Real-life Multi-hop Multi-radio Multi-channel Wireless Mesh Network Experiments: 802.11n is not any better than 802.11a!

Ivan Wang-Hei Ho*, Patrick P. Lam*, Peter Han Joo Chong† and Soung Chang Liew‡

*P2 Mobile Technologies Limited, Hong Kong, {ivanho, patricklam}@p2mt.com

†School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore, ehjchong@ntu.edu.sg

‡Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong, soung@ie.cuhk.edu.hk

***Abstract*—Wireless Mesh Network (WMN) has become a popular access network architecture in the community due to its low cost and readily deployable nature. However, it is well known that multi-hop transmission in WMN is vulnerable to bandwidth degradation, primarily due to contention and radio interference. A straightforward solution to this problem is to use mesh nodes with multiple radios and channels. In this paper, we demonstrate through real-world experiments that the use of multiple radios and channels solely cannot solve the multi-hop TCP throughput degradation problem in IEEE 802.11n mesh networks. We verify that it is because the round trip time (RTT) for wireless communication path increases significantly with the number of hops, and the TCP throughput is limited inversely by the RTT. Due to this TCP throughput limitation, we also found interestingly that the multi-hop TCP throughput (up to five hops) of 802.11a is comparable to that of 802.11n. Our results give a key message that a new generation of TCP needs to be proposed and adopted so as to support the ever advancing wireless technologies and growing demand of multi-hop communications in wireless mesh networks.**

***Keywords* – Wireless Mesh Network (WMN); IEEE 802.11n; Multi-Radio Multi-Channel (MRMC); Multi-hop TCP throughput**

I. Introduction

Due to the low cost and readily deployable nature of Wireless Mesh Network (WMN), it has been attracting a lot of attentions from the research community and industry. There are also a wide range of deployments and applications of WMNs in real world [1, 2] for video surveillance, voice communications and localization services.

However, according to [3, 4], it is well known that the multi-hop nature of WMN is vulnerable to the throughput degradation problem. This is primarily due to contention or half-duplex communication in single-radio networks. Since a node with a single wireless interface cannot transmit and receive packets at the same time, it must first receive and then transmit in order to relay a packet. The channel time required to relay a packet is thus at least twice the amount of time for sending a packet directly from the source to the destination. Moreover, the carrier-sensing mechanism in 802.11 MAC may prevent simultaneous transmissions on the same channel, and radio interference (or collision) may occur when the carrier sensing fails to prevent interfering links from transmitting simultaneously. These will all affect the multi-hop throughput performance of WMNs.

A number of previous works [5 – 7] have proposed to use multiple radios and channels in WMNs to resolve the multi-hop throughput degradation problem. With multiple radios (or wireless interfaces), a node can receive and transmit simultaneously, and full-duplex communication can be achieved. Furthermore, with subtle channel allocation, the wireless interfaces can be assigned with different (non-overlapping) channels so as to prevent radio interference, and links with different channels can transmit simultaneously without any collisions. In general, a Multi-Radio Multi-Channel (MRMC) WMN is thought to be a robust solution for the multi-hop throughput degradation problem. However, our experimental results indicate that it is not the case for TCP connections under 802.11n WMNs. For instance, in an 802.11n network, we found through real-life experiments that the multi-hop TCP throughput degrades for more than 70% after traversing five hops over an 802.11n WMN even under some nearly perfect conditions.

It is well-known that TCP throughput is limited inversely by the round trip time (RTT) for the communication path given the receive window size, RWIN (the amount of data that the receiver can accept without acknowledging the sender), as follows:

$$\text{TCP Throughput} \leq \frac{\text{RWIN}}{\text{RTT}}. \quad (1)$$

In this paper, we demonstrate with experimental results that using multiple radios and channels alone is not enough to improve the multi-hop TCP throughput in a 802.11n based WMN. We also verify through control experiments that the multi-hop TCP throughput degradation in MRMC 802.11n mesh networks is primarily due to the limiting factor described in (1) and the large RTT for multi-hop wireless path. In addition, because of (1) and the fact that the data rate of 802.11n is much larger than that of 802.11a, we found interesting in our results that the multi-hop throughput performance of 802.11a could be similar or comparable to that of 802.11n.

