Hybrid User- and Network-Initiated Flow Handoff Support for Multihomed Mobile Hosts

Q. Wang, R. Atkinson, C. Cromar, and J. Dunlop

Mobile Communications Group, Department of Electronic and Electrical Engineering,
University of Strathclyde, Glasgow G1 1XW, UK

***Abstract*—With the increasing deployment of overlay networks, a mobile host with a range of network interfaces can be connected to multiple access networks simultaneously. Such multihoming technology can be exploited to distribute (or hand off) traffic flows among the interfaces and access networks to achieve seamless, robust and even quality-of-service-aware communications for mobile hosts. At present, there is little preexisting work that has sufficiently addressed the problem of supporting both user- and network-initiated flow handoffs in a unified architecture. In this paper, we propose a hybrid approach to manage both kinds of handoffs in a flexible yet standardized way by enhancing Mobile IPv6 (MIPv6) or its variant Hierarchical MIPv6. Address management strategy in mobile host multihoming context is investigated, and the corresponding architectural and protocol choices are presented and analyzed. Particularly, the proposed comprehensive handoff management is described with signaling and operations highlighted.**

Keywords—Flow handoff, Hierarchical Mobile IPv6, Mobile IPv6, Multihoming

I. INTRODUCTION

MULTIHOMING is an attractive value-added feature for a mobile host (MH) equipped with multiple network interfaces in converged networking environments where homogeneous and/or heterogeneous networks such as Wi-Fi, WiMAX, and cellular systems are overlaid. In this context, it is desired and practical that both a mobile user and the network can trigger a handoff to switch all or *selected* application flows from one interface or access network to another from its own perspective. We refer to such an operation as a *flow handoff*. Clearly, effective flow handoff support is a key enabler to achieve the "Always Best Connected" vision [1] when coupled with intelligent network selection algorithms.

From the user's perspective, a flow handoff may be triggered by the user's preference (e.g., selection of an access technology with lower tariff when available), or by the requirements of certain applications or services with the help of "local" measurements by the MH (e.g., redirection of delay-sensitive application flows to the link with lower round-trip time). From the network's perspective, a flow handoff may be initiated to improve the network's service e.g., by load sharing, fault tolerance etc. For example, if the network detects that one access network is overloaded it has the option of distributing a subset of the flows through another access network. Or perhaps, one access network is underutilized and could share some of the loads. Detection of these kinds of "regional" events can be fulfilled by a network-side entity (other than an MH) more conveniently and accurately. Certainly, there are plenty of other examples that necessitate such flow handoffs to gain better quality of service (QoS) in a broader sense for a mobile user, the network (including network operator, service provider etc.), or even both.

Therefore, these observations justify a hybrid flow handoff approach, where both a user (MH) and the network can trigger a flow handoff, and both user- and network-triggered handoffs can be supported in a unified architecture. Furthermore, the architecture should fulfill the flow handoff management in a standardized and flexible way to facilitate implementation and deployment. So far, prior work in this field has not accomplished this objective.

The remaining of this paper is organized as follows. Section II reviews related work on multihoming support for mobile users. In Section III, we expound our proposed architecture towards a comprehensive, flexible and standardized solution. We begin with the reference network model for the following design. Then we present address management options, which can be applied to the uniform platform and directly influence the architectural and base protocol choices. Subsequently, we describe the handoff signaling and operations with an emphasis on the more complex network-initiated handoffs though both kinds of handoffs are covered. Finally, Section IV concludes this paper.

II. RELATED WORK

Mobile IPv6 (MIPv6) [2] is the de facto standard for IPv6 mobility support. Unfortunately, MIPv6 in its current form does not support advanced multihoming beyond handing off all the flows from one interface to another. Recently, the IETF MONAMI6 WG is standardizing MIPv6-based mechanisms to facilitate handoffs of selected flows for multihomed MHs. Ref. [3] allows multiple Care-of Addresses (CoAs) to be bound with a single Home Address (HoA) using Binding Update (BU) and Binding Acknowledgement (BA) messages. Ref. [4] enables a particular flow to be bound with a CoA associated with an interface. A Flow ID option accommodates the flow

identifier such as a subset of the five-tuple (source and destination addresses and port numbers, transport protocol), e.g., the well-known port numbers can identify different application flows (e.g., 80/8080 for HTTP flows). The Flow ID option, also placed in a BU/BA message, can indicate adding, replacing or deleting of a flow binding (flow ID, CoA). A default (HoA, CoA) binding exists in case of no matching. User-initiated flow handoffs have been focused on though these existing separate proposals can be exploited as building blocks towards a comprehensive unified architecture.

