On the Capacity Region of MANET: Scheduling and Routing Strategy

Michele Garetto † Paolo Giaccone * Emilio Leonardi *
† Dipartimento di Informatica, Università di Torino, Italy
* Dipartimento di Elettronica, Politecnico di Torino, Italy

***Abstract*— We characterize the capacity region of a mobile ad hoc network in which nodes employ the store-carry-forward communication scheme and move according to an arbitrary ergodic mobility process. We identify the class of scheduling policies achieving maximum throughput, and introduce a joint scheduling and routing formulation which maps the problem into a multi-commodity flow over an associated *contact graph*. Previous capacity results have been derived under the strong assumption that nodes are identical and uniformly visit the entire network area, resulting in a fully connected, homogeneous contact graph in which a simple 2-hops routing scheme is optimal. Our approach allows to extend the analysis to heterogeneous nodes with anisotropic mobility patterns, as typically encountered in realistic mobility traces. In particular, we apply our framework to an experimental network based on vehicular mobility and show that, in scenarios with inhomogeneous contact times, the 2-hops routing strategy can result significantly inefficient in terms of throughput and delay.**

***Index Terms*— Delay Tolerant Networking, Routing, Network Capacity.**

I. INTRODUCTION

In recent years Delay Tolerant Networking (DTN) has emerged as a new area of research with many promising applications [3], [4]. Delay tolerant networks are characterized by intermittent connectivity, network partitioning, long and variable delays, high error rates. Such performance-challenging conditions can be found in many different environments such as vehicular networks, sparse sensor/actuator networks, in-the-field military or disaster-relief networks, satellite and deep-space interplanetary communications [5], terrestrial networks serving remote or rural areas. Interesting real-life experiments comprise "pocket switched networks" based on human mobility [6], networks based on public transportation systems [7]-[9], wildlife tracking [10], rural kiosks providing Internet access in developing nations [11].

A typical DTN scenario consists of a sparse network of fixed or mobile devices, where most of the time there does not exist a complete path from a source to a destination, or such a path is highly unstable and may soon break. Over time, different links come up and down due to node mobility. This implies that a message can be sent over an existing link and be buffered at the receiving node for quite a long time before being forwarded on the next available link toward the destination. While being stored at a mobile node, a message is physically carried to a different location in the network area. This new communication paradigm is usually referred to as *store-carry-and-forward* or *mobility-assisted routing*. It is particularly interesting in the context of vehicular networks where a fixed communication infrastructure is not available and the cooperation among the vehicles is needed to support the communication.

Node mobility plays a fundamental role not only in providing end-to-end connectivity in the first place, but also because it can dramatically increase the overall transport capacity of interference-limited wireless networks, allowing them to scale up to a large number of nodes. Indeed, wireless networks of static nodes are known to suffer severe per-node throughput decay (in the order of $1/\sqrt{n}$) as the number of nodes n goes to infinity [12]. In contrast, Grossglauser and Tse have first shown in [13] that a simple 2-hops relay scheme can keep the per-node throughput constant, under an ideal mobility process in which each node independently and uniformly visits the entire network space.

While the work of Grossglauser and Tse suggests that delay tolerant networks can indeed scale up to large sizes, and that the 2-hops routing strategy is able to fully exploit the transport capacity offered by node mobility, it is unclear whether these results are applicable to real-life scenarios such as the ones mentioned above. One reason is that the analysis in [13] strongly relies on the assumptions that: i) all nodes are identical; ii) each node uniformly visits the entire network area according to an ergodic mobility process; iii) the trajectories of different nodes are independent and identically distributed (i.i.d.). Notice that these same assumptions have been maintained in the several papers that, after [13], have analyzed asymptotic delay-capacity trade-offs, like [14]-[16].

In many practical settings, the above assumptions on the nodes behavior do not hold, and in particular the one that the mobility process of each node uniformly covers the entire space over time, making all nodes basically indistinguishable from each other. This observation is especially true for vehicular networks, in which the mobility is restricted along the roads. In addition, both everyday life experience and campus- or city-wide traces containing spatial information (i.e., based on GPS coordinates or radio beacons from base stations and access points) [17]-[20], suggest that a node spends most of the time just in a small portion of the network area, comprising a few frequently visited "significant places" [21], and rarely goes outside this region. Although any two nodes are likely, in the long run, to eventually come in contact with each other, the impact of rare contacts on the overall network capacity

Preliminary versions of this paper have been presented at IEEE Infocom'07 [1] and IEEE ICC'07 [2].

has to be carefully investigated.

In this paper, we relax the "homogeneous mixing" assumption on the node behavior and extend the capacity analysis to heterogeneous nodes with general mobility patterns. First, we provide a general framework within which the capacity analysis of mobile ad hoc networks can be carried out. More specifically, we only require the mobility processes of the nodes to be jointly stationary and ergodic, though possibly correlated (e.g., group movements, which often occur in vehicular mobility), as defined in Sec. IV. Under this fairly general assumption, in Sec. V we formally prove that the class of scheduling policies based only on instantaneous node position information achieves the maximum network throughput. This means that the capacity region of a mobile network depends on the mobility process only through the joint-stationary distribution of the nodes, and not on the details about how nodes change their speed and direction over time. Moreover, no gain in terms of throughput can be obtained by scheduling policies using dynamical variables such as instantaneous queues lengths, age of stored information at the nodes, history of nodes behavior and past encounters [1].

Armed with this result, we map the problem of determining the transport capacity of a mobile network into a joint scheduling and routing problem which only requires knowledge of the stationary distribution of the nodes over the network area. The capacity region is then obtained by solving a multi-commodity flow over the *contact graph* defined by the scheduling policy and described in Sec. VI.

We have applied our framework to a specific example of experimental vehicular network, although its validity is more general and applies to any delay tolerant network. In the considered scenario we have observed that the contact graph is almost fully connected, however link capacities are highly inhomogeneous, spanning several orders of magnitude. This fact motivated us to investigate the performance of the 2-hops routing scheme, originally proposed under the "homogeneous mixing" assumption, in a more realistic case. Our theoretical analysis (presented in Sec. VII for the case of finite number of nodes and Sec. VIII for the case of infinite number of nodes) suggests that the 2-hops scheme can incur significant throughput losses in the presence of highly asymmetric contact graphs. Therefore, smart routing algorithms discovering links with high available capacity and possibly routing messages over paths of several hops are indeed necessary to exploit the network capacity under a realistic mobility process. In contrast, we emphasize that most of the routing protocols proposed so far for DTN are oblivious of link capacities, and aim at discovering minimum-hop paths or minimizing the instantaneous traffic congestion at the nodes (as in [22]).

In Sec. IX we compare the performance of different routing strategies in the considered vehicular network, in terms of both throughput and delays; the routing algorithms that we have compared are the following: (i) the throughput optimal algorithm, (ii) the 2-hop throughput optimal algorithm, and (iii) the delay optimal algorithm, which was previously proposed in [23].

Finally, to complete the overview of the paper, in Sec. II we discuss related work and in Sec. III we introduce our notation and discuss the main assumptions of our work. We provide some conclusions in Sec. X.

We emphasize that our results have mainly a theoretical significance, since we ignore several practical issues, such as those related to the implementability of our scheduling/routing schemes, to congestion control, etc.; nevertheless our results provide useful guidelines for the design of efficient routing/scheduling strategies.

[1]Notice that in our work we are only concerned with network throughput, not with message transfer delay; moreover we assume infinite buffer size at the nodes.

II. Related work

In the literature related to DTN, routing is the main issue that has been addressed so far. Current proposals to route messages in a DTN strongly depend on the nature of contacts among the nodes (i.e., random or predictable), the amount of information available to them, and the possibility to disseminate or not multiple copies of a message. If meeting times of nodes are predictable, intelligent routing and forwarding decisions can be made, although optimal schemes (i.e., with complete information) that account also for finite buffers lead to complex linear programming formulations [23]. In [24], the authors assume that nodes movements are not only known in advance, but can be controlled to further improve data collection and delivery across the network. More recently, [25] focuses still on periodic movements, but exploits multilevel clustering and hierarchical routing, based on minimum delay paths, to improve the scalability of the routing algorithms.

When mobility is random, nodes have to communicate during opportunistic, unscheduled contacts. Epidemic routing [26] is one of the first proposals to enable message delivery in intermittently connected networks with random mobility. Each node maintains a list of messages, whose delivery is pending. Whenever it encounters another node, the two nodes exchange all messages that they do not have in common. This way, all messages are eventually disseminated to all nodes, including their destination. Although optimal in terms of minimizing delay, epidemic routing is very wasteful of network resources.

To avoid flooding the entire network, messages can be duplicated in a limited number of copies [27], [28], or probabilistically [29], possibly exploiting history of encounters [30], [31] for a better selection of next-hops, combined with various utility functions [32]. Other approaches make use of erasure coding techniques [33] or network coding [34], [35] to cope with partial data loss and to reduce routing overhead. Practical routing issues are discussed in [36].

