Transactions on AUTOMATIC CONTROL and COMPUTER SCIENCE
Vol. 49 (63), No. 4, May 2004, ISSN 1224-600X

Delay-Centric Handover in SCTP over WLAN

Andrew Kelly*, Gabriel-Miro Muntean*, Philip Perry*, and John Murphy**

* Performance Engineering Laboratory, School of Electronic Engineering, Dublin City University, Glasnevin, Dublin 9, Ireland
Phone: (+353) 1 700-7645, E-Mail: kellya@eeng.dcu.ie, munteang@eeng.dcu.ie, perryp@eeng.dcu.ie,
WWW: http://www.perfenglab.com

** Performance Engineering Laboratory, Computer Science Department, University College Dublin, Belfield, Dublin 4, Ireland
Phone: (+353) 1 716-2921, E-Mail: murphyj@eeng.dcu.ie, WWW: http://www.perfenglab.com

Abstract ***– In the field of personal mobile communications, it has been suggested that SCTP could provide a solution to the problems encountered by the currently implemented TCP/Mobile IP scheme. By exploiting SCTP's multihoming feature to connect to several separate wireless networks concurrently, allows a Mobile Node to choose which wireless path suits the particular needs of the user application it is running. However, there is one drawback to this scenario - the current handover scheme implemented in SCTP is failure-centric in nature.***

This paper proposes an improved handover scheme for SCTP. This proposed scheme offers the benefit of performing handover based on measured path delays, thus it does not require a path failure for handover to occur. In some cases it can actually pre-empt the path failure, and handover before it occurs.

Keywords: ***SCTP, Handover, WLAN, Computer Networks, Real-time Systems***

I. INTRODUCTION

The current trend of mobile communications has facilitated the introduction of 3G, or 3rd Generation, cellular networks, such as UMTS. However, as the roll-out of these networks has, so far, been slow, a new wireless standard has emerged. IEEE 802.11 [1], the standard for Wireless Local Area Networks (WLANs), operates in the Industrial, Scientific, and Medical (ISM) band and has a number of variants. IEEE 802.11b is the most popular of these and is used in the work presented here. The use of this technology for Internet access has seen huge subscriber growth in a relatively short time period. This growth can be attributed to the relatively low cost and ease of installation of WLAN hardware, and also to the roll-out of public "hotspots": IEEE 802.11 coverage areas, in popular consumer areas, such as coffee shops, fast food restaurants, and airports.

Due to the popularity of WLAN and the transmission rates it supports (up to 11Mbps), services that were exclusive to the "wired" internet have now become available to the mobile user. Examples of these services include streaming video to a portable computer, such as a PDA, or conducting a VoIP call from a laptop. Providing these services in a mobile environment previously meant using implementations that required the user to remain quite static for the duration of the service, or employed technology such as Mobile IP to handle information forwarding as the user roamed. However these solutions were either undesirable to the user, or added an unnecessary level of complexity to the network. The introduction of the Stream Control Transmission Protocol (SCTP) [2] has opened the possibility of a mobile aware transport protocol. The multihoming feature of SCTP negates the need for a solution such as Mobile IP [3] and, as SCTP is a transport layer protocol, it adds no complexity to the network. Combining the handover procedure provided by SCTP with the increasing proliferation of WLAN hotspots can facilitate the delivery of services promised in the roll-out of 3G. As evidenced by the results obtained in [4] it would appear that not only is SCTP a feasible alternative to Mobile IP, but a practical solution for today's mobile internet.

In this paper we analyze the current handover scheme implemented in SCTP and detail some of its disadvantages. We propose a new handover scheme designed to perform handover more efficiently in a wireless environment. We outline the testing of this new handover scheme in the context of a WLAN environment and the results obtained from this testing. Section II provides a brief description of IEEE 802.11. Section III contains an overview of SCTP, including descriptions of the path monitoring and multihoming features of SCTP. Section IV describes handover in the context of SCTP. It also details the current handover scheme implemented in SCTP and provides a description of our proposed handover scheme. Our experimental setup and the results of our testing are given in Section V. Further work is outlined in Section VI. Section VII concludes the paper.

