A Dummy Segment Based Bandwidth Probing Technique to Enhance the Performance of TCP over Heterogeneous Networks

Tarik Taleb, Kenichi Kashibuchi, Nei Kato, and Yoshiaki Nemoto
Graduate School of Information Sciences - Tohoku University
Sendai, Japan 980–8579
Email: taleb@nemoto.ecei.tohoku.ac.jp, {buchiken, kato}@it.ecei.tohoku.ac.jp, nemoto@nemoto.ecei.tohoku.ac.jp

***Abstract*— Communication over wireless heterogeneous networks is still a challenging task. The underlying reason beneath this is in the waste of available bandwidth or the bursts of packet drops that may occur during handoff due to the disparity in the amount of available bandwidths among wireless cells. In mobile environments, the fundamental challenge upon a handoff phenomenon consists in an efficient probing of the availability of the new network resources and an appropriate rate adjustment in the new network cell.**

As a remedy to the above issue, this paper argues the usage of low-priority dummy packets to probe the availability of the new network resources. Indeed, when a mobile node enters a cell overlapping area and is about to change its point-of-attachment to the network, two connections are simultaneously set between the mobile node and the sender: one through the old point-of-attachment and another through the new one. The sender transmits actual data through the old connection. Meanwhile, it sends dummy segments through the new connection to verify the bandwidth availability of the new network. The proposed scheme is dubbed *Dummy Segment based Bandwidth Probing (DSBP)*.

The performance of the DSBP scheme is evaluated and compared with existing schemes through extensive simulations. The simulation results show that the DSBP scheme substantially improves the system efficiency, reduces the number of packet drops, and makes better utilization of the network bandwidth.

I. Introduction

Wireless communication systems have been gaining ground at a tremendous pace during the last few years. Due to the universality of the Internet Protocol (IP), an in-depth understanding of IP protocols and recognition of their merits and drawbacks in wireless networks are of vital importance. This understanding underpins the research work outlined in this paper.

Wireless networks exhibit different characteristics from the wired counterparts. The success of wireless networks in delivering high-speed access hinges on the ability of the underlying Internet protocols to function correctly and efficiently in wireless networks. The effect of such a communication environment on the working of Transmission Control Protocol (TCP), which forms the backbone of today's Internet protocol communication [1] [2], is the focus of this research work.

Communication over wireless networks is well characterized by frequent handoff occurrences. As mobile users perform handoff over fairly long periods of time and data packets can be lost during these periods, TCP performs poorly in mobile environments. Some propositions have been made in the recent literature to ameliorate the performance of TCP in such environments. Freeze-TCP [3] is a notable example. While the Freeze-TCP scheme enhances the performance of TCP when handoff occurs between homogeneous access media, its performance remains remarkably limited when mobile users roam over heterogeneous networks. The reason beneath this poor performance consists in the fact that upon a handoff occurrence, a Freeze-TCP sender starts transmitting data in burst after a short pause. This operation overloads the new network with packets and ultimately results in congestion, excessive queuing delays, and a large number of packet drops.

When a mobile node changes its point-of-attachment to the network, the amount of bandwidth available at the new point-of-attachment may be different than that of the old one. This bandwidth disparity can be due to difference in traffic load in both wireless cells. This mismatch in available bandwidths is further exacerbated as different wireless access techniques with different link speeds, such as WLAN, W-CDMA, and Bluetooth, are being deployed to provide ubiquitous connectivity. In general, when a mobile node performs handoff, two scenarios can be envisioned. If the mobile node moves from a higher bandwidth network (ex. Wireless LAN) to a lower bandwidth network (GPRS), and continues transmitting data without any adjustment in its sending rate, the new network will be congested and a potential number of packets will be dropped. The connection throughput will be eventually degraded. On the other hand, if the mobile node enters a higher bandwidth network, no adjustment to the sending rate of the mobile node will lead to a waste of the network bandwidth and ultimately lower network utilization.