All of the above approaches aim either at minimizing the message transfer delay or at maximizing the delivery probability, and they result, in general, to be inefficient in terms of bandwidth exploitation, especially when employing information duplication/redundancy. In contrast to previous work, in this paper we study scheduling/routing schemes that aim at maximizing the overall network throughput in the presence of heterogeneous nodes with general mobility. To the best of our knowledge, only a few results have appeared in the literature on this subject. In [37] the authors extend the results in [13], proving the asymptotic optimality of the 2-hops forwarding scheme in a network where each node independently moves along a randomly chosen great circle on the sphere of unit surface. The work in [38] studies the maximum rate at which one or more mobile nodes can relay data between a set of static sources and a set of static destinations. Under some simplifying assumptions they prove that the relay throughput depends only on the stationary distribution of the nodes' position. In particular, they assume that nodes move independently of each other and have the same stationary distribution over the area (not necessarily uniform). In this paper we extend the analysis to a more general context.

III. System assumptions and notation

We consider a mobile ad hoc network composed of n nodes moving according to a general mobility model inside a bidimensional, compact and convex region $\mathcal{A}$ of area $|\mathcal{A}|$. $X_i(t)$ denotes the position of node i at time t and $\mathbf{X}(t) = (X_1(t), X_2(t) \dots X_n(t))$ the vector of nodes positions; we define with $d_{ij}(t)$ the euclidean distance between mobile i and mobile j at time t, i.e., $d_{ij}(t) = ||X_i(t) - X_j(t)||_2$.

We assume the node mobility process to be stationary and ergodic; i.e., given any m-uple $(B_1, B_2, B_3 \dots B_m)$ of Lebesgue measurable subsets of A, it results:

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \mathbb{I}_{(\cap_i X_i(\tau)\in B_i)} d\tau = E[\mathbb{I}_{(\cap_i X_i(t)\in B_i)}] \quad \text{w.p.1} \qquad (1)$$

where $\mathbb{I}$ represents the logical indicator function. Note that stationarity and ergodicity are rather general assumptions. Most of the mobility models proposed in the literature, such as those described and discussed in [39]-[42], fits within the class considered in this paper.

Node s generates traffic for destination d according to a stationary and ergodic process with average traffic rate λ_{sd} bit/s[2]. We denote by $\mathbf{\Lambda} = [\lambda_{sd}]$ the corresponding $n \times n$ traffic matrix.

We assume that interference between simultaneous transmissions is described either by the *protocol interference model* or the *physical interference model* [12]. According to the protocol interference model, transmission from node i to node j at time t at rate r is successful only if, for any other simultaneously transmitting node k, it holds:

$$d_{kj}(t) > (1+\Delta) d_{ij}(t) \qquad (2)$$

for some guard factor $\Delta > 0$. According to the physical interference model, instead, transmission from node i to node j at time t at rate r is successful only if the associated Signal to Noise Ratio (SNR) is above a given threshold z. When signal power decays with the distance d as $d^{-\alpha}$, the SNR ratio is expressed by:

$$\frac{P_i d_{ij}^{-\alpha}(t)}{N_0 + \sum_{k\neq i} P_k d_{kj}^{-\alpha}(t)} \qquad (3)$$

being P_k the power level of transmitter k, and N_0 the thermal background noise [3]. Note that according to these interference models: (i) no node can be either origin or destination of multiple simultaneous transmissions, (ii) a node cannot be simultaneously origin and destination of transmissions.

We denote with E the set of all possible transmitter-receiver pairs (i, j) (by construction it must be $i \neq j$). Subsets π of E in which nodes appear at most once (either as transmitter or receiver), represent possible *transmission configurations*, i.e., sets of transmission-receiver pairs (i, j), which may be simultaneously enabled to communicate at time t. We denote by Π the set of all the possible transmission configurations and with $A(t) \subseteq \Pi$ the set of all non-interfering (hence, implementable) transmission configurations at time t. A given interference model induces a correspondence between the vector of instantaneous nodes positions $\mathbf{X}(t)$ and the set of non-interfering transmission configurations $A(t)$; we formalize this concept introducing function $\mathcal{I}$ mapping vectors of nodes positions into sets of non-interfering transmission configurations: $\mathcal{I}(\mathbf{X}(t)) = A(t)$.

Given any set A of implementable transmission configurations, we denote with $\mathcal{I}^{-1}(A)$ the set of node positions $\mathbf{X}$ to which A corresponds through mapping $\mathcal{I}$, i.e., $\mathcal{I}^{-1}(A) = \{\mathbf{X} : \mathcal{I}(\mathbf{X}) = A\}$. [4] For any A, we can univocally determine the probability that $A(t) = A$, i.e. the probability that configurations in A are the only implementable at time t:

$$P(A) = E[\mathbb{I}_{X(t)\in\mathcal{I}^{-1}(A)}] = \lim_{t\to\infty} \frac{1}{t} \int_0^t \mathbb{I}_{X(\tau)\in\mathcal{I}^{-1}(A)} d\tau \qquad (4)$$

w.p.1. Note that the above probability depends only on the joint stationary distribution of the node mobility process.

At last, we denote with $\{T_n\}$ the sequence of random instants at which the set of implementable transmission configurations changes; i.e., $\lim_{t\uparrow T_n} A(t) \neq A(T_n)$.

IV. Scheduling policy

The scheduling policy S dynamically selects an implementable transmission configuration $\pi^S(t)$ belonging to $A(t) = \mathcal{I}(\mathbf{X}(t))$. In this paper we restrict our investigation to stationary and ergodic scheduling policies: i.e. those policies for which:

$$E[\mathbb{I}_{(i,j)\in\pi^S(t)}] = \lim_{t\to\infty} \frac{1}{t} \int_0^t \mathbb{I}_{(i,j)\in\pi^S(\tau)} d\tau \qquad \text{w.p.1} \qquad (5)$$

In general the selection of $\pi^S(t)$ may be influenced by several dynamical parameters, including instantaneous queues lengths, age of stored information at nodes, etc. Particularly relevant are those scheduling policies driven only by $\mathbf{X}(t)$. In this paper we call *stateless and memoryless* such scheduling policies.

We also introduce the class of *simple* scheduling policies $\hat{S}$, which is a strict subclass of the *stateless and memoryless* scheduling policies characterized as follows. At each transition time T_n a transmission configuration $\pi \in A(T_n)$ is selected

[2]Defined with $\hat{\lambda}(t,\tau)$ the amount of data generated by a source within the interval $[t,\tau)$, the traffic is said stationary and ergodic with average rate λ iff: $E[\hat{\lambda}(t,t+1)] = \lambda$ for any $t > 0$ and $\lim_{t\to\infty} \hat{\lambda}(0,t)/t = \lambda$ w.p.1.

[3]More complex expressions for the SNR are obtained when considering propagation phenomena such as fading and/or shadowing, the impact of directional antennas, MIMO systems, etc

[4]It can be verified that, for any $A \in \Pi$, $\mathcal{I}^{-1}(A)$ is a convex set of $\mathbb{R}^{2n}$ for both the protocol and the physical interference models.

according to a stationary and memoryless (possibly random) rule; the selected transmission configuration is then kept constant in the whole interval $[T_n, T_{n+1})$. Simple scheduling policies are fully specified by the conditional probabilities $p_{\hat{S}}(\pi, A)$ that the transmission configurations $\pi \in A$ are selected at time T_n, given that $X(T_n) \in \mathcal{I}^{-1}(A)$:

$$p_{\hat{S}}(\pi, A) = \Pr\{\pi^{\hat{S}}(T_n) = \pi | X(T_n) \in \mathcal{I}^{-1}(A)\} \qquad (6)$$

$\forall A$ and $\pi \in A$. According to scheduling policy S, a communication link is established between nodes i and j whose average capacity expressed in bit/s is:

$$\mu_{ij}^S = rE[\mathbb{I}_{(i,j)\in\pi^S(t)}] = \lim_{t\to\infty} \frac{r}{t} \int_0^t \mathbb{I}_{(i,j)\in\pi^S(\tau)} d\tau \text{ w.p.1} \quad (7)$$

which, in case of simple scheduling policies, can be rewritten as:

$$\mu_{ij}^{\hat{S}} = r \sum_{A\in\Pi} \sum_{\pi\in A} \mathbb{I}_{(i,j)\in\pi}\, p_{\hat{S}}(\pi, A) P(A) \qquad (8)$$

The fundamental question we would like to answer is: "how can we characterize the capacity of the mobile ad hoc network under a scheduling policy S (or $\hat{S}$)?" To this end we need to consider also the routing strategy employed to transfer data through the network. The more general and abstract way to define a routing strategy is to specify quantities $f_{ij}^{sd} \in [0, 1]$ denoting the average fraction of the traffic from node s to node d, which is routed trough link (i, j), i.e. j follows i as relay node [43] (unless $j = d$); $f_{ii}^{sd} = 0$ by construction. The above quantities f_{ij}^{sd} must satisfy the following well known flow conservation constraints:

$$\sum_i f_{ij}^{sd} - \sum_k f_{jk}^{sd} = \begin{cases} 1 & \text{for } j = d \\ 0 & \text{for } j \neq d \text{ and } j \neq s \\ -1 & \text{for } j = s \end{cases} \qquad (9)$$

A routing strategy specified by a set of f_{ij}^{sd} satisfying (9) can be implemented in many ways, for example by the following hop-by-hop randomized routing algorithm $\mathcal{R}$: node i routes data from source s and destined to d by selecting node j as next hop with probability $f_{ij}^{sd} / \sum_k f_{ik}^{sd}$.

