

Impact of Transmission Power on TCP Performance in Vehicular Ad Hoc Networks

Andrew Chen¹, Behrooz Khorashadi², Dipak Ghosal³, Chen-Nee Chuah⁴
apchen@ucdavis.edu¹, bkhorashadi@ucdavis.edu², dghosal@ucdavis.edu³, chuah@ucdavis.edu⁴
Department of Computer Science^{2,3}, Department of Electrical & Computer Engineering^{1,4}
University California, Davis

Abstract— Vehicular ad hoc networks (VANET) are characterized by high velocities and relatively uniform mobility of vehicles. As a result, the dynamics of VANET are different from generic mobile ad hoc networks. The behavior of vehicles and their effect on the protocol performance has not been comprehensively investigated in the context of VANET. In this work, we analyze the performance of multi-hop TCP in a multi-lane highway environment where vehicles are configured as clients or routers trying to reach a fixed access point. In particular, we look at the effect of tuning transmission power in dense and sparse road scenarios on the TCP throughput and latency. To conduct our study, we have developed an integrated simulation tool that accurately models both the vehicular mobility and the networking protocol stack. Using our simulation tool we obtain detailed TCP statistics and correlate them with the location, velocity, and other properties of VANET nodes. We study scenarios with a single sender as well as multiple senders. In each case, results show that throughput significantly deteriorates as the number of hops from the sender to the access point increases. A higher number of hops results in higher losses due to interference and loss of connectivity, which causes TCP to throttle back the sending rate unnecessarily.

I. INTRODUCTION

The focus of this paper is to understand the dynamics of TCP with respect to the transmission power in a vehicular ad hoc networks (VANET). We examine the scenario where vehicle and roadside access points communicate through persistent multi-hop connections. Vehicles are equipped with storage and processing devices that enable them to store and process data from other vehicles and access points. Enabled vehicles can then use on-board devices to form on-demand ad hoc networks which can be used to provide a variety of useful services. This network can provide infotainment function such as in-vehicle Internet access or enable inter-vehicular peer-to-peer applications.

Many of these applications require the use of a reliable end-to-end transport protocol, the most popular of which is TCP. While there have been proposals to improve or replace TCP in wireless environments, the fact that it is so widely implemented in existing devices means that it will remain in use for quite some time. Unfortunately, studies have shown that TCP performance in MANET/VANET is very poor. Inherently, TCP treats packet loss as a sign of network congestion. In wireless networks, however, losses can occur for reasons besides congestion, such as poor channel conditions or collisions. As a result of a loss, TCP will unnecessarily reduce

congestion window size, effectively decreasing throughput. In addition, because of the ad hoc nature of VANET network, path symmetry is not guaranteed. This can also decrease TCP performance due to incorrect estimation of RTT times, resulting in unnecessary retransmissions.

In response to the shortcomings of TCP over wireless channels, many solutions have been proposed. The first category of solutions are TCP variants, which have been proposed to achieve more efficient transport in wireless networks. Second is the class of entirely new transport protocols developed to replace TCP, which specifically address the needs of wireless environments. We chose to focus, not on modifying the transport layer, but rather understanding how TCP performance can be enhanced through tuning transmission power. We look at the effects of tuning transmission power not only on TCP throughput but also its effect on TCP packet latency and congestion window size. We investigate the issue using a detailed simulation tool which accurately models both the vehicular traffic and the wireless network characteristics and the networking protocols.

In this paper we show that tuning transmission power can be a viable option in enhancing TCP performance and can also be used to increase throughput and optimize spatial reuse in VANET environments.

The remainder of this of this paper will be laid out as follows: Section 2 will discuss the Problem Statement followed by the simulation tool design and experiment setup in section 3. Results of our simulation will be presented in Section 4. Section 5 contains the conclusion and references.