The fundamental challenge in wireless communication over highly heterogeneous networks consists in adjusting the sending rate of mobile nodes to the available bandwidth in the new network. In the proposed mechanism, this adjustment is based on the transmission of a burst of low-priority dummy packets instead of actual data packets, similar in spirit to the idea of [4]. These dummy segments are used to probe the availability of the new network resources. Extensive simulations were conducted to evaluate the performance of the proposed scheme. The results reveal that the proposed scheme substantially improves the performance of TCP over heterogeneous networks;

the scheme provides higher throughput and lower packet drops rate.

The remainder of this paper is structured as follows. Section II highlights the relevance of this work to the state-of-art in the context of performance of TCP over heterogeneous wireless networks. The key design philosophy and distinct features that were incorporated in the proposed scheme are described in Section III. Section IV portrays the simulation environment and reports the simulation results. Following this, the paper concludes in Section V with a summary recapping the main advantages and achievements of the proposed scheme.

II. Related Work

After nearly two decades since its standardization, the TCP protocol has seen deployment at unforeseen scale. Despite the widespread of TCP/IP in terrestrial networks, its performance over wireless networks is still limited due to a number of reasons related to the protocol syntax and semantics [5]. The remainder of this section describes the main post-standard improvements that have been devised in recent literature to overcome the shortcomings of TCP in wireless networks.

In the recent literature, several approaches have been proposed to counter the shortcomings of TCP in wireless networks. They can be classified into two categories: wireless link errors related techniques and handoff disconnection related techniques. Methods of the first category aim to solve the problem of throughput degradation due to wireless transmission errors. As a matter of fact, TCP was initially conceived for wired networks where the link error rate is low, such that the majority of packet drops is due to network congestion. A TCP sender operates hence on the assumption that any packet losses are caused by congestion. Accordingly, the sender throttles the transmission by bringing the congestion window down to the minimum size and entering the slow-start phase. In recent years, many researchers have investigated possible solutions to this link errors related problem. [6] gives an overview of the recent research activities in this area.

On the other hand, methods of the second category tackle issues related to handoffs. Indeed, to deal with packet losses that are due to temporary disconnections that may occur when a mobile node moves to a new network, a set of mobility management techniques has been recently proposed. The major objective of most of these techniques is to reduce the packet loss during the handoffs due to the broken data path from the server to the destination. These techniques can be classified, in turn, into two categories: network layer mobility management techniques and transport layer mobility management techniques.

Network layer mobility management techniques can be classified also into two groups [7]. In the first group, when a handoff occurs, old base station caches and forwards the packets to the new base station based on a request to forward the packets. Most pioneering examples that use this technique are Fast Handovers Mobile IP [8] and HAWAII [9]. In the second group, packets are routed to multiple nearby base stations around the mobile node to ensure delivery of the packets to the mobile node. In addition to the recently proposed multipath smooth handoff scheme [10], multicast mobility support [11] and bicast used in Cellular IP [12] use this technique.

Among transport layer mobility management techniques, several TCP extensions have been proposed to provide reliable data transmission by avoiding unnecessary timeouts, and hence reducing shrinkage events of TCP window during handoffs. Two new concepts tailored to this area are worth mentioning: TCP Westwood [13] and Freeze-TCP.

TCP Westwood is a sender-side modification of the TCP congestion window algorithm. It improves upon the performance of TCP Reno in wired as well as wireless networks. The key concept of TCP Westwood is to use an estimate of the available bandwidth to set the congestion window and slow start threshold after a congestion episode. A TCP Westwood source performs end-to-end estimate of the bandwidth available along a TCP connection by measuring and averaging the rate of returning ACKs. Whenever the sender perceives a packet loss, inferred by a timeout occurrence or reception of a certain number of duplicate ACKs, the sender uses the bandwidth estimate, an approximate of the effective bandwidth at the congestion time, to select the optimum values of the congestion window and the slow start threshold. By so doing, TCP Westwood ensures faster recovery and efficient utilization of network resources. While TCP Westwood has been shown to be efficient in hybrid wired and wireless networks with high bandwidth, large propagation, yet light random loss links, it performs poorly in environments with high loss rates. TCP Westwood with Bulk Repeat [14] has been proposed as an enhancement of TCP Westwood to combat its poor performance in heavy loss environments.

