

Examining Mobile-IP Performance in Rapidly Mobile Environments: The Case of a Commuter Train.

Edwin Hernandez and Abdelsalam (Sumi) Helal

Department of Computer and Information Science and Engineering
University of Florida, Gainesville, FL 32611
helal@cise.ufl.edu

<http://www.harris.cise.ufl.edu/projects/rapmobile.htm>

ABSTRACT

Trains travel at speeds ranging from 0 to 80m/s (0 to 288 Km/hr). Providing in-train wireless Internet access to multimedia applications will require the use of a mobile networking protocol, such as Mobile-IP, to achieve uninterrupted connectivity. Although Mobile-IP represents a promising solution, its performance under “extreme” mobility is questionable. We simulated a train scenario and identified the limitations of the current mobile-IP standard in terms of throughput, handoff, and packet loss of a train moving at different velocities. We investigated the performance of UDP- and TCP-sessions, and examined the effect of different base station interleaving distances on throughput and packet loss. The results presented in this paper are part of an investigative research into adaptive mobile networking protocols in rapidly mobile networks.

1. INTRODUCTION

The dissemination of mobile networks and the increasing demand for Internet access in terrestrial vehicles makes commuter trains a suitable platform to provide wireless access to the information superhighway. Internet applications such as browsing, emailing, and audio/video streaming, currently common in wired networks, will be demanded in this *rapidly mobile* environment. This translates to QoS requirements for multimedia data from UDP- and TCP-based application sessions. Continuous connectivity, responsiveness and steady throughput are among the important QoS variables that have to be optimized throughout the train paths and trajectories.

Mobile-IP [2, 3] is the proposed standard for IP mobility support by the Internet Engineering Task Force (IETF). The standard defines three entities: mobile node, home agent, and foreign agent. Each mobile node has a permanent IP address assigned to a home network, also called *home address*. Once the mobile node decides to move away from its home network, the new location of the mobile host is determined by the *care-of-address*, which is a temporary IP address from the foreign network. In addition to the addressing procedures, the standard proposes a mobility binding method between the mobile node, the home and foreign agents. The way IP packets are routed from the correspondent host towards the mobile unit is through tunneling of the packets by the home agent to the care-of-address, bound by the mobile host. Once the packet arrives to the destination, the foreign agent proceeds to remove the encapsulated information and forwards it to the mobile node. This process is also called triangular routing.

This standard for mobility assumes relatively low speeds, which makes it very suitable for macro-mobility and nomadic environments. However, high-speed commuter trains move at speeds up to 288 Km/hr (0 to 80 m/s) which drastically reduces the effectiveness of the Mobile-IP protocol and diminishes the quality of its services.

This paper presents a performance analysis of Mobile-IP under rapid mobility conditions. We examine the effect of speed on throughput, delays, and packets drop rate of UDP and TCP transfers. In our experiments, we manufacture speed as a product of two factors: 1) the velocity of the mobile unit, and 2) the average (or fixed) interleaving distance between base stations. We base our experiments on the IEEE 802.11 wireless LAN technology [1], which represent the most appropriate MAC and physical layers available today that can deliver high-speed services to commuter train users (3G-NOW).

The paper is organized as follows. The remainder of this section summarizes related work. Section 2 presents the train simulation scenario that we developed using the ns-2 simulator tool [5]. Performance experiments examining throughput, packet loss, and effect of speed, are described in Section 3. Section 4 presents our conclusions and future work.

1.1. Related Work

Speed could have a negative effect on mobility, especially the performance of ad-hoc routing protocols. Holland [4] simulated several ad-hoc scenarios and showed that the average throughput decreases at higher speeds. The results corresponded to a simulation performed in *ns* network simulator [5] focusing only on the performance of ad-hoc networks. The same effects on throughput due to speed were measured by Gerla [6], while experimenting with tree multicast strategies in ad-hoc networks at speeds up to 100 km/hr.

Fladenmuller [7] analyzed the effect of Mobile-IP handoffs on the TCP protocol. The issues addressed in [7] involve handoff as a dual process taking effect in the wired and wireless networks. Although, no quantitative considerations on the speed of mobile nodes were presented, a conclusion was drawn in terms of the effect of pico-cells and highly frequent handoff as the main factor to modify the TCP and Mobile-IP protocols. Similarly, Caceres [8] proposed fast-retransmissions as a solution to solve the problems related to hand-off. Although his protocol improves the performance of TCP, it does not assert the issue of speed and how the retransmission timers could be modified or calculated.