II. WLAN

IEEE 802.11 stipulates two modes of operation for WLANs: Distributed Coordination Function (DCF), and Point Coordination Function (PCF) [5]. The more commonly implemented of these is DCF. DCF uses a similar medium contention scheme as IEEE 802.3 (Ethernet). However, as it is difficult to detect collision in the wireless medium, the IEEE 802.11 scheme aims to avoid collisions. Its medium contention scheme is therefore entitled Carrier Sense Medium Access with Collision Avoidance (CSMA/CA). Using this medium access control scheme, a station first senses the medium to determine if it is idle. If the medium is free, then the station is permitted to transmit. If the medium is busy, then the station defers its transmission until the medium is free, where it then waits a random interval before transmitting. This is to avoid colliding with any other station that may also be waiting to transmit.

Collisions may still occur using this scheme as two stations may sense that the medium is idle and begin transmitting simultaneously. The use of the two control frames, Request To Send (RTS) and Clear To Send (CTS), minimizes the chance of collision occurring. These control frames reserve the medium in advance of a station's transmission. All stations in the network are made aware of how long the transmission will be and to which destination a station is transmitting. This information is stored at each station in the Network Allocation Vector (NAV).

Research done previously in the area of SCTP over WLAN [6] has proved SCTP offers little benefit over TCP with regards to congestion schemes in a WLAN environment. However this research was concerned only with ad-hoc networks and didn't take into consideration SCTP's multihoming feature.

III. SCTP OVERVIEW

Traditionally, signaling in telecommunications networks had been achieved using a logically separate network, with dedicated infrastructure in place. Packet switched networks, such as IP, do not provide dedicated infrastructure for signaling, and this has led the IETF to approach signaling in IP based packet switched networks differently. The result of this approach is SCTP [2].

SCTP is a datagram transport protocol designed to provide its services at the same layer as TCP and UDP, with IP being the primary underlying network. There are however, some differences between SCTP and these two other transport protocols. UDP provides a connectionless, best effort service that does not guarantee reliable delivery of datagrams. It cannot provide error control or flow control. TCP on the other hand, provides a connection-oriented, reliable service. Its reliability is described as strictly ordered, ensuring data sequence preservation. This can lead to head-of-line blocking should a packet become lost or re-ordered in the network. SCTP is connection-oriented, like TCP, and provides the same service reliability. SCTP contains features that TCP does not, such as multihoming, multistreaming, preservation of message boundaries, and unordered reliable message delivery [7]. These features are characteristic of signaling networks. Applied to SCTP, they can provide resilience to network failures (multihoming), and can avoid the head-of-line blocking (multistreaming) that TCP is susceptible to. Multihoming is of particular interest here as it is the core concept that allows for transport layer handover and is discussed in more detail later in this document.

In SCTP terminology a connection between two endpoints is known as an association, and is identified by a source port, a destination port and a Verification Tag. An SCTP datagram contains the common SCTP header and various control or data chunks. The control chunks are used to initiate association setup (INIT, INIT-ACK, COOKIE-ECHO, COOKIEACK), and association teardown (SHUTDOWN, SHUTDOWN-ACK, SHUTDOWN-COMPLETE). Control chunks are also used to monitor the status of all destination addresses of the association, periodically, for the duration of the association. The control chunks that perform this monitoring are the HEARTBEAT and HEARTBEAT-ACK chunks. Data chunks are used to transmit user data.

A. Multihoming

One of the features of SCTP that differentiates it from both TCP and UDP is its support of multihoming. If a host may be reached by more than one IP address, it is said to be multihomed. It may be that the host has bound one or more of its IP addresses to a physical interface card, or that it has a separate physical interface card for each of its IP addresses. This latter case is referred to as simple multihoming in [8] and for the duration of this document is what is meant when referring to hosts that are multihomed. In the case of multihomed hosts, all IP addresses available to each host are exchanged in the INIT and INIT-ACK chunks during association setup. There has also been some work done on adding IP addresses to an existing association dynamically [9]. Currently this is a work-in-progress but there are some provisions in the reference implementation that would facilitate the addition of this functionality. The area of load balancing, transmitting data to two or more of the multiple addresses of an SCTP association concurrently, is also a hotbed of SCTP research [10], [11]. This idea of load balancing is not a feature of classic SCTP.

