Analysis of Path Characteristics and Transport Protocol Design in Vehicular Ad Hoc Networks

Ralf Schmitz*‡, Alain Leiggener†, Andreas Festag*, Lars Eggert† and Wolfgang Effelsberg‡
*NEC Deutschland GmbH, Email: {schmitz|festag}@netlab.nec.de
†NEC Europe Ltd., Email: {leiggener|eggert}@netlab.nec.de
‡University of Mannheim, Email: {effelsberg|schmitz}@informatik.uni-mannheim.de

***Abstract*— Vehicular *ad hoc* networks (VANETs) enable new applications by providing self-organizing vehicle-to-vehicle and vehicle-to-roadside communication. Some of these applications require reliable, in-order data delivery across end-to-end connections. The performance of a vehicular transport protocol (VTP) that provides such a service depends on its ability to adapt quickly to the varying path characteristics of highly dynamic environments. This paper studies path characteristics of VANETs in highway scenarios. An analytical evaluation derives upper bounds on the expected connectivity and disruption duration. Simulations validate these results and study further metrics, such as packet loss, packet reordering and round trip times. The paper also briefly outlines a preliminary VTP design that incorporates these analytical results.**

I. Introduction

Vehicular *ad hoc* networks (VANETs) support vehicle-to-vehicle and vehicle-to-roadside communication by providing a self-organized, wireless, multi-hop, *ad hoc* network. The *Network on Wheels (NoW)* project investigates key technical questions for VANET communication, including transport protocols or position-based routing [1].

VANETs enable a variety of new applications. They fall into two categories, broadcast applications, such as active road traffic safety or forecast services, and unicast applications, such as media transmission or email. Many unicast applications require reliable, in-order data delivery, similar to the service provided by TCP in the Internet [2]. TCP, however, performs poorly in wireless networks that have a high degree of mobility and frequent topology changes [3][4][5].

This paper investigates the path characteristics that transport protocols experience in VANETs in highway scenarios. These findings aid the design of a vehicular transport protocol (VTP). The behavior and performance of a VTP mainly depends on its ability to adapt quickly to varying path characteristics. Analytical and simulative evaluations of connectivity and disruption durations estimate the expected connectivity between communication partners for specific distances. Furthermore, the results quantify packet losses, reorderings, round trip times (RTT) and RTT jitter through simulations.

Finally, the paper presents a preliminary design of a VTP. Its key features are *(i)* utilization of statistical path characteristics for error and congestion control, *(ii)* decoupling of error and congestion control and *(iii)* congestion control via explicit signaling of available bandwidth information by intermediate hops.

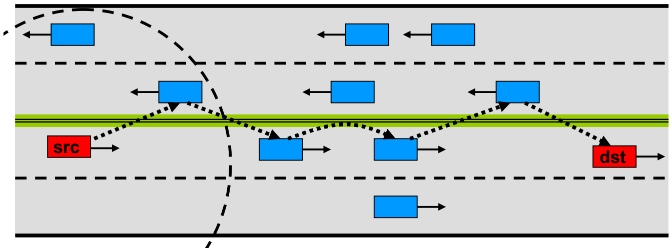

Fig. 1. Multi-hop inter-vehicle communication in the highway scenario.

Section II defines the scenario and metrics for the path characteristic evaluation in Section III. Section IV presents a preliminary design for VTP that incorporates these analytical results. Section V surveys related work and Section VI concludes the paper.

II. Scenario and Metrics

This section presents the simulation scenario that the following sections use to evaluate the path characteristics between communicating vehicles on a highway. It also defines a number of metrics for the following analysis, including connectivity and disruption durations, packet losses, RTT, RTT jitter and packet reorderings.

A. Scenario and Simulation Environment

The scenario, as illustrated in Figure 1, simulates varying numbers of vehicles on a 10km stretch of highway. The spatial distribution of the vehicles and their mobility behavior, *i.e.*, position, direction, speed, derive from validated highway mobility patterns [6][7]. The analysis considers different scenarios that have different numbers of lanes in each direction and varying numbers of vehicles per kilometer.