Additionally, Balakrishnan [9, 10] studied the protocols: Snoop, I-TCP (Indirect TCP), and SACK (Selective acknowledgments). These protocols alleviate the problems of TCP in wireless links and improve the performance during handoff. The first protocol, called Snoop, is a software agent located at each base station, which attempts to use multicast addresses to hide the location of mobile nodes. The second protocol, I-TCP, separates the wired and wireless links into two for the wireless and wired environments. SACK has shown improvements in recovering from multiple packet losses within a single transmission window. The results showed that these protocols improve the performance of TCP under high BER (Bit Error Rate) links, and consequently can avoid unnecessary activation of the congestion control protocol [11] during handoff. Again I-TCP and SACK did not address the speed factor directly.

2. SIMULATION EXPERIMENTS

The results in this paper are based on simulation experiments performed within the *ns* network simulator from Lawrence Berkeley National Laboratory (LBNL) [4], with extensions from the MONARCH project at Carnegie Mellon/Rice University [12]. We made use of the IEEE 802.11 MAC layer implementation as our physical transport layer. The mobile host and base stations were configured with the standard mobile node features defined by *ns*. The experiments were classified by traffic type: a) UDP and b) TCP. For UDP, a back-to-back randomly generated 532 bytes packet was sent at a constant rate of 0.8 Mb/sec to a destination in a wired network. While TCP transfers consisted of an FTP session executed from the mobile node to the wired network. The standard TCP implementation was used throughout the experiments. We measured the performance in the opposite direction and the differences between both experiments were minimal and only results of the former are presented in this paper.

The network topology consisted of a set of base stations located at 250, 500, 750, and 1000 meters away from each other, as depicted in Figure 1. The separation distance between the base stations was considered an important experimental variable, provided that at different speeds the mobile host will have different rendezvous periods with the cell covered by a specific base station and that coverage area gaps might be present throughout the trains trajectory. In addition, infrastructure cost might be a factor that requires certain spacing between base stations, and hence the importance of this variable.

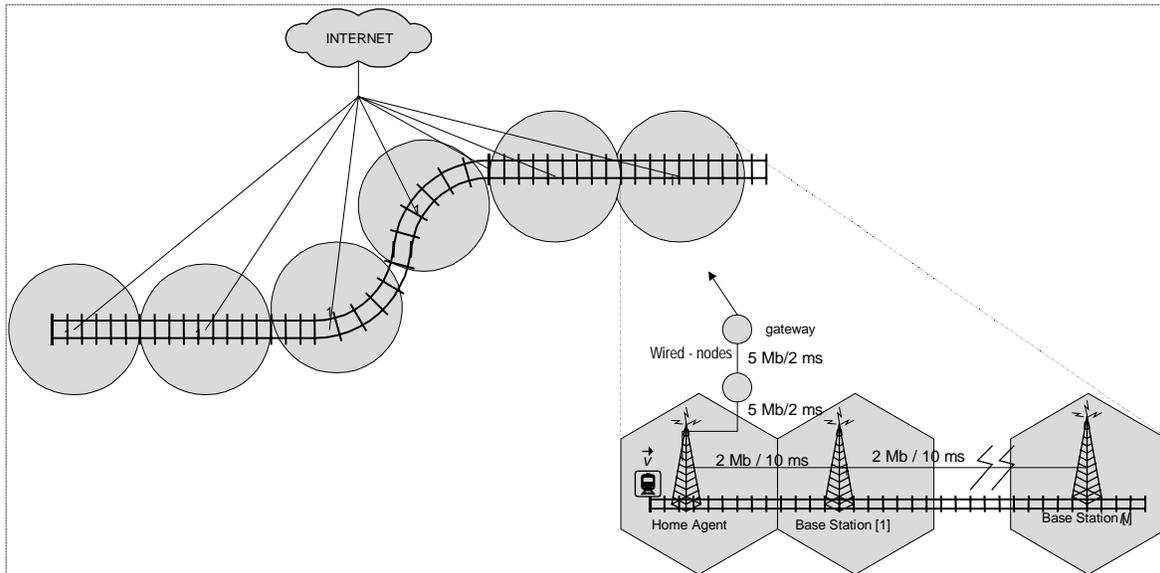


Figure 1. Network topology for a train environment simulated with *ns*.

The architecture depicted in Figure 1 represents a realistic design that may be used to cover a railroad track of a train. All foreign agents share a bus instead of a star topology, which is unrealistic. The links between base stations are of 2 Mb/s and 10ms delay, the wired links are 5 Mb/s and 2 ms delay connections. All traffic is generated from the mobile node to the farthest wired node show in the figure. We simulated one mobile host moving from the origin (the home network) towards a destination following a straight line, traversing a set of foreign nodes at a constant speed ranging from 0 to 80 m/s. It is not practical to simulate more than one mobile node, since internally in each compartment a different LAN might

coexist and the simulation only represents a mobile-bridge or a router conducting the traffic from all the mobile stations inside the train cars.

