On the Role of Network Coding in Uncoordinated Multihop Content Distribution over Ad Hoc Networks

Mark Johnson[1], Minghua Chen[2], and Kannan Ramchandran[1]

[1]Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, CA, 94720
[2]Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong
{mjohnson,kannanr}@eecs.berkeley.edu, minghua@ie.cuhk.edu.hk

***Abstract*—We consider the problem of multihop communication in ad hoc networks. This work was originally motivated by vehicular networks, and has application in numerous fields. The key constraint in such systems is the infeasibility of coordination. Because no single node has global knowledge of the network topology, centralized scheduling of transmissions and routing of packets is prohibitively expensive. We propose instead a simple distributed protocol, based on independent channel access and random network coding. This scheme does not require link level feedback, and nodes need not even track the identities of their current neighbors.**

Our main result shows that when using this scheme, the end-to-end throughput is identical to the throughput in a one hop network utilizing the same protocol. Thus, in an application requiring uncoordinated communication, network coding enables a multihop network to perform as efficiently as a single hop network. In the context of vehicular networks, in a region of congested traffic a constant throughput can be provided to vehicles arbitrarily far from the base station. In essence, every vehicle can be thought of as a *digital repeater*, despite the packet losses due to collisions and fading.

I. Introduction

The development and commercialization of low cost radio architectures is enabling many novel ad hoc wireless networks, in applications ranging from environmental monitoring to intrusion detection to energy efficient buildings. One very promising field is vehicular ad hoc networks (VANETs). Significant research and development in vehicular networking was spurred by a recent FCC spectrum allocation [1], which was targeted at vehicle safety systems but has the potential to also enable a wide variety of other services.

The key challenges in VANETs revolve around the lack of centralized coordination of scheduling and routing. Due to limitations and the power and computational complexity of the nodes, it is either technically infeasible or prohibitively expensive for any one node to collect global knowledge of the network topology. Therefore, distributed medium access control (MAC) and network layer protocols are a necessity in such ad hoc networks.

In this work, we focus on the problem of multihop communication in ad hoc networks. We propose a very simple protocol, based on random channel access and randomized network coding, which does not even require feedback at the link level. Our main result shows that the resulting throughput is the same as in a single hop network using the same protocols. Thus, a constant throughput can be provided, regardless of the number of hops between the source and destination. The intuition behind this result is that network coding, the mixing of packets at intermediate nodes, makes the uncoordinated network robust against packet losses. Without coding, in contrast, the throughput would degrade with distance due to the inability to coordinate among nodes. In effect, each vehicle functions as a *digital repeater*.

In the course of studying the throughput of uncoordinated networks, we address the problem of *starvation*, the possibility that intermediate nodes may temporarily not have any new information because of the random ordering of transmissions. We use queueing theory to demonstrate that starvation does not affect the end-to-end throughput, a result which has application in numerous problems extending beyond vehicular networks.

The rest of this paper is organized as follows. In Section II we present a qualitative discussion of the unique requirements of vehicular ad hoc networks, and then discuss the MAC and network layer protocols that we propose. In Section III we describe a model for the connectivity of a highway vehicular network. In Section IV we present a general method for finding the throughput of an uncoordinated multihop network, and in Section V we apply it to the model for vehicular ad hoc networks. Conclusions and future work are discussed in Section VI.

II. Content Distribution Problem in Ad Hoc Networks

In this work, we analyze the achievable rate for a unicast session in an ad hoc network. Figure 1 shows an example vehciular network, where the source is a fixed base station which wishes to send some information to a destination vehicle. Due to the distance between the source and destination, and limitations on the transmit power, it is not possible for the source to directly communicate with the destination. Instead, the vehicles on the road will form an ad hoc network, with the vehicles between the base station and the destination acting as

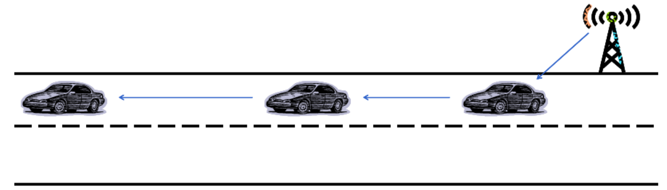

Fig. 1. We consider a unicast communication problem in a multihop vehicular ad hoc network.

relays. We assume that the data stream is short enough that the network topology can be modeled as being constant over the session.