II. PROBLEM STATEMENT

With the growing demand for convenient wireless applications on-the-go, significant effort has been put into developing vehicular-based applications. Vehicles equipped with commodity 802.11 wireless network cards can form an ad-hoc wireless network, referred to as VANET. These networks can be leveraged to perform a variety of services, for example, traffic applications such as accident alerts that warn drivers of upcoming obstructions in the road or traffic control mechanisms that regulate traffic and reduce congestion. In addition to safety and traffic applications, VANETs can be used to provide users with infotainment applications such as on-board Internet access or a platform for peer-to-peer file distribution networks. An example of this is the proposed

CarTorrent [1], which provides vehicles with a single hop and multi-hop communication network for the sharing of content such as music and videos. Vehicles can also form multi-hop mesh networks in conjunction with roadside access points in order to provide Internet access.

Many of these applications require the use of a transport layer protocol that guarantees delivery of packets. Although TCP achieves good throughput in wired networks, its performance in wireless medium has shown to be poor because TCP cannot differentiate between packets lost due to congestion and packets lost due to temporary channel conditions. The problem is exacerbated in highly mobile networks, such as vehicular mesh networks, because of the fluctuations in source to destination paths [2].

Studies such as [3] examine the performance of both UDP and TCP in vehicular ad-hoc networks and shows that TCP performance is poor in highway scenarios. A detailed study in [2] was conducted to compare the different aspects of the transport layer performance in wireless mesh environments. In general, two classes of transport layer protocols have been proposed. The first class of new transport layer protocols are TCP variants and secondly are entirely new transport layer protocols designed specifically for wireless mesh networks (WMN). The fundamental drawback of entirely new protocols such as ATP proposed in [4] is that, although they address many of the fundamental drawbacks of TCP, they inherently are not compatible with TCP. Since ATP is not TCP-compatible, it will only function in stand-alone WMNs and must be translated or proxied if communication with a TCP network is required. The alternative is TCP-variants. There are 3 primary drawbacks of TCP: 1) the standard TCP protocol cannot distinguish between congestion and interference, which is the primary reason for TCP performance degradation in wireless environments. During temporary poor channel conditions, TCP will back off, assuming a queue along the path is full. Once wireless channels are back to normal conditions, standard TCP has no mechanism to recover quickly. 2) TCP also has no mechanism to distinguish between congestion and link failure. In wireless networks, link failure is not as critical as in wired scenarios since a new link can be established fairly quickly. However, because of mobility, link failures can occur more often. 3) Finally TCP is highly-dependent on timely delivery of ACKs and is severely impacted by the asymmetry of WMNs. Because of their dynamic nature, forward paths and reverse paths can greatly differ. This causes significant problems with TCP since ACKs on one reverse path may experience much different packetloss rates, latency, and bandwidth.

Studies have shown that TCP performance over a wireless mobile network can fall to as low as 5% of the possible utilization [5]. As previously mentioned, many TCP-variants have been proposed to address these issues. These variations can be categorized into four different categories: 1. Splitting TCP connections; 2) Snooping TCP at base station [6]; 3) Notifying Cause of packet loss (ELN- explicit loss notification) [7]; and 4) Adding selective ACKs to TCP. Each TCP modification brings improvements to TCP in different scenarios but each,

in-turn, also has its drawbacks.

Splitting TCP connections enhances TCP throughput but unfortunately violates the end-to-end argument. This is because of the fact that an acknowledgment sent by the wireless gateway can reach the sender before the packet corresponding to that ACK has reached its destination. This can cause problems since a server crashing before the packet is received will cause the sender to think that the destination has really received the data packet when it actually has not. Snooping benefits are only seen in a single direction, mainly from wired to wireless environments. This is because the snooping agent caches data packets and suppresses duplicate ACKs at the wireless interface. ELN also works well but suffers from slow error recovery. Packets can only be retransmitted after the round trip time has elapsed when an acknowledgment with the notification bit is set. A more comprehensive explanation of ELN drawbacks can be found in [8]. Finally, adding selective acknowledgments also shows significant improvement in TCP throughput and can be used in conjunction with Split-TCP or end-to-end. Unfortunately recovery in the end to end systems remains very slow in the cases where paths have high delay.

