USHA: A Practical Vertical Handoff Solution

Ling-Jyh Chen, Tony Sun, and Mario Gerla
Computer Science Department, UCLA
{cclljj,tonysun,gerla}@CS.UCLA.EDU

***Abstract*— We present a seamless handoff solution, called Universal Seamless Handoff Architecture (USHA). USHA is simple and requires minimal modification to the current Internet infrastructure. Therefore, it is instantly ready for real-world deployment. Using testbed experiments, we evaluate USHA by observing TCP (FTP downloading) behavior in various vertical handoff scenarios. The results shows that USHA can successfully maintain the application connectivity in all cases. Furthermore, the results also shows that when handoff from a low capacity link to a high capacity link, there is no service latency caused by the handoff; whereas when the handoff is from high capacity link to a lower capacity link, the non-negligible latency could not be alleviated unless early handoff notification can be provided.**

I. INTRODUCTION

In view of the proliferation of mobile applications, a universal seamless handoff solution across wireless domains is becoming increasingly important. While wired connections usually provide high speed, reliable access to the Internet, wireless networking technologies enable users to access customized Internet services even when they are moving. A seamless handoff solution with both low latency and low packet loss is mandatory for mobile users who wish to receive continuous, uninterrupted Internet service while frequently switching from one network connection to another. Additionally, the handoff solution should be network-layer-transparent and infrastructure-modification-free so that existing Internet server and client applications can painlessly survive the rapid pace of wireless technology evolution.

Seamless handoff aims to provide continuous connection for an end-to-end data transmission in the presence of any link outage or handoff event. Low latency and low packet loss are the two most critical design issues. Low latency requires that the switch to the new path be completed almost instantaneously; service interruption should be minimized to provide the illusion of continuous connectivity. In the event of actual connection failure, the architecture should attempt to re-connect as soon as the service becomes available, and the packet loss due to the connection switch should also be minimized.

In this paper, we present Universal Seamless Handoff Architecture (USHA), a simple handoff solution that satisfies the requirements for seamless handoff. Compare to other seamless handoff solutions, USHA is simple and requires minimal changes to the current Internet infrastructure. Therefore, it is instantly ready for real-world deployment. Using testbed experiments, we evaluate USHA by observing TCP (FTP downloading) behavior during vertical handoff scenarios. The results shows that USHA can successfully maintain the application connectivity in all cases. Furthermore, the results also show that when handoff from a low capacity link to a high capacity link, there is no service latency caused by the handoff; whereas when the handoff is from high capacity link to a lower capacity link, the non-negligible latency could not be alleviated unless early handoff notification can be provided.

The rest of the paper is organize as follows. In section II, we summariz related work and recap the definition of vertical handoff. Section III describes the enabling mechanisms behind USHA. Section 4 presents experiment results from various vertical handoff scenarios. Section 5 discusses USHA deployment issues as well as potential enhancements that are applicable to USHA. Section 6 concludes the paper.

II. BACKGROUND AND RELATED WORK

A vertical handoff is a process of switching the ongoing network connection from one interface to the other [15], as shown in Fig. 1. With the eventual unification between cellular and IP networks, it is increasingly desirable to provide seamless vertical handoff support for mobile users. The biggest challenges in designing a seamless vertical handoff lies in maintaining an established application sessions, such as a connection-oriented flow (e.g. TCP), as well as achieving a smooth (low loss) and fast (low delay) data transmission.

