Impact of Correlated Mobility on Delay–Throughput Performance in Mobile Ad Hoc Networks

Impact of Correlated Mobility on Delay-Throughput Performance in Mobile Ad-Hoc Networks

Delia Ciullo *, Valentina Martina *, Michele Garetto †, Emilio Leonardi *
* Dipartimento di Elettronica, Politecnico di Torino, Torino, Italy
† Dipartimento di Informatica, Università di Torino, Torino, Italy

***Abstract*—We extend the analysis of the scaling laws of wireless ad hoc networks to the case of correlated nodes movements, which are commonly found in real mobility processes. We consider a simple version of the Reference Point Group Mobility model, in which nodes belonging to the same group are constrained to lie in a disc area, whose center moves uniformly across the network according to the i.i.d. model. We assume fast mobility conditions, and take as primary goal the maximization of per-node throughput. We discover that correlated node movements have huge impact on asymptotic throughput and delay, and can sometimes lead to better performance than the one achievable under independent nodes movements.**

***Index Terms*—Ad-hoc networks, asymptotic scaling laws, delay-throughput performance, correlated mobility**

I. Introduction and related work

In the last few years the *store-carry-forward* communication paradigm, which allows nodes to physically carry buffered data as they move around the network area, has opened an entire new area of research with many promising applications in the context of delay-tolerant networking [1].

In their seminal work [2], Grossglauser and Tse have shown that mobile nodes employing the *store-carry-forward* paradigm can achieve constant throughput even when the number of nodes grows to infinity, in contrast to the severe throughput decay (like $1/\sqrt{n}$) incurred in fixed networks [3]. The basic requirement of their 2-hop scheme is that nodes uniformly visit the entire network space according to an arbitrary, stationary and ergodic mobility process with *independent* trajectories.

When considering also the delay performance, the specific details about how nodes move become important. Several papers have analyzed throughput-delay trade-offs for various mobility models, ranging from the simple reshuffling model (also referred to as i.i.d. model) [4], [5], [6], to the Brownian motion [7], and variants of random walk and random waypoint [8], [9]. In [10] the authors have extended the throughput and delay scaling results of Grossglauser-Tse to more general inter-contact time distributions than the exponential distribution, allowing to account for the correlations existing in the mobility pattern of individual nodes.

The impact of limited buffers has been considered in [10], [11]. In [12] it is shown that delay-throughput trade-offs, close to those achievable in mobile networks under reshuffling mobility models, can be achieved in fixed networks by employing advanced cooperative (MIMO) transmission schemes.

A preliminary version of this paper has appeared at INFOCOM 2010

In the above-mentioned works, the mobility of the nodes has always been assumed to be uncorrelated (*i.e.*, independent from node to node) and uniform over the area.

Some authors have already considered the impact on the capacity of restricted mobility models (*i.e.*, relaxing the assumption that nodes uniformly visit the network area) [13], [14], [15], [16], still maintaining the independence assumption on the nodes mobility processes.

To the best of our knowledge, no work has been done so far to investigate the impact of correlation among nodes movements on the asymptotic throughput and delay of large mobile networks. This is rather surprising in light of the fact that real mobility processes (of pedestrians, vehicles, animals) exhibit significant degrees of correlation, as observed in several traces [17], [18], [19], [20].

The goal of our work is to study, for the first time, the scaling laws of capacity and delay for large mobile networks including correlated nodes movements. To this aim, we consider a very simple model of correlated mobility based on the popular Reference Point Group Mobility (RPGM) model introduced in [21]. Nodes are organized into several groups, and the mobility of nodes belonging to the same group is confined within a disc area. Each group has a logical center, which moves around the network dragging behind all nodes belonging to the group. Notice that in the long run each node uniformly visits the entire network space, however the trajectories of individual nodes are not independent because they are constrained to jointly follow their respective groups. By changing a few parameters, our model allows to explore various degrees of correlation in the node mobility process.

We propose novel scheduling-routing schemes whose primary goal is to maximize the per-node throughput. As a secondary goal, we also seek to minimize the packet delivery delay. Our main finding is that node correlation has a strong impact on both throughput and delay performance. Interestingly, correlated mobility can lead both to better and to worse performance with respect to the case in which node movements are independent.

Prior to our work, the impact of correlated node movements on existing and novel routing protocols has been extensively investigated by simulation. In the context of traditional *store-and-forward* networks, [22] analyzed the effect of various mobility models, including correlated movements, on classical routing protocols (DSR, AODV), while in [23] the authors have proposed a novel routing protocol, called LANMAR, which directly exploit group mobility patterns to improve routing efficiency. Similarly to our scheme, they propose a

hierarchical approach in which data are first routed at the group level, and then routed within the group containing the destination. A similar idea is proposed in [24] for the *store-carry-forward* communication paradigm. In particular, a history-based approach similar to PRoPHET [25] is adopted at the group level. Like us, the authors of [24] also employ a replication strategy to improve the delivery delay. We emphasize that previous work relied entirely on simulations to evaluate the performance of the proposed schemes, without analyzing asymptotic scaling laws nor the optimality of the proposed solutions in terms of system throughput and delay.

The rest of the paper is organized as follows. We first introduce our system assumptions in Section II. In Section III we analyze the *cluster sparse* regime, where nodes belonging to different groups (*clusters*) meet sporadically. The *cluster dense* regime, in which nodes belonging to different clusters meet frequently, is briefly discussed in Section IV. In Section V we illustrate our main findings. In Section VI we present several extensions of our scheme considering more realistic mobility models. We conclude in Section VII.

II. System assumptions

A. Mobility Model

We consider an extended network comprising n nodes moving over a square region $\mathcal{O}$ of area n with wrap-around conditions (*i.e.*, a torus), to avoid border effects. Note that, under this assumption, the node density over the area remains constant and equal to 1, as we increase n.

We assume that nodes are partitioned into m groups, with[1] $m = \Theta(n^\nu)$, $\nu \in [0, 1)$. For simplicity, we assume that each group comprises an integer number $q = n/m$ of nodes. Note, however, that our results would not change, in scaling order, if the cardinalities of the groups were not exactly the same, as long as each group contains $\Theta(n/m) = \Theta(n^{1-\nu})$ nodes.

Time is divided into slots of equal duration, which is normalized to 1. Nodes belonging to the same group move over the network area in a correlated fashion. To model this behavior, we assume that, at any given slot, all nodes of a group have to reside concurrently within a same portion, of area $o(n)$, of the total network space. In the following we will refer to such a portion as the *cluster-region* or simply the *cluster*, associated to the group.

We assume that each cluster-region has a circular shape of radius R. We can explore various degrees of correlation in the node mobility process by letting R scale with n as well, as $R = \Theta(n^\beta)$, with $\beta \in [0, 1/2)$. Notice that $\beta = 0$ corresponds to the extreme case in which each group occupies a constant fraction of the network area (just as if all nodes of a group were located at a single point), irrespective of the number of nodes in it.

We have yet to specify how nodes move over the network area from one slot to another. The mobility process of a given node i belonging to group j is described by the combination of two movements: i) a group movement (*i.e.*, the shift of the cluster-region associated to group j during a slot); ii) a node movement (*i.e.*, the change of position of node i within cluster-region j).

For what concerns the group movement, we start assuming that each cluster-region has a center point, whose position is updated at each time slot by choosing a new location uniformly at random in the network area, independently for each group. This is similar to the so-called reshuffling model, or bi-dimensional i.i.d. mobility model, considered in previous work [4], [5], [6], however note that here we adopt this model only to update the positions of the cluster centers. The mobility processes of individual nodes are not independent in our model because, once the new position of a cluster center has been selected, all nodes belonging to the corresponding group have to move to a place close to it (*i.e.*, inside a region of area $o(n)$ around the cluster center). We observe that the degree of correlation in the node mobility process increases as we either i) reduce the area of each cluster-region (smaller values of β); ii) reduce the number of groups (smaller values of ν). Table I summarizes the notation used throughout the paper.

Later on, in Section VI, we will generalize our analysis along three directions. First, we will study the case in which nodes do not change their relative positions with respect to the cluster center (we refer to this case as the *crystallized* model). Second, we will consider more realistic mobility patterns than the simple reshuffling model, assuming that both center-points and nodes within their cluster region move according to a general random walk. Third, we will allow the nodes to migrate from cluster to cluster, instead of being permanently associated to the same cluster.

B. Communication Model

To account for interference among simultaneous transmissions, we adopt the protocol model introduced in [3] and widely used in the literature[2]. According to the protocol model, nodes employ a common range r for all transmissions which occur in the same time slot (r can be different from slot to slot); equivalently, they employ a common power level in each slot. A transmission from node i to node j using transmission range r can be successfully received at node j if and only if the following two conditions hold:

1) the distance between i and j is smaller than or equal to r, *i.e.*, $d_{ij}(t) \leq r$.
2) for every other node k simultaneously transmitting, $d_{kj}(t) \geq (1 + \Delta)\, r$, being Δ a guard factor.

Transmissions occur at fixed rate which is normalized to 1. Moreover, we consider fast mobility conditions, according to which data can be transmitted over just one hop during any slot[3].

C. Traffic Model

Similarly to previous work we consider permutation traffic patterns in which every node is origin and destination of

[1]Given two functions $f(n) \geq 0$ and $g(n) \geq 0$: $f(n) = o(g(n))$ means $\lim_{n\to\infty} f(n)/g(n) = 0$; $f(n) = O(g(n))$ means $\limsup_{n\to\infty} f(n)/g(n) = c < \infty$; $f(n) = \omega(g(n))$ is equivalent to $g(n) = o(f(n))$; $f(n) = \Omega(g(n))$ is equivalent to $g(n) = O(f(n))$; $f(n) = \Theta(g(n))$ means $f(n) = O(g(n))$ and $g(n) = O(f(n))$; at last $f(n) \sim g(n)$ means $\lim_{n\to\infty} f(n)/g(n) = 1$

[2]Our results would not change under the physical model defined in [3], provided that the power loss exponent is larger than 2.

[3]We leave to future work the extension of the analysis to the slow mobility case, in which multi-hop transmissions can be performed during the same slot.

TABLE I
NOTATION

Symbol	Definition
n	number of nodes
m	number of clusters
ν	growth exponent of m: $m = \Theta(n^\nu)$, $\nu \in [0,1)$
q	number of nodes in each cluster, $q = n/m$
R	cluster radius
β	growth exponent of R: $R = \Theta(n^\beta)$, $\beta \in [0,1/2)$
λ	per-node (or per-flow) throughput
D	end-to-end packet delivery delay

a single traffic flow of rate λ. Hence there are n source-destination (S-D) pairs in the network.

Messages[4] are generated at every source according to independent memoryless Bernoullian processes.

