INTER-RELATIONSHIPS OF PERFORMANCE METRICS AND SYSTEM PARAMETERS IN MOBILE AD HOC NETWORKS

Wing Ho Yuen Roy D. Yates
WINLAB, Rutgers University
andyyuen@winlab.rutgers.edu ryates@winlab.rutgers.edu

Abstract — In this paper, the network behavior of a routing algorithm for mobile ad hoc networks is investigated. Extensive simulations are performed using *ns-2* in a variety of mobility scenarios and offered load regimes. In the literature, performance metrics (goodput, delay and path length) are often obtained through ensemble averaging of many flows. Here we advocate an alternate graphical interpretation of simulation results similar to that used by Holland and Vaidya. Performance metrics of individual monitored flows are plotted instead. By identifying the correlations between performance metrics and system parameters, inter-relationships between them are revealed. For example, we have shown that path length is dependent on system parameters such as mobility, offered load and even the node distribution. These observations often give us insights to the mechanisms that underline the network behavior. In particular we have resolved a conjecture that goodput improvement under high mobility is due to the load balancing effect. We show that at high mobility, goodput improvement for heavy offered load regimes is a consequence of the reduction of path length in the flows. Furthermore, we have introduced the concept of fraction of congested flows as a new performance metric. This and some other metrics such as fairness can be visualized from our graphs and are important in characterizing network performance.

I. INTRODUCTION

A mobile ad hoc network consists of mobile nodes which communicate with each other through multihop routing. Due to the dynamically changing topology, network routing is an important issue. Numerous routing algorithms [5] are proposed to facilitate efficient packet delivery in mobile environments. The focus is on the relative performance of routing algorithms, which are often characterized by a few performance metrics such as packet goodput, delay and path length. On the other hand, the network behavior for mobile ad hoc networks is not well understood. There are no systematic studies on the correlations between various performance metrics and system parameters such as node mobility and offered load. We show in this paper by interpreting the simulation results in an alternate graphical representation, interesting relationships are revealed.

Our graphical interpretation of simulation results is similar to that used by [2]. The performance metrics of individual *flows* (source destination pair) are plotted. Whereas [2] use this graphical representation to compare metrics on individual flows, we extend the use of these graphs to investigate the correlations between various metrics and system parameters. This has led to some new observations not reported before. Inter-relationships between the goodput, path length and node mobility and offered load are revealed, yielding insight to the mechanisms that underline the network behavior. For example, by studying the relationships of path length and goodput to speed simultaneously, we resolve a conjecture that goodput improvement under high mobility is due to the load balancing effect. Furthermore, we have introduced the concept of *fraction of congested flows* as a new performance metric. This and some other metrics such as *fairness* could be easily visualized from our graphs and are important in characterizing network performance.

Our data interpretation described in this paper is general. Many inter-relationships between performance metrics and system parameters could be obtained. Due to the space constraints, however, we focus on the inter-relationships between goodput, path length and the system parameters only. The goodput (packet delivery ratio) G denotes the fraction of packets that is correctly received. The path length L denotes the number of hops a packet travels along a route. Inter-relationships on packet delay is not discussed. The rest of the paper is organized as follows. In section 2, we discuss the inadequacies of using ensemble averaging in obtaining the performance metrics. A graphical representation similar to that used by [2] is described. Section 3 describes our simulation setup and our main observations are discussed in section 4.

II. ENSEMBLE AVERAGING IN PERFORMANCE METRICS

In the routing literature [1], [4], the performance metrics of a network are obtained using ensemble averaging. For a given node mobility, a number of mobility scenarios are created. A number of flows in each scenario is monitored for some duration. We obtain the performance metrics by averaging over all the monitored flows. The metrics are usually plotted against system parameters such as mobility [1], [4] or offered load [4].

Although an average performance metric is useful in ranking the performance of different routing algorithms, it tells little about the performance metric of individual flows in each scenario. In different mobility realizations of the network, node placement and the traffic patterns are different. The aggregate network performance in each scenario varies. Even within a mobility scenario, individual flows also exhibit heterogeneous performance characteristics due to different path lengths and the non-uniform spatial distribution of the offered load in the network. Thus when the ensemble average of a performance metric is plotted, the variation of individual flows within the network is not captured.