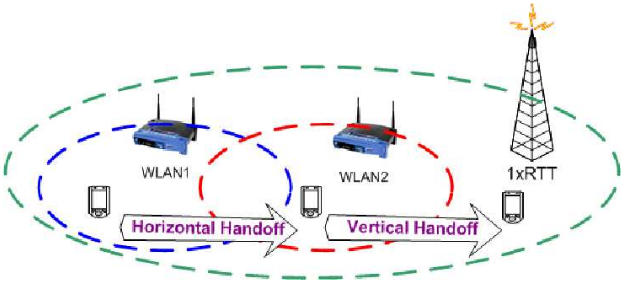

Fig. 1. Horizontal and Vertical Handoff

Mobile IP is a network layer solution for vertical handoff scenarios, and it is the current IETF standard for Internet mobility support [11]. Mobile IP ensures the delivery of packets destined to the mobile host (using home address), instead of directly sent to the mobile hosts via its current Internet address (i.e. care-of address). An address translator is deployed on the home agent, and an IP tunnel is carefully maintained from the home agent to the foreign agent; therefore, after handoff

This work is funded by NSF under NRT grant ANI-0335302.

to the foreign agent network, all ongoing traffic, which is destined to the home agent, will be forwarded to the foreign agent (and then to the mobile host) through the IP tunnel. As a result, the handoff will be transparent to the upper layer applications, and the connectivity can be kept. On the other hand, IPv6 [2], the succeeding network layer protocol, provides native mobility support. By combining home address and care-of address into a 128 bits Internet address, IPv6 can easily support vertical handoff scenarios without deploying foreign agents and address translators on the home agents.

Besides network layer solutions, there are also several handoff solutions using transport layer approaches [1], [10], [13], [16]. [1] proposed the use of the Indirect TCP (I-TCP) to deal with the handoffs by splitting the connection into two segments: one is the fixed (wired) connection, and the other is the mobile (wireless) connection. However, I-TCP suffers from the high expense in deployment and the loss of end-to-end semantics. [10], [16] use multi-homing technique to support vertical handoff. However, the deployment of these two solutions requires upgrading both transport layer and applications on both mobile hosts and Internet servers. Therefore, the deployment cost is still too high to become feasible (not very clear, replace existing TCP). [13] proposed another end-to-end approach to support host mobility via dynamically updating Domain Name System (DNS) and adding a set of Migrate options to TCP. However, it not only requires updating the mobile hosts, it also calls for an upgrade for all DNS servers in the world, which is an impracticality that can't be realized.

Additionally, several proxy based solutions are proposed to handle the network heterogeneity problem, this includes mobility and vertical handoff [4], [9], [12], [14]. Using a split-connection proxy, the mobile host can virtually maintain application sessions even with the presence of vertical handoffs. The end-to-end semantics can also be preserved, since the handoff is transparent to the transport and application layers. However, such proxy based solutions suffer from scalability problems, due to the limited resource and computation capability of individual proxy.

III. UNIVERSAL SEAMLESS HANDOFF ARCHITECTURE (USHA)

We propose Universal Seamless Handoff Architecture (USHA) to deal with both horizontal and vertical handoff scenarios with minimal changes in infrastructure (i.e., USHA only requires deployment of handoff servers on the Internet.) USHA is a mobile IP-less solution; however, instead of introducing a new session layer or a new transport protocol, it achieves seamless handoff by following the middleware design philosophy [3], integrating the middleware with existing Internet services and applications.

USHA is based on the fundamental assumption that handoff, either vertical or horizontal, only occurs on over-laid networks with multiple Internet access methods (i.e. soft handoff), which translates to zero waiting time in bringing up the target network interface when the handoff event occurs. If coverage from different access methods fails to overlap (e.g. hard handoff), it is possible for USHA to lose connectivity to the upper layer applications.

In Fig. 2, a handoff server (HS) and several mobile hosts (MHs) are shown. USHA is implemented using IP tunneling techniques (IP encapsulation), with the handoff server functioning as one end of the tunnel and the mobile host as the other. An IP tunnel is maintained between every MH and the HS such that all application layer communications are "bound" to the tunnel interface instead of any actual physical interfaces. All data packets communicated through this IP tunnel are encapsulated and transmitted using the connectionless UDP protocol.

