A Framework for Network Media Optimization in Multihomed QoS Networks *

Hiroyuki Koga[1]†, Hiroaki Haraguchi[2], Katsuyoshi Iida[3], and Yuji Oie[4]
[1]Kitakyushu JGN2 Research Center,
National Institute of Information and Communications Technology, Japan
[2]Yaskawa Information Systems Corporation, Japan
[3]Global Scientific Information and Computing Center, Tokyo Institute of Technology, Japan
[4]Department of Computer Science and Electronics, Kyushu Institute of Technology, Japan

ABSTRACT

Future pervasive networks will be constructed by various access networks, which are managed by different operators. Wireless LAN hotspot services will be available in numerous locations, and cellular services, which provide a wide coverage area, will also be available. In order to realize efficient mobile computing in such an environment, mobile hosts must have multiple network media as well as functions to select the optimum network among them based upon the current situation without the disconnecting communications. In order to provide mobile hosts with such a function, we propose an association layer, which is inserted between the transport and network layers. The association layer selects the appropriate network medium for each flow according to the network conditions and flow characteristics, providing end-to-end QoS in multihomed networks. Furthermore, we show the effectiveness of the proposed scheme through simulation results and demonstrate proof of concept based on empirical evaluations of a prototype implementation.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Management, Design, Performance

Keywords

Network media optimization, Multihoming, QoS, Mobility, Pervasive networks

1. INTRODUCTION

The expansion of mobile computing has brought about various wireless access networks. These networks, which overlap each other, have different characteristics in terms of bandwidth, delay, stability, cost, and coverage area. For example, wireless LAN provides high bandwidth and inexpensive connectivity, but its coverage area is limited, whereas cellular networks provide a large coverage area but are expensive and have low bandwidth. In order to realize efficient mobile computing in such pervasive networks, mobile hosts must have multiple network media and must select the optimum medium used for communication according to the network conditions, without disconnection of communications.

In this context, although some protocols such as Mobile IP [1, 2], Stream Control Transmission Protocol (SCTP) [3], and Host Identity Protocol (HIP) [4, 5] have been proposed, these protocols do not provide efficiency in pervasive multihomed networks. Mobile IP provides address transparency to the transport layer, so the existing applications do not need to be modified. However, Mobile IP has difficulty in supporting multihoming by itself [6, 7, 8]. On the other hand, SCTP is a novel transport protocol that natively supports multihoming. Since the multihoming feature may help to provide mobility functions, several mobility extensions of SCTP have been proposed [9, 10]. Although the mobile SCTP has the advantage of mobility functions, it does not perform network media optimization. Since SCTP is a transport protocol, it has difficulty in obtaining the information of the other transport flows. Also, HIP supports the multihoming function based on a new namespace, called a host identity, which is created by cryptography algorithms to provide trustful host identifiers other than IP addresses. Since HIP does not provide a flow assignment mechanism to different network interfaces, QoS support is problematic in HIP. Consequently, these protocols lack QoS mechanisms for choosing the appropriate network medium according to flow characteristics in terms of, for example, the available bandwidth and required delay bounds.

Therefore, we propose a novel layer, called an association layer, to achieve network media optimization in pervasive multihomed networks. The association layer is inserted between the transport and network layers. This means that major modifications to the transport layer, including socket APIs and the network layer are not necessary, so that current applications can be used without any modifications. One of the key functions of the association layer is to select appropriate network media for each flow according to the network conditions and flow characteristics. In order to achieve this function, the association layer keeps track of the network media, i.e., IP addresses, on each corresponding host. This means that the association layer provides a mobility function without any special network instances in the networks.

†koga@kyushu.jgn2.jp

*This work was supported in part by Grant-in-Aid from International Communication Foundation and Scientific Research (A)(2)(15200005) of the Ministry of Education, Culture, Sports, Science and Technology, Japan.

DIN'05, September 2, 2005, Cologne, Germany.

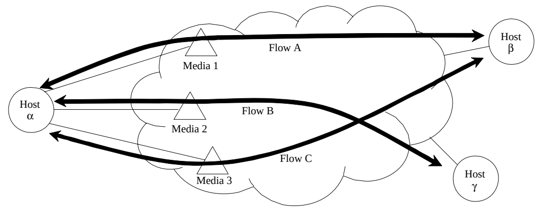

Figure 1: Network media optimization by selection of appropriate network medium for each flow

In the present paper, we demonstrate the effectiveness and validity of the proposed scheme through simulation results and empirical evaluations of the prototype implementation.