The Stream Control Transmission Protocol (SCTP) [5] is an emerging transport protocol that supports a kind of multihoming in its own right. However, instead of supporting parallel flows flexibly distributed among interfaces, the multihoming feature in SCTP is designed to enable retransmissions to alternate IP address(es) for survivability when the primary IP address becomes unavailable. More importantly, as a transport-layer protocol SCTP multihoming would only benefit applications that are based on SCTP rather than TCP or UDP, over which the dominating majority of the current IP applications are running. Therefore, a network-layer solution e.g., enhanced MIPv6 would be more appropriate.

The Host Identity Protocol (HIP) [6] is another promising proposal that could facilitate IP mobility and multihoming. HIP introduces a new "host" layer between the network and the transport layers. Consequently, flows are bound to host identities instead of IP addresses so that the change of IP addresses can be transparent to the applications as in the basic MIPv6. Regarding multihoming, similar to SCTP alternate IP addresses can be exchanged between the MH and its peer so that SCTP-like multihoming can be achieved. In addition to the similar limitation as found in SCTP-style multihoming, adding a new layer to the protocol stack is not a minor modification and this may cause an updating of numerous IP applications. Hence, HIP-based multihoming may not be preferred in the short term.

In addition, there exist some other related ad hoc proposals in the literature, although little work has been accomplished for a comprehensive, flexible, standardized and unified flow handoff management framework. It is also noted that work is underway (e.g., in the IETF SHIM WG) on site multihoming, where a site's network has connections to multiple IP service providers. For the sake of this research we concentrate on end host multihoming and thus site multihoming is disregarded.

III. Proposed Architecture

Based on the above survey, in this section we propose a MIPv6-based architecture to handle both user- and network-initiated flow handoffs in the host multihoming context.

A. Reference Network Model

A reference network model is illustrated in Fig. 1 to facilitate the following description and discussions. This model assumes a generic networking scenario:

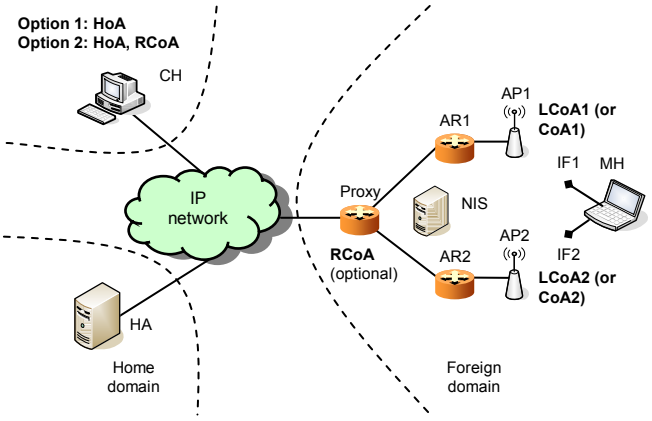

Fig. 1. Reference network model

An MH can be equipped with multiple interfaces though the two interface case (IF1 and IF2) is taken as a sufficient example. The MH is visiting a foreign domain. The Home Agent (HA) of the MH is located in the home domain. A correspondent host (CH) is stationary in a third domain. All the domains are interconnected to each other through a common IP core network.

In the foreign domain, two access points AP1 and AP2 provide wireless access to the MH. They are directly connected to or collocated with two access routers AR1 and AR2, respectively. Certainly, there can be other routers in the two access networks, e.g. a cloud can be introduced between the proxy and AR1 or AR2. For presentation clarity, additional routers are not shown in the figure. In addition, a network intelligence service (NIS) is present to monitor, collect and process QoS measurements, and determine if the network should initiate a flow handoff as the result of intelligent network selection algorithms.

Subject to the availability the local proxy is a special crossover router of the two access networks under consideration in the foreign domain. The proxy is a local mobility agent handling flow handoffs (due to intelligent network selection) and conventional movement-triggered handoffs locally. We propose that an evolved Hierarchical MIPv6 (HMIPv6) [7] Mobility Anchor Point (MAP) takes this role. It is noted that the HA can act as the proxy when the MH is in the home domain. If no such proxy functionality is available in a certain foreign domain, the architecture falls back from HMIPv6 to MIPv6 and the proxy turns out to be a standard router (e.g., the gateway). Therefore, we have two scenarios in the following discussions depending on the availability of a local proxy.

B. Address Management and Architectural Choices

Proper address management is essential to achieve effective mobility support. In a multihoming environment, a multihomed MH can have much more addresses than a conventional MH and thus a more careful address management strategy is demanded. Furthermore, the address management schemes directly influence the architectural and protocol choices.

Fig. 1 also illustrates the proposed configuration of addresses. Hierarchical address management (at least for HoAs) is proposed to simplify the complexity of multihoming. In a foreign domain, the MH acquires a Regional CoA (RCoA) if a local proxy is available and two on-link CoAs (LCoAs) for the two interfaces from the proxy's subnet and the subnets of AR1 and AR2, respectively. The MH registers the local binding (RCoA, LCoA1, LCoA2) in the proxy, and the global binding (HoA, RCoA) in the HA. When the visiting foreign domain has no local proxy, no RCoA is configured and the MH obtains two global CoAs, which must be registered with the HA. Any existing RCoA must be deregistered at the CH. If the MH is also multihomed at the home domain, the global HoA is obtained from the HA's subnet and the on-link HoAs are from the subnets of the ARs, respectively. In either scenario, the multiple address registration mechanism defined in [3] is utilized.