We use the following definitions of asymptotic throughput and delay. Let $L_i(T)$ be the number of packets delivered to the destination of node i in the time interval $[0,T]$. The delay of a packet is the time it takes to the packet to reach the destination after it leaves the source. Let $D_i(t)$ be the sum of the delays experiences by all packets successfully delivered to the destination of node i in the time interval $[0,T]$. We say that an asymptotic throughput λ and an asymptotic delay D per S-D pair are feasible if there is an n_0 such that for any $n \geq n_0$ there exists a scheduling/routing scheme for which both the following properties hold: $\lim_{T\to\infty} \Pr\left(\frac{L_i[T]}{T} \geq \lambda, \forall i\right) = 1$ and $\lim_{T\to\infty} \Pr\left(\frac{D_i[T]}{L_i[T]} \leq D, \forall i\right) = 1$. Equivalently, we say in this case that the network sustains an aggregate throughput $\Lambda = n\lambda$ at the expense of a delay D.

III. The cluster sparse regime

We can distinguish two regimes depending on the values of β and ν. We say that the system is in *cluster sparse* regime if $\nu + 2\beta < 1$ (*i.e.*, $mR^2 = o(n)$), and in *cluster dense* regime if $\nu + 2\beta > 1$[5]. In this section we consider the *cluster sparse* regime, which is more interesting and challenging to analyze. We briefly analyze the *cluster dense* regime in Section IV.

In the *cluster sparse* regime, at any time clusters cover only a negligible fraction of the entire network area. Actually they form several small, disconnected and highly dense regions (the node density within a cluster is $\frac{n}{mR^2} = \omega(1)$) floating over a huge empty space. Spatial overlaps between different clusters are sporadic.

In the following we first introduce the scheduling-routing scheme that we have developed for this case, describing the routing scheme in Section III-A and the associated scheduling scheme in Section III-B. Then, in Section III-C we analyze the performance of the proposed scheme and prove its optimality.

A. Routing scheme

We propose a multi-hop routing scheme that generalizes the 2-hop scheme introduced by Grossglauser and Tse [2].

We focus on a particular traffic stream $s \to d$. Let C_s denote the cluster containing s (*i.e.*, $s \in C_s$) and C_d the cluster containing d (*i.e.*, $d \in C_d$). We neglect the particular case in which $C_s = C_d$, since w.h.p. s and d belong to different clusters (this is also the most stressful case for the system).

[4]In this paper the terms message and packet are interchangeable.

[5]We leave for future investigations the special case in which $\nu + 2\beta = 1$

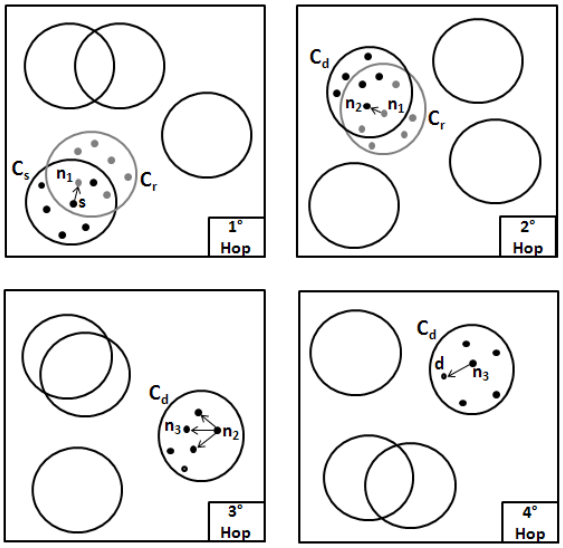

Fig. 1. Illustration of the 4-hop routing scheme

The rationale of our routing scheme is to first reach a node within the destination cluster C_d in the most efficient way, and then to forward the packets within C_d up to the final destination d. We anticipate that the system throughput is bottlenecked in the first phase of the route, in which data has to reach the destination cluster: this is due to the fact that close contacts among nodes belonging to different clusters are rare, since they occur only when two clusters overlap in space.

The same principles that inspired the 2-hop scheme of Grossglauser and Tse suggest that the most efficient way (in order to maximize the throughput) to bring a message within the destination cluster is to adopt a 2-hop relaying scheme at the cluster level, in which each packet transits through a random intermediate cluster C_r. This allows transmitters to exploit all contacts with nodes belonging to a different cluster.

Once the packet arrives within the destination cluster, we can exploit well-known schemes developed for mobile network with uniform, uncorrelated mobility patterns. Indeed, notice that, under our mobility model, each cluster can be regarded as a micro-universe of nodes forming a classical mobile network in which nodes move uniformly according to the i.i.d. model. Since the throughput is bottlenecked in the previous part of the route, it turns out that, within the destination cluster, it is convenient to adopt a replication strategy, in which the packet is first broadcasted to all nodes falling within a suitable transmission range, and then one of the copies is delivered to the final destination in one more hop. This replication strategy allows to reduce the packet delivery delay without negatively impacting the overall system throughput.

Figure 1 graphically illustrates the routing scheme outlined so far. There are 4 hops and 3 intermediate relays. In the first hop, source node s sends the message to a relay node n_1 belonging to an arbitrary cluster C_r different from C_s. In the second hop, node n_1 forwards the message to a node n_2 belonging to C_d. In the third hop, the message is replicated by n_2 through a single transmission to several nodes belonging to the same cluster C_d, exploiting the intrinsic broadcast

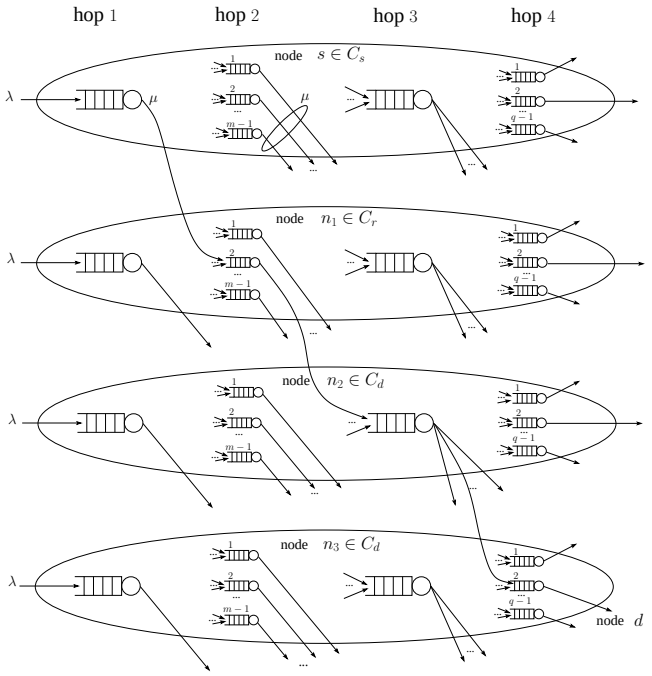

Fig. 2. System representation as an open (acyclic) network of FIFO queues. Connections among queues are drawn for flow $s \to d$.

capability of the wireless channel. In the fourth hop, one of the nodes holding a copy of the message (let it be node n_3) delivers the message to the final destination.

B. Scheduling scheme

To implement the above described routing scheme, each node is equipped with (see Figure 2): i) one queue storing its own generated packets (*i.e.*, packets at hop 1); ii) $m-1$ parallel queues, one per cluster, storing packets at hop 2; iii) one queue for packets at hop 3; iv) $q-1$ parallel queues, one for each possible destination within its own cluster, for packets at hop 4. The service discipline is First Come First Served (FCFS) at all queues.

The scheduling scheme is in charge of selecting, at any time slot, a set of transmitter-receiver pairs which can communicate successfully according to the protocol model. Recall that the protocol model requires the adoption of the same transmission range for all communications occurring in the same slot; on the other hand, it is convenient to employ different transmission ranges for the various hops of the routing scheme. For this reason, each slot is devoted only to the transmission of packets which are at the same hop of the route. This can be equivalently done in a round-robin or in a probabilistic fashion. Following a round-robin approach, we identify the slots by a sequence number t, and in the generic slot t we allow only the transmission of packets at hop $i = |t|_4 + 1$, where $|\cdot|_4$ denotes the modulus-4 operation.

One simple way to completely eliminate interference among concurrent transmissions, as required by the protocol model, is the following. Let r_i be the transmission range of packets at hop i ($i = 1, \ldots, 4$). In any slot devoted to hop i, domain $\mathcal{O}$ is divided into squarelets $\mathcal{A}_i^k$ of area A_i and edge length r_i. A subset of squarelets, regularly spaced, is selected, and at most one node is allowed to transmit in each squarelet belonging to the selected subset. Figure 3 illustrates this construction for a protocol model having $\Delta = 2$. Shaded squarelets represent one possible subset of regularly spaced squarelets. Domain $\mathcal{M}_i$ around one of the squarelet denotes the maximum-size region where we can find a receiver for an arbitrary transmitter falling in the squarelet. Domain $\mathcal{T}_i$ denotes instead the region where we cannot have any other receiver belonging to a different communication pair. By spacing the selected squarelets with step $d_i = (\Delta + 2)r_i$ we can assure that one transmitter per squarelet can be enabled to transmit without generating any conflict, irrespective of the locations of transmitters within their squarelets.

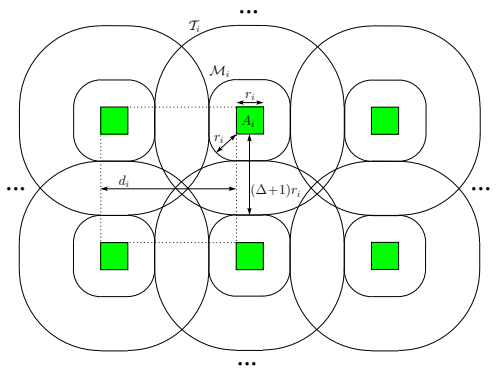

Fig. 3. Illustration of the scheduling scheme in the case of $\Delta = 2$.

C. Performance analysis

To evaluate the performance of our scheme, we proceed in four steps. In Section III-C1 we compute a simple upper bound to the throughput that any possible scheme can achieve. The obtained upper bound allows to gather some insight on the effect of the transmission range on system performance. In Section III-C2 we move a step forward computing the maximum theoretical throughput that our scheme can achieve. At this stage, we assume that all queues are constantly backlogged with packets, and we compute the maximum saturation throughput achieved by inter-cluster communications, *i.e.*, the aggregate service rate of all queues storing packets to be transmitted to nodes in different clusters. This quantity is simpler to analyze, because it requires only geometric considerations.

In Section III-C3 we take into account traffic and queuing effects. We show that nodes' queues can be loaded in such a way that the actual system throughput is in order sense the same as the saturation throughput. Given the parameters that maximize the system throughput, in Section III-C4 we compute the resulting end-to-end delivery delay. We also show that any scheduling-routing scheme that achieves the maximum throughput computed in Section III-C3 cannot incur a delay smaller than the one derived for our scheme.