In [1], [4], the performance metrics are plotted against mobility. Each point on a graph indicates the average performance metric for a specific value of node mobility. Adjacent points are obtained from mobility scenarios with different node mobility. Since these metrics are averaged over flows that have very different characteristics, the graphs in [1], [4] show zigzag patterns. The large variations in performance over different flows undermine the validity of the average metric. The trends of the performance metrics with system parameters are not obvious. This hinders our objective to find the inter-relationships between the performance metrics and system parameters.

An average performance metric sometimes yields a misleading conclusion too. A classical example is when the average packet delay is plotted. In a typical ensemble of monitored flows, some flows may have lousy routes and many packets are dropped. Since these flows have higher packet delay in general, the average packet delay will be artificially smaller since few packets from these flows reach the destination node. In the comparison of two routing algorithms, if one algorithm drops a lot of packets from lousy routes, the average delay is then artificially smaller, rendering the comparison between two algorithms meaningless. The more robust algorithm will exhibit an anomaly behavior of a higher goodput and packet delay.

In this paper, we advocate an alternate graphical representation of simulation results that is similar to the approach of [2]. In [2], the throughput for each of 50 mobility patterns for the 20m/s and 30m/s mean speeds used in the simulations are plotted. The patterns are sorted in the order of throughputs at 20m/s. It is demonstrated that for some mobility patterns the goodput improves with mobility. We use similar graphs in our data interpretation with some modifications. We assign a *pattern number* to each monitored flow. The pattern number are assigned such that the flows are ordered in the order of increasing goodput, path length, or delay. We plot the performance metric of each monitored flow versus the pattern number. Under the above assumptions, the throughput plot for each of the 50 mobility patterns for the 30m/s mean speed is different. The same pattern number no longer refers to the same flow in each of the two mean speeds. Since the throughput plot for each mean speed is an increasing function, it is obvious to observe the inter-relationship of throughput versus mobility.

III. Simulation Setup

The simulations are performed on *ns-2* [6], with its wireless extensions developed by the Monarch project [7]. The simulations consist of 50 mobile nodes that are distributed uniformly in a 1500m by 300m area. The propagation model consists of a simple path loss model with attenuation due to distance only. The path loss exponent is chosen to be $\beta = 4$. The default parameters of the wireless radios are used, such that each node has a transmit range of $250m$.

Nodes move in the network under the random waypoint mobility model. We characterize the mobility using the parameter *max_speed* and keep the *pause time* equal to 1 second in all mobility scenarios. Each node has a velocity that is uniformly distributed between 0 and *max_speed*. Four different values of *max_speed* are investigated in this numerical study, namely $v = 0, 2, 10, 20$ m/s. These values correspond to the stationary, pedestrian, slow and fast vehicular scenarios.

The traffic is generated through a CBR application over UDP [1], [4]. This simulates the routing performance of the best effort delivery paradigm. The offered load could be varied by any of the three parameters, namely packet transmission rate, packet size and the number of traffic flows in the network. We simulate four offered loads regimes, as shown in Table I. The traffic types 1 to 4 correspond to the network operating in the light, normal, heavy, and saturation load regimes respectively. Packet sizes are chosen such that fragmentation occurs on neither the network nor the MAC layer.

Traffic type	packet rate	packet size	number of flows	total load
1	5	64Byte	20	51.2Kbps
2	10	64Byte	20	102.4Kbps
3	10	512Byte	20	819.2Kbps
4	20	768Byte	20	2.458Mbps

TABLE I
Traffic parameters adopted in the numerical studies

In this paper we use the *dynamic source routing* (DSR) routing algorithm in our simulations, since DSR shares many of the salient characteristics typical to reactive routing algorithms. The DSR runs on top of the 802.11b standard with a channel reservation mechanism enabled by the use of *request to send* (RTS) and *clear to send* packets. In general, packet loss can result from contention in wireless transmissions, unavailability of route due to mobility, or buffer overflow due to congestion. Nevertheless, the RTS/CTS mechanism in the 802.11b standard is efficient in combating the hidden terminal problem. Thus, most packet loss are due to mobility or congestion.

We have selected the network size such that network partitioning does not occur in any mobility scenario. Nevertheless, the DSR routing algorithm may fail to discover a route in heavy traffic regimes or high mobility scenarios. By convention, if there is no connection for a flow for the whole simulation, we define the goodput to be 0, and the delay and path length to be infinity.