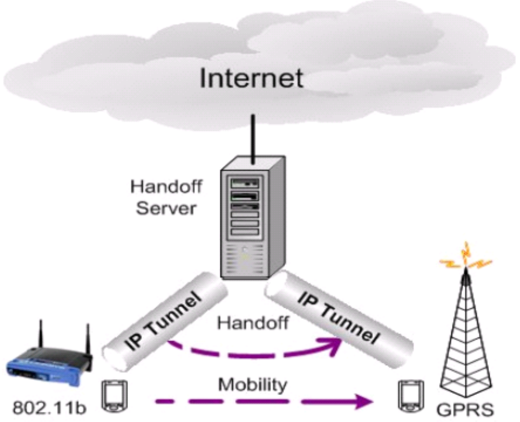

Fig. 2. Diagram of Universal Seamless Handoff Architecture

The IP tunnel above utilizes two pairs of virtual/fixed IP addresses, one on HS and one on MH. The fixed IP addresses are necessary for an MH to establish a physical connection to the HS. When the handoff event occurs and the physical connection from MH to HS changes, the MH is responsible for automatically switching the underlying physical connection of the virtual tunnel to the new interface, as well as notifying the HS of its change in physical connection. Upon handoff notification, the HS immediately updates its IP tunnel settings so that any subsequent data packets will be delivered to MH's new physical link.

Since all data packets are encapsulated and transmitted using UDP, there is no need to reset the tunnel after the handoff. Therefore, end-to-end application sessions (e.g. TCP) that are bound to the IP tunnel are kept intact. This provides handoff transparency to upper layer applications.

IV. EXPERIMENTS

In this section, we present USHA measurements in two vertical handoff scenarios. In subsection IV-A, we first evaluate the "indoor-mobility" scenario, which is common for mobile users performing handoffs between 100Mbps Ethernet (e.g. his/her working cubicle) and 802.11b (e.g. conference room) network interfaces. In subsection IV-B, we evaluate USHA in the "outdoor-mobility" scenario, where the mobile user performs vertical handoffs between different wireless technologies (e.g. 802.11b in the cafe and 1xRTT on the freeway).

A. Vertical Handoff between Ethernet and 802.11b

Fig. 3 illustrates the “indoor-mobility” scenario, where the mobile user performed a vertical handoff between Ethernet (100Mbps link capacity) and 802.11b (11Mbps mode[1]). A FTP (TCP based) download session was initiated at the beginning of each experiment and stopped after two minutes. The vertical handoff was manually triggered after the first second. The detail configuration of the experiment scenario, including link capacity (measured by CapProbe [7]) and round-trip delay (of 1500 bytes packet size), are also shown in Fig. 3.

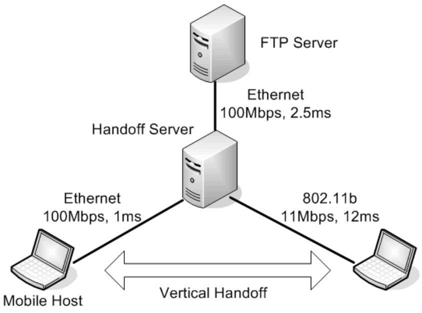

Fig. 3. Testbed configuration of the vertical handoff experiment between Ethernet and 802.11b.

1) Handoff From 802.11b to Ethernet: Fig. 4 and 5 shows the instantaneous FTP (TCP) throughput and packet sequence number during a vertical handoff (which is occurred at around 60th second) from 802.11b to Ethernet, i.e. from 11Mbps capacity link to 100Mbps link. The TCP throughput increased from 4Mbps (before the handoff) to around 35Mbps (after the handoff), and the packet sequence number increased without any additional latency due to the handoff. From the results, it’s clear that USHA can truly achieve seamless vertical handoffs.