1) Throughput upper bound: We start our performance analysis by establishing an upper bound to the network throughput. We begin with the following lemma, that characterizes the aggregate amount of data that can be transferred in one slot among nodes belonging to different clusters.

Lemma 1: Under the *cluster sparse* regime, the maximum amount of data that can be exchanged during one slot, in a single hop, by any feasible set of transmitter-receiver pairs

employing a common transmission range r (*i.e.*, tx-rx pairs which can be enabled to transmit simultaneously according to the protocol model), such that for each pair the transmitter and the receiver belong to different clusters, is $O(mR^2)$, provided that $r = O(R\sqrt{m/n})$.
The proof of Lemma 1 is reported in Appendix A.

The above lemma can be used to readily prove the following fundamental result:

Theorem 1: Under the *cluster sparse* regime, the network throughput is $O(mR^2)$. Furthermore, every scheme employing a transmission range $r = \omega(R\sqrt{m/n})$ for inter-cluster transmissions, necessarily achieves a throughput $o(mR^2)$.

Proof: First we observe that, under a permutation traffic matrix, almost all flows (*i.e.*, $\Theta(n)$ flows) are established between pairs of nodes belonging to different clusters, since the event that both the source and the destination belong to the same cluster is negligible. Then, almost all traffic requires to be exchanged at least once between a pair of nodes belonging to different clusters. An upper bound to the system throughput is then obtained assuming that flows established between different clusters (almost all flows) require just a single-hop communication between an arbitrary pair of nodes belonging to different clusters. Devoting all slots and all network resources to such arbitrary inter-cluster communications, and applying Lemma 1, we immediately conclude that the system throughput is $O(mR^2)$. Moreover, maximal throughput can be potentially achieved only by scheduling algorithms employing a transmission range $r = O(R\sqrt{m/n})$ (as computed in the proof of Lemma 1) for transmissions between nodes belonging to different clusters. ∎

2) Saturation throughput analysis: The result reported in the previous subsection represents only an upper bound to the maximum throughput achievable by any implementable scheme in the *cluster sparse* regime. Now we show that our 4-hop scheduling-routing scheme (described in Sections III-A and III-B) is able to sustain a throughput (in the first two hops), which scales with n as the upper bound established in Theorem 1. This scaling law holds provided that nodes have all their queues (see Figure 2) constantly backlogged by packets to send.

Similarly to the derivation of the upper bound, we start with a useful lemma that characterizes the aggregate amount of data that can be exchanged, under saturated conditions, among nodes belonging to different clusters, in one slot.

Lemma 2: Under the *cluster sparse* regime, assuming saturated conditions at every node, the amount of messages that can be transferred in one slot among nodes belonging to different clusters is $\Theta(mR^2)$, by employing a transmission range $\Theta(R\sqrt{m/n})$.
The proof of Lemma 2 is reported in Appendix B.

Now we observe that the network of queues modeling the system is an acyclic network of FIFO queues, since every packet traverses the queues sequentially, as illustrated in Figure 2. The aggregate service-rate obtained by queues storing messages either at the first or in the second hop can be interpreted as the amount of messages that can be transferred in one slot among nodes belonging to different clusters, as specified in Lemma 2.

As immediate consequence of the above lemma we have:

Corollary 1: By adopting a transmission range $r_i = \Theta(R\sqrt{m/n})$ at hops $i = 1, 2$, our scheduling-routing scheme provides an aggregate service rate $\mu = \Theta(mR^2/n)$ to each node, for messages either at the first or the second hop.

Proof: Consider a time slot devoted to the transmission of packets in the first or second hop. According to Lemma 2, our scheme allows $\Theta(mR^2)$ transmission opportunities (with high probability) between nodes belonging to different clusters. By symmetry, these transmission opportunities are uniformly distributed among the nodes in the network. Considering that each hop is scheduled once every 4 slots, an average number of transmission opportunities per slot (*i.e.*, a service rate) $\mu = \Theta(mR^2/n)$ is guaranteed to each node for both the first and the second hop. Note that, in the case of the second hop, the service rate of a node corresponds to the sum of the service rates of the $m - 1$ distinct queues storing packets at hop 2. Furthermore the service rates of these $m - 1$ queues are all equal by symmetry. ∎

3) Maximum achievable throughput: We are now ready to derive our main result on the maximum throughput achievable by our scheme:

Theorem 2: The maximum sustainable throughput of our scheme is $\Lambda = \Theta(mR^2)$ by employing a transmission range $r_i = \Theta(R\sqrt{m/n})$ for $i = 1, 2$ and $r_i = \Theta(1)$ for $i = 3, 4$. The corresponding per-node throughput is $\lambda = \Theta(mR^2/n)$.

Proof: First observe that, by adopting a transmission range $r_i = \Theta(R\sqrt{m/n})$ at hop $i = 1, 2$, Corollary 1 guarantees that our routing-scheduling scheme can provide an aggregate node service rate $\mu = \Theta(mR^2/n)$ for packets at either the first or the second hop.

Second, we consider that, by symmetry, our scheme uniformly distribute the traffic among all the nodes/queues, so that all queues in the network storing packets at hop i are subject to the same ingress packet arrival rate. As a consequence all queues storing packets at hop $i = 1, 2$ are jointly stable under an arrival rate that is strictly below the service rate (*i.e.*, the saturation throughput of inter-cluster communications), which is given in Corollary 1.

Once a packet arrives in its destination cluster, we can exploit well known schemes developed for networks with i.i.d. mobility. In principle, we could get an optimal per-node intra-cluster throughput $\Theta(1)$ by employing the 2-hop scheme of Grossglauser and Tse [2], using the same transmission range $r_i = \Theta(R\sqrt{m/n})$ adopted in previous hops, which is adapted to the node density within a cluster. However, this is a bad choice, because by so doing the throughput would be bottlenecked by the previous hops, and we would pay excessive intra-cluster delays for nothing. Therefore in the destination cluster the optimal choice is to exactly match the throughput achievable in the previous hop, trading off capacity and delay. Indeed it is possible to enlarge the transmission range within the destination cluster up to $r_i = \Theta(1)$ without affecting the overall system throughput. This allows to adopt a replication scheme according to which the packet is forwarded in the third hop to $\Theta(\frac{n}{mR^2})$ nodes (all nodes falling within transmission range $r_3 = \Theta(1)$ of the sender) with a single broadcast transmission. Then the first node holding a copy of the message that arrives within transmission range $r_4 = \Theta(1)$ of the destination, eventually delivers the packet in hop 4. By

so doing, the per-node throughput achievable at hop $i = 3, 4$ is reduced to $\Theta(\frac{mR^2}{n})$, for effect of the reduced spatial reuse, however the delay performance of the scheme is greatly improved thanks to replication (as better explained in the following delay analysis). ■

At last, by combining Theorem 1 and Theorem 2, it immediately descends:

Corollary 2: Our proposed scheduling-routing scheme is in order sense throughput-optimal.

4) Delay Analysis: Turning our attention to the delay performance of our scheme, we focus on a particular packet belonging to a generic flow $s \rightarrow d$ (we assume that s and d belong to different clusters) and evaluate the different components of its end-to-end delivery delay, denoted by D. Let D_i be the total delay experienced by the packet at hop i. We have $D = \sum_{i=1}^{4} D_i$.

Our first step is to compute the average service time (*i.e.*, the access delay) of the queues associated with the four hops made by the packet. Let D_i^{α} be the average service time of hop i. As shown in Appendix C, we have

$$\begin{array}{ccc} D_1^{\alpha} = \Theta\left(\frac{n}{mR^2}\right) & ; & D_2^{\alpha} = \Theta\left(\frac{n}{R^2}\right) \\ D_3^{\alpha} = \Theta\left(1\right) & ; & D_4^{\alpha} = \Theta\left(\frac{mR^4}{n}\right) \end{array} \quad (1)$$

We observe that, at each queue, the total delay would be equal to the access delay in the absence of any contention with other packets in the network (both in the same queue and in the queues of other nodes competing for the wireless medium access). Similarly to previous work, we can show that, at any hop, contention with other packets in the network does not change the order of magnitude of the total delay with respect to the access delay. As a result,

Theorem 3: In the *cluster sparse* regime, the delay performance of our scheme satisfies

$$D = \Theta\left(\max\left\{\frac{n}{R^2}, \frac{mR^4}{n}\right\}\right) \quad (2)$$

Proof: Considering the first two hops, we observe that contention among different queues (*i.e.*, different transmitter-receiver pairs) within the same squarelet can be neglected (in order sense), since by construction only a finite number of such pairs fall w.h.p. in the squarelet. Hence, both inter-arrival (for the queue storing second hop packets) and service times at queues can be bounded by a geometrically distributed number (with finite average) of cluster inter-meeting times, which, in turns, are i.i.d geometrically distributed as shown in Appendix C (due to the memoryless property of cluster positions).

This implies that both inter-arrival and service times at queues can be bounded by geometrical distributed variables and the resulting queuing delay is of the same order of the access delay, as it immediately follows from the application of the Pollaczek-Khinchine formula for discrete time queues [27].

For the analysis of the third and fourth hops, instead, we can directly apply previous results obtained for networks in which node movements are i.i.d [6] (recall once again that relative movements of nodes within a cluster are i.i.d.) and claim that $D_3 = \Theta(D_3^{\alpha})$ and $D_4 = \Theta(D_4^{\alpha})$. Since $D_i = \Theta(D_i^{\alpha})$ for all i, we have $D = \Theta(\sum_{i=1}^{4} D_i^{\alpha})$, and the result follows applying the expressions in (1). ■

We observe that similar arguments have been applied in [7], [9], [4], [6] in the case of uncorrelated i.i.d. mobility, showing that the end-to-end delay equals, in order sense, the sum of the access delays whenever the traffic injected in the network is strictly less than the saturation throughput.

Moreover we can prove the following result, which shows that our scheme achieves optimal delay performance among the class of schemes maximizing the throughput:

Theorem 4: Any scheduling-routing scheme that achieves an aggregate throughput $\Lambda = \Theta(mR^2)$ necessarily incurs a delay $D = \Omega\left(\max\left\{\frac{n}{R^2}, \frac{mR^4}{n}\right\}\right)$.