V. Traffic sustainability and capacity region

In this subsection we analyze the performance of a mobile ad hoc network comprising n users, obtaining a precise characterization of its capacity region[5]. We emphasize that our results are fairly general since only stationarity and ergodicity of traffic and mobility processes are required. We remark that, in our framework, nodes can move in a correlated fashion (like in platoons, groups, etc.), since we do not need to assume independent movements.

Definition 1: We denote with $Z(t)$ the network backlog, that is, the amount of traffic (in bits) already generated by sources that has not yet been delivered to destinations at time t.

Definition 2: Traffic $\boldsymbol{\Lambda}$ is *sustainable* if there exists a scheduling policy S and a routing strategy $\mathcal{R}$, such that: $\limsup_{t\to\infty} Z(t)/t = 0$ w.p.1.

Definition 3: Traffic $\boldsymbol{\Lambda}$ is *strongly sustainable* if there exists a *simple* scheduling policy $\hat{S}$ and a routing strategy $\mathcal{R}$, such that: $\limsup_{t\to\infty} Z(t)/t = 0$ w.p.1.

We are now in a position to state our first result:

Theorem 1: A mobile ad hoc network sustains a traffic $\boldsymbol{\Lambda}$, if a scheduling policy S and a routing strategy $\mathcal{R}$ can be found such that:

$$\sum_{sd} \lambda_{sd} f_{ij}^{sd} \leq \mu_{ij}^S \qquad \forall i, j \qquad (10)$$

Moreover, if a simple scheduling policy $\hat{S}$ and a routing strategy $\mathcal{R}$ can be found such that:

$$\sum_{sd} \lambda_{sd} f_{ij}^{sd} \leq \mu_{ij}^{\hat{S}} \qquad \forall i, j \qquad (11)$$

the mobile ad hoc network strongly sustains traffic $\boldsymbol{\Lambda}$.

Proof: The system dynamics can be described by a network of queues representing the evolution of the backlog at different nodes. We suppose that every node i is equipped with $n - 1$ separate transmission queues, each one storing data to be routed through a different node j. Upon reception, new data are immediately routed according to policy $\mathcal{R}$ and enqueued in the transmission queue associated to the next hop. Transmission queues are served at fixed rate r according to a FCFS service policy, during the periods of activity of the corresponding link (i, j). Note that, by construction, the average service rate in bit/s of the transmission queue of link (i, j) is μ_{ij}^S. The network of queues describing the system falls in the class of generalized Kelly networks, which are stable under the condition that no queues are overloaded [44]. Being, by construction, the load at the queue of link (i, j) equal to $\sum_{sd} \lambda_{sd} f_{ij}^{sd} / \mu_{ij}^S \leq 1$, the assert follows immediately [6]. ■

As a corollary, we get a strict characterization of the traffic matrices which are strongly sustainable:

Proposition 1: A traffic matrix $\boldsymbol{\Lambda} = [\lambda_{sd}]$ is strongly sustainable iff a set of $f_{ij}^{sd} \in [0, 1]$, $\forall i, j, s, d$ and $p(\pi, A) \in [0, 1]$ $\forall A, \forall \pi \in A$ can be found satisfying the following equations:

$$\begin{cases} f_{ij}^{sd} \text{ satisfies (9)} & \forall i, j, s, d \\ \sum_{\pi\in A} p(\pi, A) = 1 & \forall A \in \Pi \\ \sum_{sd} \lambda_{sd} f_{ij}^{sd} \leq r \sum_{A\in\Pi} \sum_{\pi\in A} \mathbb{I}_{(i,j)\in\pi}\, p(\pi, A) P(A) & \forall i, j \end{cases} \qquad (12)$$

Definition 4: The *capacity region* of the mobile ad hoc network is the set of all sustainable traffic matrices.

Definition 5: The *restricted capacity region* of the mobile ad hoc network is the set of all strongly sustainable traffic matrices.

Note that the restricted capacity region, by construction, depends on nodes mobility only via the joint stationary distribution of nodes. We can now state the following fundamental result:

Theorem 2: If traffic matrix $\boldsymbol{\Lambda}$ is sustainable, then it is also strongly sustainable.

[5] We consider only the net traffic (goodput), neglecting the effect of the signaling traffic necessary to implement the routing and scheduling algorithms

[6] When $\sum_{sd} \lambda_{sd} f_{ij}^{sd} / \mu_{ij}^S = 1$ the concept of stability is weak.

Proof: Let S be the stationary and ergodic scheduling policy which sustains $\mathbf{\Lambda}$. Define for every configuration A and every $\pi \in A$:

$$q_S(\pi, A) = \Pr\{\pi^S(t) = \pi | X(t) \in \mathcal{I}^{-1}(A)\} = \lim_{t\to\infty} \frac{1}{t} \int_0^t \mathbb{I}_{\pi^S(\tau)=\pi|X(\tau)\in\mathcal{I}^{-1}(A)} d\tau \quad \text{w.p.1} \quad (13)$$

Due to the ergodicity of the mobility process and of the scheduling policy, the above quantities are well defined. It is immediate to verify that: $\sum_{\pi\in A} q_S(\pi, A) = 1$. Thus considering a stationary, simple scheduling policy $\hat{S}$ such that $p_{\hat{S}}(\pi, A) = q_S(\pi, A)$, it follows, by construction: $\mu_{ij}^S = \mu_{ij}^{\hat{S}}$ $\forall i, j$. ∎

The previous result has three significant implications: i) the class of simple policies achieves maximum throughput, i.e., no gain in terms of throughput can be obtained by adopting complex scheduling policies that select transmission configurations by considering dynamical variables such as instantaneous queues lengths, age of stored information at the nodes, etc., provided that the capacities associated to the communication links are known in advance; ii) a tight characterization of the capacity region is provided by Proposition 1; iii) the capacity region depends on the mobility process only through the joint-stationary distributions of nodes. This result extends and generalizes recent findings in [38]. At last,

Corollary 1: The capacity region of an ad hoc wireless network with mobile nodes is convex.

Proof: Let $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$ be two sustainable traffic matrices. Let $\hat{S}_1$ and $\hat{S}_2$ be two simple scheduling policies which sustain $\mathbf{\Lambda}_1$ and $\mathbf{\Lambda}_2$ respectively. Any traffic pattern $\mathbf{\Lambda} = \alpha\mathbf{\Lambda}_1 + (1-\alpha)\mathbf{\Lambda}_2$, with $0 \le \alpha \le 1$, is sustainable by the simple policy $\hat{S}$ defined according to: $p_{\hat{S}}(\pi, A) = \alpha p_{\hat{S}_1}(\pi, A) + (1-\alpha)p_{\hat{S}_2}(\pi, A)$, $\forall A, \pi \in A$. Note that $\hat{S}$, at time t, with probability α emulates $\hat{S}_1$ and with probability $(1-\alpha)$ emulates $\hat{S}_2$. ∎

VI. CONTACT GRAPH: THROUGHPUT AND ROUTING

To understand the relationship between the scheduling policy and the routing strategy, we first need to characterize which traffic patterns are sustainable by employing an assigned scheduling policy S. Observe that the capacities[7] μ_{ij} associated to the communication links are univocally determined, once the scheduling policy S has been selected. Thus a (capacitated) graph $G(\mathcal{V}, \mathcal{E})$ whose vertices correspond to network nodes and capacitated edges correspond to communication links, fully characterizes the mobile ad hoc network adopting S. In the following we refer to $G(\mathcal{V}, \mathcal{E})$ with the term *contact graph*. Therefore, the routing problem through the mobile ad hoc network adopting S can be formalized in terms of a *multi-commodity flow* problem on the contact graph. Fig. 1 shows an example of contact graph construction. For simplicity, in the figure we have defined $c_{ij}^S(t) = \mathbb{I}_{(i,j)\in\pi^S(t)}$, while service rates are evaluated according to (7), i.e. by summing (or averaging) the capacity of all the transmission opportunities between the corresponding nodes.

[7]To simplify the notation we omit, in this section, the explicit dependency from the scheduling policy S.

Proposition 2: A traffic matrix $\mathbf{\Lambda} = [\lambda_{sd}]$ can be sustained employing a policy S iff the multi-commodity flow problem defined by (9) and (10), where communication link capacities are determined by S, admits a feasible solution. In such a case the set of variables f_{ij}^{sd} univocally defines the routing strategy $\mathcal{R}$.