3. EXPERIMENTS AND RESULTS

The main goal of this experimental study is to measure the effect of speed and the interleaving of the base stations on the overall performance of the Mobile-IP protocol. The experiments, which were performed according to the topology presented in Figure 1, measured the network throughput at the destination node, overall packet loss, and the overhead generated during handoff. The overhead of the registration process or binding was measured as well at the mobile node, at different speed values. In this paper, we present four experiments and their results:

- The first experiment measures the throughput of TCP transfers between the mobile unit and a destination, at different mobile unit speeds and over a range of base station interleaving distances
- Similarly, the second experiment measures throughput performance of UDP packet transfers at a constant bit rate, also at different mobile unit speeds and over a range of base station interleaving distances
- The third experiment quantifies and analyzes the packet loss error rate in TCP and UDP transfers, under rapid mobility
- Finally, the fourth experiment analyzes the delay overhead and the effect of speed on the performance of the Mobile IP protocol and on the cost-effectiveness of the wireless network topology design.

3.1. Throughput Performance for TCP Transmissions

The effect of the interleaving distance of base stations can be observed in Figure 2. The figure depicts the TCP behavior at the destination node while the mobile node is moving at a constant speed of 10 m/s. Figure 2.a represents a highly overlapped cell configuration. As expected, handoff occurs constantly, which degrades throughput performance. This represents a waste of resources since on the average, the throughput remains the same as if the cells had interleaving distances of 500 or 750m. Figures 2.b, 2.c, and 2.d describe similar experiments, where the base stations are separated by 500, 750, and 1000 m, respectively. The forwarding delay is another factor that can be easily confused with congestion delays. In fact, this behavior is observed at higher speeds, where the hopping rate among base stations increases. Therefore, there is a need to modify the retransmission times in TCP as well as the beacons according to both the speed and the base station topology pattern.

Experimental measurements showed that the coverage area of a base station in the simulator was about 506 m. Therefore at 10 m/s, the dwell time of the mobile unit in any cell is 50 sec. Our results show that the best throughput performance is achieved at a speed of 10 m/s and a distance of 500 m.