Altogether we have four mobility scenarios and four offered load regimes. For each of the sixteen *network scenarios*, ten topology realizations are simulated. Each simulation lasts for $300s$. Each flow starts at a staggered time that is uniformly distributed between 0 and $100s$. Simulation data is logged during the interval between $100s$ and $300s$ to ensure the network has reached to a steady state. In each topology realization, 5 out of the 20 flows are monitored. Thus, for each network scenario, we have logged the data of 50 monitored flows.

IV. Simulation Results

A. Dependence of Path Length L on Speed

The time averaged path length L of each flow of each offered load regime is plotted in Figure 1. In each subgraph, the path length in each mobility scenario is plotted. The average path length of a flow is obtained by averaging the number of hops each packet traverses in the flow during the experiment. In the heavy traffic regimes of Figure 1(c)(d), we consistently observe that the path length decreases as mobility increases for each pattern number. The path length difference is as large as 4 hops for some pattern numbers. This explains the prevalence of short routes in high mobility scenarios of the heavy traffic regimes. A similar trend is observed in the non heavy traffic regimes of Figure 1(a)(b). The relative difference in path length in different mobility scenarios is smaller. In particular, for short path lengths (1-2 hops) the trend is reversed and mobility increases the path length slightly.

In the DSR protocol, each node continuously snoops into every packet it receives. The length of an existing route may be *shortened* or *lengthened* as time evolves due to mobility. Whenever a node along a route detects there is a shorter path to the destination node, a route change is triggered. Thus under higher mobility, route optimization is triggered more often, leading to

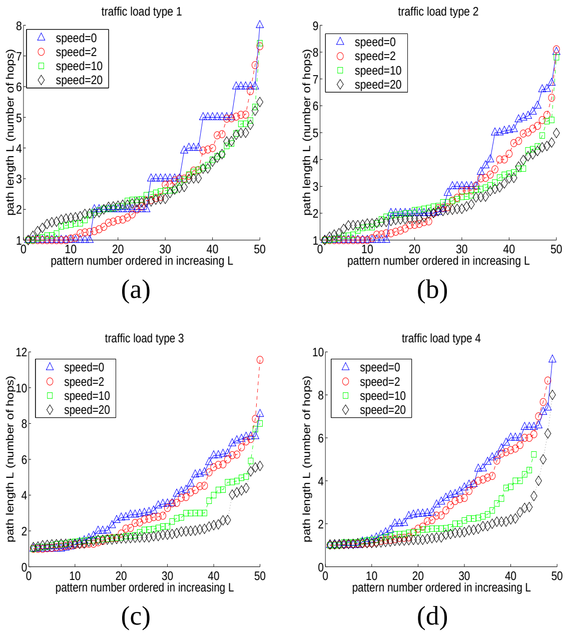

Fig. 1. Path length vs. pattern number in all mobility scenarios. (a)light offered load regime, (b)normal offered load regime, (c)heavy offered load regime, (d)network saturation regime.

shorter routes. However, for flows with a short path length of 1 or 2 hops, mobility usually *lengthens* a route. This explains that pattern 1 to 23 in Figure 1(a)(b) have longer routes in high mobility. In general, the path length of long routes (path length $L \geq 3$) decreases in mobility, whereas the path length of short routes increases in mobility. The variations of path length in high mobility scenarios is thus smaller.

In the heavy traffic regimes, there is a larger disparity of path length in different mobility scenarios. In particular we observe from Figure 1(c)(d) that the path length of 80% of the monitored flows is less than 3 hops at speed 20m/s. When the network is under the stress of *heavy traffic* and *high mobility*, only short routes (1 to 2 hops) are discovered during route discovery. In these regimes, the packet delay incurred at each hop is in the order of 10 seconds. In route discovery, if a *route request* (RREQ) packet traverses along a long route, the round trip delay is sufficiently long such that route discovery is aborted. Thus, in high mobility scenarios, the source and destination nodes are intermittently connected. When the nodes are in proximity, route discovery is successful; otherwise, route discovery fails. This explains the prevalence of short routes in high mobility scenarios of the heavy traffic regimes.