To illustrate how (1) limits the multi-hop TCP throughput of 802.11n, let us consider an example. Assume RWIN = 16 KB = 16,384 bytes and packet size of 1,500 bytes. Then, the RWIN = 16,384/1,500 = 11 packets. We found from our experiment that the effective throughput of 802.11n is about 80 Mbps. So the time needed to transmit one packet at the first hop is 1,500*8/80 μs = 0.15 ms. It will take 11*0.15 ms = 1.65 ms before the source transmits all the 11 packets in a TCP window. If at this time, the ACK for the first packet still has not returned from the destination that is multiple hops away, then the TCP source will be idling, and the bandwidth at the first hop will be wasted. Hence, the RTT for the communication path must be smaller than 1.65 ms in order to have no bottleneck. According to our experiments, the RTT for a two-hop wireless path is more than 2 ms, and this is the reason why multi-hop TCP throughput for 802.11n degrades starting from the second hop (more about the results in Section III). On the other hand, for 802.11a, the effective throughput is only about 20 Mbps. Therefore, the critical RTT that will limit the throughput is larger (1500*8*11/20 μs = 6.6 ms). For instance, the five-hop RTT for a wireless path is about 5 ms according to our experiments. Thus, the TCP throughput for MRMC 802.11a can be maintained after traversing five hops in the network.

II. Experimental setup

Two experiments were conducted, one is for measuring the multi-hop TCP throughput for 802.11n and 802.11a networks (up to five hops). The other one is a control experiment that replaces the wireless links gradually with wired links (up to three links wired) to demonstrate and verify that the large round-trip delay for wireless path is the key factor that limits the multi-hop TCP throughput performance.

Six mesh nodes are used in the experiments. The mesh nodes we used are the MeshRanger2-o mesh routers provided by P2 Mobile Technologies Limited [8]. Each of the MeshRanger2-o has two dual-band Wi-Fi radio interfaces, in which the Atheros AR7161 chipset is used as the interface controller, capable of operating in either the 2.4 GHz or 5 GHz frequency bands.

The six mesh nodes are aligned in a chain topology (with uniform separation of 1 m) and operating in non-overlapping 5 GHz channels as illustrated in Fig. 1. Note that when the high throughput (HT) mode of 802.11n is enabled, a 40 MHz-wide channel is used. So, the channel assignment needs to take this into account and needs to avoid channel overlapping and radio interference introduced by the "secondary" channel used in the HT mode. All of the nodes use a transmit power of 17 dBm (50 mW) and are equipped with two pairs of omni-directional antennas with a gain of 5 dBi.

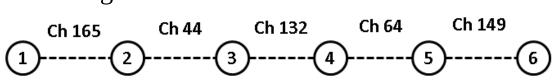

Figure 1. The alignment of the six mesh nodes and the 5 GHz channel assignment.

We measure the multi-hop data rate for 802.11a and 802.11n (with 40MHz-wide channels used in the latter case) with RTS/CTS disabled. Specifically, the network tool, Iperf [9], is used to measure the TCP and UDP throughput for connections with different number of hops, and packets originating from node 1 are transmitted to other nodes in the linear network for the measurements. At the beginning of the experiments, the throughput for every wireless link is examined to ensure that it is working properly. For instance, we found that the normal effective single-hop TCP throughput for 802.11n is about 80 to 85 Mbps, while that for 802.11a is about 20 to 25 Mbps. In the experiments, the TCP receive window size is set as 16 KB, while for UDP transmission, the offered UDP bandwidth is 100 Mbps. The ping command is used to get a rough estimation of the RTT, though the TCP RTT should theoretically be larger than the ping RTT because of the extra processing at the transport layer.

III. Experimental Results and Discussion

A. *Multi-hop throughput results and potential factors that affect the multi-hop performance*

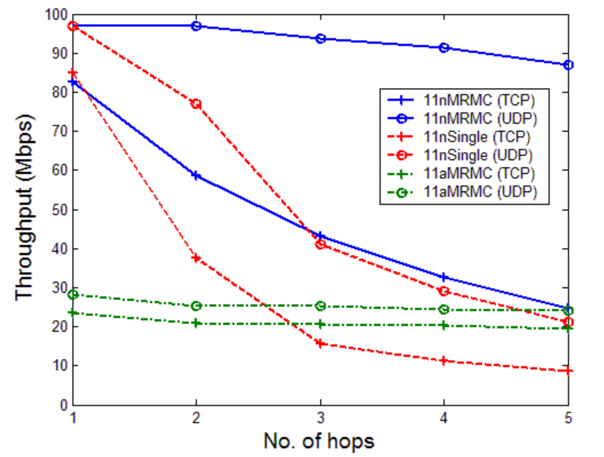

Figure 2. The TCP and UDP multi-hop throughput for 802.11n (MRMC), 802.11n (single channel), and 802.11a (MRMC).