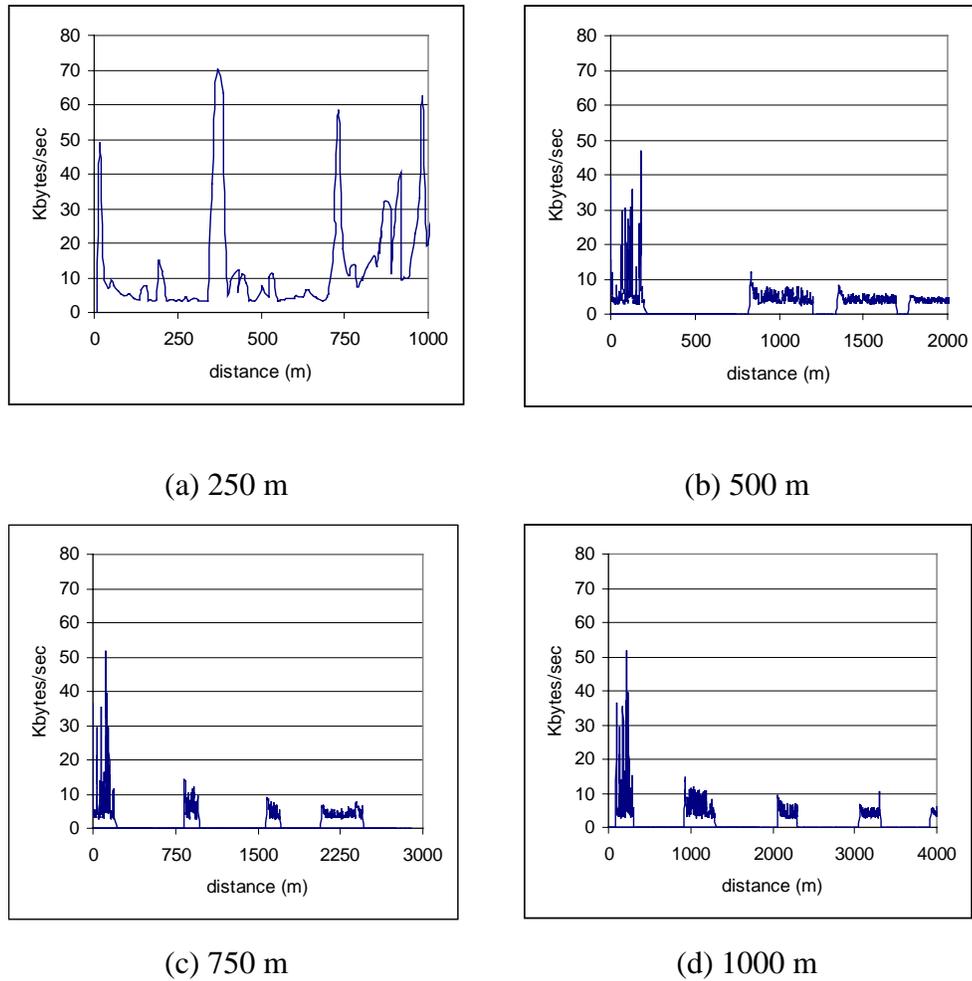


Figure 2. Throughput of FTP traffic and interleaving distances of base stations
(Speed = 10 m/s)

Figure 3 depicts the average throughput and the time spent in handoff at different network topologies and configurations. Figure 3.a shows the average TCP throughput, which in all cases decays to a minimum value, 0 Kbytes/sec, at a speed of 80 m/sec. Figure 3.b depicts the percentage of usable time at the mobile host, or the time not in handoff. As shown, a distance of 750m between base stations improves the usable time of the mobile host to transfer information at 40 m/s. Additionally, highly overlapped cells with interleaving distances of 250 m, performs much worse than interleaving cells at 500m, at speed of 20 m/s.

The worst-case scenario was observed at 80 m/s where no transfers were registered and the TCP throughput was almost zero. The main cause of error was observed in the ARP packets being dropped (45%). A second cause of error that was observed is packet loss at the queue interface (12%), since the packets can never be routed from the mobile host to the destination network, and the remainder packets were unable to find the proper route. At this speed, the mobile unit has traversed several cells before it could reach a steady state. The registration time plus the round-trip time of the link are such that the time spent in a cell ≈ 12 sec is not enough to reach the destination host. In addition to the observed causes of errors,

the transport protocol interprets the increase in the round-trip time as congestion. The adaptation process is unaware that the multi-hop topology increases the latency by hopping and due to registration and forwarding delays after every handoff event.

We conclude that the main bottleneck is the lack of awareness of the rapid mobility environment on the part of the TCP and Mobile-IP protocols. Mobile-IP and its inability to determine the proper location of the mobile unit limits the performance of the communication links at high speeds. We will show the effect of the forwarding delay in the next experiment by substituting UDP traffic for the TCP source used in this experiment.

Figure 3.d shows that at a speed of 40 m/s the usable time is the same for cells interleaved at 750m and 250m. In other words, a configuration of 20 base stations and 750m of cell interleaving distance will have the same effect of a configuration of 60 base stations separated by 250m. This observation can be used to reduce installation cost of the wireless infrastructure, assuming that the mobile unit travels at a constant speed of 40 m/s.

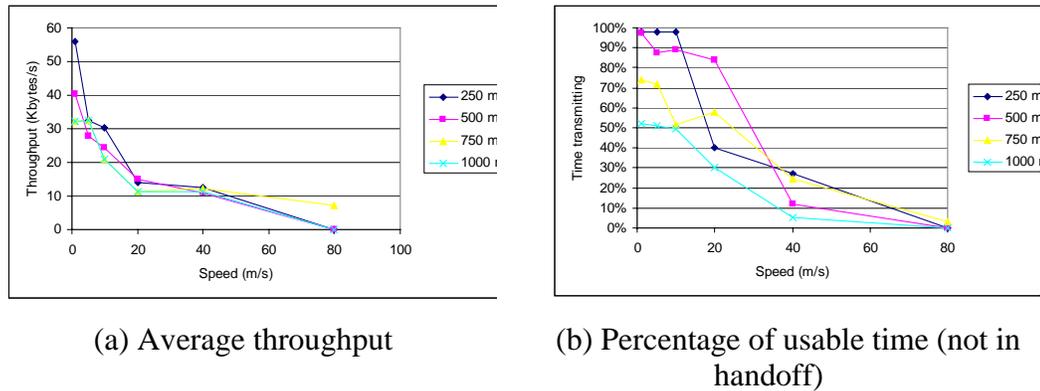


Figure 3. Throughput behavior at different speeds and base station separation distances with FTP traffic

3.2. Throughput Performance for UDP Transmissions

Similar to the experiments performed using TCP traffic, several applications require UDP to establish voice or video transfers. We used a CBR input at random intervals with an average of 200 packets/sec and a packet size of 512 bytes. Figure 4 shows the throughput at a speed of 20 m/s with highly overlapped cells, or 250 m of interleaving distances between two base stations and the non-overlapping cells at distances of 1000 m.