Multihoming is used to add resilience to network failures, providing a certain degree of network stability to critical transmission paths. This may be extremely important for some applications, e.g. signaling transport of PSTN signaling messages. It can also provide a certain level of quality of service (QoS) to applications that rely on real-time communication, such as Voice over IP (VoIP) or video streaming applications. Multihoming can also be used to facilitate handover in SCTP. During normal operation, once an association has been established, one of the destination addresses is chosen as the primary destination and the route to this destination address is set as the primary path. All other destination addresses in the association are secondary addresses and are monitored periodically to obtain their reachable status.

B. Monitoring

All paths in an SCTP association are monitored periodically. This is done to track any changes in the reachable state of a destination address, and also to update the Round Trip Time (RTT) measurement for each of the routes to these secondary addresses. The monitoring is performed using HEARTBEAT chunks. These chunks are sent periodically to all destination addresses that have not been communicated with in the previous heartbeat interval. The timing of this heartbeating is determined from [12]:

$$H_i = RTO_i + HB.Interval \times (1 + \delta) \quad (1)$$

where RTO_i is the latest RTT Time-Out value for destination i, and δ is a random value between -0.5 and 0.5

RTO_i usually has a minimal effect and δ is used to introduce some variance into the timing of HEARTBEAT chunks. According to (1) above, a heartbeat chunk is sent to destination address i every H_i seconds. This equation is used for secondary addresses whose reachable status is active. An address is marked as active when it successfully responds after being polled. An address that does not respond to the poll is marked as inactive. For an inactive address, then the RTO for the destination address is no longer meaningful and some initial value, *RTO.Initial*, must be used instead. A typical implementation value for *RTO.Initial* is 1 second and for *HB.Interval* a value of 30 seconds is chosen.

Since all data is transmitted along the primary path, this means that at random intervals of time between 16 and 46 seconds, all secondary addresses are probed and their RTT (and subsequently, their RTO) values are updated. Clearly, since these probes are substantially smaller than the traffic that will be carried on the primary path, the RTT measurement on the secondary paths is merely a guide as to the expected RTT if data traffic were to be carried on that path.

IV. SCTP HANDOVER

Handover, in general, involves transferring data transmission from one communications link to another communications link, in such a way that will cause minimal disruption to the data transmission in progress. With regards to SCTP, this involves transferring data transmission from the primary path to one of the secondary paths.

A. Current Handover Scheme

In the current SCTP implementation, the handover process only occurs in the presence of path failure, i.e. handover will only occur once the primary path has failed and the primary destination address is marked as inactive.

This happens after four consecutive timeouts have occurred on transmissions to the primary destination address. The value of four is an implementation-set parameter, *Path.Max.Retrans*, and was chosen as a tradeoff between the accuracy of detecting that a destination address has failed and the speed of detecting that failure.

The process of marking a destination address as inactive begins with a timeout on an acknowledgement of data, or the response to a polling message, sent previously to that address. When the timeout occurs, an error counter is incremented and the value of the RTO is doubled for that path. In the case where the timeout was on an acknowledgement of data, the data is retransmitted along one of the secondary paths. Later, when the next segment of data to be transmitted, it is sent along the primary path. The next time that data, or a HEARTBEAT, is sent to the failed destination address, SCTP waits this doubled RTO period before it times out. If a timeout occurs, the error counter is incremented once again, the RTO is doubled again, and the data retransmitted along one of the secondary paths. The HEARTBEAT is not retransmitted.