Fig. 2 plots the TCP and UDP multi-hop throughput for single-channel 802.11n (channel 149), Multi-Radio Multi-Channel (MRMC) 802.11n, and MRMC 802.11a. First of all, we can see from the figure that both the TCP and UDP throughput of single-channel 802.11n (red-dashed lines) drop drastically when traversing the network, which is primarily due to radio interference and collision. For its TCP throughput, it is more than halved for every increase of hops. With the use of multiple radios and channels, the TCP and UDP throughput of 802.11a, and the UDP throughput of 802.11n can be sustained over five hops (a fluctuation of less than 10% could be due to the processing of routing overheads or variation of the wireless link quality and is considered to be normal). However, the TCP throughput of MRMC 802.11n (cross-marked blue-solid line) still drops significantly (over 70%) after five hops. This result is disappointing, especially to researchers that propose the use of multiple radios and channels in wireless mesh networks for boosting the overall network capacity. It is also surprising to note that the TCP throughput of 802.11a (cross-marked green-dash-dotted line) is almost equal to that of 802.11n after five hops.

In general, there are several potential reasons to account for the multi-hop TCP throughput drop in MRMC 802.11n mesh network that are worth considering.

1. *Radio Interference and Clear Channel Assessment (CCA)*: First of all, if overlapping channels are assigned or channels are reused in the network, co-channel interference and collision might occur. Furthermore, the CCA mechanism in 802.11 consists of Physical Carrier Sensing (PCS) and Energy Detection (ED) [10]. The former detects transmissions of similar systems, while the latter provide information about medium usage. If the receive power level exceeds the corresponding thresholds, PCS or ED will indicate a busy wireless medium and prevent the device from transmitting.

2. *Limited Computation Power of Mesh Points*: The computation power of a mesh router is usually much lower than a personal computer. For instance, the CPU speed of the MeshRanger-2o is 680 MHz and the memory size is 128 MB. The limited computation power of mesh points might incur delay on handling the routing overheads in large-scale wireless mesh network.

3. *Cross-talk and Radio Leakage*: According to [5], due to the near-field effect, inter-radio board cross-talk, radio leakage, and hardware imperfections, self-interference problem may arise when two radios are in close proximity to each other. Since we have two radio cards installed in every mesh point that operate simultaneously, and they are close to each other, inter-radio board cross-talk might happen and degrade the transmission performance.

4. *End-to-end Round-trip Delay*: A large RTT is recorded for multi-hop wireless path in our experiments, which should be incurred by the MAC protocol of 802.11. The large round-trip delay restricts the TCP throughput performance as described by (1).

It is obvious that 1 is not the key factor that limits the multi-hop throughput in our case. Since we have already carefully assigned non-overlapping channels for the MRMC mesh network, and the results for 802.11a and the UDP result for 802.11n all indicate that co-channel interference and the CCA mechanism are not affecting the data transmission in our setup at all, otherwise, the multi-hop throughput cannot be sustained and will drop significantly like the single-channel 802.11n case does. In the following, we are going to validate with further calculations and a control experiment that Factors 2 and 3 cannot be the primary reason that limits the multi-hop throughput performance either. As a result, 4 becomes the most possible answer for the phenomena observed in our experiments.

B. *Control Experiment: Results and Discussion*

Let us move on to the control experiment. In this experiment, the overall topology and configurations are the same as the previous one, except that wireless links are gradually replaced with wired links. We first wire the link between nodes 1 and 2 (the first hop), then the link between nodes 3 and 4 (the third hop), and finally the link between nodes 5 and 6 (the fifth hop) as illustrated in Fig. 3. After a link is wired, the wireless interfaces that establish the wireless link are disabled to eliminate the possibility of inter-radio board cross-talk within any mesh nodes. The changes in multi-hop throughput and RTT are recorded.