2. MEDIA OPTIMIZATION NETWORK ARCHITECTURE

As described in the previous section, our primary requirements include mobility and multihoming support as well as network media optimization. In order to achieve the optimization, we propose the Media Optimization Network Architecture (MONA), which consists of the association layer and its protocol.

2.1 Overview

Before describing the proposed framework, we explain the network media optimization. Figure 1 illustrates an example of network media optimization. In this figure, there are three hosts, α, β, and γ, connected to the Internet. Host α, a multihomed host, is connected to three different media: media 1, 2, and 3. In this example, host α communicates with host β by two different flows (Flow A and B), and another flow (Flow C) communicates with host γ. Host α selects an appropriate medium for each flow based on the media condition; i.e., media 1 for flow A, media 3 for flow B, and media 2 for flow C. Of course, the appropriate media should be selected dynamically based on changes in the conditions. Since host α is assumed to be mobile host, most of the media have relatively low bandwidth, long delay, unstable links, and are expensive. Therefore, we assume the first hop from the host is the most significant link in terms of quality of service. On the other hand, the cloud illustrated in the figure is assumed to represent relatively high bandwidth, short delay, inexpensive and stable links.

Based on the above observation, we focus on the first hop and its media selection. Since hosts can easily obtain media status information from physical and datalink layers implemented in the host, it is natural to implement the selection. However, the selection requires knowledge of the instantaneous status of every flow and/or of every media. Therefore, in order to collect this status information in one place, we insert a new layer, called the *association layer*, between the transport and network layers. This information collection further allows us to allocate media for each flow. Moreover, for the down link to host α, the host should select the appropriate medium for each flow. This reverse direction of media selection requires signaling with the corresponding host.

Some of the existing protocols support the mobility function, and others support multihoming functions. However, all existing protocols have difficulty in achieving network media optimization. Lack of optimization will cause performance degradations such as high packet loss rates and long delays, as well as wasted resources. Therefore, in order to achieve optimization, we herein propose an association layer.

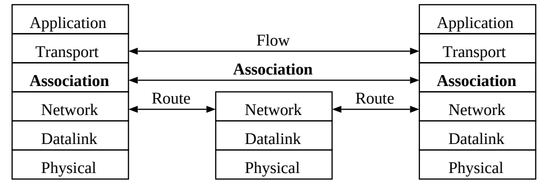

Figure 2: Association layer in network layer model

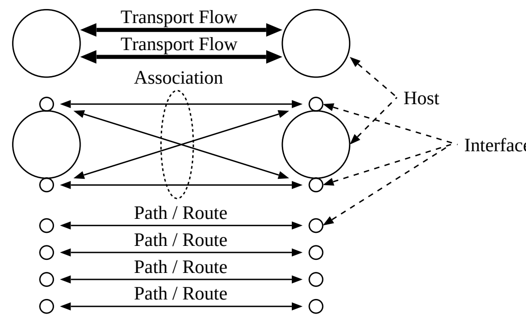

Figure 3: Relationship between transport flow, association, and route

2.2 Association layer

The core layer of MONA is the association layer introduced herein. The association layer is inserted between the transport and network layers, as illustrated in Fig. 2, and provides host-to-host abstraction service. In order to realize media optimization in down links, the IP addresses of all media in use must send notifications to peers. However, multiple IP address notifications consume a large amount of header overhead. Therefore, the address information is shared among flows between the same peers. This relationship, which is referred to as "association", is illustrated in Fig. 3. The association manages the set of paths generated by the combinations of the network media, i.e., the IP addresses, of each host. In this illustration, there are two transport flows between the same peers and two media for each peer. The information about flows and media are collected at the association layer in order to achieve network media optimization. Moreover, in order to perform the network media optimization dynamically for the case in which a mobile host moves in/out of wireless access networks, the association layer coordinates the separated identities: the host identity in the transport layer for distinguishing flows and the interface identity in the network layer for transmitting data packets. In this way, the association layer supports the mobility and multihoming functions as well as the network media optimization without any network assistance.