The key feature of vehicular networks, which distinguishes them from computer networks and static sensor networks, is the lack of global state knowledge. One strategy to deal with this challenge is for the vehicles to track the set of their current neighbors, and continuously update routing tables. However, such a scheme will consume a significant portion of the network resources, which in turn will limit the amount of data which can be carried.

We propose instead to use a set of distributed protocols, where vehicles do not schedule their channel access or route packets from end-to-end. In fact, we study here a completely uncoordinated protocol, in which vehicles do not even receive feedback as to which, if any, of their current neighbors receive each transmitted packet. We believe that distributed protocols are particularly well matched to vehicular networks.

We consider an independent medium access layer, inspired by the unslotted Aloha protocol from computer networking [2]. Relay node i will choose to transmit a packet in each time slot with probability p_i, independent of all other time slots and all other vehicles. Because the nodes are distributed, however, their clocks are not synchronized and thus the boundaries of the slots will be different at each vehicle. As in the classical unslotted Aloha protocol, a packet transmitted by one vehicle can interfere with two consecutive packets at neighboring nodes due to the offset in the time slot boundaries[1].

We propose to combine this distributed MAC protocol with intra-session network coding. When each vehicle chooses to access the channel, it will send a random linear combination of all of the packets that it has received up to that point. Nodes that receive the transmission will check whether it is independent of their set of packets. If so, they will store the new packet in their memory, and if not they can discard the dependent packet.

We emphasize that we are not using network coding to improve the capacity of the network. In this unicast problem, the maximum throughput could be achieved without coding. However, the vehicles would need full knowledge of the topology at all times. Network coding allows us to use a distributed protocol and still achieve this maximum throughput. As we will see, coding effectively provides robustness to the lack of coordination in the VANET.

[1]While this work focuses on an unslotted Aloha MAC protocol, the results can be easily adapted to networks that use CSMA.

III. Modeling Highway VANET Connectivity

The transmission power used in vehicular ad hoc networks is typically quite large, which in turn leads to a large transmission range. For example, transmitters operating in the Dedicated Short Range Communication (DSRC) band can transmit up to a range of about 300 meters [3]. Thus, the transmission range is significantly larger than the width of the road, which is no more than about 50 meters for an 8 lane highway, and clearly much smaller for a road with fewer lanes. Figure 2 shows the relative sizes of the transmission range and road width when the vehicles are making full use of the available transmit power.

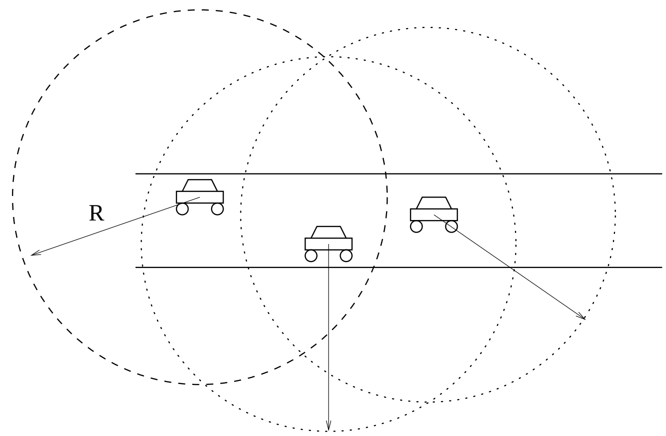

Fig. 2. The transmission range R of a vehicular network is typically much larger than the width of the road.

Because the transmission range is much larger than the width of the road, a reasonable model for the VANET connectivity is a one-dimensional line network. Although the vehicles may be located in different lanes, we will make the approximation that each vehicle can communicate with any other vehicles that are within a longitudinal distance R from itself along the road[2].