All vehicles are equipped with a single IEEE 802.11 wireless interface providing a radio transmission range of 250m. Vehicles in radio range can communicate directly. In case the distance between communication pairs exceed the radio range but a multi-hop path exists, the vehicles form a self-organizing *ad hoc* network that supports multi-hop communication.

Generally, communication in this VANET occurs between random pairs of vehicles distributed throughout the simulated area. However, some parts of the analysis restrict communication to vehicles within specific distances. The number of vehicle pairs that communicate concurrently determine the network load. The simulations investigate 5, 10 and 15 concurrent communications, representing light, medium and high loads. Each communication is a constant-bitrate bulk data transfer with a fixed packet size.

The VANET uses position-based routing (PBR) [8], because it outperforms topology-based routing in highly dynamic vehicular environments [9]. With PBR, each node selects for each packet the next reachable forwarder that is geographically closest to the destination. Thus, each packet may follow a different path due to mobility.

Mobility can also create temporary network partitions that interrupt end-to-end connectivity and cause packet loss. In order to reduce the number of network partitions, oncoming traffic is included when determining next hops [7].

B. Metrics

This section defines several metrics that describe the path characteristics experienced by a single communication instance.

A *connectivity period* denotes the existence of an end-to-end path between source and destination that enables communication. A *disruption period* denotes the absence of such a path. The *connectivity duration* hence describes the length of a connectivity period, whereas the *disruption duration* describes the length of a disruption period.

The *packet loss probability* describes the likelihood that an individual packet is lost between source and destination, independently of other packets. The *packet loss burst length* describes the number of consecutively lost packets.

The *RTT* describes the time between transmission of a packet and the reception of the first corresponding acknowledgment. The simulation takes one RTT sample at any given time. The sampling frequency is hence inversely proportional to the RTT. The *RTT jitter* describes the difference between two subsequent RTT samples. The *mean RTT* describes the mean across all RTT samples for a given communication.

Reordered packets are received in a different sequence than they were sent in. The *packet reordering probability* describes the likelihood that a packet is reordered, independently of other packets. The *reordering period* describes the time from the reception of the first reordered packet until the originally expected packet arrives. Note that lost and duplicated packets do no contribute to reordering.

III. EVALUATION RESULTS

This section presents selected simulation results for a highway scenario with two lanes per direction and on average six vehicles per lane and kilometer. Many additional results are available in [10]. Each sender generates a constant bit rate (CBR) stream of 100Kb/s. Although some path characteristics are expected to be different in the presence of a transport protocol with congestion control, the CBR streams approximate these environments. The duration of each simulation is 60s.

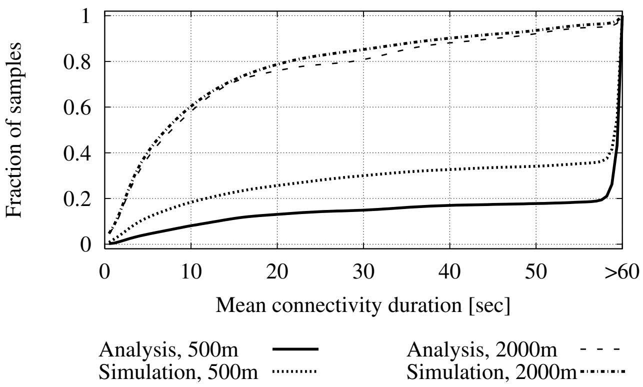

Fig. 2. CDF of connectivity duration for analysis and simulations.

A. Connectivity and Disruption Durations

This section presents the evaluation of connectivity and disruption durations and compares the analytical evaluation and simulation results for maximum source-destination distances of 500m and 2000m.