On the other hand, the general idea of Freeze-TCP is to change the TCP algorithm in the mobile host so that the base station can be prevented from sending packets during handoff. The key concept is to move the onus of signaling an impending disconnection to the mobile host side instead of the base station side [3]. This is practicable as a mobile host can easily monitor signal strengths and detect an impending handoff, and even predict a temporary disconnection. If a handoff occurs, the mobile host advertises a zero window size to force the sender into frozen mode and to prevent it from dropping its congestion window size. The sender does not change its window size when it starts retransmitting data. As a Freeze-TCP sender restarts transmission with the old window size in burst upon entering a new point-of-attachment, the performance of Freeze-TCP remains limited when handoff occurs between heterogeneous networks. To cope with such an issue, a large body of bandwidth probing techniques have been proposed to measure the available bandwidth in the new network [15]. Most of these pioneering techniques require, however, accurate measurement of propagation delay. Under heavy traffic load, an accurate estimate of propagation delay is usually not possible to obtain; a fact that ultimately leads to erroneous estimates of the available bandwidth.

III. The Dummy Segment based Bandwidth Probing Approach

This section presents an operational overview of the proposed scheme, DSBP. Before delving into details of the proposed scheme, we present an outline of the core ideas behind the scheme and its requirements.

First, it should be emphasized that the DSBP scheme exploits the fact that cells in wireless networks overlap with each other. The considered network is assumed to be end-to-end Quality-of-Service (QoS) enabled (e.g., wireless ATM and IntServ supported networks). In fact, the proposed scheme requires that all network elements in the connection path support some priority disciplines. Currently, most networks are best effort and most routers in the Internet do not apply any priority policy. However, in the near future, through the Differentiated Service Model (DiffServ) [16], routers will be able to support multiple service classes. As inter-networking over wireless systems is still a virgin field, implementing DiffServ on wireless networks should not be an issue. To allow a mobile node to simultaneously register multiple Care-of-Addresses (CoAs), mobile IP simultaneous binding option [17] [18] is used. On the other hand, to keep senders always informed of these CoA registrations directly from the mobile nodes, the route optimization option [19] is used. Recall that each CoA identifies a connection path from the sender to the mobile node and each new CoA binding update reports either the set-up or the loss of a CoA address. It should be noted also that the new CoA of the mobile node in the new cell is different from its old CoA used in the old cell.

The key concept behind the proposed scheme is that as a mobile node moves into a cell overlapping area, a new connection between the mobile node and the sender through the base station (BS) of the new network cell is established (Fig. 1). Obviously, the sender becomes aware of the new BS based on a new CoA binding update message from the Mobile IP protocol. Before reaching the middle point of the overlapping area (where the handoff usually takes place), the mobile node keeps on receiving data from the sender using the old connection through the old BS. Simultaneously, the sender sends *rwnd* dummy segments to the mobile node through the new BS, where *rwnd* is the receiver window size that limits the maximum value of congestion window. The value of *rwnd* indicates the rate at which the sender transmits dummy segments to the mobile node. The algorithm of the proposed scheme is based on the concept of using these dummy segments to probe the availability of network resources without carrying any new information to the sender. This concept was first proposed in [20] and has been used since then in several researches in the recent literature. Notable examples are TCP-Peach [21] and TP-Planet [22].

Dummy segments are generated by the sender as a copy of the last transmitted data packet. They are treated as low priority segments and accordingly they do not affect the delivery of the actual data traffic. Indeed, when a router on the connection path is congested, IP packets carrying dummy segments are discarded first. Overhead of these dummy segments in terms of bandwidth consumption is, therefore, not an issue.

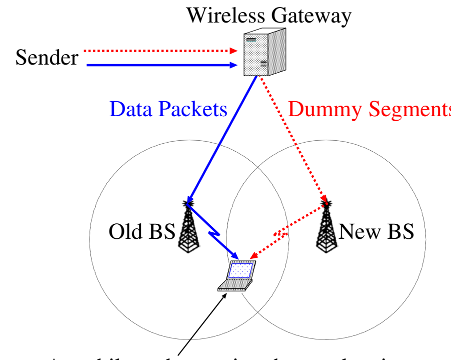

Fig. 1. Upon a handoff occurrence, two connections are simultaneously set between the mobile node and the sender upon its entrance to the cell overlapping area: One through the old BS and the other through the BS of the new network cell.