SCTP continues in this fashion until either it receives a reply from the destination address, in which case all timers and error counters are reset for that address, or the error counter reaches *Path.Max.Retrans*. In this latter case the destination address is marked as inactive and a new primary destination address is chosen. This is called handover. From the example in [12], it can be shown that the time taken to mark a destination address as inactive, F_t, is given by:

$$F_t = RTO + 2 \times RTO + 4 \times RTO + 8 \times RTO \quad (2)$$

SCTP calculates a new RTO for a destination address each time the RTT for that address is updated. In the reference implementation of SCTP, this calculation of the RTO is lower bounded to 1 second. Using this value for the RTO in (2), it will take at least 15 seconds for SCTP to handover from the current (failed) primary address to one of the secondary addresses of the association. This is not acceptable in the case of real-time traffic. As noted in [13], when averaged over several trials, the detection of the failure and handing over to a secondary path took approx. 15 seconds. This serves as confirmation of (2).

B. Proposed Handover Scheme

This current scheme does not take into account any of the characteristics of the secondary links when performing the handover. The handover scheme being proposed periodically measures the RTT of each path and makes a handover decision based on the measurements obtained. The RTT was chosen as it is well known that it is one of the factors upon which the throughput of a TCP connection is dependant [14]. As SCTP was designed to use very similar congestion control schemes as TCP it can be inferred that SCTP's throughput is similar to TCP's. There is still research being carried out in this area [15].

However, as the scheme is being evaluated in a mobile environment (WLAN), some provisions had to be made with regards to delay spikes that are characteristic of such environments. The Smoothed Round Trip Time (SRTT) was chosen as the metric upon which the handover decision was to be made. The SRTT is calculated from the RTT but also uses a previous value as a baseline. This acts as a low pass filter to minimize the effect spikes have on the RTT. In SCTP the SRTT is calculated from (3):

$$SRTT = (1-\alpha)SRTT + \alpha RTT \qquad (3)$$

where

$$\alpha = \frac{1}{8}$$

Using this proposed scheme, SCTP will always handover to the path with the shortest delay between the two endpoints. This can provide real-time applications with a certain level of quality of service. As handover is performed based on a delay metric it does not incur the penalty that is associated with the current scheme, namely the four timeouts before the destination address is marked as inactive. From this it can be inferred that the proposed handover scheme will result in fewer retransmissions and a seamless handover between networks. These characteristics make a transport layer handover scheme particularly suitable for real-time applications.

V. RESULTS

The test-bed used to verify our proposed scheme is as in Fig. 1. Two multihomed SCTP hosts running the Linux operating system were used as the endpoints of the association. The implementation of SCTP installed on the hosts was a version of the reference implementation, available from [12], with the necessary modifications to implement our proposed delay-centric handover scheme. It is worth noting that the modifications were made to one SCTP host only. As the modifications adhere to the SCTP specification as defined in [2], no changes needed to be made to the second host. This provides a form of backwards compatibility with hosts that have not implemented the proposed scheme.

From Fig. 1, Host A is connected to Host B through two separate wireless networks. The initial primary path is the path from interface 1.2 on Host A to interface 3.2 on Host B. This is referred to as Path 0. The second path is from interface 2.2 on Host A to interface 4.2 on Host B. This is the secondary path, and is referred to as Path 1.

Once the association between Host A and Host B had been established, the amount of traffic in the first WLAN cell was increased. This was achieved by having several wireless stations transmit background UDP