The sustainable region achievable by scheduling policy S (i.e. the set of $\mathbf{\Lambda} = [\lambda_{sd}]$ that can be sustained employing S) can be related to the capacities associated to cuts of the contact graph as follows:

Proposition 3: Traffic $\mathbf{\Lambda} = [\lambda_{sd}]$ is sustainable by policy S only if, for any partition $(D, \bar{D})$ of the nodes, it results:

$$\sum_{s\in\bar{D}} \sum_{d\in D} \lambda_{sd} \le \sum_{s\in\bar{D}} \sum_{d\in D} \mu_{sd} \quad (14)$$

Consider a network adopting a scheduling policy S and a routing strategy $\mathcal{R}$, under a sustainable traffic pattern $\mathbf{\Lambda}$. Let ν be the network throughput, equal, by definition, to the offered network load $\nu = \sum_{sd} \lambda_{sd}$, and $|\pi(t)|$ be the size of $\pi(t)$ (i.e., the number of parallel transmissions enabled by S at time t); it results:

$$E[|\pi(t)|] = E\Big[\sum_{ij} \mathbb{I}_{(i,j)\in\pi(t)}\Big] = \sum_{ij} \frac{\mu_{ij}}{r} \quad (15)$$

Let C be the average aggregate transmission rate over all of the links; by construction:

$$C = \sum_{sd} \lambda_{sd} \sum_{ij} f_{ij}^{sd} \le rE[|\pi(t)|] \quad (16)$$

The ratio $h_{ave} = C/\nu$ represents the average number of times that data are transmitted in the network; thus h_{ave} is the average length of the paths followed by information flows, expressed in number of hops. The following relationship is of immediate verification: $\nu h_{ave} \le rE[|\pi(t))|]$.

In general, to efficiently exploit the network bandwidth, the routing strategy should minimize h_{ave}. This consideration justifies the fact that shortest path routing approaches have been widely used in several application contexts related to computer communications. In the context of mobile ad hoc networks, the 2-hops routing strategy proposed in [13] has gained a wide popularity; according to this strategy, data are delivered from source s to destination d either through the direct communication link, or through routes $s \to k \to d$, using every other node k of the network as relay. Although the 2-hops routing strategy is appealing because of its simplicity, in general it does not allow to optimally exploit the network bandwidth, possibly causing a reduction of the sustainability region achievable by the scheduling policy, as we show in the following section.

VII. PERFORMANCE OF THE 2-HOPS ROUTING STRATEGY

Necessary and sufficient conditions for traffic Λ to be sustainable under a 2-hops routing strategy are provided by the following statement:

Proposition 4: A traffic pattern $\Lambda = [\lambda_{sd}]$ is sustainable by the 2-hops routing scheme if the multi-commodity flow problem on the contact graph defined by (9) and (10) with the

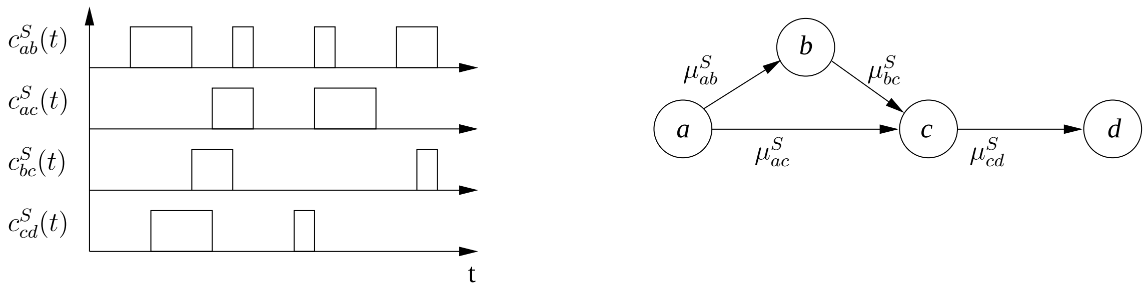

Fig. 1. Contact graph construction: on the left the temporal sequence of contacts among node pairs, and on the right the corresponding capacitated contact graph.

extra constraint: $f_{ij}^{sd} = 0$ when both $i \neq s$ and $j \neq d$ admits a feasible solution.
The set of all sustainable traffic matrices under the 2-hops routing strategy represents the *2-hops capacity region* of the mobile ad hoc network. In the following we are interested to better understand the relationship between the *capacity region* and *2-hops capacity region* on realistic contact graphs associated to mobile ad hoc networks.

To simplify the discussion, in analogy to what has been done in previous works [13], [12], [37], we assume that every node in the network is source and destination of a single information flow (i.e. Λ contains only one non-null element per row and per column). In the following we refer to the above class of traffic patterns with the term "permutation traffic patterns".

Of course, the performance of the 2-hops routing strategy can be arbitrarily bad on a general contact graph. For example, considering graphs with maximum nodal degree d, the average distance between nodes is (using standard Landau asymptotic notation) $\Omega(\log_d(n))$. Thus, for sufficiently large n, if traffic is exchanged only between node pairs whose distance on the contact graph is greater than two, the throughput sustained by the 2-hops strategy turns out to be equal to zero. As we will see in the next section, contact graphs associated to mobile ad hoc networks are typically strongly connected (i.e., $d = \Theta(n)$), and often fully connected (all edges exist). In the latter cases the contact graph diameter is expected to be very short. On this regard we recall that condition $d = \Theta(n)$ has been proved to be sufficient to asymptotically guarantee a diameter equal to two [8] in several families of random graphs, such as Erdős-Rényi graphs and graphs with regular degree [45]. Nevertheless, we will show that, even when either $d = \Theta(n)$ or $d = n - 1$, 2-hops routing may result strongly inefficient, provided that capacities associated to the edges of the contact graph are highly unbalanced.

Under a permutation traffic pattern, let us focus on a particular traffic relation (s, d) and consider all paths from node s to node d followed by data according to a 2-hops routing strategy. The aggregate maximum transmission capacity obtainable by flow $s \to d$ on all parallel 1-hop and 2-hops paths of the graph is given by:

$$\mu_{sd}^{2h} = \mu_{sd} + \sum_{k \neq s,d} \min(\mu_{sk}, \mu_{kd}) \qquad (17)$$

[8]the above property holds with a probability that tends to 1 when $n \to \infty$

Moreover, the direct link from s to d is used exclusively by flow (s, d). Link (s, k) instead must be shared by flow $s \to d$ and the flow generated by $s' \neq s, k$ and directed to k, whereas link (k, d) is shared by flow $s \to d$ and the flow generated by k and destined to $d' \neq k, d$. Hence, the capacity of each edge can be shared by at most two flows.

Sufficient conditions for traffic sustainability can be easily derived under the 2-hops routing strategy, assuming that bandwidth of links is evenly partitioned among the flows.

Proposition 5: Under the 2-hops routing strategy, a permutation traffic pattern is guaranteed to be sustainable if, for every source and destination pair (s, d), the following constraint is met:

$$\lambda_{sd} \leq \mu_{sd} + \frac{1}{2} \sum_{k \neq s,d} \min(\mu_{sk}, \mu_{kd}) \qquad (18)$$

It is interesting to understand under which conditions the 2-hops routing strategy is guaranteed to efficiently exploit the system bandwidth. The following theorem provides an answer to this question.

Theorem 3: If for every pairs of nodes (s, d), μ_{sd}^{2h} equals the minimum-cut capacity between (s, d), then the 2-hops routing scheme is $1/2$-efficient in terms of throughput, i.e., if $\Lambda = [\lambda_{sd}]$ cannot be sustained by the 2-hops routing strategy, no other routing strategy can sustain traffic $2\Lambda = [2\lambda_{sd}]$.

We remind that, given any partition of graph vertices (S, D) with $s \in S$ and $d \in D$, the capacity $\mu_{(S,D)}$ associated to the cut (S, D) is defined as: $\mu_{(S,D)} = \sum_{i \in S} \sum_{j \in D} \mu_{ij}$. The min-cut capacity $\hat{\mu}_{sd}$ between (s, d) is defined as $\min_{(S,D)} \mu_{(S,D)}$.

Proof: Consider any source destination pair (s, d). Under a general routing strategy the amount of traffic λ_{sd} that can be sustained by the network cannot exceed the min-cut between s and d on the contact graph. On the other side the 2-hops routing strategy sustains, between s and d, a traffic $\lambda_{sd} \geq \frac{1}{2}\mu_{sd}^{2h}$ according to Proposition 5. The assert follows immediately, since $\mu_{sd}^{2h} = \hat{\mu}_{sd}$ for all s and d. ■

Since $\sum_k \mu_{sk}$ and $\sum_k \mu_{kd}$ are two cut capacities, the condition:

$$\mu_{sd}^{2h} = \min\Big(\sum_{k \neq d} \mu_{sk}, \sum_{k \neq s} \mu_{kd}\Big) \qquad (19)$$

is sufficient to guarantee that at least half of the optimal throughput is delivered by the 2-hops routing strategy, under any permutation traffic pattern.

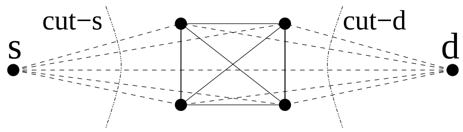

Fig. 2. Contact graph for $n = 6$.

Note that, when nodes move uniformly over the network area $\mathcal{A}$, by symmetry all link capacities μ_{ij} are equal, thus $\mu^{2h}_{sd} = \hat{\mu}_{sd}$ for all s and d. Furthermore in this specific case the 2-hops strategy can be easily shown to be 1-efficient, under uniform permutation traffic patterns (i.e., when all nodes generate the same amount of traffic). In this case, indeed, the adoption of the 2-hops routing scheme permits to completely saturate the capacity of every link, while minimizing the expected number of hops. For a general contact graph $G(\mathcal{V}, \mathcal{E})$, define

$$\eta^{2h} = \min_{sd} \frac{\mu^{2h}_{sd}}{\hat{\mu}_{sd}} \tag{20}$$

The statement of Theorem 3 can be generalized to the case $\eta^{2h} < 1$:

Theorem 4: Provided that $\eta^{2h} = \alpha < 1$, the 2-hops scheme is $\alpha/2$-efficient in terms of throughput, i.e., if $\Lambda = [\lambda_{sd}]$ cannot be sustained by the 2-hops scheme, no other routing strategy can sustain traffic $\frac{2}{\alpha}\Lambda = [\frac{2}{\alpha}\lambda_{sd}]$.

