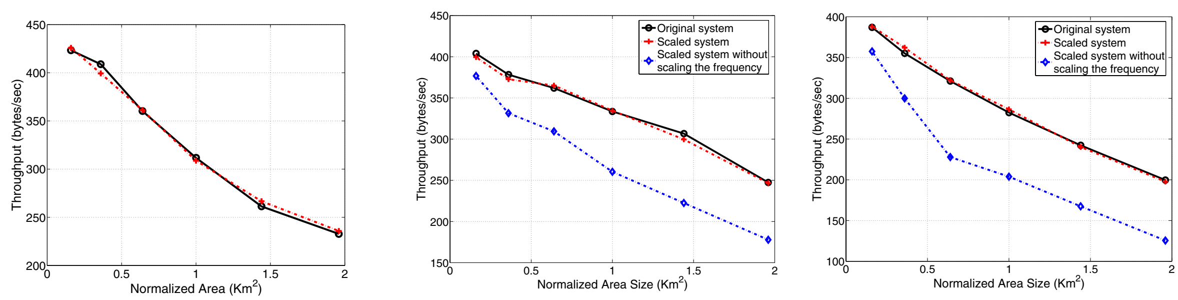

(i) (ii) (iii)

Fig. 4. System throughput under (i) log-normal shadowing, (ii) Ricean fading, and (iii) Rayleigh fading. (Space-downscaling only.)

original system is not scaled according to our procedure. Also, notice that the plots show the *normalized* area size for the scaled system, where the normalization is done as follows: if αA is the area size of the scaled system, then $\frac{\alpha A}{\alpha} = A$ is its normalized area size.

It is again evident from the plots that the scaled system can predict the performance of the original system quite accurately, and that if the carrier frequency is not scaled according to our procedure, performance prediction is inaccurate, as expected. Next, we show results, where in addition to space-downscaling, we have also performed time-downscaling.

B. Simulation results with space- and time-downscaling

Figures 5...7 show how accurately the scaled system (in both space and time) can predict the performance of the original system, under the log-normal shadowing and the Ricean fading models. (The results for Rayleigh fading are similar. We do not present those due to space limitations.). Recall that now, the scaled replica runs δ times faster than the original system. For the results presented here, we have run the original system for 2000sec of simulation time, and the scaled replica for 200sec only, since $\delta = 10$.

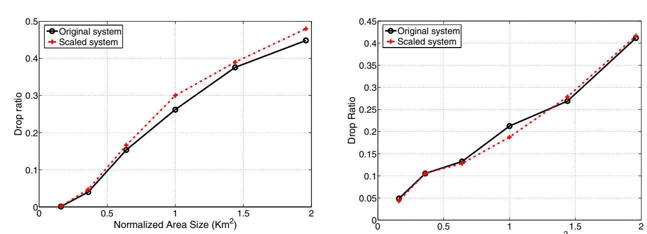

(i) (ii)

Fig. 5. Packet drop ratio (i) under log-normal shadowing, and (ii) under Ricean fading. (Space- and Time- downscaling.)

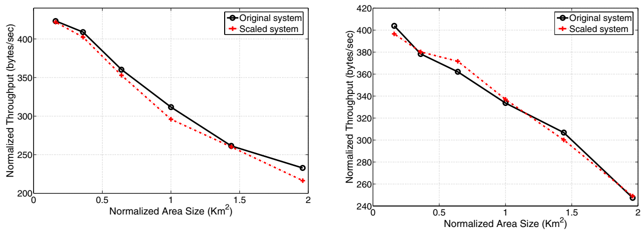

(i) (ii)

Fig. 6. Normalized system throughput (i) under log-normal shadowing, and (ii) under Ricean fading. (Space- and Time- downscaling.)

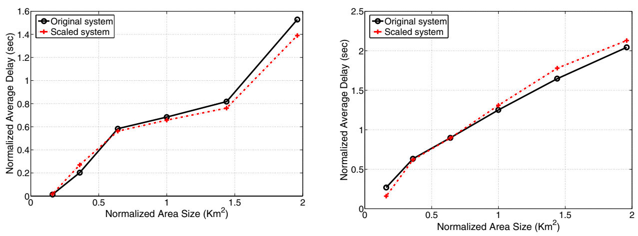

(i) (ii)

Fig. 7. Normalized average end-to-end packet delays (i) under log-normal shadowing, and (ii) under Ricean fading. (Space- and Time- downscaling.)

Notice that the plots show the *normalized* throughput and delay for the scaled system, where the normalization is done as follows: (normalized throughput)=(throughput)/δ, and (normalized delay)=(delay)$\cdot\delta$. These normalizations can be easily justified by the construction of the scaled-in-time system, as all time-intervals were multiplied by $\frac{1}{\delta}$. Also, the normalization of the area size is performed as described earlier. From the plots we observe again that the scaled replica can predict performance quite accurately.