The number of neighbors of each node depends on the transmission range and the inter-vehicle spacing. In heavily congested scenarios, the spacing is nearly constant and thus the connectivity can be modeled by a graph where every node has a fixed number of neighbors. An example where each node has 2 neighbors in both directions is shown in Figure 3. In scenarios where traffic is flowing freely, on the other hand, the spacing is variable. For example, the vehicle locations might be modeled with a Poisson process.

IV. Throughput of Multihop Networks

The end-to-end throughput of a network using any randomized MAC protocol and network coding can be computed in

[2]When vehicles use a low transmit power, this approximation is no longer valid. In the case of extremely low transmit power and a heavily congested road, where all lanes are at the maximum vehicle density, the connectivity can be modeled as a two dimensional grid lattice. We will study the throughput of the low power vehicular network in the sequel to this work.

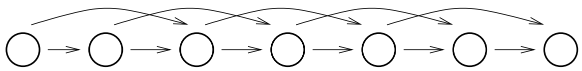

Fig. 3. The connectivity of a VANET with uniformly spaced nodes. Each vehicle has 2 neighbors in both the forward and reverse directions.

two steps. First, the capacities of the individual links, which depend of the MAC protocol, are computed. Then, it follows from the max-flow min-cut theorem and Ahlswede [4] that the throughput from source to destination is equal to the min-cut of the link capacity graph. The key technical detail, which differentiates this from prior work, is the problem of *starvation*. Because the randomized MAC protocol induces a random ordering of packet transmissions, the intermediate relays may be temporarily out of new information to forward to their neighbors.

A. Link Capacities

The link capacities can be found via a recently proposed method from the CSMA literature. The throughput of links was shown to be computable from a time-reversible Markov chain in [5]. An efficiently computable approximation was given in [6] and the approach has been further developed in [7].

As a concrete example, we will consider the linear network shown in Figure 4, with links labeled 1 through 4. We begin by listing the *independent sets*, the sets of links that can be used concurrently, without interfering with each other. In this example, links 1 and 4 can be used at the same time. However, if any other set of multiple links are used at the same time, a collision will occur. Thus, there are five independent sets in this network: $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$, and $\{1, 4\}$.

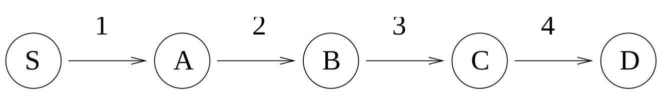

Fig. 4. An example network with four links. The independent sets are $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$, and $\{1, 4\}$.

Next, we form a Markov chain whose state is the set of links that are carrying information at the current time. We refer to this set as the *active* links. Clearly, the state space of the Markov chain is given by the independent sets plus the null set. This Markov analysis is valid for networks using Aloha, and CSMA networks without hidden nodes[3].

The valid transitions of this Markov chain, and the transition probabilities, depend on the specific MAC protocol being used. For Aloha protocols, the Markov chain is a fully connected graph, since the set of active links in two different timeslots is independent. The network can transition from any independent set to any other independent set. If the vehicular network used CSMA, on the other hand, time dependencies are introduced because of the backoff counter that is started when a node senses that the channel is in use. The transition diagram for the network in Figure 4 will have a form as shown in Figure 5.

[3]Any CSMA network can be made free of hidden nodes by appropriately adjusting the carrier sensing range.

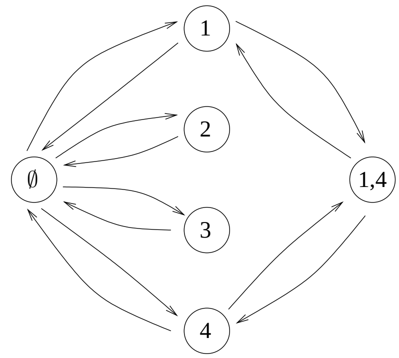

Fig. 5. The Markov chain corresponding to the independent set of the network in Figure 4 under a carrier sensing MAC protocol.