Furthermore there exists a permutation traffic pattern Λ which is sustainable by a properly defined routing scheme and such that $(\alpha + \epsilon)\Lambda$, for any $\epsilon > 0$, is not sustainable by the 2-hops strategy.

Proof: The first statement can be proved using exactly the same arguments of Theorem 3.

For what concerns the second statement, consider the source-destination pair (s, d) in correspondence to which $\mu^{2h}_{sd}/\hat{\mu}_{sd} = \alpha$. If only (s, d) exchanges traffic in the network, the maximum amount of traffic that can be sustained in the network by an arbitrary routing strategy equals the min-cut $\hat{\mu}_{sd}$; on the other hand, according to the 2-hops scheme the maximum amount of traffic that can be sustained cannot exceed μ^{2h}_{sd}. ■

In conclusion, if $\eta^{2h} = \alpha$ we should be prepared to get a throughput reduction of order $2/\alpha$ in case the 2-hops routing strategy is adopted.

VIII. Some asymptotic considerations on the throughput loss of the 2-hops routing strategy

In this section we focus on fully connected contact graphs with inhomogeneous capacities, such as those resulting from experimental traces. We assume that edge capacities μ_{ij} are i.i.d. random variables with assigned distribution. In this case we are able to show analytically that, for $n \to \infty$, the 2-hops routing strategy can be strongly inefficient, and that the throughput loss depends on the particular edge capacity distribution.

Essentially, we need to evaluate parameter η^{2h} introduced in (20). Due to the symmetry of the contact graph, we consider a randomly chosen source-destination pair (s, d) and estimate the ratio between the aggregate capacity of 2-hops paths and the minimum capacity among all cuts between s and d, as n grows large. In Fig. 2 the contact graph is represented for $n = 6$.

As n increases, to keep finite the capacities of the cuts to be evaluated, edge capacities are scaled by n, i.e., their distribution is given by $f^n_\mu(x) = ng(nx)$, being $g(x)$ an assigned distribution function with average $\hat{\mu} = \int xg(x)dx$. As a consequence, the average capacity of each edge is $E[\mu_{ij}] = \hat{\mu}/n$. Note that the scaling effect of edge capacities is intrinsically obtained when contact capacities are computed according to (7) (see [13]). Given two independent random variables X and Y distributed according to $g(x)$, we denote with $\mu_{\min} = E[\min(X, Y)]$.

Observe in Fig. 2 that the cut corresponding to node partition (S, D) with $|S| = k$ and $D = n - k$, being $k = 1, \ldots, n-1$, contains $k(n-k)$ edges; thus the cut around the source ("cut-s") (i.e, the cut for which $S = \{s\}$) and the cut around the destination ("cut-d") (i.e., $D = \{d\}$), are minimal in terms of number of edges traversing them $(n - 1)$; this suggests that they are minimal also in terms of capacity.

Indeed, under mild assumptions, standard concentration arguments allow to confirm this intuition for large n. In particular we can claim that with high probability (w.h.p.) (i.e., with probability that tends to 1 when $n \to \infty$) the minimum-capacity cut corresponds to either "cut-s" or "cut-d". Hence, the maximum achievable flow from s to d corresponds to the minimum between cut-s and cut-d:

$$\hat{\lambda}_{sd} = \min\left(\sum_{k \neq s} \mu_{sk}, \sum_{k \neq d} \mu_{kd}\right) \to \hat{\mu} \qquad \text{w.h.p.} \tag{21}$$

when $n \to \infty$. If we allow only 2-hops routing, the maximum achievable throughput λ^{2h}_{sd} is given by the aggregate capacity on all $n-1$ parallel 1-hop and 2-hops paths between s and d:

$$\lambda^{2h}_{sd} = \mu_{sd} + \sum_{k \neq s,d} \min\{\mu_{sk}, \mu_{kd}\} \to \mu_{\min} \qquad \text{w.h.p.} \tag{22}$$

when $n \to \infty$. Hence, asymptotically for $n \to \infty$, η^{2h} is given by:

$$\eta^{2h} = \min_{sd} \frac{\lambda^{2h}_{sd}}{\hat{\lambda}_{sd}} = \frac{\mu_{\min}}{\hat{\mu}} \qquad \text{w.h.p.} \tag{23}$$

It is interesting to compute $\mu_{\min}/\hat{\mu}$ for specific distributions of edge capacities. We consider two cases, both simple to evaluate analytically:

- $g(x)$ is an exponential distribution with mean a (i.e., $g(x) = e^{-x/a}/a$, for $x \geq 0$); then $\min\{\mu_{sk}, \mu_{kd}\}$ is still exponential with mean $a/2$: $\eta^{2h} \leq 1/2$.
- $g(x)$ is a Pareto distribution with parameters a and b (i.e., $g(x) = ba^b/x^{b+1}$, for $x \geq a$ and $b > 1$), having mean $ab/(b-1)$; then $\min\{\mu_{sk}, \mu_{kd}\}$ is still Pareto with parameters a and $2b$: $\eta^{2h} \leq 2(b-1)/(2b-1)$. Hence, η^{2h} is upper bounded by an increasing function starting from 0 for $b \to 1$ and growing to 1 only for $b \to \infty$. Note that $b \in (1, 2]$ corresponds to capacities with finite mean but infinite variance.

Hence, for exponentially distributed capacities, under a generic permutation traffic pattern we can expect 2-hops routing to be slightly inefficient, leading to throughput loss within a factor

4 (by Theorem 4). For Pareto-distributed capacities, instead, 2-hops routing can be very inefficient leading to unbounded throughput loss when $b \to 1$.

In conclusion, in case of inhomogeneous edge capacities, paths of several hops are needed to efficiently exploit the system bandwidth, even when the associated contact graph is fully connected. This is especially true when the distribution of link capacities exhibits a long tail, as typically found in experimental traces.

IX. Experimental mobile networks

We have analyzed an experimental vehicular mobile ad hoc network using the data obtained from the DieselNet [7] network running on Umass campus during the Spring quarter in 2005. This network consisted of 30 buses offering the campus transportation service and carrying some short-range radio devices. The publicly available traces [7] refer to the radio contacts among buses, measured in terms of data transferred through TCP connections. The traces provide a sequence of records in the form (t_c, s_c, d_c, b_c), meaning that during the c-th contact, started at time t_c, bus s_c was able to transfer b_c bytes of data to bus d_c. If t_{tot} is the total duration of the trace, it is possible to evaluate directly the service rate between two buses i and j, as

$$\mu^S_{ij} = \frac{1}{t_{tot}} \sum_c \mathbb{I}_{s_c=i} \, \mathbb{I}_{d_c=j} \, b_c$$

by summing the data over any contact c with source i (i.e., $s_c = i$) and destination j (i.e., $d_c = j$). Note that these records depends also on the IEEE 802.11 MAC protocol adopted in the experimentation. Furthermore, we restricted our analysis to the first three weeks, excluding week-ends and nights (in total, 14 working days). This permitted to guarantee a good degree of stationarity in the inter-contact times.

From the traces, one can immediately derive the corresponding contact graph. This turns out to be almost fully connected. The number of its edges is approximately the 86% of a complete graph; in other words, each bus got in contact with almost all of the other buses. It is interesting to note that the resulting contact graph contains significant inhomogeneous capacities; Fig. 3 visually shows the contact graph partitioned into three subgraphs: the left subgraph contains edges with small capacities, the central subgraph (the richest subgraph) contains medium capacity edges, whereas high capacity edges are shown in the right subgraph. In this last subgraph we have highlighted in bold two edges with very high capacity (> 260 Mbytes/14 days). Observe that even if less than the 40% of the overall edges have high capacity, however, high capacity edges provide the main contribution to the transport capability.

We compared the following routing policies:

- **TH-OPT**: the algorithm maximizes the aggregate transmission rate (as defined in (16)) by solving the multi-commodity flow problem on the contact graph defined by (9) and (10) using standard linear programming tools.
- **TH-2H-OPT**: the algorithm maximizes the aggregate transmission rate like TH-OPT, but with the additional constraint that only 1 and 2 hops routes are allowed (i.e., $f^{sd}_{ij} = 0$ when both $i \neq s$ and $j \neq d$).