VII. Conclusion and Future Work

In this paper we have studied ways of constructing suitably scaled-down replicas for efficient performance prediction of large-scale mobile ad hoc networks. We have identified necessary and sufficient conditions for the downscaling to be possible, and we have supported our results using both theoretical arguments and realistic simulations.

We have several interesting directions for future work. First, we believe that it is very interesting to investigate whether one could build a suitable scaled-in-space replica, under small-scale fading, without the need of scaling the carrier frequency. If such a downscaling is possible, it would alleviate the problems that may arise when attempting to scale the carrier frequency, and hence would allow us to build even smaller replicas.

Second, we are planning to investigate whether one can build scaled-down replicas for the case of ad hoc networks with significant propagation delays, and in such cases to study the impact of downscaling on traffic with feedback, *e.g.* such as TCP traffic. The downscaling in this case could be used to facilitate efficient testbed experimentation for emerging classes of such kind of networks, *e.g.* such as underwater networks.

Third, we plan to perform experiments on actual testbeds to further validate our findings under more realistic settings.

And finally, note that in this study the number of nodes between the original and scaled systems was left unaltered. This made it feasible to build scaled-down replicas having the same connectivity properties and carry the same amount of traffic, as the original network. Another interesting, yet challenging, future work direction is to investigate whether we could build scaled-down replicas consisting of only a fraction of nodes from the original network. In this case, it is unclear how one can maintain the same connectivity and traffic between the original and scaled systems. However, this kind of downscaling could be used to reduce the computational requirements of simulations and expedite experiments, since we would need to work with fewer nodes. While this has been done for the case of wireline networks that resemble the Internet [8], no such study exists for the case of mobile wireless networks.

References

[1] P. De, A. Raniwala, S. Sharma, and T. Chiueh, "MiNT: A miniaturized network testbed for mobile wireless research," in *Proc. of INFOCOM*, 2005.
[2] V. Naik, E. Ertin, H. Zhang, and A. Arora, "Wireless testbed Bonsai," in *Proc. of WiNMee*, 2006.
[3] D. Maltz, J. Broch, and D. Johnson, "Experiences designing and building a multi-hop wireless ad-hoc network testbed," *CMU TR99-116*, 1999.
[4] H. Lunndgren, D. Lundberg, J. Nielsen, E. Nordstrom, and C. Tscudin, "A large-scale testbed for reproducible ad hoc protocol evaluations," in *Proc. of WCNC*, 2002.
[5] B. A. Chambers, "The grid roofnet: A rooftop ad hoc wireless network," *MIT Masters Thesis, Tech. Rep.*, 2002.
[6] J . Lei, R. Yates, L. Greenstein, and H. Liu, "Wireless link snr mapping onto an indoor testbed," in *Proc. of Tridentcom*, 2005.
[7] K. Psounis, R. Pan, B. Prabhakar, and D. Wischik, "The scaling hypothesis: Simplifying the prediction of network performance using scaled-down simulations," in *Proc. of ACM HOTNETS*, 2002.
[8] F. Papadopoulos, K. Psounis, and R. Govindan, "Performance preserving topological downscaling of internet-like networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 12, 2006.
[9] E. Ertin, A. Arora, R. Ramnath, M. Nesterenko, V. Naik, S. Bapat, V. Kulathumani, M. Sridharan, Hongwei Zhang, and H. Cao, "Kansei: A testbed for sensing at scale," in *Proc. of IPSN/SPOTS*, 2006.
[10] N. H. Vaidya, J. Bernhard, V. Veeravalli, P. R. Kumar, and R. Iyer, "Illinois wireless wind tunnel: A testbed for experimental evaluation of wireless networks," in *Proc. of ACM SIGCOMM E-wind Workshop*, 2005.
[11] J. Kong, J. Cui, D. Wu, and M. Gerla, "Building Underwater Ad-hoc Networks and Sensor Networks for Large Scale Real-time Aquatic Applications," in *Proc. of MILCOM*, 2005.
[12] T. S. Rappaport, *Wireless Communications: Principles and Practice (2nd Edition)*, Prentice Hall PTR, 2001.
[13] R. J. Punnoose, P. V. Nikitin, and D. D. Stancil, "Efficient Simulation of Ricean Fading within a Packet Simulator," in *Proc. of Vehicular Technology Conference*, 2000.
[14] "IEEE 802 LAN/MAN standards committee," http://www.ieee802.org/22/Meeting_documents/2005_July/22-05-0048-00-0000_WRAN_Multipath_Channel_Considerations.doc.
[15] Charles E. Perkins and Pravin Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," in *Proc. of SIGCOMM*, 1994.
[16] David B. Johnson and David A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, 1996.
[17] "Network simulator," http://www.isi.edu/nsnam/ns.
[18] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *Mobile Computing and Networking*, 1998.