All of the previously mentioned improvements focus on the network layer (studying routing layer protocols and ways to improve them) or the transport layer (improving TCP or suggesting new transport layer protocols). Instead we chose to focus on the effects and benefits of modulating transmission power. By studying the effect of tuning transmission power and its effect on TCP we may be able to find an optimal setting for transmission power for vehicular ad-hoc mobile networks. Since wireless 802.11 devices are capable of transmitting using different power levels, we propose to analyze the local and global effects of transmission power modulation on the TCP protocol in a wireless mobile vehicular network. In a vehicular mesh network, vehicles all transmitting with high power can cause a high degree of interference. On the other hand, reducing power will cause less channel contention and interference but will result in an increase in the number of hops between vehicle and access point, which also has a dramatic impact on vehicle throughput.

III. SIMULATION TOOL AND EXPERIMENT SETUP

In order to accurately simulate both vehicular movement and a wireless network environment, we developed a comprehensive simulation tool. Our simulator builds upon the SWANS network simulator and the JiST (Java in Simulation Time) discrete event driven simulation tool [9]. Vehicular traffic flow is based on the cellular automata model implemented by the Nagel and Schreckenberg model [10]. The simulation granularity is $\Delta t = 1s$, i.e., the new vehicle positions are calculated using the N-S model every second. We have extended the basic N-S model to more accurately reflect real-world traffic based on the work in [4] and by adding lane-changing capability.

In this paper we focus solely on the TCP protocol. To accomplish this, we modified the network stack in SWANS in order to extract network statistics. SWANS also provides routing protocols such as AODV and DSR. However, we




Fig. 1. A single source vehicle communicates with a fixed AP, using other vehicles as routers. The road is divided into bins consisting of 100 cells each.

choose to implement a position-based routing protocol based on [11], which leverages the positioning information provided by GPS devices that are already present in vehicles. SWANS also implements the 802.11 MAC layer protocol and several path loss and fading models.

For our simulation, we consider a 4-lane, 1.5km road (divided into 1000 cells) with a fixed access point (AP) located at cell 500. Each vehicle occupies 5 cells and the road is also divided into 10 "bins" of 100 cells each, which we will use to correlate network statistics with vehicle position. This is illustrated in Fig. 1.

We examine two scenarios: single-source and multi-source (two and three vehicles). In the first case, a single source vehicle establishes a TCP connection with the AP and continuously sends data until it exits the road, at which point another source vehicle enters the road. We choose to examine this scenario in order to simplify the problem, allowing us to obtain the best-case TCP performance. In the second scenario, we repeat the experiment but with two and three vehicles spaced approximately evenly over the road.

The goal of this study and quantify the effects of transmission on the TCP transport protocol. The key parameters analyzed in the simulation are as follows:

- Traffic Density: The percentage of cells occupied by vehicles over the total number of cells.
- Transmission Range: The maximum distance two nodes are able to communicate in optimal channel conditions, i.e. no background noise or interference.
- Vehicle position: The location of vehicles on the road. The road is divided into "bins" of 100 cells (150 meters) each, with cell and bin 0 starting at the beginning of the road.

TABLE I
VEHICLE AND TRAFFIC PARAMETERS

Traffic Density	10% and 30%
Transmission Range	100 – 1000 meters
Maximum Speed	108 kph

A. Position Based Routing

The position-based routing is based on the greedy forwarding concept from GPSR [11], in which packets are forwarded through nodes geographically closer to the destination than previous node. That is, the position of the next hop will