2) Handoff From Ethernet to 802.11b: In addition to the handoff from LOW to HIGH capacity link, we also evaluated FTP downloading during a vertical handoff from Ethernet to 802.11b link, i.e. from 100Mbps to 11Mbps capacity link. Fig. 6 and 7 show the instantaneous throughput and the corresponding packet sequence numbers. Here, there is a non-negligible latency of around 5 seconds during the vertical handoff. This is due the fact that TCP suffers multiple packet losses caused by the drastic reduction in link capacity. When multiple packets are dropped, TCP reduces its congestion window and probes the new available bandwidth with its *Slow Start* mechanism. As identified by [5], such latency could not be alleviated unless early explicit handoff notification were provided.

B. Vertical Handoff between 1xRTT and 802.11b

The second vertical handoff scenario is the “outdoor-mobility” scenario, where the mobile users perform vertical handoffs between wireless technologies that are designed to

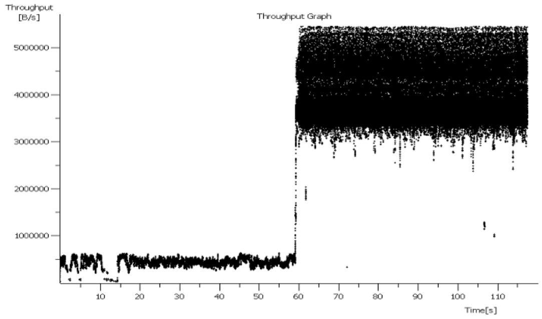

Fig. 4. Instantaneous throughout results of one TCP flow during a vertical handoff from 802.11b (11Mbps) to Ethernet (100Mbps).

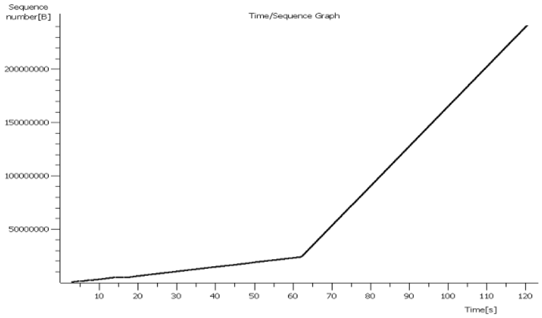

Fig. 5. Sequence number of one TCP flow during a vertical handoff from 802.11b (11Mbps) to Ethernet (100Mbps).

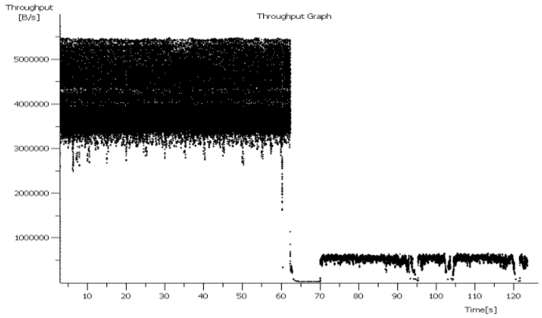

Fig. 6. Instantaneous throughout results of one TCP flow during a vertical handoff from Ethernet (100Mbps) to 802.11b (11Mbps).

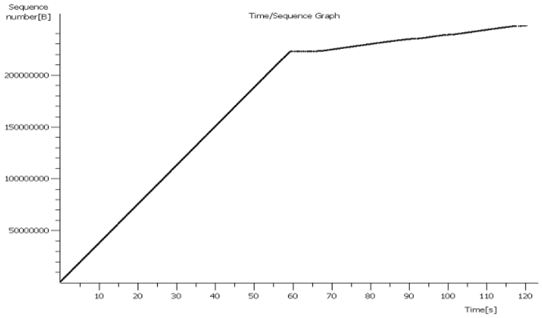

Fig. 7. Sequence number of one TCP flow during a vertical handoff from Ethernet (100Mbps) to 802.11b (11Mbps).