Upon reception of a dummy segment, the mobile node can recognize it. In response to each dummy segment, the mobile node transmits a dummy acknowledgment (ACK) to the sender. Dummy ACK packets indicate the availability of network resources in the new cell. In response to each dummy ACK, the sender transmits, in turn, an actual data packet to the mobile node. To distinguish dummy segments from actual data packets, dummy segments and their acknowledgments (ACKs) are marked using one or more of the six unused bits in their TCP headers. A simple modification of the TCP implementation is thus required in the end-terminals. ACKs for dummy segments are used to provide an efficient probing of the bandwidth availability in the new network. As a result, the senders can adjust their sending rates to the most appropriate value within one round trip time (RTT). They either increase their transmission rates to make full utilization of the new network resources or decrease their transmission rates to avoid overloading the new network with bursty traffic, and thus causing higher packet drops and significant throughput degradation.

Admittedly, reception of dummy segments and transmission of dummy ACKs by mobile nodes result in additional energy consumption. The DSBP scheme may be seen, thus, as costly in terms of reducing the battery life of mobile nodes. However, the performance gains achieved by the proposed scheme in Section IV-B, are worthwhile and can be used to advocate for this additional cost. Indeed, the obtained higher throughput and lower packet loss rates lead to significant reduction in the overall transmission time of a given data file. This intuitively reduces the overall usage time of the mobile node battery and ultimately saves its energy. Moreover, apart from the rare case of mobile nodes *flip-flopping* over the cell overlapping area, the additional energy consumption due to dummy segments remains, by and large, minimal.

IV. Performance Evaluation

A. Simulation Set-up

This section gives a detailed description of the simulation environment, justifying the choices made along the way. The

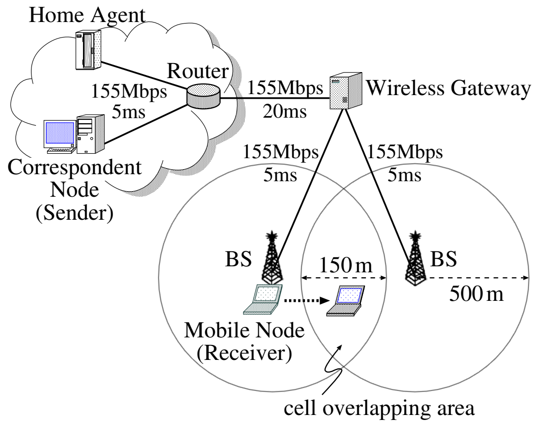

Fig. 2. Simulation network topology

design of the simulation setup relies on Network Simulator (NS) [23]. Particular attention is, thus, paid to the design of an accurate and realistic one. Unless otherwise noted, the parameters specified below are those used in all the experiments throughout the paper.

While it is more general to consider a sequence of handoffs, the behavior of the proposed scheme is best understood by considering individual handoffs. In all simulations, focus is thus on analyzing a single handoff instance. Handoffs are assumed to occur over the overlapping area of two wireless cells; which is the most common case. In the performance evaluation, TCP NewReno [24], TCP Westwood-NR[1], and Freeze-TCP are used as comparison terms. The reason behind the choice of the TCP Westwood-NR and TCP NewReno among other TCP implementations underlies beneath the fact that TCP NewReno achieves faster recover from multiple losses within the same window and has the potential of significantly improving TCP's performance in the case of bursty losses. Furthermore, since we consider a wireless environment with light random loss, comparison of the DSBP scheme to the TCP Westwood-NR, rather than to the TCP Westwood with Bulk Repeat, is satisfactory. The data packet size is fixed to 1 kB. Due mostly to its simplicity and its wide usage in today's switches and routers, all routers use Drop-Tail as their packet-discarding policy. Simulations were all run for 180 s, a duration long enough to ensure that the system has reached a consistent behavior. All results are an average of multiple simulation runs.

In the proposed scheme, mobile nodes are assumed to issue a CoA binding update message as soon as they reach the cell overlapping area. They are assumed to perform handoff when they reach the middle line of the overlapping area. Difference in the performance of the proposed DSBP scheme, TCP NewReno, TCP Westwood-NR, and Freeze-TCP is more interesting when mobile nodes travel over the cell overlapping area. Focus is thus on the performance of the four protocols during the handoff period; the time period during which a mobile node stays in the cell overlapping area.