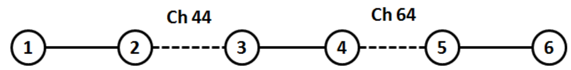

Figure 3. The control experiment with three links wired (the 1st, 3rd, and 5th hops).

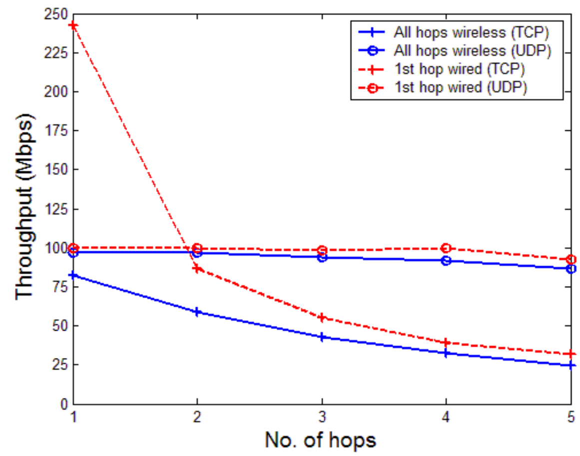

Figure 4. The TCP and UDP multi-hop throughput for 802.11n (MRMC) with all wireless hops and with the first hop wired.

Fig. 4 shows the TCP and UDP throughput across multiple hops when the first hop (or link) is wired (red dashed lines). We can see that the TCP bandwidth for the first hop is very high (243 Mbps) since nodes 1 and 2 are connected with a wire. The TCP throughput then drops starting from the second hop. It is interesting to note that the line for the wired case (cross-marked red-dashed line) appears to lag behind the wireless case (cross-marked blue-solid line) by one hop. For example, the (two-hop) TCP throughput from nodes 1 to 3 for the wired case is similar to the (single-hop) TCP throughput from nodes 1 to 2 for the wireless case.

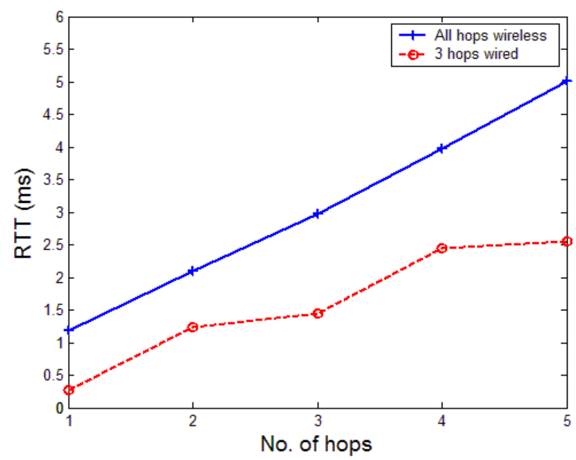

Figure 5. The ping RTT for multi-hop connections with all wireless hops and with three hops wired (the 1st, 3rd, and 5th hops).

Other than the data rate difference, wired and wireless links also differ in the RTT. Fig. 5 illustrates the differences in multi-hop RTT for the wireless and wired control experiments (with the first, third and fifth hops wired). We can see that the RTT for the wireless case increases almost linearly with the number of hops. For the wired case, there is only a slight increase in RTT at the first, third and fifth hops, which indicates that the RTT for wired links is much smaller than that for wireless links.

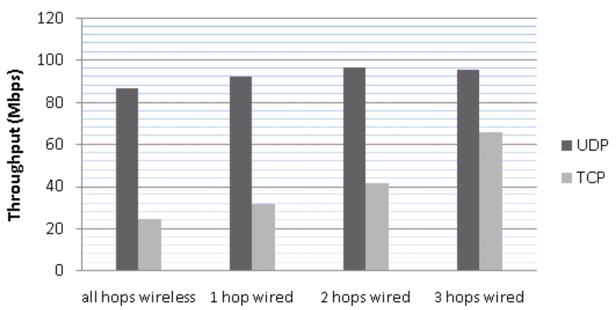

Figure 6. The five-hop TCP and UDP throughput with all wireless hops, one hop wired (the 1st hop), two hops wired (the 1st and 3rd hops), and three hops wired (the 1st, 3rd, and 5th hops).