First, an analytical evaluation examines the connectivity and disruption durations by determining the theoretical availability of an end-to-end path in discrete time intervals using global knowledge. These results represent an upper bound for the expected connectivity durations, because they do not consider MAC and physical effects. Simulations that use an ideal MAC validate the analytical evaluation. Further simulations evaluate the connectivity duration using the IEEE 802.11 MAC.

Figure 2 shows the cumulative distribution function (CDF) of the normalized connectivity durations. When the communicating nodes remain within a distance of 500m, 9% of the communications are interrupted within 10s and 20% are interrupted within the duration of the simulation. Consequently, 91% of the communications remain uninterrupted for 10s and 80% continue for the complete duration of the simulation.

The results for the 802.11 MAC in the 500m scenario show a decrease in connectivity durations. After 10s, 20% of the communications are interrupted and 38% of the communications are interrupted up to the end of the evaluation. These differences between analysis and simulations are mainly due to inaccurate location information in the latter case.

The connectivity durations significantly decrease for longer distances, as illustrated by the curve for 2000m maximum distance in Figure 2. The analysis and simulation curves converge for 2000m maximum distance, because the interruptions due to routing errors dominate in this case. The exemplary description of the simulation curve shows that after 10s, 60%, after 30s, 84% of the communications, and after 60s, 96% of the communications are interrupted

A similar evaluation of disruption durations cannot be shown due to space restrictions, but is available in [10]. The main result is that the average disruption duration is short. For 500m maximum distance, 92% of the disruptions end after 3s. In comparison, 75% of the communications over 2000m distance resume after 3s disruption.

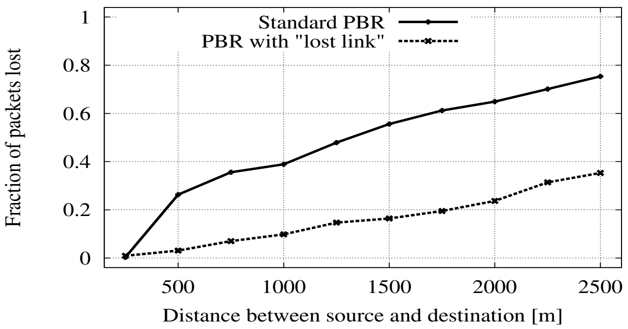

Fig. 3. Loss probability over distance (standard and lost-link enhanced PBR).

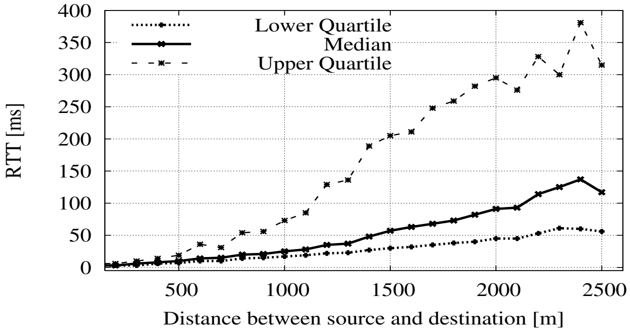

Fig. 4. Median RTT and quartiles over distance.

B. Packet Loss Probability and Distribution

Packet losses are frequent in vehicular environments, because of high mobility and the resulting topology changes. Figure 3 illustrates loss probability over distance for standard PBR and PBR with *lost link* enhancement. With standard PBR, neighbor table entries time out periodically and can become stale. The *lost link* enhancement aims to reduce packet loss by cross-layer integration, keeping neighbor tables updated based on link-layer feedback [11].

Figure 3 shows already a significant loss probability in a scenario with light network load. For both curves, the loss probability up to 250m is below 1%, due to wireless packet loss in single-hop communication. With standard PBR, loss probability increases to 26% for 500m distance. Beyond 500m distance, multi-hop communication is required and the loss probability increases linearly with longer distances.

The second curve shows a linear increase of packet loss probability for PBR with the *lost link* enhancement. Packet loss is significantly reduced, down to 3% for 500m distance. However, the reduction of packet loss comes at the cost of increased RTT and RTT jitter, because the probing of different neighbors is time consuming, as analyzed in [10].