Figure 4 shows a state diagram per peer host in the association layer. This is similar to the three-way handshake of TCP. If there has been no communication activity on a peer, then the state for the peer is "Idle". Once the upper layers, say the application and transport layers, are going to connect the peer, the association layer sends an association request "Init" to the peer, then the status becomes "Init sent". If the peer accepts the request, it will send back an acknowledgment. After receiving the acknowledgment, the status will be "Associated". If the peer does not accept the request, then the status will be "No association", which indicates that either the peer has no ability to handle the association layer or the peer rejects the association due to its policy. In this case, the peers communicate using

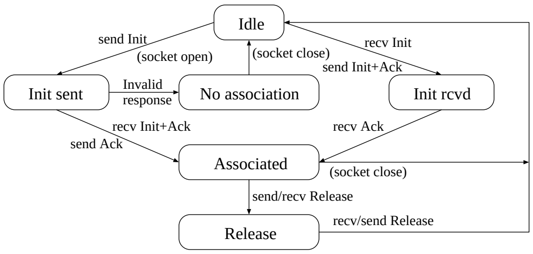

Figure 4: State diagram of the association layer

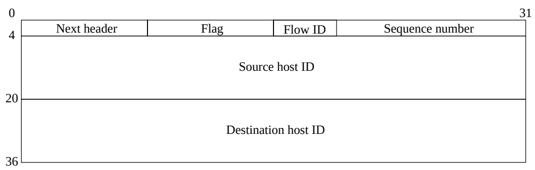

Figure 5: Association management protocol header

normal IP without association extension. After the communication, each peer can release the association through explicit signaling at the termination of the final flow. Note that the association layer controls one state, i.e., Idle, Associated and so on, sharing flows among the same peers to enable media optimization.

2.3 Association management protocol

The Association Management Protocol (AMP) is an association layer protocol, which establishes and manages the association. The basic function is to exchange the Association Identification (AID), source interface addresses, and interface preference for each flow. The AID is embedded as an AMP header, which is inserted between the IP header and its payload in all datagrams using the association, as shown in Fig. 5. An 8-bit next-header field indicates the type of the next header including an association option header, which will be explained later, or an upper-layer header such as TCP and UDP. A flag field of 8 bits consists of bits that represent the type of AMP header such as Init, Acknowledgment, and Release. A flow ID field of 4 bits is used to distinguish transport flow groups. The Source and Destination host IDs of 128 bits are AID in order to distinguish associations, which are determined at the association establishment phase.

The source interface addresses and interface preference are notified by the association option header to the peer during the establishment phase and when status changes occur. The address option header and flow preference option header are shown in Figs. 6 and 7, respectively. The 8-bit next header field is the type of next header including other following association option header if exists or an upper-layer header. The type and length fields of 4 bits represent the type and length of this association header. IP addresses of the address option header consist of the IP addresses currently used for communication with a priority order. For example, the peer of the receiver side uses the notified IP address (1), the highest priority, as the destination address for flows that not belong to the flow ID indicated by the flow preference option. The flow preference option header is used to notify the peer that a flow group, which is indicated by the flow ID, would like to use a different address from the default IP address (1). The flow preference option header includes each flow ID and its selected addresses with a priority order. The address fields store the identification number in the IP address list informed by the address option.

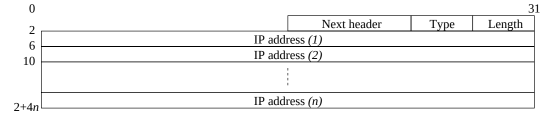

Figure 6: Address option header

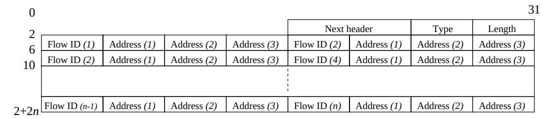

Figure 7: Flow preference option header

These association option headers and association management packets, including Init, Ack, and Release are sent as an association packet that is independent of transport flows. Namely, the peers immediately send these packets to each other as the need arises and retransmit them if needed. This enables us to prevent an increase in header overhead as well as communication with traditional hosts using normal IP. The host without an association layer simply drops the association Init packet, but can receive transport packets that have no embedded association header.