Incidentally, in the stationary scenario of the light offered regime of Figure 1(a), we observe that the time averaged path length of each flow is an integer. In general, route changes are due to mobility or congestion. Since there is no mobility and congestion in the stationary scenario of Figure 1(a), there are no route changes over the duration of simulation. Thus the time averaged path length of individual flow must be an integer. A staircase pattern is also observed in the stationary scenario of the normal offered load regime. In Figure 1(b), we observe that for pattern 1 to 35, the path length follows a staircase pattern. This shows that there are few route changes when the path length are short. For patterns 36 onwards, the staircase pattern disappears. This indicates that route changes due to congestion are common for these flows. We observe that these flows have path lengths of more than 4 hops. Thus we could also infer that flows with longer path lengths are more susceptible to route changes due to congestion.

B. Dependence of Path Length L on Node Distribution

Referring to Figure 1(a) the stationary scenario of the light offered load regime, there are 14 flows with a path length of 1 hop. As the path length L increases, the number of flows with path length L decreases. In this scenario, there is no congestion or route changes due to mobility. Consider a scenario in which an infinite number of nodes are uniformly distributed on a line of length 1. Suppose we define the distance between two arbitrary nodes as Z. It is straightforward to compute the probability distribution of Z, which is:

$$Pr[Z \leq z] = \int_0^1 \min(y+z,1) - \max(y-z,0)\, dy \quad (1)$$

$$= 2(1-z) \qquad z \in [0,1] \quad (2)$$

Equation 2 indicates that the prevalence of short routes is a direct consequence of uniform node distributions in the network. Since path length is roughly proportional to route distance, a route with short path length is more probable. The same trend of path length is also observed in other mobility scenarios and offered load regimes since all the path length curves L in Figure 1 are concave upwards.

The mean path length $E[Z]$ could be derived from Equation 2 to be $\frac{1}{3}$, which is one third of the network dimension. In our simulation, we have used a scenario size of 1500mX300m. Since each node has a nominal range of 250m, the scenario resembles a one dimensional network. The mean route length is thus 500m. Consider the stationary scenario in the light offered load regime. The mean path length of all 50 flows is found to be 2.95 hops. This agrees with our computations for a 1 dimension network. Most routes require a minimum path length of 3 hops to traverse a distance of 500m.

C. Improved Goodput G due to Load Balancing

In Figure 2, the goodput G of each offered load regime is plotted. In each subgraph, the goodput in all mobility scenarios is plotted versus pattern number. In the light offered load regime of Figure 2(a), we observe that mobility leads to a slight deterioration of goodput. Due to the light offered load, few packets in transit are lost due to buffer overflow in the node preceding a broken link. Most packets are queued in the node buffers during route maintenance. Although goodput degrades with mobility, the discrepancy is small because packet loss is uncommon.

Similarly, the goodput also deteriorates in the vehicular scenarios of the normal offered load regime of Figure 2(b). At normal load, packet loss is due to both mobility and congestion. During route failure, most packets in transit are lost due to buffer overflow. Thus, goodput is very sensitive to mobility. However, the goodput in the pedestrian scenario is higher than that in the stationary scenario. When nodes are stationary, packet loss is due to local congestion along some flows. For convenience, we define a flow as congested if the goodput of the flow is smaller than 0.8. Thus there are 9 congested flows in the stationary scenario, compared with 5 for the pedestrian scenario. This is consistent with observations in literature [1], [4], in which goodput

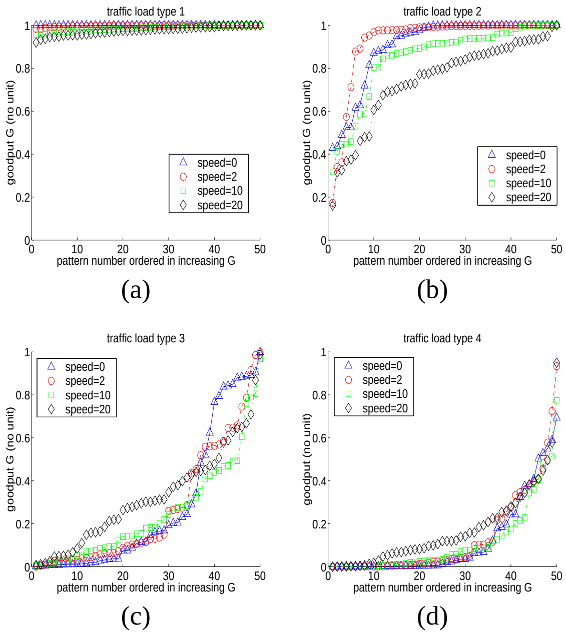

Fig. 2. Goodput vs. pattern number in all mobility scenarios. (a)light offered load regime, (b)normal offered load regime, (c) heavy offered load regime, (d)network saturation regime.