The evaluation of consecutive loss in [10] shows that over all distances in the light-load scenario, with a probability of 31% losses occur as single packet losses, with 54% probability, three consecutive packets are lost, and with 82% probability up to ten packets are lost subsequently. In comparison, the results for the high-load scenario show 29% loss of single packets and 78% of ten subsequently lost packets. Thus, for high network loads, the number of consecutively lost packets increases because additional queue drops occur as shown in the evaluation of drop reasons in [10]. However, consecutive packet loss of more than ten packets is mainly due to network partitions. The UDP communication continues transmitting during a disruption, whereas VTP should decrease the transmission rate to probing packets.

C. RTT and RTT Jitter

Transport protocols commonly use an estimation of the RTT, *e.g.*, to determine retransmission timeouts. This evaluation focuses on RTT and RTT jitter (without lost-link enhancement) in order to determine whether this metric is appropriate in VANET environments. Figure 4 shows the medium RTT over distance for 15 streams. For a 500m source-destination distance, the RTT median is 10ms and the upper 75% quartile is 19ms. However, the upper quartile increases significantly for longer distances, *e.g.*, the median for 2000m distance is 91ms and the respective upper quartile is 295ms.

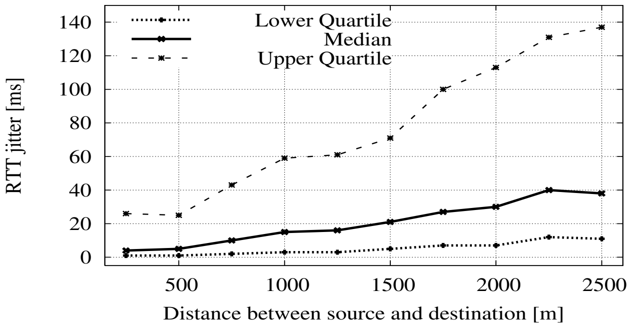

Fig. 5. Median RTT jitter and quartiles over distance.

Figure 5 shows the evaluation results of RTT jitter for consecutive samples over distance. The median RTT jitter for 500m distance is 5ms and the upper quartile is 25ms. For 2000m, the median RTT jitter is 30ms and the upper quartile increases to 113ms. The differences between the median and the upper quartile show that the RTT for consecutive packets differs significantly.

D. Packet Reordering Probability and Period

VANETs have a significant probability of packet reordering, which mainly depends on the network load and source-destination distance. The reordering probability in light-load scenarios with 5 parallel streams is below 1%, whereas the reordering probability for middle load scenarios is 15% beyond a distance of 1500m [10].

In addition to the reordering probability, the number of subsequent reordered packets and the duration of reordering are important for the VTP design. Figure 6 compares the reordering period for different network loads in a CDF graph. In the light-load scenario, 2% of the samples remain in the reordering period for 10ms, 60% for 100ms and 96% for 1000ms. In comparison, in the high-load scenario shows 7%, 50% and 93% of the samples remain in the reordering period

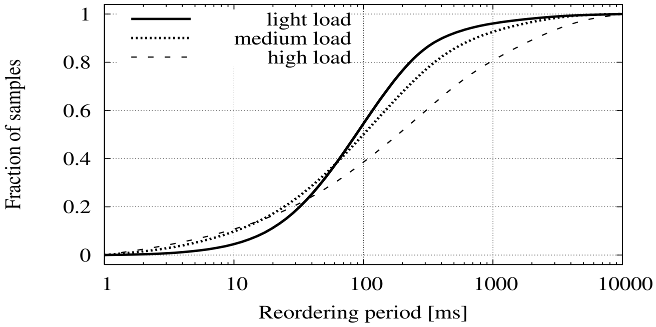

Fig. 6. CDF of reordering period for different network loads.

for 10ms, 100ms and 1000ms, respectively. Note that the reordering probability affects the form of the curves and causes e.g., the crossing of the curves.