Routing policy	Maximum throughput	Minimum delay	Number of hops
TH-OPT	yes	no	any
TH-2H-OPT	yes	no	≤ 2
DL-OPT [23]	no	yes	any
random 2-hop relay[13]	only asympt.	no	2

TABLE I
Comparison of the metrics optimized by different routing policies

- **DL-OPT**: the algorithm computes the minimum-delay routes as proposed in [23]. It adopts a modified version of Dijkstra's algorithm to find the paths with the minimum delay. The delay is defined as the time taken by a data unit to be transferred across the sequence of relays up to reaching the destination. To find the complete set of routing paths between a source s and a destination d, the algorithms iterates among the following phases:
  1) compute the minimum delay path[9] $\mathcal{P}_{sd}$;
  2) compute the maximum bandwidth across $\mathcal{P}_{sd}$: $b_{sd} = \min_{(i,j)\in\mathcal{P}_{sd}}\{\mu_{ij}\}$;
  3) allocate b_{sd} across $\mathcal{P}_{sd}$ and remove all the edges corresponding to fully allocated edges; the edges capacities will be updated as $\mu_{ij} - b_{sd}$;
  4) on the residual graph, start again with phase 1) until no bandwidth can be allocated anymore.

  This algorithm can be extended to the multi-flows scenario by just modifying the two initial steps in the following way:
  1) compute the set of minimum delay paths $\{\mathcal{P}_{sd}\}$ for all active flows;
  2) compute the bandwidth among the paths included in $\{\mathcal{P}_{sd}\}$ according to a max-min fairness algorithm [46].

Table I provides a synoptic view of the algorithms; note that the random 2-hop policy in [13] is asymptotically optimal in terms of throughput only for infinite size networks with uniform contact capacities, while TH-2H-OPT can be considered to be its optimal version for finite networks.

For a meaningful comparison, we considered the two following traffic scenarios:

- *single-flow*: just one source-destination pair (i.e. a pair of buses corresponds to a single traffic flow) is active in the network and all the nodes cooperate to transfer the data of such flow. For an exhaustive analysis, we considered all possible $n(n-1)$ flows, one for each source-destination pair; in the case of Umass trace, this corresponds to 870 single-flow cases.
- *multi-flows*: each node is both source and destination of exactly one traffic flow, whose generation rate is set for all flows equal to λ (i.e. $\lambda_{sd} = \lambda, \ \forall\, (s,d)$). All the nodes

[9] The contact graph used by DL-OPT is slightly different from the contact graph described in Sec. VI, since it is a multi-graph storing a detailed information regarding the time and the capacity of all the contact events; the reader can directly refer to [23] for the details.

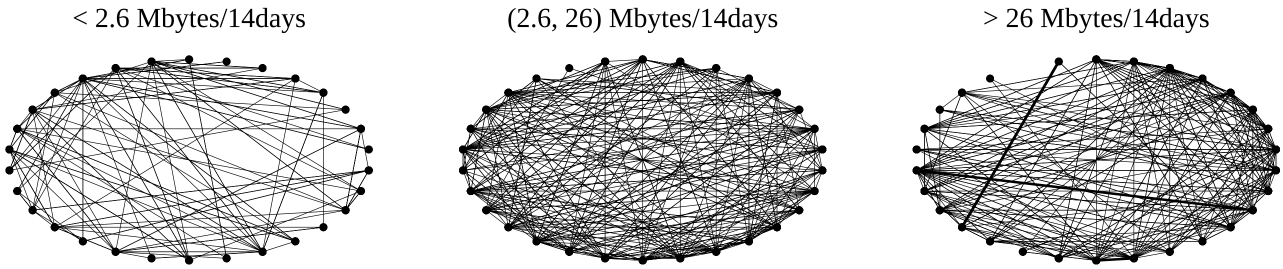

Fig. 3. Contact graphs of Umass buses divided in different classes of edge capacities. Larger capacity edges 3-4 contribute for the 85% of the overall transport capacity.

Routing algorithm	Average capacity per flow	Average delay
TH-OPT	299 Mbytes/14days	105.7 hours
TH-2H-OPT	134 Mbytes/14days	52.6 hours
DL-OPT	216 Mbytes/14days	30.4 hours

TABLE II
PERFORMANCE ACHIEVABLE UNDER SINGLE-FLOW SCENARIO, AVERAGED AMONG ALL THE POSSIBLE FLOWS, FOR UMASS BUSES

cooperate to transfer the n flows active in the network. Since there are $n!$ ways to match sources and destinations, we considered only the following two extreme cases:

- *mDC* (minimum Direct Capacity): nodes pairs are selected in such a way that the sum of the capacities associated to the direct communication links between sources and destinations is minimized; as a consequence, direct communications between sources and destinations are inefficient and multihop transmission should be exploited to maximize the throughput.
- *MDC* (Maximum Direct Capacity): nodes pairs are selected to maximize the same sum as above; in this case, direct communications between a source and its destination are quite efficient and we expect that a small number of hops will be enough for maximizing throughput.

The performance metrics we have considered are: (i) aggregate throughput, (ii) per-flow throughput and (iii) delay; the delay is averaged over the different routing paths, weighted with the corresponding throughput along the path.

For the single-flow scenario, Table II reports the performance achievable by the three considered routing algorithms. As expected, TH-OPT obtains the highest per-flow throughput at the cost of the largest delay. By reducing the maximum number of hops, TH-2H-OPT is able to improve the delay, but at the expense of a lower throughput. To better understand the possible inefficiencies of 2-hop routing in this vehicular scenario, we evaluated the 2-hop throughput efficiency $\eta^{2h}_{sd} = \mu^{2h}_{sd}/\hat{\mu}_{sd}$ for each single flow. Fig. 4 shows the corresponding values for all 870 flows, displayed in increasing order. The average value of η^{2h}_{sd} is 0.603, suggesting that

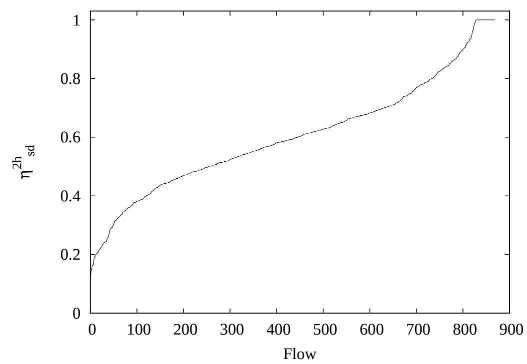

Fig. 4. Two-hops efficiencies η^{2h}_{sd} under single-flow scenario for Umass buses

moderate throughput degradation is experienced employing a 2-hops routing strategy under most traffic scenarios, like the ones previously selected. Nevertheless, the value of $\eta^{2h} = \min_{sd} \eta^{2h}_{sd}$ is rather small, equal to 0.1234. As a consequence, significant throughput degradation is experienced for the worst case source-destination pair.

By optimizing the delay, DL-OPT obtains significantly better delays than TH-OPT and TH-2H-OPT; the corresponding throughput is smaller (as expected) than TH-OPT but larger than TH-2H-OPT. That means that the best throughput-delay tradeoff cannot be achieved by TH-2H-OPT.

Fig.s 5(a)-5(c) provide a more detailed view of the performance achievable under the single-flow scenario, showing the optimal throughput-delay region achievable for each possible flow (represented by a point in the graph), under the three different routing policies. Fig. 5(a) shows that many groups of flows are characterized by identical throughputs but different mean overall delays. These groups share either the same source or the same destination node, where a capacity bottleneck can arise because the node gets sporadically in contact with just another node; note that this effect especially occurs for flows with lower throughput. The throughput on these paths is limited by the bottleneck capacity, whereas delays depend on the specific path. Fig. 5(b) shows that TH-2H-OPT experiences better delays but at the expense of a smaller throughput; the absence of the previous phenomenon can be explained by the fact that TH-2H-OPT is not capable to saturate the bottleneck

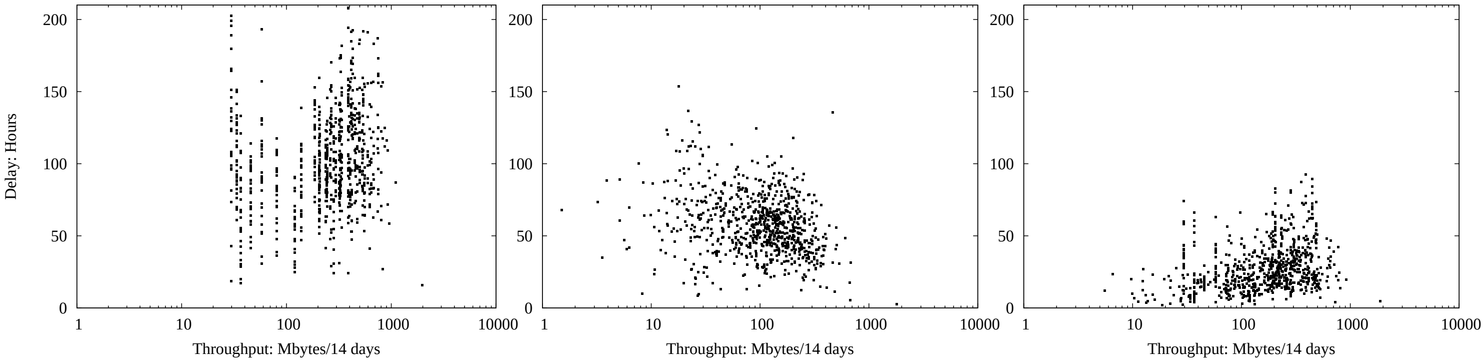

(a) TH-OPT (b) TH-2H-OPT (c) DL-OPT

Fig. 5. Optimal throughput-delay region under single-flow scenario for Umass buses

Traffic scenario	Routing algorithm	Maximum total capacity	Average delay
MDC	TH-OPT	8.36 Gbytes/14days	44.1 hours
	TH-2H-OPT	7.02 Gbytes/14days	22.3 hours
	DL-OPT	7.31 Gbytes/14days	13.6 hours
mDC	TH-OPT	4.09 Gbytes/14days	109.8 hours
	TH-2H-OPT	2.63 Gbytes/14days	65.6 hours
	DL-OPT	2.64 Gbytes/14days	39.1 hours

TABLE III
CAPACITY OBTAINED BY OPTIMAL ROUTING UNDER MULTI-FLOWS SCENARIO FOR UMASS BUSES

capacity, as for TH-OPT. At last, Fig. 5(c) shows the minimum delays achievable under DL-OPT routing: we still observe groups of flows with the same throughput, due to the saturation of bottleneck edges as in TH-OPT.