The five-hop TCP and UDP throughput from nodes 1 to 6 for the wireless, 1-hop wired, 2-hop wired, and 3-hop wired scenarios are shown in Fig. 6. There is a small increment in the UDP throughput when the links are wired, it is because wired link is more stable and has less fluctuation than wireless link in general. For the TCP throughput, it increases significantly with the number of links wired, and is improved by about three times when three hops are wired. According to Fig. 6, it again appears that the "effect" of introducing wired links to the multi-hop network is very similar to reducing the number of hops (or effectively reducing the RTT of the wireless path. For instance, the five-hop path from nodes 1 to 6 in Fig. 3 has similar throughput performance as a two-hop wireless path.

Let us consider the scenario in Fig. 3 with three links wired (the first, third, and fifth hops) and revisit the potential reasons that account for the poor performance proposed in the previous subsection. If the key factor that affects the multi-hop throughput is Factor 2 (limited computation power of mesh points), then the wired control experiment should also suffer from throughput degradation, since both the wired and wireless cases require the processing of routing overheads. But it appears that it is not the case according to the result in Fig. 6. In addition, if Factor 3 (cross-talk and radio leakage) is the key factor, then the five-hop throughput for the three links wired scenario should be able to be maintained at around 80 Mbps (but not 60 Mbps as shown in Fig. 6). Since in that scenario, none of the mesh points has more than one radio card operating simultaneously, hence, the possibility of inter-radio board cross-talk within a mesh node is eliminated.

As a result, it appears that Factor 4 (end-to-end round-trip delay) is the most possible answer to account for the multi-hop TCP throughput degradation phenomena. The latency (or increase in RTT) should be incurred by the MAC protocol of 802.11. The large RTT for multi-hop wireless path triggers spurious timeouts and retransmissions in TCP protocol and hence leads to poor multi-hop performance.

Table 1. The multi-hop TCP and UDP throughput and ping RTT for the MRMC 802.11n mesh network.

# hops	1	2	3	4	5
UDP (Mbps)	97.0	96.8	93.6	91.4	86.9
TCP (Mbps)	82.6	58.5	43.0	32.5	24.7
RTT (ms)	1.19	2.10	2.98	3.97	5.01
Estimated TCP (Mbps)	N/A	N/A	41.2	30.9	24.5

Table 1 shows the multi-hop throughput and RTT results for the MRMC 802.11n network. Let us verify if the results match with the relationship given in (1). Assume that the two-hop TCP throughput is limited by (1), that is, RWIN/RTT = 58.5 Mbps. If this assumption is true, we should be able to predict the throughput for the remaining hops. For example, the estimated throughput for the third hop is 58.5*RTT(2-hop)/RTT(3-hop) = (58.5*2.1)/2.98 = 41.2 Mbps. Our measurement is 43 Mbps, which is pretty close. By the same token, we can estimate the throughput for the fourth and fifth hops, which are (58.5*2.1)/3.97 = 30.9 Mbps and (58.5*2.1)/5.01 = 24.5 Mbps respectively. Again, these predictions match well with our measurements (32.5 and 24.7 Mbps), which further confirm that the multi-hop TCP throughput degradation problem is due to the large RTT of wireless path, and the relationship between the throughput and RTT is described by (1). Furthermore, the phenomenon that the TCP throughput for 802.11a does not drop with the number of hops in our experiment can be explained with (1). This is because the effective data rate of 802.11a (about 20 Mbps) is much lower than that of 802.11n (about 80 Mbps with two spatial streams), thus, 802.11a can sustain a much larger RTT until the TCP throughput is limited by (1) (as discussed in the last paragraph of Section I).

IV. CONCLUSION

TCP was initially designed and optimized for wired network. Therefore, it is well-known that the fluctuating wireless link quality and user mobility will result in highly variable RTT and delay spike, and thus incur spurious timeout [11]. However, no previous work has addressed the problem incurred by the RTT effect in multi-hop MRMC wireless network, primarily because most of the previous studies on wireless TCP are based on single-radio wireless networks (in which bandwidth degradation in multiple hops is expected). Our work here is the first in the literature to point out that the dramatic increase of RTT in multi-hop wireless path could hinder the harness of the high throughput feature (up to 300Mbps with two spatial streams) in IEEE 802.11n mesh networks, and hence it is important to develop efficient techniques to further improve the performance of TCP in multi-hop wireless mesh networks.