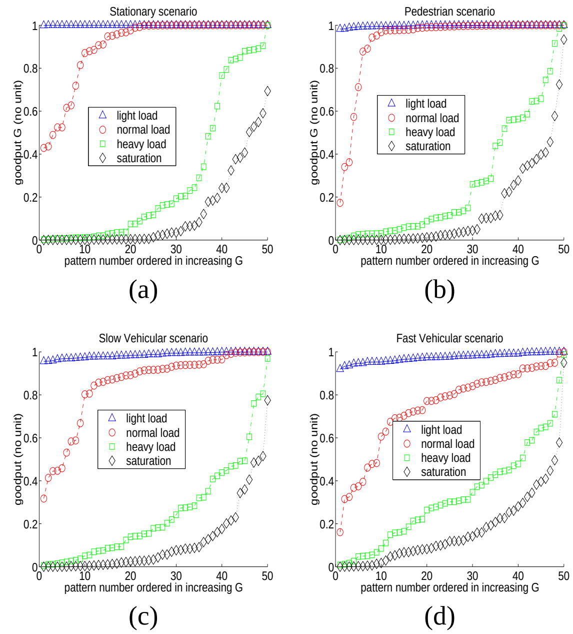

Fig. 3. Goodput vs. pattern number in all offered load regimes. (a)stationary scenario, (b)pedestrian scenario, (c)slow vehicular scenario, (d)fast vehicular scenario.

is shown to improve with speed. It is conjectured that the improved goodput in mobility is due to the load balancing effect. Some flows that pass through the congestion hot spots achieve improved goodput through rerouting brought about by node mobility.

D. Improved Goodput G due to Reduced Effective Load

The load balancing effect could correctly explain the goodput improvement with speed when localized congestion occurs. When we consider the heavy traffic regimes, the load balancing argument is inapplicable since network-wide congestion is experienced at all nodes. However, the goodput improvement with speed is still observed, albeit for a different network mechanism.

In the heavy traffic regimes of Figure 2(c)(d), we observe that many flows have better goodput in higher mobility. This is remarkably different from other traffic regimes of Figure 2(a)(b), where mobility degrades the goodput performance.

In heavy traffic, congestion is a network-wide phenomenon. All nodes are backlogged with packets. Rerouting due to mobility should not bring about any goodput improvement at all. We claim that the improved goodput in mobility is due to the decreased effective network load. As discussed earlier, most flows have shorter routes in high mobility. The total number of transmissions to forward a packet to the destination node is dramatically reduced. This effectively decreases the total network traffic, enabling a higher goodput for all flows. To see this, we consider the network saturation regime. We compare the total number of transmissions to deliver one packet for each monitored flow in each mobility scenario. Refer to Figure 1(d), by computing the area under the path length paths in Figure 1(d), we find the total number of transmissions to deliver 1 packet for each of the 50 flows is 171.9184 hops in the stationary scenario. The average path length for each flow is then 3.4384 hops. In the fast vehicular scenario, the total number transmission for 50 flows is 93.3077 hops. Thus the average path length for each flow is 1.8662 hops. Compared the normalized path length of each flow, the effective network load in the fast vehicular scenario is only 54.27% of that in the stationary scenario.

The corresponding number of delivered packets in a fixed duration is proportional to the sum of the areas under the goodput graphs in Figure 2(d) weighed by the packet rate. Suppose T_{intarr} is the mean packet interarrival time. In the stationary scenario, the expected number of delivered packets for 50 flows is 6.0541 packets. The normalized number of delivered packets for a flow in a time T_{intarr} is thus 0.1211. In the network saturation scenario, the expected number of delivered packets for 50 flows is 8.6093 packets. The normalized number of delivered packets for a flow in a time T_{intarr} is thus 0.1722. This amounts to a 42.21% increase in goodput.

Similarly, we also compare the total number of transmissions to deliver one packet for each monitored flow in the heavy offered load regime. Refer to Figure 1(c), the average number of transmissions is 3.6 hops per flow in the stationary scenario, and 2.0986 hops per flow in the fast vehicular scenario. The effective load in the fast vehicular scenario is only 58.29% of that in the stationary scenario. From the goodput graph in Figure 2(c), the expected number of delivered packets for a flow in a time T_{intarr} is 0.2880 in the stationary scenario. compared with 0.3277 in the fast vehicular scenario. This amounts to a modest 13.79% increase in goodput. In general, in high mobility scenarios, the reduction of effective network load has a more prominent effect to packet loss due to mobility. Thus we expect to obtain even better network goodput in higher mobility scenarios.