The abstract configuration of the considered network is depicted in Fig. 2. The wireless part of the network consists of two neighboring wireless cells. The coverage radius of each wireless cell is set to 500 meters. To have the longest distance across the overlapping area equal to 150 meters, the distance between the two neighboring base stations is fixed to 850 meters. These parameters are chosen with no specific purpose in mind and do not change any of the fundamental observations about the simulation results. The wireless domain is connected to the wired network through a single wireless gateway. The choice of a single wireless gateway serving the old and new cells represents a general and simple case. Indeed, considering a topology where two cells are served by two different wireless gateways will simply increase the connection RTT and shall has no influence on the overall performance of the four considered schemes. All wired links are given a capacity equal to 155 Mbps (e.g. OC3). As for the wireless links capacity, a number of test scenarios were created by setting their capacity to different values. All links are presumed to be error-free throughout this paper. This assumption is made so as to avoid any possible confusion between throughput degradation due to packet drops and that due to wireless channel errors. Unless otherwise stated, the round trip time of the connection is set to 60 ms. Table I shows a complete list of the simulation parameters and the range of values studied.

Factor	Simulation Parameters and Range of Values
Cell Coverage Radius	500 m
Overlapping Area Longest Distance	150 m
Wired Links Capacity	155 Mbps
Old Wireless Network Capacity	6 Mbps
New Wireless Network Capacity	1 Mbps – 11 Mbps
Connection RTT	60 ms
Mobile Node Speed	10 km/h – 100 km/h
Packet size	1024 Bytes

TABLE I
SIMULATION PARAMETERS

Two quantifying parameters are used to evaluate the performance of the proposed scheme: average throughput and loss rate. Throughput indicates the number of bytes received by a mobile node during the handoff period. The loss rate is the ratio of the dropped packets to the aggregate sent packets over the handoff period. It should be emphasized that dropped dummy segments were not considered in the loss rate computation.

B. Simulation Results

Having described the simulation parameters, we now direct our focus to evaluating the performance of the proposed scheme through extensive simulation experiments. We first investigate the impact of the bandwidth availability in the new cell on the performance of the proposed scheme. The bandwidth of the old cell is set to 6 Mbps, whereas the bandwidth of the new cell is varied from 1 Mbps to 11 Mbps. The moving speed of the mobile node is set to 50 km/h.

Fig. 3(a) graphs the throughput of the four schemes as a function of the difference between the available bandwidths in the new and old networks. The figure shows that the DSBP

[1]The NewReno based version of TCP Westwood

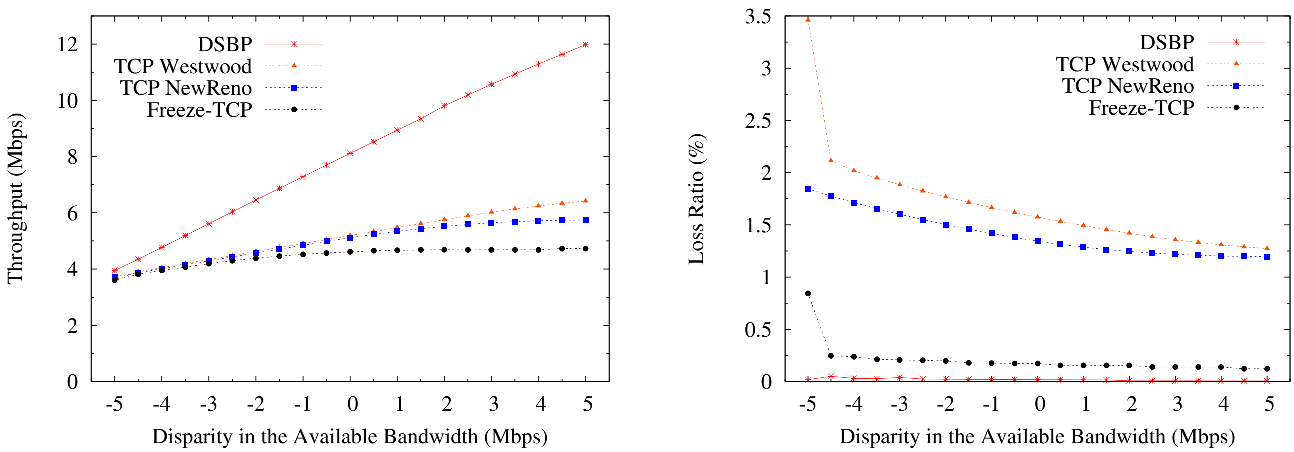

(a) Throughput

(b) Packet drops rate

Fig. 3. Transmission efficiency vs disparity in the available bandwidth (Mobile node speed = 50 km/h)