With any finite number of links in the graph, there will be a finite number of independent sets. The Markov chain will then have a steady state distribution, which can be computed from the transition probabilities. Finally, the link capacities are given by the fraction of time that a link is active, which is the sum of the steady state probabilities of independent sets containing that link. In this example, C_1, the capacity of link 1, is equal to $\pi_1 + \pi_{14}$. Therefore, we can produce a graph with edge labels given by the capacities of the wireless links, as shown in Figure 6. We will refer to this as the *link capacity graph*.

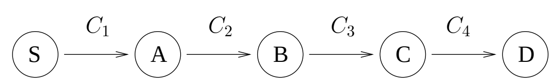

Fig. 6. The link capacity graph of the example network in Figure 4.

When random network coding is used, and the coding is done over a sufficiently large field, the end-to-end throughput is equal to the min-cut of the link capacity graph with high probability. However, this requires the assumption that all nodes are in a *saturated* state, meaning they always have new information to transmit. Consider the network in Figure 4, and suppose that the following sequence of independent sets are active in four consecutive timeslots: $1, 2, 2, 1$. Even though nodes S and A both successfully transmitted two packets, node B has only received one independent packet. This is because when A transmits the second time, it has not yet received a new packet from S, and hence must retransmit the same packet that it sent previously, or a scalar multiple of it. Hence, one of the transmissions from A is wasted because of the ordering.

In the next section, we prove that the performance loss caused by starvation at intermediate nodes is a second order effect, and does not reduce the rate below the min-cut. This result has applications beyond just vehicular networks. In the CSMA literature, for example, it is common to assume that all nodes are saturated for ease of analysis.

B. End to End Throughput

Prior work on multihop wireless networks, and in particular on carrier sense MAC protocols, often assumes that starvation does not matter, and the throughput will not been changed when all nodes are assumed to be in a saturated state [8], [9]. We use a result from queueing theory to show that removing this assumption does not decrease the throughput. Even when starvation is taken into account, the end-to-end throughput of the network is still equal to the min-cut of the link capacity graph.

For illustrative purposes, we will consider the example network shown in Figure 7, with two relays between the source and destination. The corresponding link capacity graph is shown in Figure 8, where the C_i are dependent on the MAC protocol.

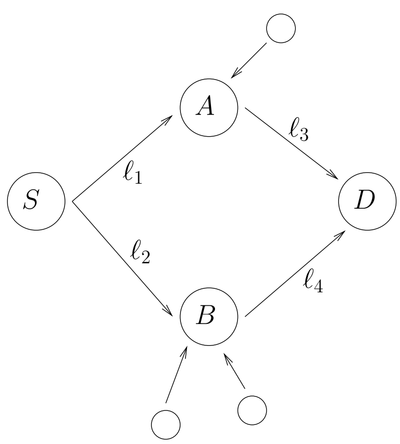

Fig. 7. An example network with source S, destination D, and two relays A and B. Transmissions from S may be received by both of the relay nodes.

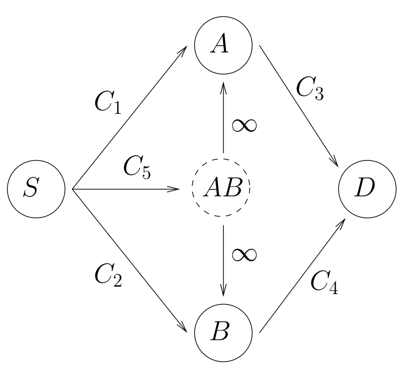

Fig. 8. The link capacity graph of a network with a source S, destination D, and two relays A and B. The dashed node represents a virtual node, which accounts for packets that are received by both of the relays.

Theorem 1: *The maximum throughput of a multihop wireless network using random linear network coding is given by the min-cut of the link capacity graph corresponding to the MAC protocol.*

Proof: We present a sketch of the proof of this theorem. The technical details will be presented in the sequel. First, we find the maximum flow through the link capacity graph, and one rate allocation that achieves this flow. We denote the rate assigned to each link by R_i^*, and the value of the flow by λ^*. By the max-flow min-cut theorem, λ^* is equal to the min-cut of the graph. Figure 9 shows the *rate allocation graph* corresponding to the the link capacity graph in Figure 8. Because this is a valid flow, $R_i^* < C_i$ for all i, and the total rate entering each intermediate node equals the total rate exiting. We further observe that for any $\lambda < \lambda^*$, we can write $\lambda = \gamma\lambda^*$, where $0 < \gamma < 1$. Then, $R_i = \gamma R_i^*$ is a valid flow that achieves rate λ.