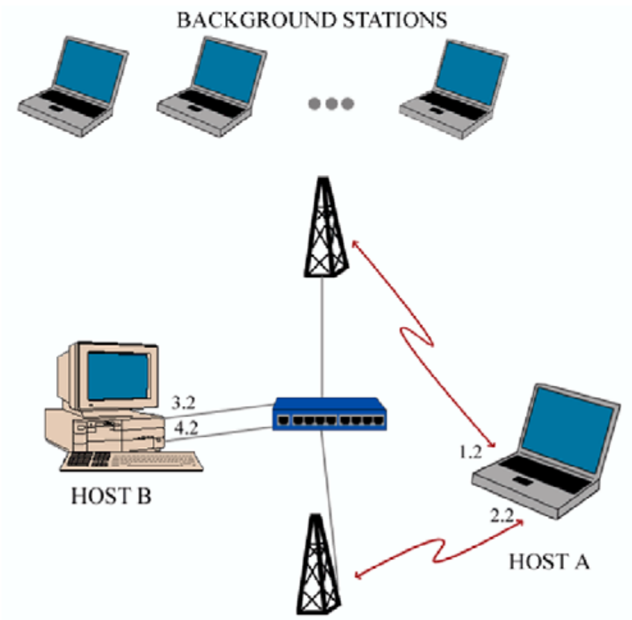

Fig. 1. Equipment Setup

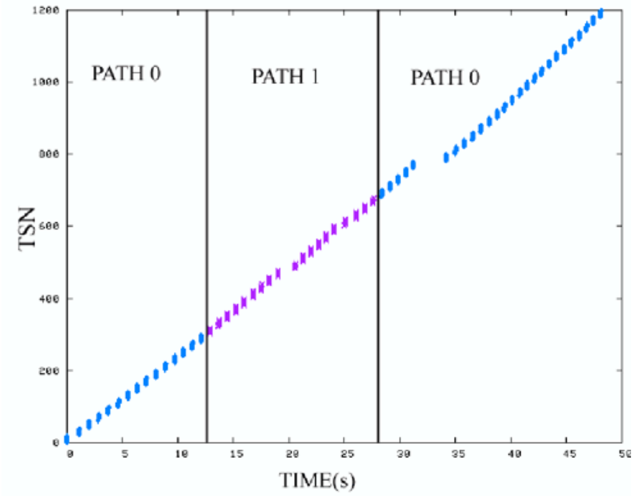

Fig. 2. Handover results

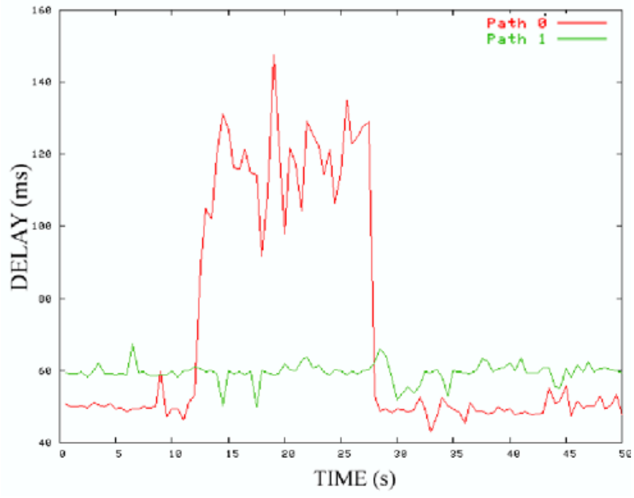

Fig. 3. Path delays for the association

traffic using the Iperf package [16]. This caused congestion in the cell and lead to an increase in the delay experienced by the SCTP association between Host A and Host B. It is at this time that the association performed handover from Path 0, the initial Primary Path, to Path 1, the secondary path. Subsequently the amount of background traffic in the first WLAN cell was decreased, reducing the congestion. This had the effect of reducing the delay that the SCTP association experienced, causing it to perform hand back to Path 0 from Path 1. Fig. 2 shows a trace of TSNs (SCTP data packets) that were transmitted during this test sequence. It can be seen, when correlated with the path delays shown in Fig. 3, that handover to Path 1 occurred when the delay on Path 0 increased, due to congestion in the cell, and handed back to Path 0, when the congestion had passed.

This behaviour emulates the expected network delay as a user moves away from one AP and experiences lower access speeds and increased channel errors. If another AP is available, then the delay experienced to this AP should reduce as the user moves closer to it. For optimal performance, the application needs to be able to automatically select the new AP as its primary point of connection.

These results show that in spite of the limited accuracy of the simple probing strategy deployed here, the modified version of SCTP was able to successfully complete a vertical handover from a congested network to a less congested network and then back again after the congestion had passed.