Proof: To achieve throughput $\Lambda = mR^2$ it is necessary to employ a transmission range $\Theta(R\sqrt{m/n})$ for inter-cluster communications, as a consequence of Theorem 2. Notice that using this transmission range it is not possible to get any delay gain (in order sense) by employing packet replication during inter-cluster communications (only $\Theta(1)$ nodes can simultaneously receive the message). Thus necessarily a delay $D = \Omega(D_2^{\alpha})$ must be suffered by a message to reach the destination cluster C_d, since the last relay node along its path, that does not belong to C_d, must come in contact with some node in C_d before it can transmit the message. Within the destination cluster we can apply the general trade-off $D = \Omega(n\lambda^2)$ derived in [5], [6] for networks with i.i.d mobility. Using $q = n/m$ in place of n in this trade-off formula, and plugging in $\lambda = mR^2/n$, we obtain a delay $D = \Omega(\frac{mR^4}{n})$ due to intra-cluster communications. Combining the above two constraints on D due to inter- and intra- cluster communications, we get the assertion. ■

IV. The cluster dense regime

In the *cluster dense* regime, which occurs when $\nu + 2\beta > 1$, clusters are highly overlapped at any point of the network area. Indeed, applying standard results borrowed from the theory of random geometric graphs [28], it can be shown that every point of the network area is w.h.p covered by a number of clusters $\Theta(mR^2/n) = \Theta(n^{\nu+2\beta-1})$. This implies that nodes are almost uniformly distributed over the network domain, hence the typical distance at which one node finds the node closest to it is $\Theta(1)$. Such closest node, however, belongs w.h.p. to a different cluster. To see this, note that the density of nodes within a cluster is $\frac{q}{\pi R^2} = o(1)$, resulting into a typical distance $\omega(1)$ between nodes belonging to the same cluster.

This fact dramatically limits the degree of freedom that we have in the design of a scheduling-routing scheme specifically targeted at maximizing the system throughput. Indeed, any scheme requiring at some stage that packets are transferred between nodes belonging to the same cluster must adopt a transmission range $\omega(1)$ for such intra-cluster communications, resulting into a per-node throughput $\lambda = o(1)$.

On the contrary, a simple 2-hop scheme similar to the one proposed by Grossglauser-Tse [2], according to which packets are sent from source s to destination d though a single relay node that does not belong neither to C_s nor to C_d, can effectively employ a transmission range as short as $\Theta(1)$ (only

TABLE II
SUMMARY OF RESULTS

	cluster dense	cluster sparse
λ	$\Theta(1)$	$\Theta(mR^2/n)$
D	$\Theta(n)$	$\Theta\left(\max\left\{\frac{n}{R^2}, \frac{mR^4}{n}\right\}\right)$

inter-cluster communications are required), thus achieving a per-node throughput $\lambda = \Theta(1)$, at the expense of a delivery delay $D = \Theta(n)$.

V. SUMMARY OF RESULTS

In this section, we summarize and graphically present the main results derived in the previous section. We recall that the primary goal of our schemes is to maximize the system throughput. Only as a secondary goal we seek to minimize delay. Hence our schemes do not explore the full range of possible capacity-delay trade-offs.

Table II summarizes throughput and delay results obtained so far. In the cluster dense regime, since clusters are largely overlapped and the overall node density is constant in order sense, we get the same results as the original Grossglauser-Tse scenario with independent node movements. In the cluster sparse regime, the correlation among node trajectories may have a significant impact on both throughput and delay performance. Throughput is reduced for effect of clustering w.r.t to the original Grossglauser-Tse scenario: the maximum number of parallel transmissions is reduced from $\Theta(n)$ to $\Theta(mR^2)$ (when tx-rx node pairs are required to belong to different clusters; see Lemma 1 and Theorem 1). On the other hand, delay performance is improved by the effect of clustering: nodes belonging to the destination cluster, which meet the destination at the highest frequency, can be effectively exploited as last hop relays. Consequently, the overall delay is reduced, since it is dominated by the last-hop delay in the original Grossglauser-Tse scenario.

Figure 4 provides a graphical representation of our results on the throughput-delay plane. Note that the x-axis and y-axis of the figure report the scaling exponents $e(D) = \log_n(D)$ and $e(\lambda) = \log_n(\lambda)$ respectively. The shaded region denotes all optimal operating points that are obtained as we vary the parameters β and ν of our model. The gray scale is related to the correlation degree in the mobility process: lower correlation, which results from increasing either β or ν, corresponds to darker gray.

We have reported on the plot the line $D = n\lambda^2$ which denotes (neglecting logarithmic factors) the best possible throughput-delay trade-offs that can be obtained under independent reshuffling of all nodes (and fast-mobility conditions), according to [5], [6]. Notice, however, that the above scaling law is achievable for the i.i.d. reshuffling model only at the cost of scaling the transmission range (or, equivalently, the transmission power) to infinite as n increases; instead, the line $D = n\lambda$ (neglecting poly-log terms) corresponds to the best achievable trade-offs when the transmission range is constrained to be $O(1)$.

We observe that there are operating points of our system above the line $D = n\lambda^2$. This means that under our correlated mobility model it is possible to obtain significant better performance than that achievable under uncorrelated mobility, even without scaling up the transmission range.

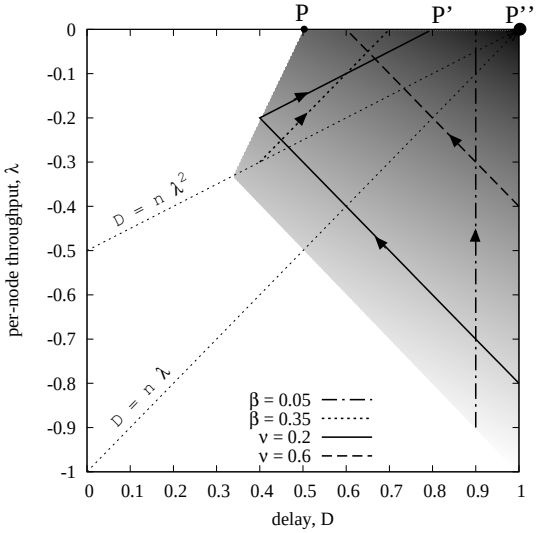

Fig. 4. Throughput-delay scaling exponents (The marks on the axes represent the orders asymptotically in n). Arrows on the trajectories correspond to increasing values of ν (for fixed β) or increasing values of β (for fixed ν).

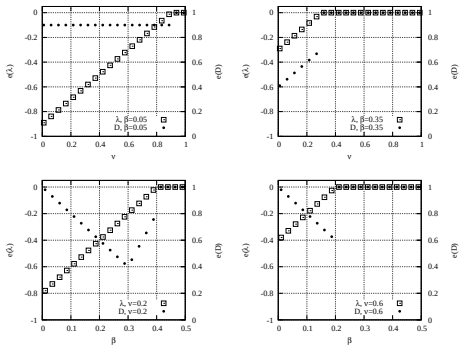

Fig. 5. Throughput and delay scaling exponents versus ν and β (The marks on the y-axes represent the orders asymptotically in n).

We have also reported on the plot of Fig. 4 a few trajectories obtained when we fix one of the parameters of the model (either β or ν), letting the other one vary. To help the reader, we separately report the throughput and delay scaling exponents curves in Fig. 5, considering the same values of ν and β for the trajectories shown in Fig. 4.

We observe a quite complex range of possible behavior. For large values of β and/or ν, the system operates in the *cluster dense* regime, with $\lambda = \Theta(1)$ and $D = \Theta(n)$ (point P'' in Fig. 4). As we increase the degree of correlation, by reducing either β or ν, at some point the system shifts to the *cluster sparse* regime.

For example, let us examine in detail the trajectory with $\nu = 0.2$ in Fig. 4 (and the corresponding curve in the bottom left plot of Fig. 5). If $\beta > 0.4$, the system is in *cluster dense* regime (point P''); as soon as β becomes smaller than 0.4 the system switches to the *cluster sparse* regime, jumping to the operating point P'. Note, indeed, that, while the throughput

scaling exponent exhibits a continuous behavior w.r.t to the parameters β and ν, the delay exponent is discontinuous at the transition between the two regimes (Fig. 5). As β decreases from 0.4 to 0.3, significant better delays are obtained paying a moderate penalty in terms of throughput. Below 0.3, the introduction of additional correlation (by further reducing β) in the mobility model leads to increasing delays (and decreasing throughput).

The overall best operating point, *i.e.*, the point characterized by the maximum throughput and minimum delay, can be approached when both $\beta \to 1/4$ and $\nu \to 1/2$ (point P in Fig. 4).

VI. Extensions under the cluster sparse regime

In this section we present and analyze several variations of our model, focusing on the *cluster sparse* regime. In Section VI-A we discuss what happens when nodes remain still within their clusters, *i.e.*, when they maintain their relative positions within the cluster-region indefinitely, starting from an initial configuration in which they are placed uniformly at random in the cluster-region. We refer to this extension as the *crystallized model*. In Section VI-B we show how results can be extended to more general mobility models than the simple i.i.d. reshuffling, for both the mobility of cluster centers and the mobility of individual nodes within their cluster. Finally, in Section VI-C we present a variation of the system in which nodes are allowed to migrate to a different cluster upon contact between clusters.

A. The crystallized model

We describe our scheduling-routing strategy for the *crystallized* model by adapting the scheme proposed in Section III. In particular, we replace the 2-hop replication technique previously adopted within the destination cluster (*i.e.*, 3rd and 4th hop in Fig.1) with a multi-hop communication similar to the one developed for static nodes by Gupta-Kumar [3]. Indeed, notice that in the *crystallized* model each cluster can be regarded as a micro-universe in which nodes are still (the relative positions of nodes within a cluster region are fixed).

We first analyze the achievable throughput and delay in the last two hops. Let λ_M and D_M be the throughput and the delay achievable by the multi-hop communication phase performed within the destination cluster. Applying standard results for a random network of q static nodes ([3], [7]), a maximum per-node throughput $\hat{\lambda}_M = \sqrt{1/(q \log q)}$ can be sustained within the destination cluster (as long as sources and destinations are chosen irrespective of their locations in the cluster area), using a transmission range $\hat{r}_M = R\sqrt{\log q/q}$, at the expense of a delay[6] $\hat{D}_M = \sqrt{q/\log q}$. Moreover, by increasing the transmission range it is possible to achieve capacity-delay trade-offs characterized by the law $D_M = \Theta(q\lambda_M)$ [7].

Similarly to the reshuffling model, the optimal choice is to match the throughput achievable within the destination cluster with that provided by the inter-cluster communications performed in the first two hops (computed in Corollary 1), *i.e.*, $\lambda_M = mR^2/n$.