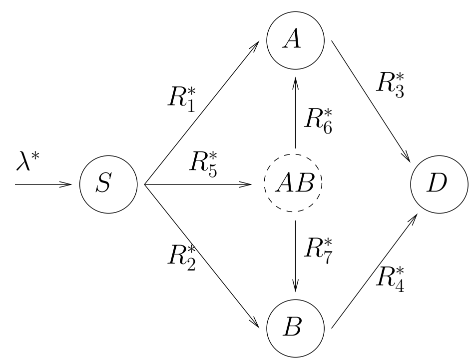

Fig. 9. A flow assignment that achieves the max-flow from S to D while respecting the link capacities.

Then, we construct a queueing network corresponding to the rate λ and allocation R_i, by introducing a genie into the original wireless network. This genie alters the physical network in two ways. First, each node randomly drops incoming packets with probability $1 - \frac{R_i}{C_i}$. Thus, the rate of each link is now R_i, instead of the link capacity C_i. Second, packets entering a virtual node are assigned to exactly one physical node, in a ratio proportional to the outgoing rates from the virtual node. Lemma 1, in Section IV-C, demonstrates that these modifications can only decrease the end-to-end throughput. Thus, the *genie-degraded* network provides a lower bound on the actual system.

With the genie inserted, there is now only one copy of each packet in the network, even if a transmission was physically received by multiple nodes. Therefore, we can represent the flow of information in the genie-degraded system by a queueing network. The queue length at each node represents how many more packets a node has received than it has successfully transmitted to one of its neighbors[4].

The key observation is that because of interference constraints in the underlying wireless network, it is not possible

[4]Recall that each node stores all packets that it receives, and therefore the physical buffer at each node only increases in length. Because there is no feedback in the network-layer protocol, nodes do not know which of their transmissions were lost due to collisions. The queue length is thus a virtual quantity, of which nodes do not have any knowledge. One consequence is that it is not possible to implement schemes where the transmission probability is modulated by queue length, such as [8], [10].

for adjacent servers in the queueing network to process packets at the same time. Therefore, the server processes are correlated across the nodes, and the queueing network is a *generalized* Jackson network, as opposed to a standard Jackson network with independent service processes.

The stability of the queueing network is the key property for calculating the maximum throughput of the system. If the queueing network is stable, the queue length at each node is finite and there are only a finite number of packets at all of the intermediate nodes at any given time. As time gets large, the rate of packets arriving at the destination must approach λ, the rate of packets injected into the network from the source. On the other hand, if the queueing network is unstable, one or more queue lengths becomes infinite and the rate of packets arriving at the destination is less than λ. Therefore, the maximum end-to-end throughput is equal to the maximum rate λ for which the queueing network is stable.

The stability of generalized Jackson networks is considered by Lelarge [11], which is equivalent to earlier results due to Baccelli and Foss [12], [13]. They first define π_i to be the average number of times that a packet visits each of the M nodes, and show that π_i satisfies the following system of equations:

$$\pi_i = p_{0,i} + \sum_{j=1}^{M} p_{ji}\pi_j \qquad\qquad i = 1, \ldots, M$$

where p_{ji} gives the probability that a packet which has been served by j is routed to i, and $p_{0,i}$ represents the fraction of external arrivals which are first served by i. Defining $\mu^{(i)}$ to be the mean service rate at server i, Lelarge shows that the generalized Jackson network is stable if

$$\lambda \frac{\pi_i}{\mu^{(i)}} < 1 \qquad \forall i$$

We now apply this theorem to the queueing system defined by our genie-degraded wireless network. The mean service rate of node i is equal to the total capacity out of that node:

$$\mu^{(i)} = \sum_{e:Out(i)} R_i^*$$

Changing notation so that R_{ij} denotes the rate along the link from node i to node j, the transition probability p_{ij} is defined to be the fraction of outgoing rate from i that is assigned to the link from i to j:

$$p_{ij} = \frac{R_{ij}}{\sum_{e:Out(i)} R_e}$$

For any $\lambda < \lambda^*$, we can write the flow balance equation at node i as

$$\sum_{e:Out(i)} R_e = \sum_{e:In(i)} R_e$$

where $e : Out(i)$ denotes the set of all edges e that flow out of node i. Defining the total outflow of node i as $V_i = \sum_{e:Out(i)} R_e$, we have

$$V_i = \sum_{e:In(i)} R_e = \sum_{e=(j,i)} \frac{R_e}{V_j} V_j = \sum_{e=(j,i)} p_{ji} V_j$$

or

$$V_i = \sum_{e=(j,i)} p_{ji} V_j$$

Now, if we set $V_i = \lambda \pi_i$ and observe that $p_{0,i} = 0$ for all nodes except the source in our problem, we recover the same system of equations in π_i as in Lelarge's theorem. This is not a trivial point. By defining $\pi_i = \frac{1}{\lambda} V_i$, we have in fact recovered the correct steady state distribution on the number of times a packet visits each server, and can then apply the theorem for the stability of generalized Jackson networks. It follows from the flow balance equations and the definitions above that

$$\gamma \mu^{(i)} = \lambda \pi_i \qquad \forall i$$

or

$$\gamma = \lambda \frac{\pi_i}{\mu^{(i)}} \qquad \forall i$$

Because $\gamma < 1$ by definition, the generalized Jackson network is stable. Therefore, when the source generates packets at rate $\lambda < \lambda^*$ the destination will receive new packets at average rate λ, and the possibility of starvation at the intermediate nodes does not reduce the throughput of the network. □

C. Genie-Degraded Wireless Networks

Lemma 1: *The unicast capacity of a network using random network coding cannot increase when nodes discard a subset of packets that were successfully received.*

Proof: We begin by examining a time-space representation of the system, shown in Figure 10. All of the nodes in a row represent the same physical node, with each column corresponding to a separate time step. Thus, we can think of the graph as "unwrapping" the time axis. We are interested in the max-flow from the source at the first time step (the top left node) to the destination at the final time step (the bottom right node). Observe that the horizontal edges represent each node's internal buffer, i.e., each physical node stores all packets that it receives and can code over them in future timeslots. The diagonal edges represent successful transmissions between nodes.

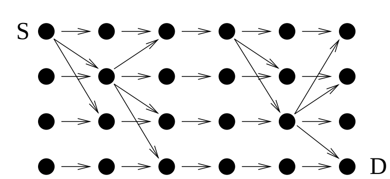

Fig. 10. A time-space representation of the network. All of the nodes in a single row represent the same physical node at different time steps.

A node ignoring a packet that it received is equivalent to removing one of the diagonal edges in the time-space graph, which cannot increase the min-cut. Thus, when coding over an infinite field, where the max-flow always equals the min-cut, the effect of removing some edges is that the max-flow either is unchanged or decreases.

When coding over a finite field, the probability that the destination can decode the source data is at least $(1 - 1/q)^\eta$ where q is the field size and η is the number of links carrying coded packets [14]. Removing edges will reduce η, which increases this lower bound on the probability of decoding. Thus, removing links has not increased the min-cut, but we must increase the field size q if we want to guarantee that the decoding probability in the original system meets a desired target. □

V. Throughput of VANETs

The results of the previous section can now be applied to the problem of finding the throughput of a vehicular ad hoc network. Utilizing the model presented in Section III, and assuming that the road is in a heavily congested state where the vehicles are approximately evenly spaced, we will assume that each vehicle has K neighbors both in front and behind itself. Thus, the throughput is governed by the minimum cut, as shown in Figure 11.