E. Determination of the Fraction of Congested Flows

In Figure 3, the goodput G of all mobility scenarios is plotted. In each subgraph, the goodput in all offered load regimes are plotted against the pattern number. We note that Figure 3 and

Figure 2 are plotted from the same results. Whereas the goodput in different mobility scenarios are compared in Figure 2, goodput in different offered load regimes are compared here in Figure 3.

Consider the stationary scenario in Figure 3(a). Recall that we defined a flow as congested if the goodput was less than 0.8. Thus, in the normal offered load regime, roughly 20% of all flows experience congestion. In the heavy offered load and the network saturation regimes, the fraction of congested flows is respectively 80% and 100%. Nevertheless, the above definition arbitrary. We show below that it is possible to classify the flows into the congested or uncongested regimes independent of a specific threshold.

Consider the light offered load regime of Figure 3. In this regime, packet loss is due to mobility in every flow. It is clear that the goodput curve for all patterns could be fitted by a straight line as a function of pattern number. Thus the variations of goodput could be modeled by an uniform distribution with some mean and variance. As mobility increases, the mean goodput drops slightly and the variance increases.

More generally, packet loss is due to a combination of mobility and congestion. Consider the normal offered load regime in the following. In the slow vehicular scenario of Figure 3(c), the goodput curve could be fitted into a piecewise linear function. From patterns 1 to 12, the goodput variations with pattern number have a steeper slope. For other patterns, the goodput variations follow a more gentle slope. The observation of two regimes could be explained by the cause of packet loss. For patterns 11 to 50, packet loss is dominated by mobility. Similar to the light offered load regime, the goodput variations of each flow could be modeled by an uniform distribution. For patterns 1 to 10, congestion dominates over mobility. Congestion is more detrimental to goodput than mobility. Also, there are more variations in goodput depending on the extent of congestion. The goodput for the congested flows is thus modeled by an uniform distribution with a smaller mean and larger variance. Similarly, the goodput variations in the fast vehicular scenario of Figure 3(d) could also be fitted by a piecewise linear function. By observing the intersection of the fitted lines, we infer that patterns 1 to 11 is in the regime where congestion is the main reason for packet loss.

In the heavy offered load and the network saturation regimes, the goodput could not be fitted nicely by a piecewise linear function. In these regimes, congestion is a network-wide phenomenon. Thus there are flows in which goodput deterioration due to mobility and congestion is both prominent.

In our simulations, the delay in uncongested flows typically does not exceed 1 second. The corresponding delay in congested flows is in the order of 10 seconds. Most applications in ad hoc network today have tight delay constraints. It is highly undesirable to deliver packets over some flows that experience congestion. The fraction of congested flows that is therefore an important performance metric to consider.

F. Dependence of Fairness on Offered Load, Speed and Path Length L

In homogeneous ad hoc networks, all nodes are peers and they cooperate to forward packets for each other. Applications in military and rescue operations fall into this context. In these networks, it is desirable for each flow to attain the same goodput independent of the offered load, mobility and path length.

Referring to Figure 3, we observe that fairness deteriorates quickly with the offered load in each mobility scenario. When the offered load is light, all flows have a goodput close to 1 irrespective of mobility. In the normal offered load regime, local congestion occurs for some flows. We discussed earlier that the variable goodput could be modeled by a uniform distribution with some mean and variance depending on the cause of packet loss. If packet loss is dominated by congestion, the goodput exhibits randomness with a larger variance. Thus, in general, fairness is very sensitive to the presence of congestion. In particular, when congestion dominates in the heavy traffic regimes, there is a high asymmetry in the goodput of the monitored flows.

We also observe the sensitivity of fairness to offered load decreases in higher mobility. This is easily visualized by comparing the stationary and the fast vehicular scenarios in Figure 3(a)(d). At high mobility the goodput curves of different load regimes are more closely spaced. Heavy load regimes have improved goodput due to the decreased effective network load while light load regimes suffer from packet loss due to mobility.