---

[1] The *effective capacity* of 11Mbps transmission rate mode is around 5∼6 Mbps. [7]

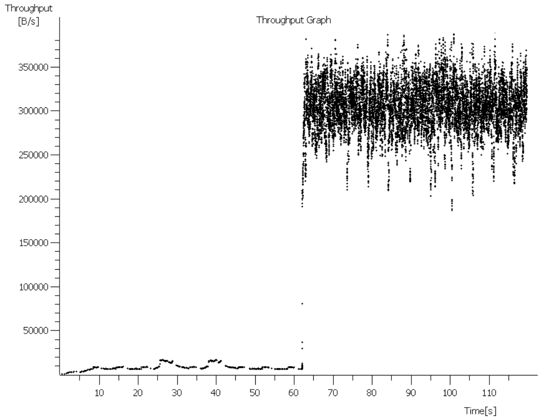

Fig. 8. Instantaneous throughout results of one TCP flow during a vertical handoff from 1xRTT (150Kbps) to 802.11b (5.5Mbps).

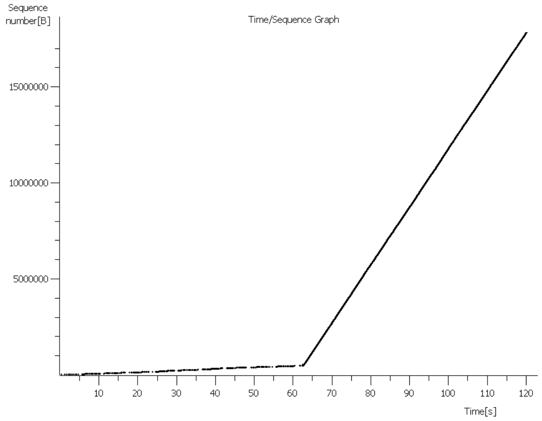

Fig. 9. Sequence number of one TCP flow during a vertical handoff from 1xRTT (150Kbps) to 802.11b (5.5Mbps).

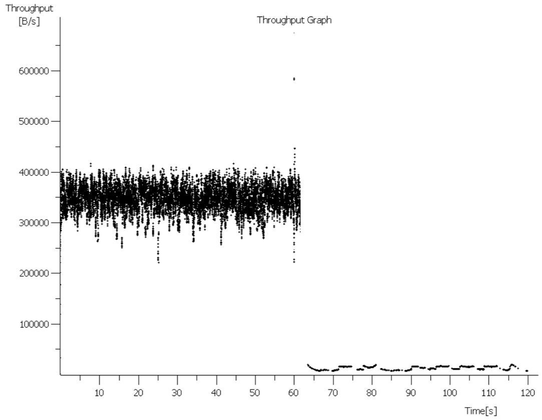

Fig. 10. Instantaneous throughout results of one TCP flow during a vertical handoff from 802.11b (5.5Mbps) to 1xRTT (150Kbps).

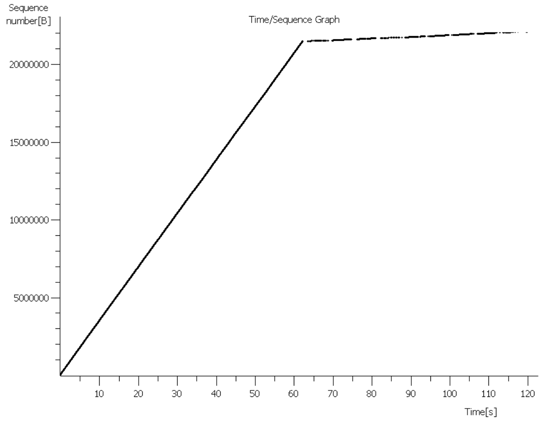

Fig. 11. Sequence number of one TCP flow during a vertical handoff from 802.11b (5.5Mbps) to 1xRTT (150Kbps).

support greater user mobility. For simplicity, we decide to use two popular wireless technologies, namely 1xRTT and 802.11b in our experiments. While 1xRTT provides greater coverage, 802.11b is able to achieve higher data throughput.