Depending on the system parameters, two cases are possible: 1) the potential per-node throughput $\hat{\lambda}_M$ achievable in the last multi-hop phase exceeds (in order sense) the maximum data rate provided by the first two hops $\lambda = mR^2/n$. In this case, as for the reshuffling case, by properly selecting the transmission range within the destination cluster (*i.e.*, increasing it with respect to the minimum value) we reduce the throughput achievable within the destination cluster so as to match $\lambda = mR^2/n$. This does not affect the overall throughput, while significantly reducing the packet delay experienced during the last multi-hop phase; 2) the potential per-node throughput $\hat{\lambda}_M$ achievable in the last multi-hop phase is smaller (in order sense) than $\lambda = mR^2/n$. In this case the optimal scheme is a bit trickier, as it requires to modify also the forwarding strategy of the second hop. Indeed, notice that we can increase the throughput λ_M beyond $\hat{\lambda}_M$ by reducing the distances that messages have to traverse within the destination cluster (and thus the number of multiple hops to be performed within the destination cluster). This can be done in such a way that the resulting intra-cluster throughput perfectly matches the throughput achievable in the previous hops. To obtain this, we need to modify the forwarding rule of the second hop, forcing the relay node n_1 to send messages destined to d only to those nodes $n_2 \in C_d$ falling within a proper distance $R_M = o(R)$ from d.

Case 1) occurs for $\beta \leq (1-\nu)/4$ (condition under which $\hat{\lambda}_M = \Omega(mR^2/n)$). In this case, by properly selecting the transmission range we reduce the achievable throughput in the destination cluster so as to achieve $\lambda_M = mR^2/n$, and a corresponding delay $D_M = \Theta(R^2)$ (recall the trade-off law $D_M = \Theta(q\lambda_M)$). Since $D_2^{\alpha} = \Theta(n/R^2) = \omega(R^2)$ for the considered range of values of β, the overall end-to-end delay is dominated by the second hop, and we have $D = \Theta(n/R^2) = \Theta(m/\lambda)$.

Case 2) occurs for $(1-\nu)/4 < \beta < (1-\nu)/2$ (recall that we are in the *cluster sparse* regime, in which $\beta < (1-\nu)/2$). We achieve a throughput $\lambda_M = \lambda = mR^2/n$, by selecting[7] $R_M = \Theta\left(\sqrt{\frac{q}{R^2 \log q}}\right)$. We observe that this change in the forwarding rule of the second hop requires to modify also the internal architecture of the nodes, providing each node with $n-m$ different FIFO queues, one for each destination belonging to a different cluster, in which to store packets at hop 2. Indeed, in this way we can still guarantee (in saturated traffic conditions) that, in slots devoted to hop 2, whenever a node $n_1 \in C_a$ comes in proximity of a node $n_2 \in C_b$, with $C_a \neq C_b$, it can always find a packet at the head of one queue devoted to hop 2, whose corresponding destination d lies within C_b at a distance not greater than R_M from n_2. In this way nodes can exploit all contacts with other nodes belonging to a different cluster, hence no throughput reduction occurs at hop 2 due to the modified forwarding rule.

Turning our attention to the delay of this modified scheme, the access delay of the second hop is increased to $D_2^{\alpha} = \Theta(n/R_M^2)$, because a tagged packet can be forwarded only

[6]Recently, it has been shown that a throughput $\hat{\lambda}_M = \sqrt{1/q}$ and a delay $\hat{D}_M = \sqrt{q}$ can be achieved adopting a more sophisticated scheduling-routing scheme [29]. However, for our scopes we resort to the original scheme analyzed in [3], [7].

[7]To sustain a throughput $mR^2/n = \omega(1/\log n)$, it is necessary that relay n_1 delivers the packet directly to the final destination d.

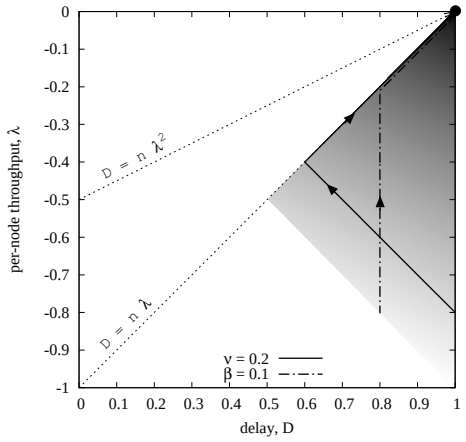

Fig. 6. Throughput-delay scaling for the *crystallized* model. The marks on the axes represent the orders asymptotically in n.

to those nodes within the destination cluster C_d which lie in a circle of radius R_M centered at the destination[8]. The delay component due to the multi-hop phase is instead equal to the number of hops $R_M/\hat{r}_M = \Theta\big(q/(R^2 \log q)\big)$. In the considered range of values for β, the end-to-end delay is always dominated by the second hop, hence $D = \Theta(mR^2 \log n) = \Theta(n\lambda \log n)$.

At last, we would like to emphasize that using similar arguments as for the reshuffling model, it can be proved that no scheme can achieve better delay performance in order sense, while guaranteeing the optimal throughput $\lambda = mR^2/n$.

Figure 6 provides a graphical representation of our results for the *crystallized* model, analogous to Figure 4 for the reshuffling model. Again, the shaded region denote all optimal operating points that are obtained as we vary the parameters β and ν of our model. Differently from the reshuffling model, we observe that all points in the *crystallized* model are below the line $D = n\lambda$, meaning that in this case it is not possible to achieve better performance with respect to the case of uncorrelated mobility.

B. Extension to more general mobility models

In this section we show how results in Section III can be generalized to the case in which the positions of cluster centers do not regenerate at each time slot. The analysis can be similarly extended to the case in which the positions of nodes within their cluster region do not regenerate at every slot.

In particular, we consider the following general class of random walks. Let $X(t)$ be the position of a cluster center (or the relative position of a node inside a cluster) at time slot t; $X(t)$ is updated according to the law $X(t) = X(t-1) + Y_t$, where Y_t is a sequence of i.i.d., rotationally invariant random vectors describing the individual movements accomplished by the moving entity during each slot. $Y(t)$ is often referred to as flight. Let moreover define the average length of the flights $L_f = E[|Y_t|]$. Random-walk mobility patterns are important since they have been widely used to describe human and vehicular mobility.

[8]The access delay D_2^α is $\Theta(n/\hat{r}_M^2)$ when packets have to be delivered directly to the final destination at hop 2.

While considering cluster centers trajectories, we further assume that flights cover distances at least of the order of the cluster radius R, *i.e.*, $L_f = \Omega(R)$. When we consider the relative position of a node within its cluster region, we instead assume that flights cover at least the typical transmission range employed within a cluster (for intra-cluster transmissions), *i.e.*, $L_f = \Omega(1)$. With these assumptions the movements accomplished by cluster centers (or individual nodes) during a single slot become appreciable.

First, we observe that our analysis of the system throughput does not require that cluster centers positions (or relative positions of nodes within a cluster area) regenerate at every slot. Indeed, our analysis only requires that the following two properties are satisfied at any time slot: 1) positions of cluster centers are independently and uniformly distributed over the area; 2) positions of individual nodes are independently and uniformly distributed within their cluster area. Thus the throughput results obtained in this paper can be immediately extended to any mobility model, including the considered class of random walks, that satisfies assumption 1) and 2). This is not surprising, in light of the fact that all throughput/capacity results obtained so far in the literature for mobile ad-hoc networks depend only on the steady-state spatial distribution of nodes over the area, and not on the specific details about how nodes move from one slot to another [2], [14], [16].

Instead, when we turn our attention to the delay, things become more involved and performance comes to depend on the specific mobility pattern ([4], [5], [6], [7]). However, following an approach similar to the one proposed in [10], we can easily evaluate the scaling laws of the delay performance achievable by our schemes under the considered class of random walks.

Focusing on the *cluster-sparse* regime (the analysis in the *cluster-dense* regime proceeds along the same lines) first recall that D_2^α in (1) (which dominates D_1^α) represents the average time for a node a belonging to cluster C_a to come in contact with a node belonging to the destination cluster $C_d \neq C_a$; moreover D_2^α can be tightly approximated (in order sense) by the average residual life time until clusters C_a and C_d meet (*i.e.*, their cluster regions overlap). This because node a has a non-negligible probability of being selected for transmission every time C_a and C_d meet. Meeting times between two clusters form a renewal process for a large class of mobility models including random-walks. Denoting with T_{ad} the inter-meeting time between C_a and C_d, standard renewal arguments provide the expression $D_2^\alpha = \frac{E[T_{ad}^2]}{2E^2[T_{ad}]}$.

Our analysis leverages results from [7] and [30]. By Theorem 3 and Proposition 1 in [30], the average inter-meeting time between clusters is invariant with respect to the considered mobility model. Hence, $E[T_{ad}] = \Theta(n/R^2)$. By Theorem 1 and Proposition 2 in [30], the second moment of the inter-meeting time increases for random walks as we decrease the flight lengths (*i.e.*, as we increase the degree of time correlation in the mobility pattern of the mobile). By Lemma 4 in [7], for random walks with minimum flight length $L_f^{\min} = \Theta(R)$, we have $E[T_{ad}^2] = \Theta(n \log n/R^2)$. In conclusion D_2^α is jointly $\Omega(n/R^2)$ and $O(n \log n/R^2)$ whenever cluster centers move according to random walks with flight length $\Omega(R)$.

Now we analyze the delay incurred by a message at the

4-th hop. First we focus on the average time $\tilde{D}_4^\alpha$ that it takes to a node belonging to the destination cluster, and holding a copy of the message, to hit the destination. Quantity $\tilde{D}_4^\alpha$ can be tightly approximated (in order sense) by the average residual life time until two specific nodes within a cluster meet (*i.e.*, their distance becomes smaller than $r = \Theta(1)$), and with arguments analogous to the ones used in the analysis of D_2^α we conclude that $\tilde{D}_4^\alpha$ is jointly $\Omega(R^2)$ and $O(R^2 \log n)$, provided that the flight size of nodes within their cluster is $\Omega(1)$.

Now observe that, after the third hop, there are $z = \Theta(\frac{n}{mR^2})$ different nodes belonging to the destination cluster, that hold a copy of the tagged message. When the first copy hits the destination, the message is delivered.

In the ideal case in which the nodes holding different copies of the message were initially uniformly (and independently) spread within the destination cluster, the time needed for the first copy to hit the destination would be $O(\tilde{D}_4^\alpha / z)$. This descends from the fact that the time taken by the first copy to hit the destination can be modeled as the minimum of z independent, identically distributed continuous random variables with mean $\tilde{D}_4^\alpha$ (and whose density support starts from 0). Unfortunately, the copies of the tagged message are initially not uniformly distributed in the cluster, since they all lie within the transmission range of the broadcasting node n_2. This reduces the effectiveness of the replication mechanism.