TABLE II
NETWORK PARAMETERS

Parameter	Description	Value
<i>BEACON_INTERVAL</i>	Delay between position beacons	1 second
<i>RETRANSMIT_DELAY</i>	Delay between forwarding attempts for a packet	1 second
<i>MESSAGE_TTL</i>	Number of times a node will attempt to forward a packet	5
<i>INFO_TIMEOUT</i>	Lifetime of a position beacon message	5 seconds

always be closer to the destination node than that of the current hop. We do not consider the "perimeter routing" mode of GPSR that searches for alternate routes that may not be geographically closer, since in a highway scenario, the width of the road is generally smaller than the transmission range. In this scenario, there is no way for a route to move away from the destination and still find its way back. Therefore, if a node cannot find a route to the destination, it will simply queue the message and check every *RETRANSMIT_DELAY* seconds to see if a route has been found. If this fails *MESSAGE_TTL* times, the message is dropped. The operation of position-based routing depends on the timely dissemination of the location of nodes in the network. This is accomplished through periodic beacons. Every *BEACON_INTERVAL* seconds, every vehicle will broadcast its own location to its neighbors. Using this information, all vehicles in range will be able to build a map of one-hop neighbors. Since vehicles can leave the road, nodes are removed from this map if no update has been received in more than *INFO_TIMEOUT* seconds.

In our simulation, we assume that the location of the access point is known a priori. When a vehicle wishes to send to the access point, it will first add an additional routing header to each outgoing packet with source position. This enables the destination of data packets to learn the location of the source node and can send return messages. The position encoding in packet headers allows for position-based routing to work in the reverse direction without requiring every node to have a complete map of the entire road.

IV. RESULTS

A. Single Source

In order to gauge TCP performance, we look at three metrics. TCP throughput is our primary concern but we also look at latency and the total number of collisions in the system as a whole. TCP throughput is collected by averaging the throughput of all source vehicles over several runs. This is done across the entire road as well as correlated to road position for each transmission range. The same analysis is done for packet latency as well.

Fig. 2 shows the TCP throughput as a function of the transmission range for different vehicular traffic densities. We find that the lower density scenario outperforms the high-density scenario at all transmission ranges. Even at 100m transmission range and 10% vehicle density, there is sufficient connectivity to allow successful communication. This is




Fig. 2. Average throughput from source to destination at 10% and 30% vehicle densities with single source.




Fig. 3. Number of collisions observed for single source 10% and 30% scenario.

verified by looking at the number of routing drops, which we define as packets lost due to the vehicle not having a route. For all combinations of density and transmission range zero packets are dropped due to routing failure. This means that vehicles never lose route paths. We expect there is a crossover point at very low transmission ranges where the higher density case would offer greater connectivity. However the purpose of this study was not finding this crossover point of connectivity. We consider 100m a reasonable minimum range for common wireless hardware.

Furthermore, even though vehicle speeds are faster at low densities, the relative position of vehicles to other vehicles does not change often since everyone is moving at approximately the same speed. This implies that, in free-flow situations, higher vehicle speed may not have a significant effect on connection lifetime. This may not be true for highly-congested highways where stop-and-go traffic can create changes in topology often, but in those cases, we expect density to be high and speeds to be lower.

The throughput performance gap between the two road densities is due to the increase in number of packet collisions at high densities. Increasing traffic density from 10% to 30% increases the number of vehicles by three-fold. This, in turn, increases the number of beaconing packets by three times. Recall that although there is only one source vehicle, the routing protocol employed by all vehicles requires periodic broadcasts of their current location. As a result, the number of collisions increases by more than an order of magnitude as seen in Fig. 3. The 30% density line also begins to diverge from the 10% line as transmission range increases, again due to the increasing number of collisions. Note also that the 10% density curve is smoother than the 30% density curve. The variance of the high-density results is significantly higher than that of the low-density, primarily due to the randomness of the MAC-layer back-off algorithms. Fig. 4, the latency plot shows that there is little difference between densities at high ranges since the number of hops requires to reach the




Fig. 4. Average latency from source to destination at 10% and 30% vehicle densities for single source.