Contrary to TCP, the UDP protocol allows us to determine the forwarding delay and how the mobile-IP affects the performance of mobile computing at high speeds.

Figure 4 depicts the throughput at the mobile host, which shows that when it reaches the farthest base station, the average throughput remains almost unchanged, and only the handoff time is affected. This differs from the behavior presented by the TCP protocol in the same figure. In addition, UDP presented a much higher packet loss ranging from 8 to 64%. Being a connection-less protocol, UDP shows no evidence of throughput degradation, and it consistently converged to a constant throughput value, for the same speed and different cell interleaving distances.

Furthermore, Figure 4.a shows that there are at least 3 base stations with no use during the registration process and first handoff period. The horizontal axis indicates the distance in meters, and the points presented denotes the location of each base station. It is shown that at any cell interleaving distance the percentage of use of the cell is a little bit more than 50%, in other words, half of the cell is not used properly and this is mainly due the forwarding delay of the Mobile-IP protocol. For example in Figure 4.d, a base station located at a distance of 2000m results in UDP transfer to start at 1900 m instead of 1500 m. The covered area is wasted on no less than 40%. This phenomenon acknowledges the fact that the *home agent's* lack of awareness of the base station interleaving pattern and speed of the mobile hosts could be costly in terms of delays and installation costs.

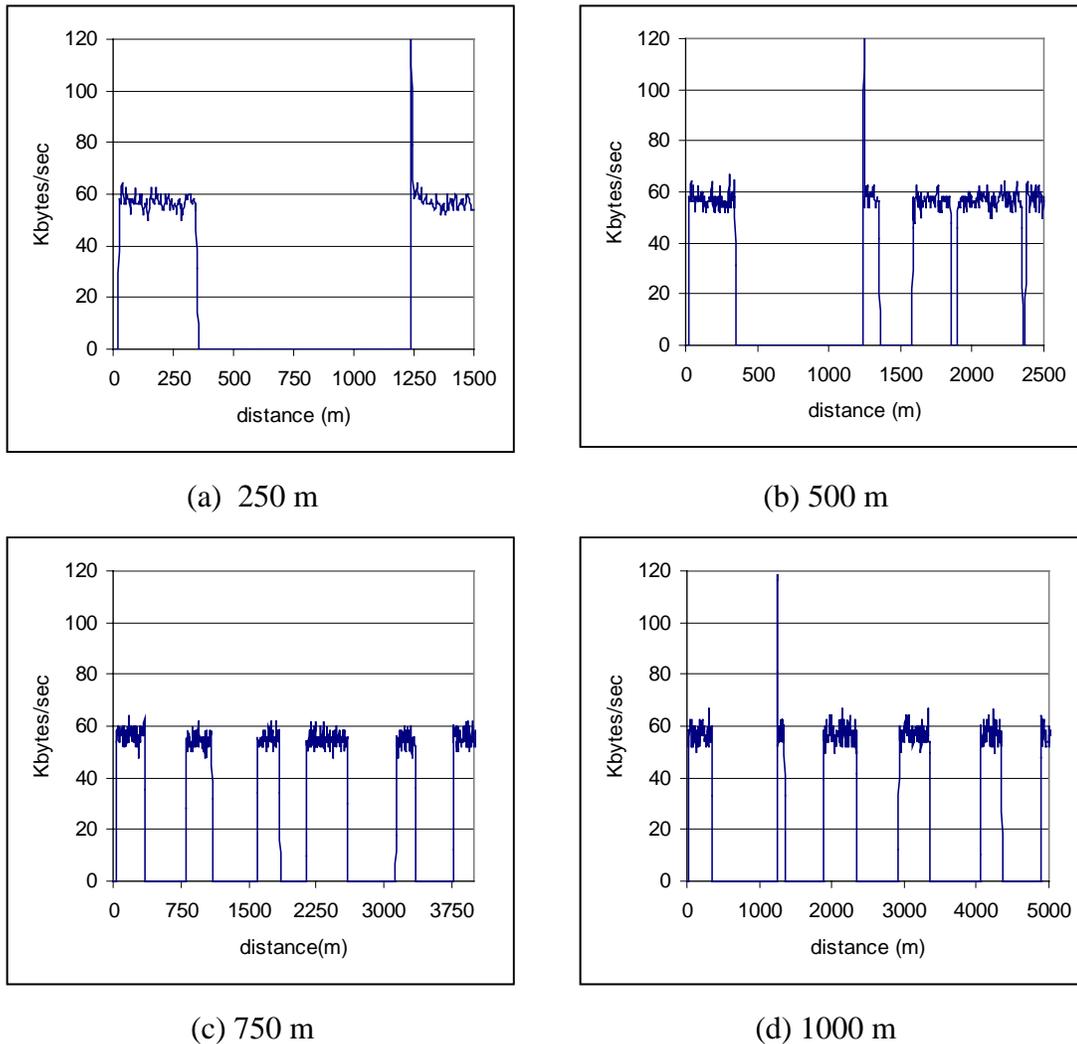


Figure 4. Throughput of UDP traffic and interleaving distances of base stations
(Speed = 20 m/s)