3. PERFORMANCE EVALUATION

In order to demonstrate the effectiveness of the proposed scheme, we perform computer simulations. We used VINT Network Simulator NS Version 2 [11] after adding the association layer and the association management protocol for the network media optimization.

3.1 Simulation model

In our simulation, a mobile host (MH) communicates with the correspond host (CH) through a cellular network and wireless LAN. The cellular network has a bandwidth of 384 Kb/s and a propagation delay of 1 ms. The wireless LAN has a bandwidth of 11 Mb/s and a propagation delay of 1 ms. In order to simplify the simulation model, we ignore the effect of interference and transmission errors on wireless links. The link between a base station (BS) and a gateway (GW) has a bandwidth of 100 Mb/s and a propagation delay of 5 ms. The link between an access point (AP) and the GW has the same characteristics. In addition, the link between the GW and the CH has a bandwidth of 100 Mb/s and a delay of d ms, which ranges from 5 ms to 100 ms. Each node is equipped with a buffer of 50 packets.

We assume that the MH accepts two flows from the CH. One is the UDP flow with a constant bit rate of 64 Kb/s, such as IP telephony. The other is the TCP flow that is used for greedy file transfer, which continues to infinitely transmit data from the sender. The UDP packet size is set to 200 bytes, the TCP packet size is set to 1000 bytes, and the TCP variant employed herein is SACK.

The MH moves into wireless LAN 30 s after the start of simulation and then moves out 30 s later. Namely, the MH can use a multihomed environment for 30 s.

3.2 Simulation results

In this section, we will show the characteristics of TCP throughput performance, UDP loss probability, and UDP delay performance during the multihomed period (30 - 60 s) by means of simulation results. First, in Fig. 9, we show the average throughput performance of TCP when the link delay

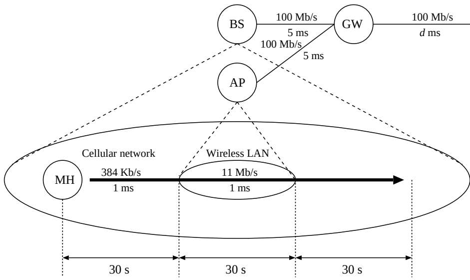

Figure 8: Simulation model

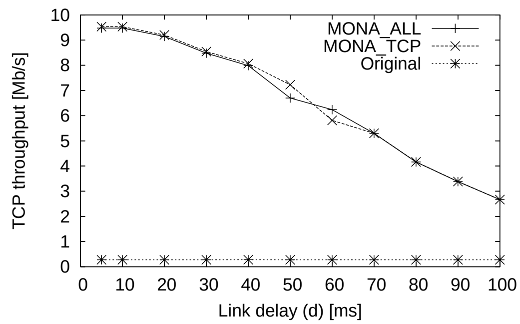

Figure 9: Average TCP throughput: multihomed period

d varies from 5 ms to 100 ms. In a multihomed environment, there are two options for the use of multihomed links. One is the case in which all flows use wireless LAN, which is indicated as "MONA_ALL" in the figure. The other is the case in which only TCP flows use wireless LAN, which is indicated as "MONA_TCP" because TCP flow requires high bandwidth, whereas UDP flow requires an uncongested link rather than a high bandwidth link in this scenario. "Original" indicates no association extension, so that all flows use cellular networks at all time. From this figure, we see that MONA schemes effectively use a high-bandwidth wireless LAN.

Next, we show the loss probability of UDP in Fig. 10. This figure shows that the MONA_ALL scheme causes packet loss of UDP, although packet loss does not occur in the other schemes because the TCP flow increases its transmission rate in order to use a wireless LAN of higher available bandwidth than the cellular networks. This leads to packet losses of the UDP flow. Furthermore, we show the UDP performance for transmission delay time in Figs. 11 (a) and (b), which indicate the average and coefficient of variation, respectively, of delay time. MONA schemes achieve the reduction of UDP transmission delay time. However, the MONA_ALL scheme worsens jitter performance although the MONA_TCP scheme attains excellent jitter performance for the same reason described above. From the above results, MONA schemes, in particular MONA_TCP, improve UDP and TCP performances due to network media optimization in multihomed networks.