AP depends only on the distance to the AP, provided there is sufficient connectivity (which we have shown there is). At low ranges, however, there is a marked change. In both scenarios, latency sharply increases at 100m and 200m. This is due to the increase in hop count, which increases delay due to processing and queuing at each router. For these reasons, we note high variance in the latency at low ranges. We will also see later that the number of hops between source and destination play a significant role in determining throughput. In Fig. 5, we correlate throughput to vehicle positions on the road. We consider several transmission ranges at a fixed 30% density. Intuitively, as the source approaches the AP, throughput should increase as the number of hops decreases, leading to a bell-like distribution. This does occur at the 100m range, with approximately 0 throughput at the edges of the road and a maximum of nearly 200KBps directly adjacent to the AP. However, as range increases, we see that throughput plateaus at approximately 250KBps or 2Mbps. While this is far below the maximum of 11Mbps, the MAC layer and TCP




Fig. 5. Throughput for 30% density, correlated to vehicle position on the road. Position 0 in the plot corresponds to packets sent by the source when it is in cells 0-99(0-150 meters); position 100 corresponds to cells 100-199(150-300 meters).




Fig. 7. Congestion window size in relationship to vehicle position for 30% density.




Fig. 6. Congestion window size in relationship to vehicle position for 10% density.

overhead significantly reduce the actual maximum throughput. Increasing transmission range further only widens the plateau.

In order to better understand the dynamics of the TCP protocol we also choose to observe the fluctuations in TCP window size with respect to vehicle position. The graphs for both 30% and 10% densities are shown in Fig. 6- 7. For the 100m range, there is a sharp peak in the window size at around 750 meters corresponding to the peak throughput. At this range, there is little interference from surrounding vehicle that would cause packets to be lost, allowing the congestion avoidance algorithm to ramp up the window size very high. As expected the congestion window size closely follows TCP throughput. The graphs also show the role congestion plays on TCP window sizes. We found that the congestion window size for ranges 300-700 was at maximum throughout the duration of the simulation for 10% density. Looking at the window size for 30% density we see that window size for 300 meter range

reaches maximum size only after 650 meters which is only within 100 meters of the AP. Similarly the congestion window size begins to decrease much more quickly (after 1000 meter marker) than in the 10% case.

B. Multi-Source (2 Vehicles and 3 Vehicles)

In the multi-source scenario we found introducing additional vehicles has no significant impact on TCP performance. Fig. 8 illustrates system throughput is the very similar for all scenarios at high transmission ranges. At lower transmission ranges the multi-source scenario achieves a higher system throughput. This is a result of the multiple vehicles being approximately evenly distributed over the road. In the multi-source scenario, at any given time, it is more likely that a vehicle is close to the access point. Vehicles closer to the access point require fewer hops and, as we've shown, can achieve a higher throughput. This is the primary reason that the multi-source scenario achieves a higher throughput at lower transmission ranges. At around the 900 meter range we begin to see the system throughput saturate at 250 KBps.

Intuitively, based on observations from the single sources scenario, we expected that throughput would increase with transmission range. We also expected a diminishing gain in throughput with multiple sources due to an increase in interference between the sources. This expectation was not borne out in the experiments. Results show that, at higher transmission ranges for 2 and 3 vehicle scenarios, system throughput was nearly identical to single source. We attribute this to the congestion avoidance algorithms inherent to the 802.11 MAC layer. As transmission range increases all source vehicles are able to sense each other's transmission. As a result collisions are avoided and the channel is shared between the contending source vehicles. Each vehicle transmits with throughput proportional to the number of sources. With 2 vehicles each source vehicle transmits with $250\text{KBps}/2=125\text{Kbps}$ and with 3 vehicles $250\text{KBps}/3=83.3\text{KBps}$, where 250KBps is the max throughput we found during simulation of the single




Fig. 8. System throughput of both single source multi-source with respect to transmission range in 10% traffic density scenario.




Fig. 10. Throughput of 2-source scenario at various transmission ranges. Vehicle density for this graph is 10%.




Fig. 9. Congestion window size as a function of time for a one of the 2 vehicle in the multi-source scenario. Vehicle density is 10% and transmission range is 100 meters.

source case.