This section evaluated connectivity and disruption durations, packet loss, RTT, RTT jitter and reordering probabilities for a highway scenario. The next section presents a preliminary VTP design based on these statistical results.

IV. PRELIMINARY VTP DESIGN

The objectives of VTP include the establishment and release of an end-to-end connection, reliable delivery of data packets and flow and congestion control. The design of selected VTP functions is directly influenced by the path characteristics results, as described below. The major design goals are to:

- maximize the throughput of a connection,
- preserve fairness among contending flows and adapt to the available bandwidth,
- reliably transmit data, cope with frequent packet loss rate and high end-to-end RTT and RTT jitter,
- deliver data in-order, coping with high ratio of reordered packets.

A VTP instance can either be in a *connected* or *disrupted* state. The arrival of acknowledgments (ACKs) indicates a connected state. In absence of ACKs, the sender calculates the expected remaining connectivity duration, using the previously determined statistical results for the given source-destination distance. If the result is lower than a threshold, VTP switches to the disrupted state. The arrival of an ACK triggers the transition from disrupted to connected state. The VTP sender immediately resumes transmission at the maximum allowed data rate, as an ACK contains the minimal available bandwidth along the path.

In a connected state, a VTP sender uses explicit, per-packet feedback as collected by intermediate nodes along the multi-hop path to adapt its data rate to the current path characteristics. In a disrupted state, a sender periodically probes whether connectivity has resumed. The probing interval derives from the statistical results based on the source-destination distances.

Congestion control in VTP uses explicitly signaled information from intermediate nodes. Decoupling of error and congestion control avoids unnecessary transmission rate reductions due to non-congestion packet loss, such as routing errors or wireless effects. The feedback information includes local and environmental data of intermediate nodes, such as the available bandwidth based on bandwidth utilization, packet queue length and number of vehicles in transmission range.

Intermediate nodes compute the locally available minimum bandwidth to the current value in the packet header and update it if the local minimum is below the header value. The VTP receiver piggybacks the minimum of the available bandwidth to reverse-path ACKs. This feedback allows the VTP sender to calculate the bandwidth-delay product using the bandwidth information from the ACK and to estimate the RTT to pace the next data packets. The sender estimates the RTT based on RTT measurements and the statistical deviation for the source-destination distance. Assuming that the determined available bandwidth is valid in the spatial vicinity of the intermediate node, it is therefore not restricted to a specific route (in contrast to the existing ATP approach [12]). Though the next packet might follow a different path (as position-based routing allows for every packet), the VTP source assumes the same data rate until a new feedback is received.

The high packet loss rate, the distribution and packet loss burst length in combination with reordering demand for selective acknowledgments (SACK) for efficient retransmission that avoids spurious retransmissions of already received packets.

An intermediate VTP node distributes the available bandwidth based on the accumulated number and size of packets per flow in a feedback interval, similar to XCP [13]. This packet-based bandwidth distribution provides fairness among the contending flows without maintaining flow information on the intermediate hops, because packets carry all required flow information in the header.

V. RELATED WORK

PATHS [14] provides an analysis of path duration statistics and their impact on reactive routing protocols. In contrast to this paper, PATHS investigates AODV and DSR routing protocols. The analysis of connectivity and delivery ratio of the PBR routing protocol is given in [6][7][11]. The evaluation of path characteristics in this paper extends their metrics and observes the performance over time and distance.

Numerous enhancements to TCP aim to improve performance in wireless environments. [15] classifies single-hop TCP modifications in link-layer [16], end-to-end [17] and split connection proposals [18]. A variety of proposals focuses on multi-hop TCP enhancements [19][20][21]. Other proposals focus on the more general problem of improving TCP behavior across arbitrary intermittently connected links [22]. However, a transport protocol that is tailored to the specific path characteristics of VANETs is expected to outperform such generic TCP enhancements, because they are not tuned to the characteristics of highly dynamic vehicular environments.