For the multi-flows scenario, Table III shows the aggregate performance for the two considered traffic patterns, MDC and mDC. The global capacity is computed by summing the throughput of all active flows in the network. As in the single-flow case, TH-OPT achieves the best capacity, but now TH-2H-OPT and DL-OPT obtain almost the same throughput since the multi-flow scenario tends to distribute the load among the nodes and TH-2H-OPT is able to saturate all the edges as DL-OPT. The fact that DL-OPT achieves better delays indicates that there are many paths with similar capacities but different delays; DL-OPT is more able than TH-2H-OPT to exploit these minimum delay paths. In conclusion, the best throughput-delay tradeoff is not obtained by TH-2H-OPT also in the multi-flows scenario.

Between the two traffic patterns, mDC is the one highlighting the throughput performance differences obtained by the three routing algorithms. This is due to the fact that several optimal paths longer than 2-hops are ignored by TH-2H-OPT; Such long paths are likely to be ignored also by the greedy algorithm adopted in DL-OP, since they often lead to large delays.

Fig.s 6 and 7 report the detailed performance obtained by each flow, under MDC and mDC traffic patterns. Results in these plots have been sorted in increasing order, by choosing the appropriate flow sequence. Fig. 6(a) shows that TH-OPT, as expected, always outperforms the other two algorithms in terms of throughput, whereas TH-2H-OPT and TH-DL-OPT exhibit very similar behavior. Furthermore, because of the inhomogeneous nature of the contact graph, throughput performance of different flows is very different (ranging over more than two orders of magnitude), even if the set of flows was chosen to maximize the direct connection capacity. Interestingly, the delays of DL-OPT are better then TH-2H-OPT even if they are associated to the same throughput. At last, Fig. 6(b) suggests that delays for each single flow follow the same relative performance than the aggregate delays.

Under the mDC traffic pattern, Fig. 7(a) presents the throughput achieved by the different flows. Large discrepancies are experienced in this scenario as well. Also the delays experienced by the flows, presented in Fig. 7(b), are significantly different among them.

As a final comment on the numerical results discussed in this section, we note that the absolute performance indexes considered so far show that the DTN exploiting the Umass bus transport service achieve quite poor performance, in terms of both throughput and delays. This is not due to routing inefficiencies, since all the considered routing schemes are optimal with respect to some performance metric and rely on centralized algorithms exploiting complete knowledge of future contacts[10]. Instead, in the considered vehicular scenario, poor performance can be explained easily by considering the test-bed nature of the network, exploiting (in the available traces) the mobility of only few nodes in a large area. The nature of the transportation service itself tends to reduce contact times between buses so as to minimize the delays experienced by the users at the bus-stops.

X. CONCLUSIONS

In this paper we have considered an ad hoc wireless network composed of n heterogeneous mobile nodes and proposed a general methodology that allows to precisely characterize its

[10]Notice that the absolute performance that we have computed provides an upper bound on the best performance achievable by any distributed algorithm running without full knowledge of the contacts

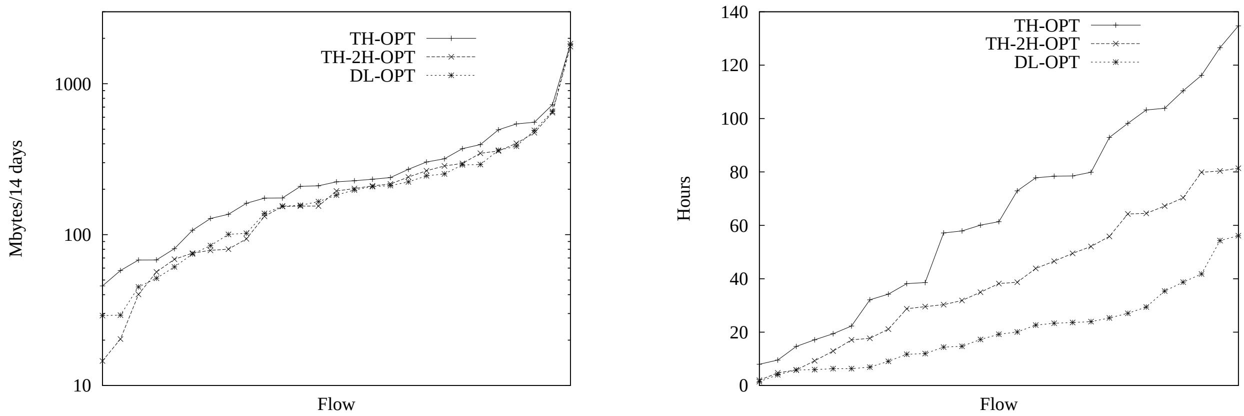

(a) Throughput (b) Average delay

Fig. 6. Performance achievable for each flow, under MDC multi-flows scenario for Umass buses

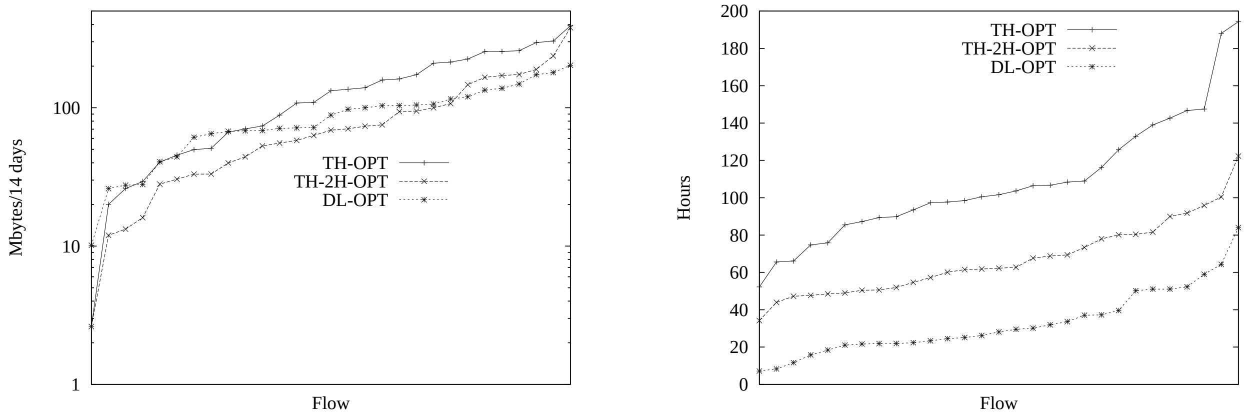

(a) Throughput (b) Average delay

Fig. 7. Performance achievable for each flow, under mDC multi-flows scenario for Umass buses

capacity region by considering the associated *contact graph*, highlighting several important structural properties of the system. Then we have turned our attention to the throughput performance of the 2-hops routing scheme. We have related the effectiveness of the 2-hops routing strategy to the structural properties of the contact graph associated with the network. We have then presented experimental evidence that contact times between nodes are highly inhomogeneous in a real environment. At last we have shown that, in realistic networks, the 2-hops routing strategy may result strongly inefficient in terms of aggregate throughput.

In a real vehicular scenario, we have assessed the capacity-delay region achievable by routing algorithms that are optimal from the throughput and delay point of views. The performance indexes that we have obtained in the considered vehicular scenario are not compatible with the QoS constraints of realistic applications. However we emphasize that our framework is general and can be applied to any vehicular network, provided that the contact history is known or can easily be estimated (as in the case of scheduled mobility for public transportation systems). Our methodology provides an upper bound for the performance achievable by any routing algorithm, and for this reason can be considered as a reference model to test the performance of future implemented routing protocols for delay tolerant networks.

References

[1] M. Garetto, P. Giaccone, E. Leonardi, "On the Capacity of Ad Hoc Wireless Networks Under General Node Mobility", *IEEE INFOCOM'07*, Anchorage, AK, May, 2007

[2] M. Garetto, P. Giaccone, E. Leonardi, "On the Effectiveness of the 2-hop Routing Strategy", *IEEE ICC'07*, Glasgow, UK, June, 2007

[3] Delay Tolerant Network Research Group (DTNRG): `www.dtnrg.org`

[4] K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets," *ACM SIGCOMM '03*, 2003.