VI. FURTHER WORK

All the current work has involved performing handover from one WLAN cell to another less-congested WLAN cell. It is our aim to next evaluate our delay-centric handover scheme in a heterogeneous wireless environment. To achieve this it is proposed that an association is setup between two endpoints, with one path through a WLAN network and a second path through a GPRS network. It is expected that when in the presence of the WLAN cell handover will occur from the GPRS path to the WLAN path. As the mobile station moves out of coverage of the WLAN Access Point, the delay associated with the WLAN path, as perceived from the SCTP host's point of view, should increase to the point where the delay on the GPRS path is more favourable. At this point the association will hand back to the GPRS link. Coupling this handover ability with the dynamic IP address addition proposed in [9] would allow for a fully mobile aware transport layer protocol.

VII. CONCLUSION

In this paper we have given a brief discussion of IEEE 802.11b and introduced SCTP. We have detailed the current handover scheme implemented in SCTP and described some of its failings. We have proposed a new handover scheme that would be more suitable to a mobile environment such as WLAN. It has been shown through testing that this proposed handover scheme does offer a feasible alternative to the current failover-centric scheme. It has been surmised that by performing handover on the basis of delay as opposed to interface failure, fewer retransmissions will occur. In an environment such as a WLAN this will offer a better utilization of the wireless medium, increasing the efficiency of the WLAN.

More significantly, this new scheme will provide a seamless handover mechanism to real-time applications such as VoIP or video streaming. We believe that this work represents the first published results of a vertical handover implemented in the transport layer that can provide such functionality.

ACKNOWLEDGMENTS

The support of both the Research Innovation Fund and the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

REFERENCES

[1] 802.11 Working Group, http://grouper.ieee.org/groups/802/11/

[2] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, V. Paxson, "Stream Control Transmission Protocol", RFC2960, October 2000

[3] C. Perkins (Editor), "IP Mobility Support for IPv4", RFC3220, 2002

[4] W. Xing, H. Karl, A. Wolisz, and H. Müller, "M-SCTP: Design and Prototypical Implementation of an End-to-End Mobility Concept", Proc. of 5th Intl. Workshop The Internet Challenge: Technology and Applications, Berlin, Germany, October 2002

[5] LAN/MAN Standards Committee, IEEE Computer Society, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification", IEEE, New York, NY, USA, IEEE Std 802.11, 1997

[6] G. Ye, T. Saadawi, M. Lee, "SCTP Congestion Control Performance In Wireless Multi-Hop Networks", MILCOM, California, USA, October 2002

[7] L. Coene (Siemens), "Stream Control Transmission Protocol Applicability Statement", RFC3257, 2002

[8] R. Braden (Editor), "Requirements for Internet hosts - communication layers", RFC1122, 1989

[9] R. Stewart, M. Ramalho, Q. Xie, M. Tuexen, I. Rytina, M. Belinchon, P. Conrad, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", http://www.ietf.org/internet-drafts/draft-ietf-tsvwgaddip-sctp-08.txt, Expired, work-in-progress

[10] A. Caro, J. Iyengar, P. Amer, G. Heinz, "Using SCTP Multihoming for Fault Tolerance and Load Balancing", SIGCOMM 2002 Student Poster - ACM Computer Communication Review, vol. 32, no. 3, July 2002

[11] J. Iyengar, K. Shah, P. Amer, R. Stewart, "Concurrent Multipath Transfer Using SCTP Multihoming", Technical Report 2004-02, CISC Dept, Univ of Delaware, 2003

[12] R. Stewart and Q. Xie, "Stream Control Transmission Protocol (SCTP): A Reference Guide", Addison Wesley, New York, NY, USA, 2001

[13] A. Kelly, P. Perry, and J. Murphy, "A Modified SCTP Handover Scheme for Real Time Traffic", HETNETs Working Conference, Ilkley, England, July 2003

[14] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modelling TCP Throughput: A Simple Model and its Empirical Validation", IEEE/ACM Trans. on Networking, vol.8, no. 2. April 2000

[15] R. Rajamani, S. Kumar, N. Gupta, "SCTP versus TCP: Comparing the Performance of Transport Protocols for Web Traffic", University of Wisconsin-Madison, May 2002

[16] Iperf, Traffic Generator, available from http://dast.nlanr.net/Projects/Iperf/