Whereas fairness is sensitive to the offered load, it is insensitive to mobility. Refer to Figure 2(a), all flows have a goodput close to 1 in all mobility scenarios. The goodput is insensitive to mobility in the light offered load regime. Similarly, in the normal offered load regime of Figure 2(b), packet loss due to mobility leads to a slight deterioration of goodput. In the heavy traffic regimes, the reduction of effective network load in high mobility scenarios leads to a slight improvement in goodput. Thus, for each offered load regime, the disparity of goodput among all flows in different mobility scenarios is small. This is intuitively plausible since fairness depends on the resource allocation of each flow in the network. Although mobility impacts the efficiency of route discovery, the resource allocation of each flow does not depend on mobility once a route is found. Thus, in all node mobility of interest in this study, mobility does not impact the fairness. In extremely high mobility, however, route discovery is inefficient. The flow will be intermittently connected, leading to poor goodput performance.

We have argued that fairness is intimatedly connected to the resource allocation in each flow. Since the path length of a route determines the amount of forwarding and thus the resource requirement of a flow, it is instructive to investigate the dependence of fairness to path length L. Refer to Figure 4, the goodput G for each mobility scenario is plotted. In each subgraph, the goodput of all patterns in all offered load regimes are plotted ordered in increasing path length. This allows us to inspect the effect of path length L on fairness in each offered load regime and mobility scenario.

As shown in the figure for all mobility scenarios, we observe some correlation between the path length L and goodput G. When the path length is long, the goodput is likely to be smaller. Consider the light offered load regime. Although routes with long path lengths have lower goodput, fairness is not a problem in this regime since the goodput of the worst flow in each mobility scenario is close to 1. Consider the normal offered load regime. In the stationary and pedestrian scenarios of Figure 4(a)(b), only the flows with long path lengths have small goodput. In these scenarios, packet loss is dominated by congestion. We could infer that long routes are more susceptible to congestion. This is intuitively reasonable since it is more likely to route through some local congestion hot spots if the path length is long. In the slow and fast vehicular scenarios, we also observe that goodput is likely to be smaller for longer path lengths.

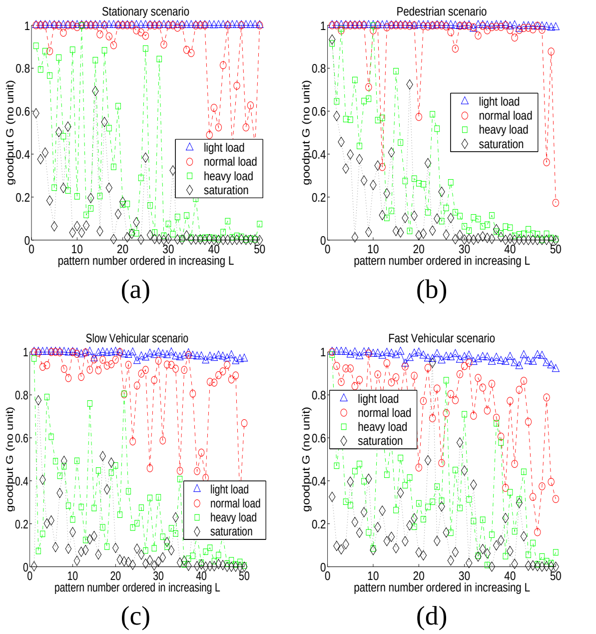

Fig. 4. Goodput vs. pattern number in all offered load regimes. (a)stationary scenario, (b)pedestrian scenario, (c)slow vehicular scenario, (d)fast vehicular scenario.

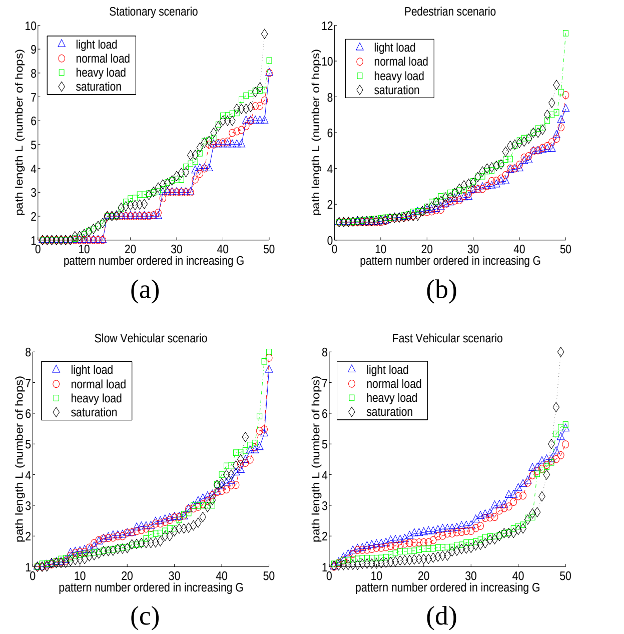

Fig. 5. Path length vs. pattern number in all traffic regimes. (a)stationary scenario, (b)pedestrian scenario, (c)slow vehicular scenario, (d)fast vehicular scenario.