In contrast, *non-TCP* approaches, such as ATP [12] or the approach in [23], assume that the window-based concept of TCP is not appropriate for mobile wireless networks. ATP proposes explicit signaling of queue length and average transmission delay by intermediate hops, similar to the VTP design. However, their approach assumes the establishment of routes as in traditional topological-based routing protocols. Thus, mechanisms such as static feedback or route failure feedback (*i.e.*, assuming a single route per communication pair) are not applicable for VANETs using PBR.

Furthermore, explicit signaling for mobile wireless networks [24][25] uses a comprehensive flow control via explicit signaling and additional mechanisms for dynamic bandwidth estimation, safety window and route probing. In order to regulate the flows, each router has to keep per-flow state. However, PBR forwards packets *on the fly* via different paths and flow state maintenance at intermediate hops is not possible.

XCP [13] proposes explicit signaling for high bandwidth-delay-product networks. Though XCP is not designed for wireless networks, the explicit signaling concept appears suitable for VANETs.

VI. CONCLUSION AND FUTURE WORK

This paper evaluates the communication path characteristics for VANETs in typical highway scenarios, namely connectivity and disruption duration, packet loss, packet reordering, RTT and RTT jitter.

The connectivity evaluation results show that steady communication is feasible for distances of up to 2000m. For a distance of 2000m, about 40% of the connections remain uninterrupted for 10s in average. With decreasing distance, the connectivity duration even increases. Disruptions resume latest after 3s, only marginally dependent on the distance. The packet loss ratio for a constant packet stream is, however, huge: For a distance of 2000m, standard PBR shows a packet loss rate of almost two thirds, which can be significantly reduced to 22% when using cross-layer integration. Although the RTT and RTT jitter are acceptably small for source-destination distances below 700m, higher distances result in extreme fluctuation in RTT, *e.g.*, up to 300% for a 2000m distance. Finally, reordering ratios for light loads are small (below 1%), but increases to 15% even for medium data loads.

Based on the results, this paper outlines a preliminary design for a unicast transport protocol in VANETs. The protocol aims to maximize throughput with reliable and in-order delivery of data and to provide flow and congestion control. Specific functions of VTP use knowledge about the distance between communication partners, elapsed connectivity duration and statistical knowledge about the connectivity between vehicles on highways to predict the current and future connectivity state and to cope with the extreme conditions in VANETs.

In the future work, we will extend the evaluation to city scenarios, provide a detailed VTP specification and analyze the approach through simulation and real-world measurements.

ACKNOWLEDGMENTS

R. Schmitz and A. Festag acknowledge the support of the German Ministry of Education and Research (BMB+F) for the project *NoW – Network on Wheels* under contract number 01AK064F. L. Eggert was partially supported by *Ambient Networks*, a research project under the European Commission's Sixth Framework Program.

REFERENCES

[1] "NoW - Network on Wheels," http://www.network-on-wheels.de/.

[2] J. Postel, "Transmission Control Protocol," September 1981, RFC 793.

[3] G. Holland and N. Vaidya, "Analysis of TCP over Mobile Ad Hoc Networks," *Proc. ACM MobiCom*, August 1999.

[4] Z. Fu, X. Meng, and S. Lu, "How bad TCP Can Perform In Mobile Ad Hoc Networks," *Proc. Symposium on Computers and Communications (ISCC)*, July 2002.

[5] M. Gerla, K. Tang, and R. Bagrodia, "TCP Performance in Wireless Multi-hop Networks," *Proc. IEEE Workshop on Mobile Computer Systems and Applications (WMCSA)*, February 1999.

[6] R. Krüger, H. Füßler, M. Torrent-Moreno, M. Transier, H. Hartenstein, and W. Effelsberg, "Statistical Analysis of the FleetNet Highway Movement Patterns," *TR-2005-004, Department of Computer Science, University of Mannheim, Germany*, July 2005.