The testbed configuration is similar to Fig. 3, except here we replace Ethernet link with 1xRTT link, and use 5.5Mbps transmission mode for the 802.11b link[2]. The 1xRTT link is with around 150Kbps link capacity[3] and 350 ms round-trip delay. A single FTP download session is initiated from the MH to the FTP server, and the vertical handoff is manually triggered after the first minute.

1) Handoff From 1xRTT to 802.11b: Fig. 8 and 9 shows the instantaneous FTP (TCP) throughput and packet sequence number during the vertical handoff process (which occurred around 60th second) from 1xRTT to 802.11b. The TCP throughput increased from 80Kbps to around 3Mbps after the handoff, and the packet sequence number kept on increasing without additional latency. The results further proves that USHA can provide continuous connectivity for various vertical handoff scenarios

2) Handoff From 802.11b to 1xRTT: Fig. 10 and 11 show another measurement results, where the vertical handoff is from 802.11b to 1xRTT. Similar to the results of the handoff from 802.11b to Ethernet, there is a non-negligible latency of roughly 2 seconds during the vertical handoff. Again, such latency is due to the drastic capacity reduction, because TCP needed time to recover from multiple packet losses.

V. DISCUSSION

We have presented the technical details of USHA approach and evaluated USHA in various vertical handoff scenarios. In this section, we discuss deployment issues associated with USHA, as well as potential enhancements applicable to USHA.

A. Choosing the "Best" Handoff Server (HS)

To deploy an USHA system in the real world, a MH requires registration with at least one HS, allowing the HS to manage

---

[2] The *effective capacity* of 5.5Mbps transmission rate mode is around 3∼4 Mbps.

[3] The ISP claims 150Kbps link capacity, but CapProbe [7] only measures around 90Kbps capacity.

an IP tunnel from itself to the MH. In the case where there are multiple HS available, it is necessary for a MH to determine the "best" HS to register to. The decision process can be modelled as follows.

First of all, the MH should collect a set of candidate HS in accordance to some pre-defined policy and cost. For instance, the MH should filter out the HSs that are not accessible to the MH's network interfaces(i.e. behind the firewall) .

Secondly, the MH should select the HS with tolerable delay and highest capacity, as the "best" HS. Specifically, suppose there are k network interfaces on the mobile host, and there are n candidate HS. $D_{i,j}$ and $C_{i,j}$ denote the round-trip delay and link capacity of the ith network interface to the jth HS. D_{thresh} is the maximum tolerable delay for the mobile user. Suppose the rth HS is the "best" HS for the mobile, the following two equations should be satisfied.

$$\max_{i=1...k} D_{i,r} \leq D_{thresh} \tag{1}$$

$$\min_{i=1...k} C_{i,r} \geq \min_{i=1...k} C_{i,j} \quad ; \quad j = 1...n \tag{2}$$

B. USHA Enhancements

In addition its support for continuous connectivity, USHA could be easily extended by inserting additional '*plug-in*s'. For instance, on the MH side, one can implement *policy-enabled handoff system* [18] to determine the "best" target interface at any given moment through a well-designed cost function. Therefore, by continuously updating the cost function, MH can automatically trigger a vertical handoff if necessary.

Moreover, the HS can be extended by incorporating the *transcoding* functionalities. More specifically, when the available bandwidth[4] of the access link (i.e. from HS to MH) is lower than the required bit rate required by applications (e.g. video streaming application), it would be of great help if HS could transcode application data to a lower bit rate. Resulting in better utilization of available network resources to provide best possible services.

Additionally, since data packets are all encapsulated and transmitted via UDP between HS and MH, one can enhance the connection security by encrypting the encapsulated UDP packets (e.g. using IPsec [6]) and improve the end-to-end data goodput by compressing all UDP packets through the IP tunnel in real-time (e.g. using LZO algorithm [8]). A more detailed account management and access control schemes are also possible with USHA, since both functionalities can be easily handled and deployed on the HS.