To account for this fact, we need to consider also the additional delay before the copies spread within the destination cluster, which is $\Theta(\frac{R^2}{L_f^2})$ (observe that $\frac{R^2}{L_f^2}$ can also be regarded as the time taken by one copy to cover a distance $\Theta(R)$, for which we can apply standard results of random walks in [31]). As a result we have:

$$D_4^\alpha = O\left(\frac{R^2}{L_f^2} + \frac{\tilde{D}_4^\alpha}{z}\right) = O\left(\frac{R^2}{L_f^2} + \frac{m\,R^4 \log n}{n}\right) \qquad (3)$$

At last, when considering the queueing delay, we can resort to arguments similar to those developed in [7]. Individual queues at nodes can be modeled as GI/GI/1-FCFS queues (recall that the whole system is an acyclic network of queues), in which both inter-arrival and inter-departure time distributions of packets match cluster/nodes inter-meeting time distributions. Then, applying Kingman's upper bound[9] to the average delay of a GI/GI/1-FCFS queue we can conclude that the delay experienced by a packet in the i-th hop is $D^i = \Theta(D_i^\alpha)$. We conclude that the packet delivery delay is given by the maximum between D_2^α and D_4^α (which dominates both D_1^α and D_3^α) and we finally obtain:

$$D = O\left(\max\left\{\frac{n}{R^2}, \frac{R^2}{L_f^2}, \frac{m\,R^4 \log n}{n}\right\}\right). \qquad (4)$$

Figure 7 provides a graphical representation of our results on the throughput-delay plane, for the case $L_f = \Theta(1)$. The shaded region denotes all the operating points of our scheme that are obtained as we vary the parameters β and ν of the model. Comparing with the shaded region in Figure 4, whose perimeter is shown in Figure 7 by a thick solid line, we observe that some operating points obtained under the reshuffling model are no longer feasible under the random walk mobility model, yet we can still go above the line $D = n\lambda^2$ which characterizes delay-throughput trade-offs in the absence of correlated movements.

[9]Kingman's upper bound states that the average delay in a GI/GI/1-FCFS queue is $O(\frac{E[a^2]+E[d^2]}{E[a]})$ where $E[a]$ and $E[a^2]$ are, respectively, the first and second moment of the inter-arrival time, whereas $E[d^2]$ is the second moment of the inter-departure time.

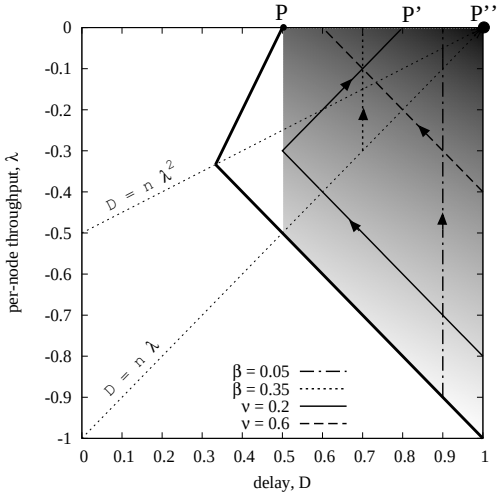

Fig. 7. Throughput-delay scaling for the extension to the Random walk mobility model, in the case of flight length $L_f = \Theta(1)$.

C. Node migrations among clusters

In the class of networks considered so far, one limitation is related to the fact that we assume the nodes to remain indefinitely associated with the same cluster. In realistic, correlated mobility processes, it could happen that nodes decide to follow different groups of nodes with the passing of time, switching from one cluster to another.

In this section we present a variation of our model in which nodes are allowed to migrate from one cluster to another when two clusters come in contact which each other (*i.e.*, when their cluster regions overlap). To make the model more general, we assume that clusters are partitioned into $M = n^\gamma$ super-clusters, with $0 \leq \gamma \leq \nu$, each containing $K = m/M = n^{\nu-\gamma}$ clusters analogous to the ones defined in our original model. When two clusters belonging to the same super-cluster overlap in space, their nodes independently migrate from one cluster to the other with some fixed constant probability. As γ varies, we obtain a wide gamut of systems, comprising as special cases our original model without nodes' migrations (for $\gamma = \nu$), and the other extreme case (for $\gamma = 0$) in which all nodes can migrate to any cluster in the network. In the following analysis, we assume that nodes know the identities of all other nodes in their super-cluster (such information does not change over time, thus it can be distributed once and for all at system start up). Instead, we assume for simplicity that nodes do not know the identity of the other nodes currently associated to the same cluster (this information changes over time).

For this more general class of networks, the first goal of our performance optimization (*i.e.*, throughput maximization) is very simple to analyze. Indeed, for any $0 < \gamma \leq \nu$, the majority of flows (*i.e.*, all flows established between nodes belonging to different super-clusters) still require at least

one inter-cluster communication, hence we can directly apply Theorem 1 and conclude that in this case the network throughput is necessarily $O(mR^2)$. The above bound is actually achievable by a scheduling scheme similar to the one proposed in the absence of node migrations, still based on a four hop scheduling/routing strategy. The only algorithmic difference with respect to the previous case is that now, in the second hop, we allow a message to be transferred to the first contacted node belonging to the super-cluster of the destination (we recall that nodes ignore the identity of the cluster in which the destination is). Observe also that different copies of a message generated after the third hop, will migrate from one cluster to another (as nodes holding them migrate) within the destination super-cluster. The message will be then delivered when the first copy hits the destination. The throughput analysis of the above scheme can be easily carried out along the same lines as before. It turns out that the achievable network throughput is $\Theta(mR^2)$, as in the non-migration case.

The special case $\gamma = 0$ leads instead to the condition in which nodes uniformly come in contact with any other node in the network, hence in this case we can apply the classic Grossglauser-Tse 2-hop scheme and conclude that we can achieve per-node throughput[10] $\lambda = \Theta(1)$. Notice that there is a sharp discontinuity in the throughput when the value of γ steps away from zero.

As far as the delay performance with $\gamma > 0$, the delays D_1^α and D_3^α for the first and third hop of the new scheme are exactly the same as in the previous scheme. The delay of the second hop, instead, reduces by a factor K, since now a message can be transferred into any one of the K clusters belonging to the super-cluster of the destination, instead of just into the cluster of the destination. By the same arguments in Appendix C, we get $D_2^\alpha = \Theta\left(\frac{n}{KR^2}\right)$ which dominates both D_1^α and D_3^α.

The delay of the fourth hop is slightly more difficult to analyze. Recall that the message gets replicated in the third hop, generating a number of copies that are initially confined within a unique cluster belonging to the super-cluster of the destination. With the passing of time, these copies distribute themselves among the clusters of their super-cluster, as in an epidemic process, as a result of nodes' migrations. Note that the average inter-contact time between two clusters belonging to the same super-cluster equals D_2^α, the delay of the second hop. Moreover, it can be shown that the spreading phase of the copies over the super-cluster of the destination takes at most a duration $\Theta(\log n \cdot D_2^\alpha)$ before we can assume that replicas are uniformly distributed in the super-cluster. At this point, all replicas are equally likely to fall in communication range of the destination node, thus the additional delay before the first one of them hits the destination is $\Theta(\frac{mKR^4}{n})$, by arguments analogous to those in Appendix C. Considering (pessimistically) that the spreading phase must always be completed before the message can reach the destination (such pessimistic assumption introduces at most a $\log n$ factor in the overall delay), we obtain that $D_4^\alpha = O\left(\log n\, D_2^\alpha + \frac{mKR^4}{n}\right)$. In conclusion:

[10]Moreover, it can be shown that in this case the delivery delay is $D = \Theta(n)$ for any $\beta > 0$.

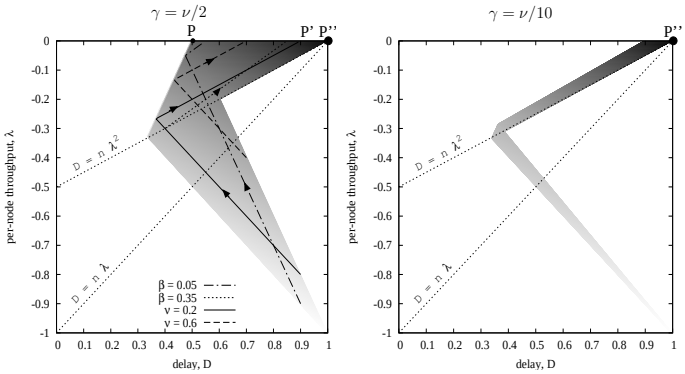

Fig. 8. Throughput-delay scaling for the extension to node migrations among clusters, for $\gamma = \nu/2$ (left plot) and $\gamma = \nu/10$ (right plot).

Proposition 1: In the case of nodes' migrations among clusters, restricted to occur among $K = n^{\nu-\gamma}$ clusters, with $0 < \gamma \leq \nu$, it is possible to achieve, under the cluster sparse regime, a maximum per-node throughput $\lambda = \Theta(mR^2/n)$, at the expense of an end-to-end delay

$$D = O\left(\log n \frac{n}{KR^2} + \frac{mKR^4}{n}\right) \tag{5}$$

Fig. 8 shows the scaling exponents of per-node throughput and delay in the case of $\gamma = \nu/2$ (left plot) and $\gamma = \nu/10$ (right plot). We observe that, for decreasing values of γ (i.e., smaller number of super-clusters), the region of feasible operating points shrinks, yet we still have points above the $D = n\lambda^2$ line.

VII. Conclusions

Correlated nodes movements have huge impact on the throughput and delay performance of mobile ad hoc networks. In this paper we have provided a first characterization of the scaling laws of networks with correlated node mobility, devising novel scheduling-routing schemes which maximize the per-node throughput as primary goal. Being the first analysis of this kind, we have considered a simplified group mobility model, yet flexible enough to explore various degrees of correlation in the nodes mobility process. Our study reveals the existence of a wide range of correlated node movements which can lead to significant better performance than the one achievable under independent nodes movements.

References

[1] Delay Tolerant Network Research Group, www.dtnrg.org
[2] M. Grossglauser, D.N.C. Tse, "Mobility increases the capacity of ad hoc wireless networks", *IEEE/ACM Trans. on Networking*, vol. 10, no. 2, August 2002.
[3] P. Gupta, P.R. Kumar, "The capacity of wireless networks", *IEEE Trans. on Information Theory*, 46(2), pp. 388–404, 2000.
[4] M. J. Neely, E. Modiano "Capacity and delay trade offs for ad-hoc mobile networks" *IEEE Trans. on Inf. Theory*, 51(6), pp. 1917–1937, 2005.
[5] S. Toumpis, A. J. Goldsmith, "Large Wireless Networks under Fading, Mobility, and Delay Constraints," *in Proc. INFOCOM '04*.
[6] X. Lin, N. B. Shroff, "The Fundamental Capacity-Delay Tradeoff in Large Mobile Ad Hoc Networks," *in Proc. MedHoc'04*.
[7] A. El Gamal, J. Mammen, B. Prabhakar and D. Shah, "Throughput and Delay in wireless networks – Part I: Fluid case," Special joint issue of *IEEE Transaction on Information Theory* and *IEEE/ACM Transaction on Networking*, Vol. 52, No. 6, page. 2568- 2592, 2006.