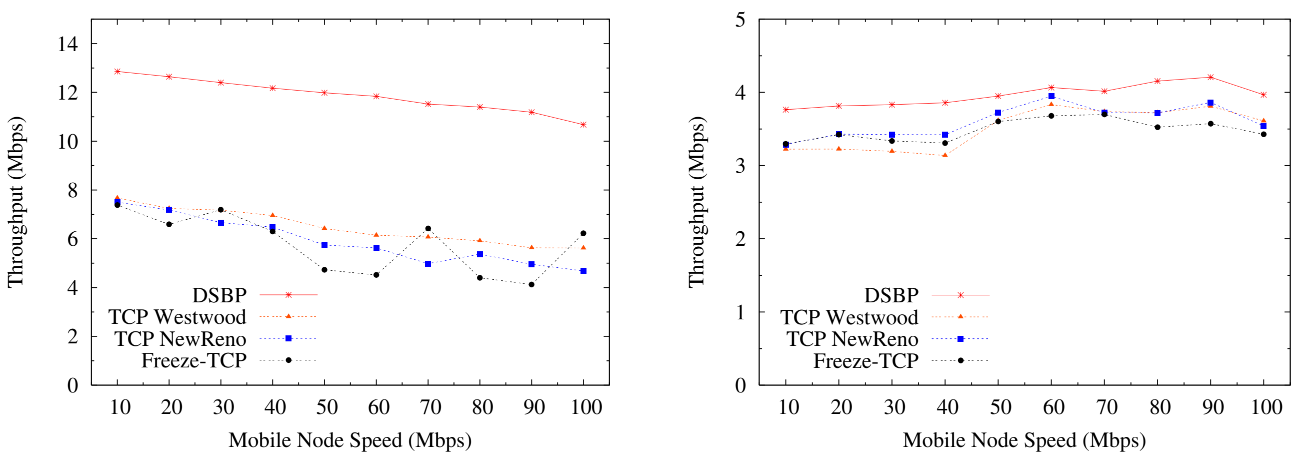

(a) Case I: Handoff from a low bandwidth cell (6 Mbps) to a high bandwidth cell (11 Mbps)

(b) Case II: Handoff from a high bandwidth cell (6 Mbps) to a low bandwidth cell (1 Mbps)

Fig. 4. Throughput variation for different mobile node speeds

scheme achieves the highest throughput compared to TCP Westwood-NR, Freeze-TCP, and TCP NewReno. Furthermore, its throughput exhibits an abrupt increase. When the available bandwidth in the new network is lower (the range of negative values on the x-axis), the four simulated schemes exhibit smaller throughputs. This is simply due to the fact that the bandwidth in the new network becomes less available. On the other hand, when the new cell has a higher bandwidth (the range of positive values on the x-axis), the proposed scheme gains up to more than 200% over the three other schemes. This significant gain is mainly due to the fact that the dummy segments inform the sender of the extra-bandwidth within a single round trip time and stimulate it accordingly to increase its sending rate.