Fig. 9 plots the congestion window size versus time. At the edges of the road, the congestion window size never grows very large, even though there may be sufficient bandwidth available. The graph shows TCPs inability to increase window size to a significant level as packet drops due to poor channel condition cause to window to cut back transmission rates. This is a commonly known problem with TCP performance. However, we choose to include this result in order to validate that VANET also suffers from this as well.

Interestingly, our results for multi-source show that TCP throughput near the access point is higher for lower transmission ranges. Fig. 10 and 11 show that within 250 meters of the access point, both 100 and 300 meter transmission ranges achieve higher throughput rates than 500 or 700 meter transmission ranges. Higher throughput is achieved with lower transmission ranges near the access point because at low transmission ranges the source vehicles farther away from the AP




Fig. 11. Throughput of 3-source scenario at various transmission ranges. Vehicle density for this graph is 10%.

require multiple hops to reach the AP. The instability caused from the multi-hop communication restricts TCP throughput causing those vehicles to transmit at very low rates. This leaves the channel free for vehicles closer to the AP. In contrast, at high power all vehicles are within one hop of the access point and will transmit at their maximum rates, which results in the source vehicles all contending for the channel.

V. CONCLUSION

Our study found that TCP performance is highly susceptible to the number of hops between source and destination. In order to maximize TCP throughput vehicles should reduce the number of hops to a minimum. This can be achieved through increasing transmission power to the maximum. We also found that interference caused by increasing transmission range does not play a significant role in system throughput due to congestion avoidance algorithms in 802.11 MAC layer protocol. As a result, in such a scenario, vehicles should set their individual transmission power to the maximum in order

to maximize both individual and system throughput regardless of what other vehicles are doing. However, this applies only to scenarios with a few source vehicles.

VI. FUTURE WORK

In the future we plan to expand the experiment setup to include additional source vehicles. Currently we look at the effects of no more than 3 vehicles. With significantly more vehicles in contention for the channel, there will be a point at which the 802.11 collision avoidance algorithm breaks is unable to fairly share the channel. We would then see more transmission collisions in the system.

We also plan to look at TCP behavior with specific wireless applications. In this study our focus was vehicle to roadside node communication. As part of our future work we plan to look at inter-vehicular or vehicle-to-vehicle communication and how to modulate transmission power dynamically to maximize system throughput.

REFERENCES

- [1] S. D. A. Nandan, "Cooperative downloading in vehicular ad hoc networks," *In Proc. of Wireless On-Demand Network and Services(WONS)*, 2005.
- [2] W. W. I. Akyildiz, X. Wang, "Wireless mesh networks: A survey," *Science Direct*, pp. 445–487, 2005.
- [3] P. M. Fay Hui, "Experimental characterization of multi-hop communications in vehicular ad hoc network," *Vanet*, 2005.
- [4] H.-Y. H. E. K. Sundaresan, V. Anantharaman, "ATP: A reliable transport protocol for adhoc networks," *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, pp. pp. 64–75, 2003.
- [5] S. L. Zhenghua Fu, Xiaoqiao Meng, "How bad TCP can perform in mobile ad hoc networks," *Proceedings of the Seventh International Symposium on Computers and Communications*, 2002.
- [6] R. H. K. H. Balakrishnan, S. Seshan, "Improving reliabletransport and handoff performance in cellular wireless networks," *ACM Wireless Networks*, 1995.
- [7] H. Balakrishnan and R. Katz, "Explicit loss notification and wireless web performance," *Proceedings of the IEEE Globecom Internet Mini-Conference*, 1998.
- [8] A. J. Wenqing Ding, "Delay performance of the new explicit loss notification TCP technique for wireless networks," *IEEE*, 2001.
- [9] "Jist / swans. <http://jist.ece.cornell.edu/>."
- [10] K. Nagel and M. Schreckenberg, "A cellular automaton model for freeway traffic," *Journal de Physique*, vol. 2, pp. 2221–2229, 1992.
- [11] H. T. K. Brad Karp, "GPSR: Greedy perimeter stateless routing for wireless networks," *MOBICOM*, 2000.