[8] N. Bansal, Z. Liu, "Capacity, Delay and Mobility in Wireless Ad-Hoc Networks," *in Proc. IEEE INFOCOM '03*.

[9] G. Sharma, R. R. Mazumdar and N. B. Shroff, "Delay and Capacity Trade-offs in Mobile Ad Hoc Networks: A Global Perspective" *in Proc. IEEE INFOCOM '06*.

[10] U. Lee, S. Y. Oh, K. W. Lee, M. Gerla "Scaling properties of delay tolerant networks with correlated motion patterns", *ACM CHANTS '09*.

[11] J. D. Herdtner, E. K. P. Chong, "Throughput-storage tradeoff in ad hoc networks", *in Proc. INFOCOM '05*.

[12] A. Ozgur, O. Leveque, "Throughput-Delay Trade offs for Hierarchical Cooperations", IEEE Trans. on Information Theory, vol. 56, no. 3, pp. 1389-1277, Mar. 2010.

[13] S.N. Diggavi, M. Grossglauser, D.N.C. Tse, "Even one-dimensional mobility increases ad hoc wireless capacity", IEEE Trans. on Information" Theory, vol. 51, no. 11, pp. 3947-3954, Nov. 2005.

[14] M. Garetto, P. Giaccone, E. Leonardi, "Capacity Scaling in Delay Tolerant Networks with Heterogeneous Mobile Nodes" *in Proc. ACM MobiHoc '07*.

[15] J. Mammen, D. Shah, "Throughput and Delay in Random Wireless Networks With Restricted Mobility," *IEEE Trans. on Inf. Theory*, 53(3), pp. 1108-1116, 2007.

[16] M. Garetto, P. Giaccone, E. Leonardi, "Capacity Scaling of Sparse Mobile Ad Hoc Networks", *in Proc. INFOCOM '08*.

[17] J. Yoon, B. D. Noble, M. Liu, M. Kim, "Building realistic mobility models from coarse-grained traces", *in Proc. MobiSys '06*.

[18] V. Naumov, R. Baumann, T. Gross, "An evaluation of inter-vehicle ad hoc networks based on realistic vehicular traces," *in Proc. MobiHoc '06*.

[19] M. Musolesi, C. Mascolo, "Designing mobility models based on social network theory," *Mob. Comput. Commun. Rev.*, 11(3), pp. 59–70, 2007.

[20] P. Hui, J. Crowcroft, "Human mobility models and opportunistic communications system design," *Phil. Trans. R. Soc. A*, no. 366, pp. 2005-2016, 2008.

[21] X. Hong, M. Gerla, G. Pei, and C. Chiang, "A group mobility model for ad hoc wireless networks", *in Proc. ACM MSWiM '99*.

[22] F. Bai, N. Sadagopan, A. Helmy, "The IMPORTANT framework for analyzing the Impact of Mobility on Performance Of RouTing protocols for Adhoc NeTworks", *Ad Hoc Networks*, vol. 1, pp. 383–403, 2003.

[23] G. Pei, M. Gerla, X. Hong, "LANMAR: landmark routing for large scale wireless ad hoc networks with group mobility", *in Proc. MobiHoc '00*.

[24] M. Thomas, S. Phand, A. Gupta, "Using group structures for efficient routing in delay tolerant networks", *Ad Hoc Networks*, 7(2), pp. 344–362, 2009.

[25] A. Lindgren, A. Doria, O. Schelén, "Probabilistic Routing in Intermittently Connected Networks, *LNCS*, Vol. 3126, Springer, 2004.

[26] K. Blakely, B. Lowekamp, "A structured group mobility model for the simulation of mobile ad hoc networks", *in Proc. MobiWac '04*.

[27] H. Takagi, Queueing Analysis : Discrete-Time Systems, North Holland, 1993.

[28] M. Penrose, Random Geometric Graphs, Oxford University Press, 2003.

[29] M. Franceschetti, O. Dousse, D. N. C. Tse, and P. Thiran, "Closing the gap in the capacity of wireless networks via percolation theory," *IEEE Trans. Inf. Theory*, vol. 53, no. 3, pp. 1009-1018, Mar. 2007.

[30] H. Cai, D.Y. Eun, "Toward stochastic anatomy of inter-meeting time distribution under general mobility models", *in Proc. MobiHoc '08*.

[31] D. Aldous and J. Fill, "Reversible Markov Chains and Random Walks on Graph," Monograph in preparation, available at http://www.stat.berkeley.edu/users/aldous/RWG/book.html.

APPENDIX A
PROOF OF LEMMA 1

Recall that the protocol model requires the adoption of the same transmission range for all communications occurring in the same slot. Therefore, let r be the common transmission range by all transmitters during the considered slot. The proof consists in showing that the number of simultaneous transmissions between nodes belonging to different clusters is necessarily $O(mR^2)$, for any value of r and under any possible scheduling policy.

First, we recall a basic result proved in [3], which states that, according to the protocol model, the distance between any pair of transmitting nodes cannot be smaller than Δr where r is the transmission range and Δ the guard factor defined in Section II-B. This implies that, considering discs of radius $\frac{\Delta r}{2}$ around simultaneously transmitting nodes, all discs are necessarily disjoint.

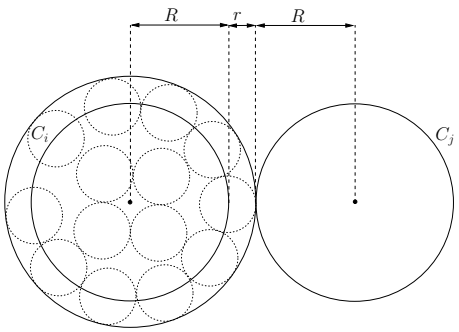

Fig. 9. Geometry used to upper bound the number of simultaneous transmissions among nodes belonging to different clusters

Focusing on a particular cluster C_i, we observe that any transmitting node belonging to cluster C_i lies within a distance R from the cluster center; then all discs associated to transmitting nodes in C_i are constrained to lie within a circle of radius $R + r$ centered at the cluster center.

From the above argument we can obtain an upper bound H_1 to the number of transmitting nodes belonging to C_i, by bounding the maximum number of discs of radius $\Delta r/2$ that can fit within a circle of radius $r + R$. A simple upper bound to this number is $\pi(R + r)^2 / \left(\pi \frac{\Delta^2 r^2}{4}\right)$. By noticing that, in any case, the number of transmissions originated by nodes in C_i can not exceed the number of nodes q in C_i, we can refine the previous upper-bound as:

$$H_1 = \min\left\{q, \frac{\pi(R+r)^2}{\pi\frac{\Delta^2 r^2}{4}}\right\} \tag{6}$$

Next we observe that any transmission between nodes in C_i and nodes belonging to a different cluster requires that at least one center of a cluster $C_j \neq C_i$ falls within a distance $2R + r$ from the center of cluster C_i (see Figure 9). This event occurs with probability

$$H_2 = 1 - \left(1 - \frac{\pi(2R+r)^2}{n}\right)^{m-1} \tag{7}$$

and optimistically we can say that, upon its occurrence, all potential transmitters in C_i can find a receiver belonging to a different cluster using transmission range r.

At last, summing over all clusters we obtain an upper bound H to the number of simultaneous transmissions among nodes belonging to different clusters that any policy employing a transmission range r can achieve $H = mH_1H_2$:

$$H = m \min\left\{q, \frac{\pi(R+r)^2}{\pi\frac{\Delta^2 r^2}{4}}\right\}\left[1 - \left(1 - \frac{\pi(2R+r)^2}{n}\right)^{m-1}\right]$$

It remains to show that the scaling order of the above quantity H is $O(mR^2)$, or equivalently, that $H_1H_2 = O(R^2)$. To this purpose, we separately analyze the scaling order of H_1 and H_2, whose behavior is qualitatively represented in Figure 10 as function of r.

We observe that H_1 saturates to its maximum theoretical value $q = \frac{n}{m}$ when we select a transmission range $r =$

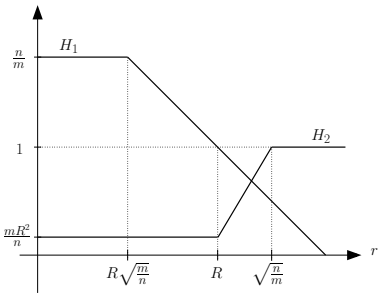

Fig. 10. The (qualitative) scaling behavior of terms H_1 and H_2 as function of r

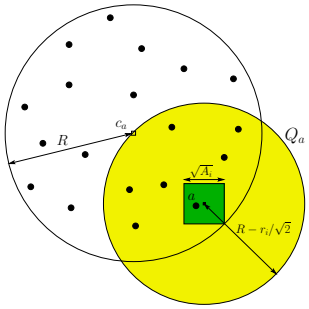

Fig. 11. The shaded disk denotes the region where a cluster center must fall so that the selected squarelet is completely covered by the cluster containing transmitting node a.

$O\left(R\sqrt{\frac{m}{n}}\right)$, which allows all nodes in a cluster to (potentially) transmit concurrently without interfering with each other. For $r = \omega\left(R\sqrt{\frac{m}{n}}\right)$, interference among concurrent transmitters progressively reduces the number of feasible transmitters within a cluster, as we further increase r.

On the contrary, H_2 is an increasing function of r. For $r = O(R)$, we can neglect r with respect to R, hence the scaling behavior of H_2 is insensitive to the selected transmission range. Furthermore, since in the cluster sparse regime $m\frac{R^2}{n} = o(1)$, we obtain:

$$H_2 \sim 1 - \left(1 - \frac{\pi(2R)^2}{n}\right)^m \sim m\frac{\pi(2R)^2}{n} = \Theta\left(\frac{mR^2}{n}\right)$$

The same approximation can be applied also when $r = \omega(R)$ and $r = O(\sqrt{\frac{n}{m}})$, obtaining the scaling order $\Theta\left(\frac{mr^2}{n}\right)$, which increases with r. For $r = \omega(\sqrt{\frac{n}{m}})$, the scaling behavior of H_2 saturates to $\Theta(1)$, since in this case the transmission range is large enough to cover the typical distance $\sqrt{\frac{n}{m}}$ between cluster centers, allowing a transmitter in a cluster to find a receiver belonging to a different cluster with finite probability (at the cost of being, typically, the unique transmitter enabled in its cluster).