Figure 5.a shows that the mobile unit presents lower average throughput at higher speeds, while Figure 5.b depicts the percentage of usable time at the mobile host while covered by a

base station. As shown, base stations separated by 750m presented 10% less usable time than the ones at 500m, which is very similar to the behavior of TCP transfers at 40 m/s.

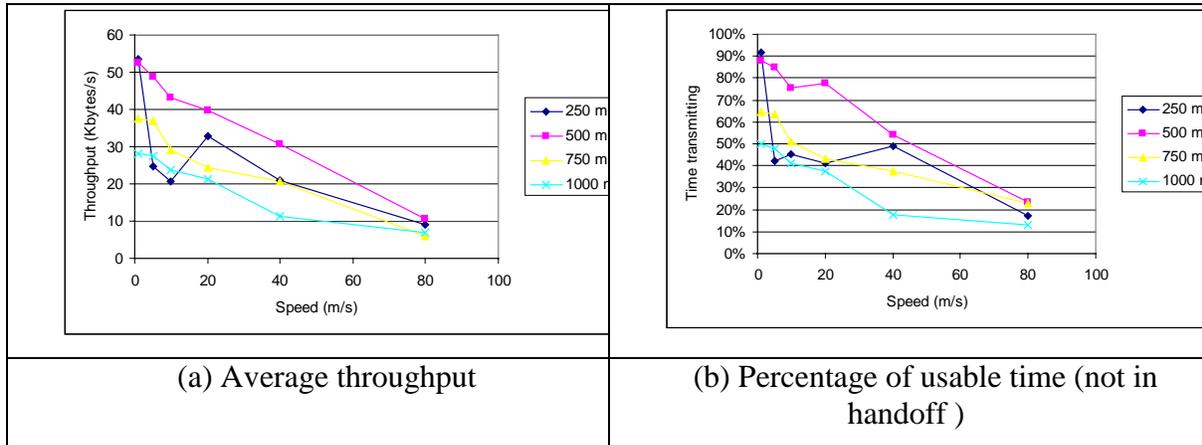


Figure 5. Throughput behavior at different speeds and base station interleaving using UDP traffic

3.3. Packet Loss Errors

Packets are continuously dropped during handoff, and sporadically dropped during data transfers performed through the wireless media. Figure 6 illustrates the packet loss in TCP and UDP transmissions. It is expected that for TCP the packet loss is lower than UDP since the former is a connection-oriented protocol. As shown, packet loss for TCP at 80 m/s was almost 0% since the throughput was zero as well. The percentage of packets dropped during the UDP session ranged from 0 to 60%. As shown, the error is reduced to 20-30% at 80 m/s which might be a consequence of the reduction in the amount of packets transferred during 20% of usable time, and not 40% as in the 40 m/s case (Figure 5.b).