VI. CONCLUSIONS

In this paper, we described a simple and practical vertical handoff solution, called USHA. USHA is seamless and with minimal changes on the current Internet infrastructure. Using testbed experiments, we evaluated USHA in both indoor and outdoor mobility scenarios. The results show that USHA is able to maintain the application connectivity (i.e. TCP session) and introduce no additional latency during a vertical handoff. Moreover, we presented the deployment issues and potential enhancements of USHA. The ongoing work is to investigate the application performance in vertical handoff scenarios, where the link capacity and round-trip delay may change dramatically after a handoff, and design appropriate "adaptation" schemes for applications to more intelligently and faster adapt themselves in mobile computing scenarios.

REFERENCES

[1] A. Bakre and B.R. Badrinath, "Handoff and System Support for Indirect TCP/IP," In Proc. of 2nd USENIX Symposium on Mobile and Location-Independent Computing, 1995.

[2] S. Deering, and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 2460, IETF, Dec. 1998.

[3] A. Fox, S. Gribble, Y. Chawathe, E. Brewer, and P. Gauthier, "Cluster-based Scalable Network Services," SOSP'97, 1997.

[4] V. Ghini, G. Pau, P. Salomoni, M. Roccetti, and M. Gerla, "Smart Download on the Go: A Wireless Internet Application for Music Distribution over Heterogeneous Networks," In Proc. IEEE ICC, 2004.

[5] A. Gurtov and J. Korhonen, "Measurement and Analysis of TCP-Friendly Rate Control for Vertical Handovers," ACM Mobile Computing and Communications Review, 8(3):73-87, July 2004

[6] "IP Security Protocol (ipsec) Charter," http://www.ietf.org/html.charters/ipsec-charter.html

[7] R. Kapoor, L.-J. Chen, L. Lao, M. Gerla, M. Y. Sanadidi, "CapProbe: A Simple and Accurate Capacity Estimation Technique," In Proc. of ACM SIGCOMM 2004.

[8] "LZO real-time data compression library," http://www.oberhumer.com/opensource/lzo/

[9] D. Maltz, and P. Bhagwat, "MSOCKS: An architecture for transport layer mobility," In Proc. of IEEE Infocom, p.p. 1037-1045, March 1998.

[10] A. Matsumoto, M. Kozuka, K. Fujikawa, and Y. Okabe, "TCP Multi-Home Options," draft-arifumi-tcp-mh-00.txt, IETF Internet draft, Oct. 2003.

[11] C. Perkins, Ed. "IP Mobility Support for IPv4," RFC 3344, IETF, Aug. 2002.

[12] M. Schlaeger, B. Rathke, S. Bodenstein, and A. Wolisz, "Advocating a Remote Socket Architecture for Internet Access using Wireless LANs," Mobile Networks & Applications (Special Issue on Wireless Internet and Intranet Access), vol. 6, no. 1, pp. 23-42, January 2001.

[13] A. C. Snoeren, "A Session-Based Approach to Internet Mobility," PhD Thesis, Massachusetts Institute of Technology, December 2002.

[14] A. C. Snoeren and H. Balakrishnan. "An End-to-End Approach to Host Mobility," In Proc. of MobiCom 2000.

[15] Mark Stemm, and Randy H. Katz. "Vertical Handoffs in Wireless Overlay Networks," ACM Mobile Networking (MONET), 1998.

[16] R. Stewart et al. "Stream Control Transmission Protocol," RFC 2960, IETF, Oct. 2000.

[17] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," In Proc. of ACM IMC 2003.

[18] H.J. Wang, R. H. Katz, and J. Giese, "Policy-Enabled Handoffs across Heterogeneous Wireless Networks," Proc. of ACM WMCSA, 1999.

---

[4] For instance, we can use *Spruce* [17] to quickly estimate the available bandwidth of a given path.