To illustrate the performance of the four schemes in terms of packet drops, we plot in Fig. 3(b) the packet loss rate as a function of the difference between the available bandwidths in the new and old networks. The results show that the proposed DSBP scheme and Freeze TCP achieve the lowest packet drop rate. The proposed scheme outperforms further the Freeze TCP scheme and achieves almost *zero* drops regardless of the available bandwidth of the new cell. The main reason beneath this performance is in the intrinsic characteristic of the proposed scheme. Indeed, DSBP uses dummy segments to estimate the optimum rate the sender should send data at. Accordingly, the sender avoids overloading the network with data packets that would ultimately be dropped otherwise. The results show also that TCP Westwood-NR and TCP NewReno experience the highest packet drop rate. This poor performance is mainly due to their bursty nature. In fact, both schemes keep transmitting data with window sizes that can not be accommodated by the new network. This leads to congestion, higher drops, and ultimately throughput degradation. It is also worth noting that although TCP Westwood-NR exhibits a slightly higher throughput than TCP NewReno (Fig. 3(a)), it experiences more packet drops. This performance is mostly attributable to the failure of TCP Westwood-NR in obtaining a good estimate of the available bandwidth in the new cell based on only the arrival rate of ACK packets. In summary, since the proposed scheme uses the dummy segments to exploit the maximum available bandwidth of the new cell, the proposed scheme achieves the highest throughput and the lowest drop rate compared to the three other schemes.

To investigate the impact of the length of the handoff period on the throughput of the proposed scheme, we vary the mobile node speed from 10 km/h to 100 km/h. It should be noted that the length of the handoff period becomes longer as the mobile node speed decreases, and vice versa, and hence the influence of the mobile node velocity on the system performance. In this experiment, two scenarios are considered. The first scenario is when the mobile node moves from a low bandwidth cell (6 Mbps) to a high bandwidth cell (11 Mbps). The second scenario considers the case of a handoff from a high bandwidth cell (6 Mbps) to a low bandwidth cell (1 Mbps). Figures 4(a) & 4(b) show the throughput of the four schemes for different speeds of the node. We notice that the throughput of the four schemes decreases as the mobile node speed increases. The reason behind this performance underlies beneath the fact that the handoff period becomes shorter and the time required for the probing phase becomes accordingly shorter. In the first scenario, we observe that the DSBP scheme significantly outperforms the three other schemes and that is for all considered speeds. We remark also that its throughput remains the highest in both scenarios.

V. Conclusion

In this paper, we proposed a dummy segment based bandwidth probing technique to enhance the performance of TCP over heterogeneous wireless networks. The basic idea behind the proposed scheme is to use dummy segments to probe for bandwidth availability in wireless networks. When a mobile node enters a cell overlapping area and is about to perform handoff, the sender establishes a new connection with the mobile node through its future point-of-attachment to the network. Dummy segments are sent through this connection to probe the bandwidth of the new cell. These dummy segments are low-priority packets and do not affect the actual delivery of data. To guarantee a smooth transmission of data, the sender meanwhile preserves the transmission of actual data through the old connection. The proposed scheme has the potential of efficiently adjusting the sending rates of mobile nodes in new network cells.

The performance of the proposed scheme was investigated through extensive simulations. Performance evaluation relied on computer simulation and a set of scenarios was considered. Obtained performance measures included transmission efficiency related metrics; throughput and packet loss ratio. The obtained results were encouraging and elucidated both the high throughput and low packet drops rate of the proposed scheme.

Acknowledgment

The authors take great pleasure in acknowledging the financial aid from the Koden Electronics Co., LTD, to conduct a part of this research work.

References

[1] K. Thompson, G. J. Miller, and R. Wilder, "Wide-Area Internet Traffic Patterns and Characteristics," *IEEE Network Magazine*, Vol. 11, No. 6, pp. 10-23, November/December 1997.

[2] N. Brownlee and K. C. Claffy, "Understanding Internet Traffic Streams: Dragonflies and Tortoises," *IEEE Commun. Magazine*, Vol. 40, No. 10, pp. 110-117, October 2002.

[3] T. Goff, J. Moronski, D. S. Phatak, and V. Gupta, "Freeze-TCP: A True End-to-End TCP Enhancement Mechanism for Mobile Environments," in *Proc. of the IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, Vol. 3, pp. 1537-1545.

[4] G. Bianchi, A. Capone, and C. Petrioli, "Throughput Analysis of End-to-End Measurement-Based Admission Control in IP," in *Proc. of the IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, Vol. 3, pp. 1461-1470.

[5] V. Tsaoussidis and I. Matta, "Open Issues on TCP for Mobile Computing," *The Journal of Wireless Communications and Mobile Computing*, Wiley Academic Publishers, Vol. 2, No. 1, pp. 3-20, February 2002.