[7] M. Torrent-Moreno, F. Schmidt-Eisenlohr, H. Füßler, and H. Hartenstein, "Effects of a Realistic Channel Model on Packet Forwarding in Vehicular Ad Hoc Networks," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Las Vegas, NV, USA, April 2006.

[8] M. Mauve, J. Widmer, and H. Hartenstein, "A Survey on Position-Based Routing in Mobile Ad-Hoc Networks," *IEEE Network*, vol. 5, no. 6, November 2001.

[9] H. Füßler, M. Mauve, H. Hartenstein, M. Käseman, and D. Vollmer, "A Comparison of Routing Strategies for Vehicular Ad Hoc Networks," *TR-02-003, Department of Computer Science, University of Mannheim, Germany*, July 2002.

[10] A. Leiggener, "Evaluation of Path Characteristics in Vehicular Ad Hoc Networks," *Master's Thesis, Institute EURECOM, Sophia Antipolis, France*, September 2005.

[11] H. Hartenstein, H. Füßler, M. Mauve, and W. Franz, "Simulation Results and Proof-of-Concept Implementation of the FleetNet Position-Based Router," *Proc. Conference on Personal Wireless Communications (PWC)*, September 2003.

[12] K. Sundaresan, V. Anantharaman, H.-Y. Hsieh, and R. Sivakumar, "ATP: A Reliable Transport Protocol for Ad-hoc Networks," *Proc. ACM MobiHoc*, June 2003.

[13] D. Katabi, M. Handley, and C. Rohrs, "Congestion Control for High Bandwidth-Delay Product Networks," *Proc. ACM SIGCOMM*, August 2002.

[14] N. Sadagopan, F. Bai, B. Krishnamachari, and A. Helmy, "PATHS: Analysis of PATH Duration Statistics and their Impact on Reactive MANET Routing Protocols," *Proc. ACM MobiHoc*, June 2003.

[15] H. Elaarag, "Improving TCP Performance over Mobile Networks," *ACM Computing Surveys*, vol. 34, no. 3, Sept. 2002.

[16] H. Balakrishnan, S. Seshan, and R. Katz, "Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks," *Wireless Networks*, vol. 1, no. 4, December 1995.

[17] B. Bakshi, P. Krishna, N. Vaidya, and D. Pradhan, "Improving Performance of TCP over Wireless Networks," *Proc. International Conference on Distributed Computing Systems (ICDCS)*, May 1997.

[18] A. Bakre and B. Badrinath, "I-TCP: Indirect TCP for Mobile Hosts," *Proc. International Conference on Distributed Computing Systems (ICDCS)*, June 1995.

[19] M. Gerla, R. Bagrodia, L. Zhang, K. Tang, and L. Wang, "TCP over Multi-Hop Protocols: Simulation and Experiments," *Proc. Int. Conference on Communications (ICC)*, June 1999.

[20] Z. Fu, B. Greenstein, X. Meng, and S. Lu, "Design and Implementation of a TCP-Friendly Transport Protocol for Ad Hoc Wireless Networks," *Proc. IEEE International Conference on Network Protocols (ICNP)*, November 2002.

[21] J. Liu and S. Singh, "ATCP: TCP for Mobile Ad Hoc Networks," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 7, July 2001.

[22] S. Schütz, L. Eggert, S. Schmid, and M. Brunner, "Protocol Enhancements for Intermittently Connected Hosts," *ACM Computer Communication Review (CCR)*, vol. 35, no. 3, July 2005.

[23] G. Anastasia and A. Passarella, "Towards a Novel Transport Protocol for Ad Hoc," *Proc. Personal Wireless Communications (PWC)*, Sept. 2003.

[24] K. Chen and K. Nahrstedt, "EXACT: An Explicit Rate-based Flow Control Framework in MANET," *Technical Report UIUCDCS-R-2002-2286/UILU-ENG-2002-1730*, July 2002.

[25] K. Chen, K. Nahrstedt, and N. Vaidya, "The Utility of Explicit Rate-Based Flow Control in Mobile Ad Hoc Networks," *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, March 2004.