At last, it is easy to see that the product H_1H_2 attains its maximum when $r = O\left(R\sqrt{\frac{m}{n}}\right)$, for which $H_1H_2 = \Theta(R^2)$. We conclude that, in the *cluster sparse* regime, the maximum amount of data that can be transferred in one slot among nodes belonging to different clusters is $O(mR^2)$.

Appendix B
Proof of Lemma 2

Let r_i be the transmission range used in a generic slot i, devoted to inter-cluster communications among arbitrary nodes belonging to different clusters. Without lack of generality in light of Theorem 1, we assume $r_i = O(\sqrt{R\frac{m}{n}})$.

We first observe that, according to the scheduling scheme illustrated in Figure 3, at most one communication can be enabled in each square of area d_i^2. Hence we can express the average number $\mathbb{E}[N_i]$ of packets that can be transmitted over the entire network during a slot devoted to inter-cluster communications as

$$\mathbb{E}[N_i] = \frac{n}{d_i^2}\mathbb{P}(\text{active squarelet} \mid i) \tag{8}$$

where $\frac{n}{d_i^2}$ is the number of squarelets and $\mathbb{P}(\text{active squarelet} \mid i)$ is the probability that for a generic squarelet $\mathcal{A}_i^k$ we can find: i) a transmitting node a residing in the considered squarelet; ii) a receiving node b at distance at most r_i from a (given the location of a), and belonging to a different cluster than the one of a. Let $\mathbb{P}(a|i)$ and $\mathbb{P}(b|a,i)$ denote the occurrence probability of the two events above, respectively. Since the positions of nodes belonging to different clusters are independent, we have $\mathbb{P}(\text{active squarelet} \mid i) = \mathbb{P}(a|i)\mathbb{P}(b|a,i)$.

Probability $\mathbb{P}(a|i)$ can be approximated by the joint occurrence of the following two events: i) $\mathcal{A}_i^k$ is entirely covered by one cluster[11]; ii) given the occurrence of event i), at least one node belonging to the covering cluster is found in $\mathcal{A}_i^k$. Condition i) above occurs when at least one cluster center falls within a disk of radius $R - r_i/\sqrt{2}$ around the squarelet (see disk Q_a in Figure 11). While, condition ii) above occurs with probability $1 - \left(1 - \frac{r_i^2}{R^2}\right)^q$. It follows,

$$\mathbb{P}(a|i) = \left[1 - \left(1 - \frac{(R - r_i/\sqrt{2})^2}{n}\right)^m\right]\left[1 - \left(1 - \frac{r_i^2}{R^2}\right)^q\right]$$

The expression above can be approximated as $\mathbb{P}(a|i) \sim qr_i^2$ given that we assume $qr_i^2 = O(R^2)$. Observe that $\mathbb{P}(a|i)$ increases linearly with r_i^2 and for $r_i = \sqrt{\frac{m}{n}R^2}$ it reaches the saturation value $\Theta(1)$.

Probability $\mathbb{P}(b|a,i)$ can be approximated[12] by the joint occurrence of the following two events: i) the disk of radius r_i centered at a is entirely covered by a cluster different from the one of a; ii) given the occurrence of event i), at least one node belonging to the covering cluster is found in $\mathcal{A}_i^k$. Condition i) above occurs when at least one out of $m-1$ cluster centers falls within a disk of radius $R - r_i$ centered at a (see disk Q_b in Figure 12). Instead, condition ii) above occurs with probability $1 - \left(1 - \frac{r_i^2}{R^2}\right)^q$. We obtain

$$\mathbb{P}(b|a,i) = \left[1 - \left(1 - \frac{(R - r_i)^2}{n}\right)^{m-1}\right] \cdot \cdot \left[1 - \left(1 - \frac{r_i^2}{R^2}\right)^q\right] \tag{9}$$

[11]A transmitter could be found in $\mathcal{A}_i^k$ even if the squarelet were partially covered by a cluster. This approximation does not affect the results, in order sense, as one can show by considering $\mathcal{A}_i^k$ entirely covered by a cluster even if just a corner of it is touched by the cluster.

[12]The approximation does not affect the result, in order sense, for reasons analogous to our approximation of $\mathbb{P}(a|i)$.

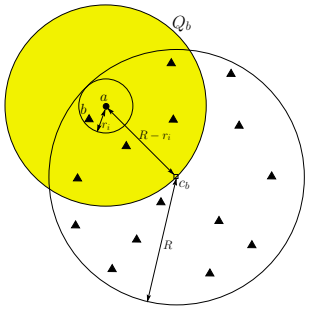

Fig. 12. The shaded disk denotes the region where a cluster center must fall so that the disk of radius r_i around transmitter a is completely covered by the cluster containing receiver b.

The expression above can be again approximated as $\mathbb{P}(b|a,i) \sim r_i^2$ since $qr_i^2 = O(R^2)$. Observe that also $\mathbb{P}(b|a,i)$ increases linearly with r_i^2 and for $r_i = \sqrt{\frac{m}{n}R^2}$ it reaches the saturation value $\Theta(1)$. Given that $d_i^2 = \Theta(r_i^2)$, putting things together we obtain:

$$\mathbb{E}[N_i] = nr_i^2 \qquad \text{for} \quad r_i = O(R\sqrt{m/n}) \tag{10}$$

From (10) we conclude that it is possible to achieve $\mathbb{E}[N_i] = mR^2$ by selecting $r_i = R\sqrt{m/n}$. This corresponds to using a transmission range which is strictly related to the density of nodes within clusters, which is equal to $n/(mR^2)$. In particular, with this choice r_i is equal to the typical distance between nodes belonging to the same cluster.

Appendix C
Computation of the access delays

Since we are interested to an order sense evaluation of the access delay of each hop, we can ignore all factors whose effect on the access delay can be bounded by a multiplicative constant, such as: i) the fact that only one slot out of four is devoted to transmission of packets at a given hop; ii) only a subset of squarelets can be activated in a given slot.

In the first hop, the tagged packet has to wait until the source node s gets in contact with a node n_1 belonging to an arbitrary different cluster C_r (Figure 1). In a slot devoted to hop 1, the probability $\mathbb{P}(n_1|s,1)$ that a generic node n_1 belonging to a cluster C_r different from C_s gets in contact with s (*i.e.*, lies at distance at most r_1 from s) is analogous to quantity $\mathbb{P}(b|a,i)$ in (9) (see Figure 12). Since we use $r_1 = R\sqrt{m/n}$, we have $\mathbb{P}(n_1|s,1) \sim R^2 m/n$.

The packet access delay at s, expressed in number of slots, follows a geometric distribution Geom($\mathbb{P}(n_1|s,1)$), since the positions of all nodes regenerate from slot to slot. Hence the average access delay at the first hop is: $D_1^\alpha = \Theta\left(\frac{n}{mR^2}\right)$.

The access delay in the second hop is similar to the one of the first hop, however in this case n_1 can transmit the tagged packet only when it gets in contact with a node n_2 belonging to the *specific* cluster containing the destination. In a slot devoted to hop 2, the probability $\mathbb{P}(n_2|n_1,2)$ that n_1 gets in contact with a generic node n_2 belonging to cluster C_d can be computed as:

$$\mathbb{P}(n_2|n_1,2) = \frac{(R-r_2)^2}{n}\left[1-\left(1-\frac{r_2^2}{R^2}\right)^q\right]$$

Since we use $r_2 = R\sqrt{m/n}$, we have $\mathbb{P}(n_2|n_1,2) \sim R^2/n$. Again, the packet access delay at n_1, expressed in number of slots, follows a geometric distribution Geom($\mathbb{P}(n_2|n_1,2)$), thus: $D_2^\alpha = \Theta\left(\frac{n}{R^2}\right)$. Note that D_2^α always dominate D_1^α.

To compute D_3^α and D_4^α we can apply standard results [7], [9], [6], [5] obtained for the i.i.d mobility model, since relative movements of nodes within each cluster area are i.i.d.

In particular, we have $D_3^\alpha = \Theta(1)$, since n_2 can broadcast the packet in any slot devoted to the third hop without any other requirement.

After the third hop, a number $\Theta(\frac{n}{mR^2})$ of nodes within the destination cluster hold a copy of the tagged packet, hence D_4^α corresponds to the average time that it takes before the first one of these nodes arrive at distance $r_4 = \Theta(1)$ from the destination. In a slot devoted to hop 4, the probability $\mathbb{P}(n_3|d,4)$ that at least one node holding a copy of the tagged packet falls within transmission range r_4 from d is given by

$$\mathbb{P}(n_3|d,4) = 1-\left(1-\frac{r_4^2}{R^2}\right)^{\frac{n}{mR^2}} = \Theta\left(\frac{n}{mR^4}\right)$$

Since the access delay of the fourth hop follows a geometric distribution Geom($\mathbb{P}(n_3|d,4)$), we have: $D_4^\alpha = \Theta\left(\frac{mR^4}{n}\right)$.

Delia Ciullo received the Master degree in Telecommunication Engineering and the Ph.D. degree in Electronics and Communication Engineering, both from Politecnico di Torino in 2007 and 2011, respectively. In 2009, she has been a visiting student at the CNRG group of MIT, under the supervision of Prof. Eytan Modiano. She is currently a post-doc at Politecnico di Torino. Her research interests are in the fields of energy-aware networks, scaling properties in wireless networks, and P2P systems.

Valentina Martina received the Master degree in Mathematical modeling in Engineering and the Ph.D. degree in Electronics and Communication Engineering, both from Politecnico di Torino in 2007 and 2011, respectively. In 2010, she has been a visiting student at the Technicolor Paris Research Lab. She is currently a post-doc at Politecnico di Torino. Her research interests are in the fields of scaling properties in Wireless Networks, mobility models, and P2P systems.

Michele Garetto (M'04) received the Dr.Ing. degree in Telecommunication Engineering and the Ph.D. degree in Electronic and Telecommunication Engineering, both from Politecnico di Torino, Italy, in 2000 and 2004, respectively. In 2002, he was a visiting scholar with the Networks Group of the University of Massachusetts, Amherst, and in 2004 he held a postdoctoral position at the ECE department of Rice University, Houston. He is currently assistant professor at the University of Torino, Italy.

Emilio Leonardi (M'99, SM'09) is an Associate Professor at the Dipartimento di Elettronica of Politecnico di Torino. He received a Dr.Ing degree in Electronics Engineering in 1991 and a Ph.D. in Telecommunications Engineering in 1995 both from Politecnico di Torino. In 1995, he visited the Computer Science Department of UCLA, Los Angeles; in 1999 he joined the High Speed Networks Research Group, at Bell Laboratories/Lucent Technologies, NJ; in 2001, the Electrical Engineering Department of the Stanford University, and finally in 2003, the IP Group at Sprint, Advanced Technologies Laboratories, CA. His research interests are in the field of performance evaluation of wireless networks, P2P systems, packet switching.