However, not all flows with long path lengths have small goodput. This could be explained by load balancing effect due to path rerouting for these flows.

Consider the heavy offered load and network saturation regimes. In these regimes, long routes are shut down completely. With reference to Figure 4(a)(b), 40% of all flows with the longest path length have a goodput close to 0. In the slow and fast vehicular scenarios the fraction of flows that are shut down decreases to 20% and 12% respectively due to the decreased effective network offered load. In general, flows with long path length are shut down completely in heavy traffic regimes. Only local communication is possible. We conclude that fairness is sensitive to the path length when the offered load is large.

G. *Dependence of Path Length L on Offered Load*

In Figure 5, the path length L of all mobility scenarios is plotted. In each subgraph, the path length of all patterns in all offered load regimes are plotted. These are the same results of Figure 1. Path length in different offered load regimes are compared here, whereas path length in different mobility scenarios are compared in Figure 1.

Consider the stationary and pedestrian scenarios of Figure 5(a)(b). We observe that L generally increases when the offered load increases. This phenomenon implies that at high traffic intensity, either routes with long range could not be used, or successful transmissions are limited to small ranges only. The latter argument is flawed since the RTS/CTS channel reservation mechanism in 802.11 is effective in resolving contention, even if the network is under the stress of excessive traffic. The paradox could be explained as follows. The RTS/CTS mechanism in 802.11 effectively prevents collisions of data packets. Route request (RREQ) packets however, can't use the RTS and CTS mechanisms since they are broadcast packets. If heavy congestion occurs, collisions between RREQ packets are likely. Since collisions are more likely to occur for longer hops, RREQ packets may never reach the destination node if a route consists of hops with large distance progress. Only routes with large number of hops and small per hop distance progress are discovered in route discovery. In general, as the offered load increases, congestion is prominent and may impede the transmissions of the broadcast RTS and CTS packets. Thus path lengths are longer due to congestion.

Consider the slow and fast vehicular scenarios of Figure 5(c)(d). At high mobility, the trend of our observations is reversed. We observe as offered load increases, the path length decreases. In high mobility scenarios, mobility has a more significant impact compared to congestion. Recall in our discussion for path length L vs. speed. At high mobility, the path length will decrease more drastically for heavy load regimes. In general, congestion and mobility have opposite effects on path length. At high mobility, the effect of mobility dominates. Therefore, path length is smaller for higher offered load.

References

[1] J. Broch, D.A. Maltz, D.B. Johnson, Y.C. Hu, J. Jetcheva, “A performance comparison of multi-hop wireless ad hoc network routing protocols,” Proc. Mobicom, 1998, pp.85-97.

[2] G. Holland, N. Vaidya, “Analysis of TCP performance over mobile ad hoc networks,” Proceedings of the fifth annual ACM/IEEE international conference on Mobile computing and networking August 1999.

[3] D.B. Johnson, D.A. Maltz, “Dynamic Source Routing in Ad Hoc Wireless Networks,” Mobile Computing, Chapter 5, pp.153-181. Kluwer Academic Publishers, 1996.

[4] C.E. Perkins, E.M. Royer, S.R. Das, M.K. Marina, “Performance comparison of two on-demand routing protocols for ad hoc networks,” IEEE Personal Communications , Volume: 8 Issue: 1 , Feb. 2001, pp.16-28.

[5] E.M. Royer, C.-K. Toh “A review of current routing protocols for ad hoc mobile wireless networks”, IEEE Personal Communications, Volume: 6 2 , pp.46-55, 1999.

[6] Network Simulator Notes and Documentation, UCB/LBNL, http://www.isi.edu/nsnam/ns/

[7] The Monarch Project of Carnegie Mellon Univ., http://monarch.cs.cmu.edu/