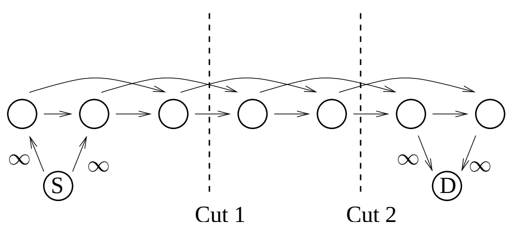

Fig. 11. In the uniformly spaced vehicular network, all cuts that do not pass through the infinite capacity edges near the source and destination will have the same (minimum) value.

When each vehicle has a constant number of neighbors, the topology is regular. Hence, the minimum cut does not change as the distance between the source and destination increases. Instead, adding more hops to the network simply adds more cuts with the same value, and the multihop capacity is identical to the one-hop capacity between the source and its immediate neighbor.

We emphasize that this is not an intuitively obvious result. For example, in prior work [15] we proposed a *synchronous* protocol based on random access and network coding, where vehicles are grouped into spatial blocks, and transmissions are coordinated at the block level. That result showed that the throughput under their scheme is constant if the number of blocks grows subexponentially in the vehicle density, but decays to zero otherwise. Thus, in practical scenarios where the vehicle density is limited by the physical dimensions of the cars, there is a limited distance over which their protocol can support a non-zero rate.

VI. Conclusions and Future Work

In this work, we have considered the problem of multihop content distribution in a vehicular ad hoc network. We have proposed distributed MAC and network layer protocols that are compatible with the rapidly changing network topology, and analyzed the resulting throughput. Our fundamental result shows that by using random network coding, the throughput is independent of the distance between the source and destination.

In the future, we intend to study the effect of practical channel models on the system performance, as well as extend our results beyond the single unicast communication problem here to consider multicast and multiple overlapping unicast sessions.

VII. Acknowledgments

The authors would like to thank Prof. Francois Baccelli of INRIA and the Ecole Normale Superieure for discussions on the connections between wireless networks and queueing theory.

References

[1] http://grouper.ieee.org/groups/scc32/dsrc/index.html.

[2] N. Abramson, "The ALOHA system - another alternative for computer communications," in *1970 Fall Joint Comput. Conf., AFIPS Conf. Proc.*, Montvale, NJ, 1970, pp. 281–285.

[3] Q. Xu, T. Mak, J. Ko, and R. Sengupta, "Medium access control protocol design for vehicle-vehicle safety messages," *IEEE Transactions on Vehicular Technology*, vol. 56, no. 2, pp. 499–518, March 2007.

[4] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. on Information Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.

[5] X. Wang and K. Kar, "Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks," in *IEEE INFOCOM*, Miami, FL, March 2005.

[6] S. C. Liew, C. Kai, J. Leung, and B. Wong, "Back-of-the-envelope computation of throughput distributions in CSMA wireless networks," in *IEEE International Conference on Communications (ICC)*, Dresden, Germany, June 2009.

[7] M. Chen, S. Liew, Z. Shao, and C. Kai, "Markov approximation for combinatorial network optimization," in *IEEE INFOCOM*, San Diego, CA, March 2010.

[8] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," in *Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2008.

[9] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. Poor, "Towards utility-optimal random access without message passing," *Special issue in Wiley Journal of Wireless Communications and Mobile Computing*, December 2009.

[10] J. Ni, B. Tan, and R. Srikant, "Q-CSMA: Queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks," in *IEEE INFOCOM Mini-Conference*, San Diego, CA, March 2010.

[11] M. Lelarge, "Fluid limit of generalized Jackson queueing networks with stationary and ergodic arrivals and service times," *Journal of Applied Probability*, vol. 42, pp. 491–512, 2005.

[12] F. Baccelli and S. Foss, "Ergodicity of Jackson-type queueing networks," *Queueing Systems (QUESTA)*, no. 17, pp. 5 – 72, 1994.

[13] ——, "On the saturation rule for the stability of queues," *Journal of Applied Probability*, vol. 32, pp. 494–507, 1995.

[14] T. Ho, M. Medard, J. Shi, M. Effros, and D. Karger, "On randomized network coding," in *Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2003.

[15] M. Johnson, L. D. Nardis, and K. Ramchandran, "Collaborative content distribution for vehicular ad hoc network," in *Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2006.