There have been a number of methods proposed to improve TCP performance in wireless networks. Generally, they can be classified into two groups. The first group requires new protocol mechanism or modifications to the existing TCP protocol, while the second group of methods do not. TCP Eifel [12] is a typical example in the first group that detects spurious timeouts and retransmissions via implementing time stamping. On the other hand, Split-TCP [13] is a solution in the second group that divides the end-to-end TCP connection into two independent connections with a proxy serving as a common point. However, the major drawback of the first group methods is that devices with the newly modified TCP protocols may not be compatible with generic devices. Furthermore, a criticism of the Split-TCP method in the second group is that the TCP ACK received by an application is not actually coming from the end receiver. Thus, it destroys the original semantic that a TCP ACK means the receiver has already received the corresponding TCP packet, and violates the end-to-end security protection.

As a conclusion, we have investigated the multi-hop throughput performance of a multi-radio multi-channel (MRMC) IEEE 802.11n mesh network via experimental studies. Specifically, we have shown that though the UDP throughput can sustain after traversing multiple hops, the TCP throughput degrades dramatically even with the use of MRMC mesh points. Our results indicate that the use of multiple radios and channels is not enough to improve multi-hop TCP throughput performance in mesh networks. This is especially prominent in 802.11n with the high throughput feature. We have verified that it is primarily because of the congestion control mechanism in TCP and the large round-trip delay of wireless communication path, which makes the bandwidth of the latest WiFi technologies (e.g., 802.11n) not able to be fully utilized. Therefore, a new generation of TCP should be proposed and adopted so as to support the ever advancing technologies and growing demand of multi-hop communications in wireless mesh networks. Otherwise, no matter how sophisticated wireless modules are invented and how fast the physical data rate can be achieved, the multi-hop performance of wireless mesh networks is still bottlenecked by the TCP protocol.

REFERENCES

[1] J. Ishmael, S. Bury and D. Pezaros, N. Race, "Deploying Rural Community Wireless Mesh Networks," *IEEE Internet Computing Magazine*, vol. 12, no. 4, 2008.

[2] D. Johnson, K. Matthee, D. Sokoya, L. Mboweni, A. Makan and H. Kotze, "Building a Rural Wireless Mesh Network: A do-it-yourself guide to planning and building a Freifunk based mesh network," Meraka Institutue, 2007.

[3] P. C. Ng and S. C. Liew, "Throughput Analysis of IEEE802.11 Multi-hop Ad-hoc Networks," *IEEE/ACM Trans. on Networking*, vol. 15, no. 2, pp. 309-322, 2007.

[4] Strix Systems, "Solving the Wireless Mesh Multi-hop Dilemma," White Paper, 2005.

[5] K. Ramachandran, I. Sheriff, E. Belding-Royer and K. Almeroth, "A Multi-radio 802.11 Mesh Network Architecture," *Mobile Networks and Applications*, vol. 13, no. 1-2, 2008.

[6] Y. Liu, Y. Xiong, Y. Yang, P. Xu and Q. Zhang, "An Experimental Study on Multi-channel Multi-radio Multi-hop Wireless Networks," IEEE Global Communications Conference (GLOBECOM), 2005.

[7] C. Zhang, K. Kowalik and M. Davis, "An Experimental Study of the Impact of Using Multi-radio in WLAN Mesh Networks," Wireless Communications, Networking and Mobile Computing (WiCom), 2009.

[8] P2 Mobile Technologies Limited, http://www.p2mt.com/.

[9] Iperf, http://sourceforge.net/projects/iperf/.

[10] G. R. Hiertz, S. Max, Y. Zang, T. Junge and D. Denteneer, "IEEE 802.11s MAC Fundamentals," IEEE International Workshop on Enabling Technologies and Standards for Wireless Mesh Networking (MeshTech), 2007.

[11] K. K. Leung, T. E. Klein, C. F. Mooney and M. Haner, "Methods to Improve TCP Throughput in Wireless Networks with High Delay Variability," IEEE Vehicular Technology Conference (VTC), 2004.

[12] R. Ludwig and R. Katz, "The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions," *ACM Computer Communication Review*, vol. 30, no. 1, 2000.

[13] F. Xie, J. L. Hammond and D. L. Noneaker, "Evaluation of a Split-Connection Mobile Transport Protocol," *Wireless Networks*, vol. 9, pp. 593 – 603, 2003.