Finally, we show the results for the entire simulation time, which indicate the effect of switching the currently used network medium. Figure 12 shows the performance of the UDP transmission delay time. MONA schemes reduce the average delay time, while increasing the jitter of the UDP flow. Therefore, a trade-off exists between the jitter performance and the other performance metrics. The impact on the jitter performance is not negligible in the case the network

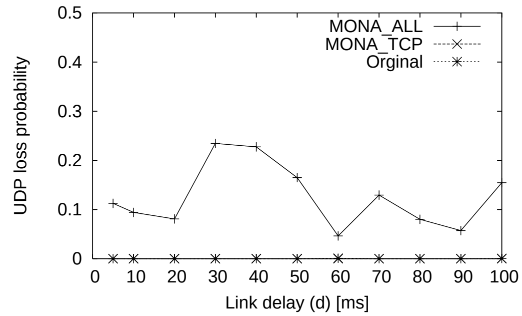

Figure 10: UDP loss probability: multihomed period

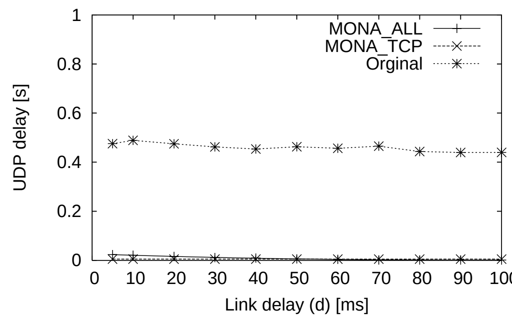

(a) average

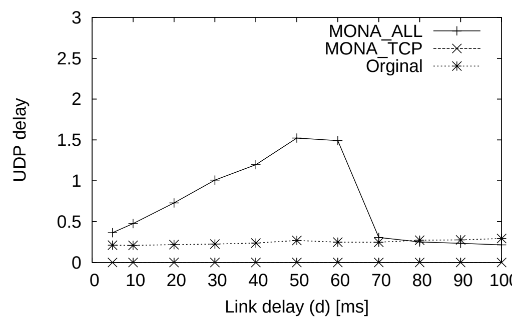

(b) coefficient of variation

Figure 11: UDP transmission delay time: multihomed period

medium is changed frequently although it is small as compared with the effect of MONA on the UDP delay performance and the TCP throughput performance in this scenario. Accordingly, in future studies, we should determine how to switch the network media while taking this trade-off into account.

4. PROTOTYPE IMPLEMENTATION

In order to demonstrate the validity of the proposed scheme, we perform an empirical evaluation of the prototype implementation. We implement basic functions of the association layer and protocols in Linux-2.4 kernel, which provide the mechanism by which to maintain transport flow without disconnection for the case in which the IP address of network interfaces and the interfaces currently used for communication change.

Using this implementation, we performed the following experiments. Figure 13 shows the model used for performance evaluation. Each host has two network media connected to