[5] S. Burleigh et. al.,"Delay-Tolerant Networking: An Approach to Interplanetary Internet", *IEEE Communications Magazine*, pp. 128–136, June 2003.

[6] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, "Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms", *IEEE INFOCOM '06*, Barcelona, Spain, April 2006.

[7] John Burgess, Brian Gallagher, David Jensen, Brian Neil Levine, "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networking," *IEEE INFOCOM '06*, Barcelona, Spain, April 2006.

[8] H.Y. Huang, P.E. Luo, M. Li, Da Li, Xu Li, W. Shu, M.Y. Wu, "Performance Evaluation of SUVnet With Real-Time Traffic Data", *IEEE Trans. on Vehicular Technology*, vol. 56, n. 6, pp. 3381-3396, Nov. 2007

[9] M. Sede, Xu Li, Da Li, Da; M.Y. Wu, M. Li, W. Shu, "Routing in Large-Scale Buses Ad Hoc Networks", *IEEE WCNC 2008*, Apr. 2008

[10] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.–S. Peh, D. Rubenstein, “Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet”, *ASPLOS-X*, October, 2002. San Jose, CA.
[11] A. Seth, D. Kroeker, M. Zaharia, S. Guo, S. Keshav, “Low-cost Communication for Rural Internet Kiosks Using Mechanical Backhauls,” *ACM MobiCom ’06*, Los Angeles, CA, September 2006.
[12] P. Gupta, P.R. Kumar, “The capacity of wireless networks”, *IEEE Trans. on Information Theory*, vol. 46, n.2, pp. 388–404, Mar. 2000
[13] M. Grossglauser, D.N.C. Tse, “Mobility increases the capacity of ad hoc wireless networks”, *IEEE Trans. on Networking*, vol. 10, n.2, pp. 477–486, Aug. 2002
[14] S. Toumpis and A. Goldsmith, “Large wireless networks under fading, mobility, and delay constraints”, *IEEE INFOCOM ’04*, Hong Kong, China, Mar. 2004
[15] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah, “Throughput-Delay Trade-off in Wireless Networks”, *IEEE INFOCOM ’04*, Hong Kong, China, Mar. 2004
[16] G. Sharma, R. R. Mazumdar and N. B. Shroff, “Delay and Capacity Trade-offs in Mobile Ad Hoc Networks: A Global Perspective” *IEEE INFOCOM ’06*, Barcelona, Spain, Apr. 2006
[17] N. Sarafijanovic-Djukic, M. Piorkowski, and M. Grossglauser, “Island Hopping: Efficient Mobility-Assisted Forwarding in Partitioned Networks”, *IEEE SECON 2006*, Reston, VA, Sep. 2006.
[18] W.-J. Hsu and A. Helmy, “On Nodal Encounter Patterns in Wireless LAN Traces”, *WiNMee ’06*, Boston, MA, 2006.
[19] J. Leguay, T. Friedman, V. Conan, “Evaluating Mobility Pattern Space Routing for DTNs”, *IEEE INFOCOM ’06*, Barcelona, Spain, April 2006.
[20] M. Balazinska, P. Castro, “Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network”, *ACM MobiSys ’03*, San Francisco, CA, May 2003.
[21] J. H. Kang, W. Welbourne, B. Stewart, G. Borriello, “Extracting Places from Traces of Locations”, *ACM Mobile Computing and Communications Review*, 9(3), July 2005.
[22] A. Boukerche, S.K. Das, A. Fabbri, “Analysis of a Randomized Congestion Control Scheme with DSDV Routing in ad Hoc Wireless Networks”, *Journal of Parallel and Distributed Computing*, vol. 61, n. 7, pp. 967-995, July 2001
[23] S. Jain, K. Fall, and R. Patra, “Routing in a delay tolerant network”, *ACM SIGCOMM ’04*, 2004.
[24] W. Zhao, M. Ammar, and E. Zegura, “A message ferrying approach for data delivery in sparse mobile ad hoc networks”, *ACM MobiHoc ’04*, 2004.
[25] C. Liu, J. Wu, “Scalable routing in delay tolerant networks”, *ACM MobiHoc ’07*, 2007
[26] A. Vahdat and D. Becker, “Epidemic routing for partially connected ad hoc networks”, Technical Report CS-200006, Duke University, April 2000.
[27] T. Small and Z. J. Haas, “Resource and performance tradeoffs in delay-tolerant wireless networks,”. *ACM workshop on Delay Tolerant Networking*, August 2005.
[28] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, “Spray and wait: an efficient routing scheme for intermittently connected mobile networks”, *ACM workshop on Delay-tolerant networking*, August 2005.
[29] A. Lindgren, A. Doria, and O. Schelen, “Probabilistic routing in intermittently connected networks”, *Mobile Computing and Communications Review*, 7(3):19–20, 2003.
[30] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, “Age matters: efficient route discovery in mobile ad hoc networks using encounter ages”, *ACM MobiHoc ’03*, Annapolis, MD, June 2003.
[31] M. Grossglauser and M. Vetterli, “Locating Mobile Nodes with EASE: Learning Efficient Routes from Encounter Histories Alone”, *IEEE/ACM Trans. on Networking*, vol 14, no 3, June 2006.
[32] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, “Single-copy routing in intermittently connected mobile networks,” *IEEE SECON ’04*, Santa Clara, CA, October 2004.
[33] S. Jain, M. Demmer, R. Patra, K. Fall, “Using redundancy to cope with failures in a delay tolerant network”, *ACM SIGCOMM ’05*, 2005.
[34] X. Zhang, G. Neglia, J. Kurose, D. Towsley, “On the Benefits of Random Linear Coding for Unicast Applications in Disruption Tolerant Networks”, *IEEE NETCOD Workshop*, April 2006.
[35] J. Widmer, J.-Y. Le Boudec, “Network Coding for Efficient Communication in Extreme Networks”, *ACM workshop on Delay-tolerant networking*, August 2005.
[36] E.P.C. Jones, L. Li, J.K. Schmidtke, P.A.S. Ward, “Practical Routing in Delay-Tolerant Networks”, *IEEE Trans. on Mobile Computing*, vol. 6, n. 8, pp. 943-959, Aug. 2007
[37] 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, n. 11, pp. 3947–3954, Nov. 2005
[38] A. Al Hanbali, A. A. Kherani, R. Groenevelt, P. Nain, E. Altman, “Impact of Mobility on the Performance of Relaying in Ad Hoc Networks,” *IEEE INFOCOM ’06*, Barcelona, Spain, Apr. 2006
[39] T. Camp, J. Boleng, and V. Davies, “A Survey of Mobility Models for Ad Hoc Network Research”, Wireless Communications & Mobile Computing (WCMC), Vol 2 n 5, pp. 483-502.
[40] A. Jardosh, E. Belding-Royer, K. Almeroth, S. Suri. “Towards Realistic Mobility Models for Mobile Ad hoc Networks”, *ACM MobiCom 2003*, San Diego, CA.
[41] A. Boukerche, *Handbook on Algorithms for Wireless Networking and Mobile Computing*, CRC, 2005
[42] A. Saha, D. Johnson. “Modeling mobility for vehicular ad-hoc networks”, *1-st ACM international workshop on Vehicular ad hoc networks*, Philadelphia, PA, 2004
[43] L. Fratta, M. Gerla, L. Kleinrock, “The Flow Deviation Method - An Approach to the Store-and-Forward Communication Network Design,” *Networks*, vol. 3, pp. 97-133, 1973
[44] M. Bramson, “Convergence to Equilibrium for Fluid Models of FIFO Queueing Networks”, *Queueing Systems*, vol. 22, pp. 5-45, 1996
[45] B. Bollobas, *Random Graphs*, Cambridge University Press, Cambridge, UK, 2001
[46] D. Bertsekas, R. Gallager, *Data Networks*, Prentice Hall, 1987

PLACE PHOTO HERE

Michele Garetto (S’01–M’04) received the Dr.Ing. degree in Te lecommunication 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. His research interests are in the field of performance evaluation of wired and wireless communication networks.

PLACE PHOTO HERE

Paolo Giaccone (M’01) received the Dr.Ing. and Ph.D. degrees in telecommunications engineering from the Politecnico di Torino, Torino, Italy, in 1998 and 2001, respectively. He is currently an Assistant Professor in the Department of Electronics, Politecnico di Torino. During the summer of 1998, he was with the High Speed Networks Research Group, Lucent Technology-Bell Labs, Holmdel, NJ. During 2000-2001, he was with the Department of Electrical Engineering, Stanford University, Stanford, CA. His main area of interest is the design of scheduling policies for high-performance routers and for wireless networks.

PLACE PHOTO HERE

Emilio Leonardi (M’99) received the Dr.Ing. degree in electronics engineering and the Ph.D. degree in telecommunications engineering from the Politecnico di Torino, Italy, in 1991 and 1995, respectively, where he is currently an Associate Professor. In 1995, he was with the Department of Computer Science, UCLA. In the summer of 1999, he was with the High Speed Networks Research Group, Lucent Technology-Bell Labs, Holmdel, NJ, and in the summer of 2001, he was with the Department of Electrical Engineering, Stanford University, Stanford, CA. His areas of interest are all-optical networks, queueing theory, and scheduling policies for high-speed switches.