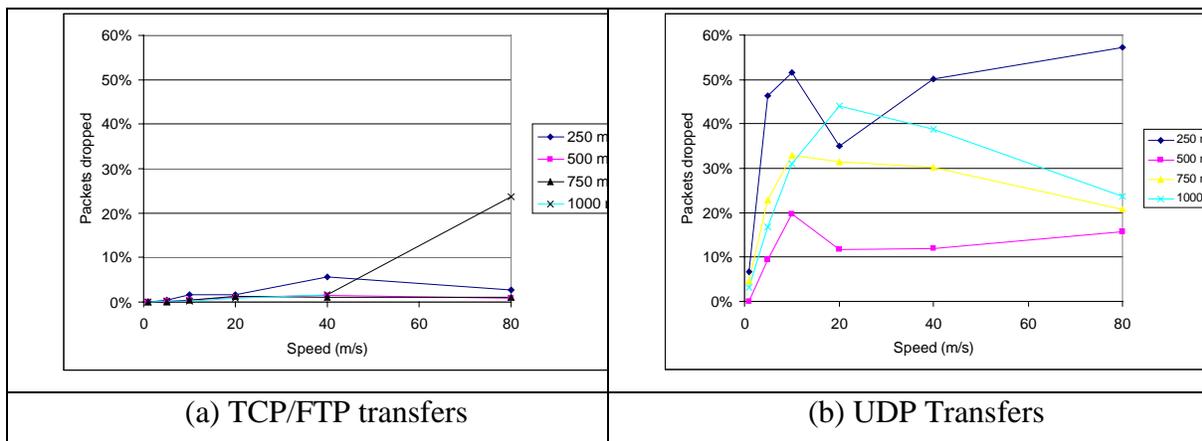


Figure 6. Percentage of packets dropped using UDP and TCP traffic

3.4. Speed and Handoff Overhead

The speed of a mobile unit and the forwarding delay from the home agent determines how the cell interleaving can be reduced. Empirically, a methodology to establish metrics for locating the base stations can be established. Figure 7 presents a mobile host moving at a constant speeds ranging from 1 to 80 m/s while involved in UDP transfers. Many cells are not active during the communication and handoff at higher speeds, therefore, the proper setting for positioning the cells is very important. The designer could use knowledge of the expected speed to reduce base station installation cost or to power off cells at high speeds, and consequently reduce power consumption.

The measurements presented in this section correspond to the UDP traffic employed for section 4.2. As shown in Table 1, the registration time takes from 37 to 47s. This time translated to distance is shown in Figure 7. At lower speeds, less than 10 m/s (36 km/hr), this is not a problem and registration occurs in the adjacent cell. However, at higher speeds, greater than 20 m/s, there are a number of cells traversed by the mobile unit that are not being used during the registration process. This observation also matches many cells that are “not seen” by the mobile-unit and this number depends upon the speed and mobility pattern of the train. In fact, if the cell size and the speed values range between, 20 and 80 m/s, the registration occurs at the 3rd, 4th, and 10th cells, respectively. For example, in Figure 7.f, the registration process occurs at 4000 m from the home agent, or the 10th cell. Therefore the spacing between cells and the awareness of the mobile-IP protocol to execute this routing is required for both the TCP and the UDP protocols.

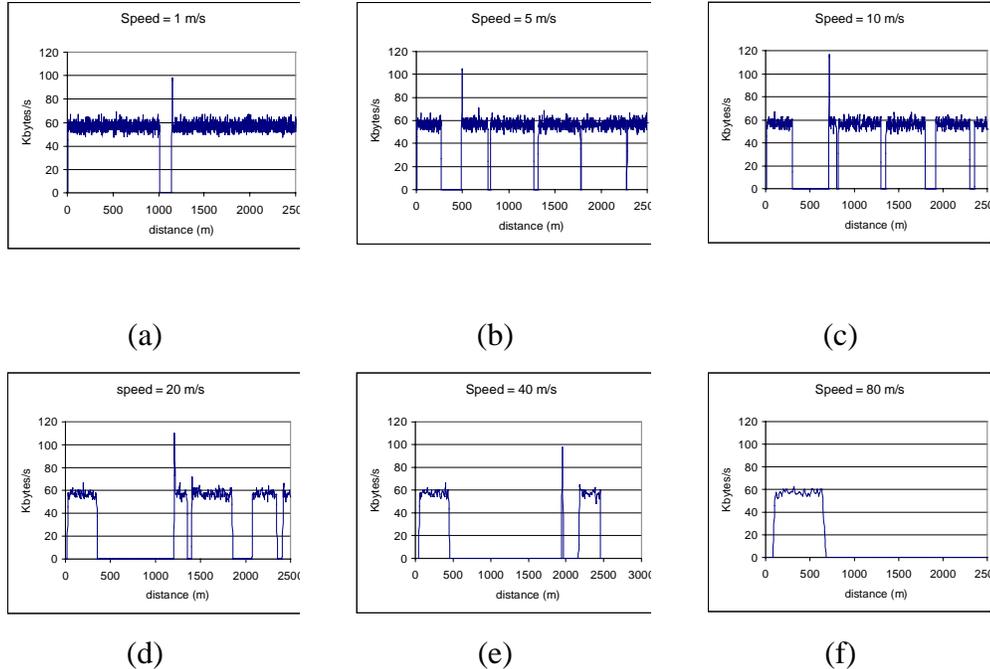


Figure 7. Registration delay and the effect of speed with cell interleaving of 500m. (UDP traffic)

Table 1 also shows the forwarding delay at the cell where the first rendezvous occurs between the mobile unit and the base station. At a speed of 80 m/s, the forwarding delay corresponds to 248 ms. This forwarding delay increases when all the foreign agents are in a multi-hop network. Under the assumptions presented in Figure 1, the delay increase is a consequence of the forwarding process, which affects the topology throughout the trajectory of the train.