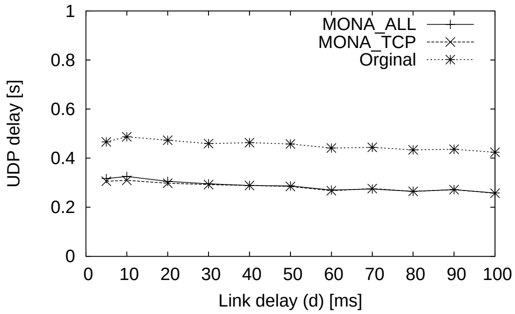

(a) average

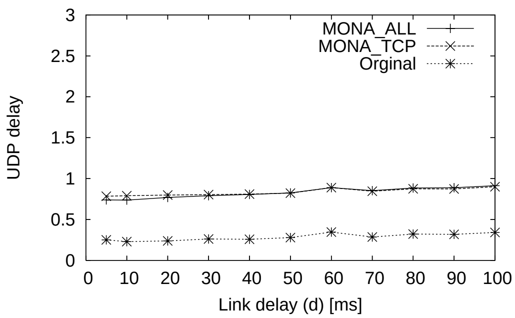

(b) coefficient of variation

Figure 12: UDP transmission delay time: entire simulation time

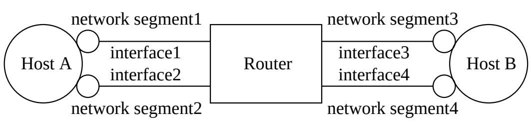

Figure 13: Experimentation model

different network segments and communicates with the corresponding host through the router. All of the links have a bandwidth of 10 Mb/s and a propagation delay of 10 ms. Figure 14 shows the sequence number of transmitting packets on each interface of host A when the host moves between network segments 1 and 2 at intervals of 1 s. This figure shows that the proposed scheme detects changes in the network media conditions and selects a suitable medium by which to continue communication without disconnection.

5. CONCLUSION

In the present paper, we proposed a novel association layer to achieve an efficient mobile computing in pervasive multihomed networks. The association layer supports mobility and multihoming functions as well as network media optimization. One of the key functions of the association layer is to select the appropriate network medium for each flow according to the network conditions and flow characteristics, providing end-to-end QoS in multihomed networks. Furthermore, we demonstrated the effectiveness and validity of the proposed scheme based on the simulation results and experimental results of a prototype implementation. The simulation results have shown that MONA improves the loss probability and delay performance of UDP and the TCP throughput performance by selecting the appropriate network medium for each flow in multihomed networks. Moreover, from the empirical evaluation, we have seen that MONA detects changes in the network media conditions and

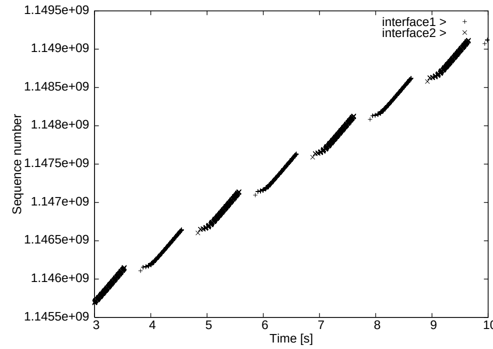

Figure 14: Experimental result of switching network medium

selects a suitable medium by which to continue communication without disconnection.

In the future, we intend to investigate a mechanism by which to reduce the jitter of the UDP flow when the currently used network medium is switched.

6. REFERENCES

[1] C. Perkins. IP mobility support for IPv4. *IETF RFC3344*, August 2002.

[2] D. Johnson, C. Perkins, and J. Arkko. Mobility support in IPv6. *IETF RFC3775*, June 2004.

[3] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson. Stream control transmission protocol. *IETF RFC2960*, October 2000.

[4] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson. Host identity protocol. *IETF Internet-Draft*, draft-ietf-hip-base-02.txt, February 2005.

[5] P. Nikander, J. Arkko, and T. Henderson. End-host mobility and multi-homing with the host identity protocol. *IETF Internet-Draft*, draft-ietf-hip-mm-01.txt, February 2005.

[6] N. Montavont, T. Noel, and M. Kassi. Mobile IPv6 for multiple interfaces. *IETF Internet-Draft*, draft-montavont-mobileip-mmi-01.txt, October 2003.

[7] N. Montavont, R. Wakikawa, T. Ernst, T. Noel, and C. Ng. Analysis of multihoming in mobile IPv6. *IETF Internet-Draft*, draft-montavont-mobileip-multihoming-pb-statement-03.txt, January 2005.

[8] R. Wakikawa, K. Uehara, T. Ernst, and K. Nagami. Multiple care-of-address registration. *IETF Internet-Draft*, draft-wakikawa-mobileip-multiplecoa-02.txt, September 2003.

[9] S. Koh, M. J. Lee, M. Riegel, M. L. Ma, and M. Tuexen. Mobile SCTP for transport layer mobility. *IETF Internet-Draft*, draft-sjkoh-sctp-mobility-03.txt, February 2004.

[10] S. Kashihara, K. Iida, H. Koga, Y. Kadobayashi, and S. Yamaguchi. Multi-path transmission algorithm for end-to-end seamless handover across heterogeneous wireless access networks. *IEICE Transactions on Communications*, E87-B(3):490–496, March 2004.

[11] VINT Project, Network Simulator ns-2, http://www.isi.edu/nsnam/ns/.