[6] F. Hu and N. K. Sharma, "Enhancing Wireless Internet Performance," *IEEE Commun. Surveys & Tutorials*, Vol. 4, No. 1, pp. 1-15, Third Quarter 2002.

[7] A. T. Campbell, J. Gomez, S. Kim, Z. Turanyi, C-Y. Wan, and A. Valko, "Comparison of IP Micro-Mobility Protocols," *IEEE Wireless Commun. Magazine*, Vol. 9, No. 1, pp. 72-82, February 2002.

[8] R. Koodli, "Fast Handovers for Mobile IPv6," work in Progress, Internet Draft, Internet Engineering Task Force, 2002.

[9] R. Ramjee, K. Varadhan, L. Salgarelli, S. Thuel, S.Y. Wang, and T. La Porta, "HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-Area Wireless Networks," *IEEE/ACM Trans. on Networking*, Vol. 10, No. 3, pp. 396-410, June 2002.

[10] Y. Pan, M. Lee, J. B. Kim, and T. Suda, "An End-to-End Multi-Path Smooth Handoff Scheme for Stream Media," in *Proc. of the* 1^{st} *ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots*, San Diego, CA, USA, September 2003, pp. 64-74.

[11] A. Helmy, "A Multicast-Based Protocol for IP Mobility Support," in *Proc. of the ACM SIGCOMM* 2^{nd} *International Workshop on Networked Group Commun.*, Palo Alto, CA, USA, November 2000, pp. 49-58.

[12] A. T. Campbell, J. Gomez, S. Kim, A. G. Valko, C-Y. Wan, and Z. R. Turanyi, "Design, Implementation and Evaluation of Cellular IP," *IEEE Personal Commun.*, Vol. 7, No. 4, pp. 42-49, August 2000.

[13] S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, and R. Wang, "TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links," in *Proc. of ACM SIGMOBILE 2001*, Rome, Italy, July 2001, pp. 287-297.

[14] G. Yang, R. Wang, M. Y. Sanadidi, and M. Gerla, "Performance of TCPW BR in Next Generation Wireless and Satellite Networks," in *Proc. of IEEE Int. Conf. Commun., ICC 2003*, Anchorage, Alaska, USA, May 2003.

[15] M. Jain and C. Dovrolis, "End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput," in *Proc. of the ACM SIGCOMM 2002*, Pittsburgh, PA, August 2002, pp. 295-308.

[16] A. T. Campbell, "QoS-Aware Middleware for Mobile Multimedia Communications," *Multimedia Tools and Applications*, Kluwer Academic Publishers, Vol. 7, No. 1-2, pp. 67-82, July 1998.

[17] W. Fritsche and F. Heissenhuber, "Mobile IPv6: Mobility Support for Next Generation Internet," *IPv6 Forum*, white paper, 2000.

[18] C. Perkins, "IP Mobility Support for IPv4," RFC 3344, August 2002.

[19] R. Vadali, J. Li, Y. Wu, and G. Cao, "Agent-Based Route Optimization for Mobile IP," in *Proc. of the IEEE Vehicular Technology Conference 2001 Fall*, Atlantic City, NJ, USA, October 2001, Vol.4, pp. 2731-2735.

[20] V. N. Padmanabhan and R. Katz, "TCP Fast Start: A Technique for Speeding up Web Transfer," in *Proc. of the IEEE GLOBECOM '98 Internet Mini-Conference*, Sydney, Australia, November 1998.

[21] I. F. Akyildiz, G. Morabito, and S. Palazzo, "TCP-Peach: A New Congestion Control Scheme for Satellite IP Networks," *IEEE/ACM Trans. on Networking*, Vol. 9, No. 3, pp. 307-321, June 2001.

[22] O. B. Akan, J. Fang, and I. F. Akyildiz, "TP-Planet: A Reliable Transport Protocol for Interplanetary Internet," *IEEE J. on Selected Areas in Commun.*, Vol. 22, No. 2, pp. 348-361, February 2004.

[23] VINT Project, "UCB/LBNL/VINT Network Simulator - ns (version 2)," `http://www.isi.edu/nsnam/ns/`

[24] S. Floyd and T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm," RFC 2582, April 1999.