*Table 1. Important factors during Mobile-IP handoff and registration
(UDP traffic, $d = 500$ m)*

Factor	Speed					
	1 m/s	5 m/s	10 m/s	20 m/s	40 m/s	80 m/s
Time required for first registration	45 sec	41 sec	45 sec	44 sec	36 sec	47 sec
Minimum handoff time	7 sec	5 sec	1 sec	1 sec	1 sec	3 sec
Minimum Forwarding delay	22 ms	28 ms	40 ms	60 ms	98 ms	248 ms
Distance from the home agent to achieve registration	298 m	459 m	705 m	113 8 m	170 3 m	446 0 m

4. CONCLUSIONS AND FUTURE WORK

Our main conclusion is that the Mobile-IP protocol requires awareness of infrastructure configuration as well as mobile unit speed in order to function properly. Registration and triangular routing generate a large overhead, which affects the communication process at speeds greater than 20 m/s (72 Km/hr). The design of the wireless infrastructure requires a-priori knowledge of the protocols employed as well as speed characteristics of the mobile hosts. Cells can be interleaved at different distances and configurations depending on the speed and mobility behavior of the mobile units. We observed that providing full wireless coverage is unnecessary and a waste of resources. The positioning of cells depends upon the train trajectory and speed. At speed of 40 m/s the location of cells at 250m or 750 m produced the same effect in throughput and the percentage of usable time within a cell. Therefore cost of the wireless infrastructure can be optimized or the topology adapted such that cells can be turned off when not required.

In fact, a wireless, adaptive topology of cells would be ideal to use in this case of a commuter train that moves at different speeds. Maximum cells (or cell interleaving) can be used at low speed (entering or departing a station). Less number of cells can be used at medium speeds (going through construction and semaphore zones). Finally, minimum cells can be used at high speed (inter-city crossing).

Mobile networking protocols, such as Mobile-IP, are not designed to handle high-speed gracefully. Such protocols produce considerable overhead and high forwarding delay. We found out that protocols based on registration and non-aware packet re-routing are not appropriate for speeds higher than 20 m/s.

Future work can be aimed at creating an adaptive topology environment, which configures according to the speed and mobility conditions of vehicles such as trains. Improvements can be made to the Mobile-IP protocol, as well as the TCP implementation. Predictability of speed and location of a mobile unit should be included in the protocol design as well in the re-configurable architecture. This will enable better scheduling of packet transfers, thus improving the usable time of cells (currently about 50%), under rapid mobility.

5. REFERENCES

- [1] B. Bing. "High-speed wireless ATM and LANs", Artech, Boston, MA, 2000.
- [2] J. Solomon. Mobile-IP. Prentice Hall, New Jersey, 1998.
- [3] C. Perkins, IP mobility support, RFC 2002, IBM, October 1996.
- [4] G. Holland and N. Vaidya. Analysis of TCP performance over Mobile Ad-Hoc networks, ACM Mobicom'99, Seattle, WA, 1999, pp 219-230.
- [5] K. Fall, K. Varadhan, editors. NS notes and documentation. The VINT project, LBL, February 2000. <http://www.isi.edu/nsnam/ns/>
- [6] M. Gerla, C. Chiang, and L. Zhang. "Tree multicast strategies in mobile, multihop wireless networks", Mobile Networks and Applications, Vol.4, No. 3, pp. 193-207, October 1999
- [7] A. Flademuller and R. De Silva. "The effect of Mobile IP handoffs on the performance of TCP", Mobile Networks and Applications, Vol. 4, No. 2, pp 131-135, May 1999.
- [8] R. Caceres and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environments. IEEE Journal of Selected Areas in Communications, Vol 13, No. 5, November 1996.
- [9] H. Balakrishnan, S. Seshan, and R. Katz. "Improving reliable transport and handoff performance in cellular wireless networks", ACM Wireless Networks Vol. 1, No. 4, December 1995.
- [10] H. Balakrishnan, V. Padmanabhan, S. Seshan, and R. Katz. A comparison of mechanisms for improving TCP performance over wireless links. Proc ACM SIGCOMM, Stanford, CA, August 1996.
- [11] V. Jacobson, "Congestion avoidance and control", Computer Communications Review, Vol. 18, No. 4, pp 314 – 329 , August 1988.
- [12] J. Broch, D. B. Johnson, D. A. Maltz, et.al." A performance comparison of multi-hop wireless ad-hoc network routing protocols", ACM/IEEE Int. Conf. On Mobile Computing and Networking, pp. 85-97, October. 1998.
- [13] J. Widmer. Network simulations for a mobile network architecture for vehicles. "TR-00-09, International Computer Science Institute", Berkeley, CA, May 2000.