Accordingly, there are options for managing the MH's address(es) at the CH. In Option 1, only the MH's HoA is available at the CH, and thus the CH can be unaware of the MH's movements. The advantage of using this option is that the CH can be a non-MIPv6 host whereas the disadvantage is that triangular routing via the HA is incurred. In Option 2, the MH registers the binding (HoA, RCoA) at the CH, and route optimization can be achieved. Since built-in route optimization is a highlighted feature in MIPv6, Option 2 is emphasized.

The hierarchical HoA management results in a single HoA of the MH as a stable global identifier. Consequently, the upper layers (layers above the network layer including applications) may safely assume an unchanged IP address as if the MH were a fixed host in Option 1. In the case of Option 2, the upper layers may choose to use the RCoA as the stable IP address of the MH as indicated in the HMIPv6 specification [7]. Furthermore, the hierarchical address management also enables a graceful transparency of the MH's multihoming at the CH – the CH is only aware of a single HoA in Option 1 or a single HoA with a single CoA (RCoA) in Option 2. In both cases, the CH is released from the burden to maintain a list of MH's addresses and to select destination addresses to initiate or resume a session as found in SCTP- or HIP-based multihoming support. From the CH's perspective, Option 1 uses one HoA to identify the MH instead of one HoA per interface (thus multiple HoAs to the CH). Using one HoA per interface introduces additional complexity for TCP flow handoffs. Option 2 further employs hierarchical CoA management in interested foreign domains so that a multihomed MH appears as a non-multihomed MH globally with route optimization supported. Moreover, when the MH is in foreign domains the local proxy takes care of the flow handoffs instead of the HA so that inter-domain signaling to the remote HA and single point failure for flow handoffs found in Option 1 are avoided. Note that inter-domain signaling would generate more overhead and delay, plus tighter requirements on security. In addition, with HMIPv6 is introduced, micro-mobility (movement within a domain) of the MH can be supported as the HMIPv6 was designed for. On the other hand, Option 2 demands more deployment complexity since a local proxy is introduced, and HMIPv6 is necessitated. In Option 1, the HA alone is in charge of both home and away cases and HMIPv6 in the home domain may not be necessarily deployed: the MH may register its on-link HoAs as CoAs bound to the single HoA obtained from the HA's subnet.

In contrast, if no such hierarchical address management is employed, the use of multiple HoAs (and/or CoAs) may generate ambiguity to applications running over the MH and/or the CH for selecting source/destination addresses and thus extra run-time management overhead. One possible solution to alleviating these problems is to set one of the HoAs (and/or CoAs when away) as the "well-known" HoA (and/or CoA) to a CH.

TABLE I summarizes the architectural and protocols configurations in the two proposed options, and lists their pros and cons.

TABLE I
ARCHITECTURAL AND PROTOCOL OPTIONS

			Option 1	Option 2
Architecture	Hierarchical address management		At home domain only	At home and foreign domains
	Flow distributor	Downlink	HA	HA at home domain, and MAP at foreign domain
		Uplink	MH	MH
	MH's IP address(es) know to CH		HoA	(HoA, RCoA)
	Communications between MH and CH		Bi-directional tunneling (via HA)	Route Optimization (via local proxy)
Base Protocol	MIP protocol at home domain		HMIPv6 with RO disabled; or MIPv6	HMIPv6
	MIP protocol at foreign domain		MIPv6 with RO disabled	HMIPv6 with RO enabled
Advantages			An unchanged HoA is referred to globally, regardless of multiple HoAs and CoAs; session continuity ensured; no local proxy when away needed	An unchanged HoA and a single CoA simplify multihoming to base MIPv6 largely; route optimization supported; more prompt response to flow handoffs; micro-mobility supported
Disadvantages			Triangular routing and global signaling overhead caused	A local proxy needed

C. Handoff Signaling and Operations

In the proposed hybrid architecture, a flow handoff due to intelligent network selection can be triggered by either an MH or the NIS on behalf of the user and the network, respectively. A set of specific intelligent network selection algorithms for this purpose is to be designed and is beyond the scope of this paper. We assume that the triggers are available as a result of these algorithms, which can be designed as another independent building block in a modularized way. In addition, we take the scenario where a proxy is available in a foreign domain as an example in the illustrations. As a comprehensive framework, both user- and network-triggered handoffs are handled as follows.

1) Support for User-Initiated Handoffs: As an example, we assume that two flows (flow 1 and flow 2) have been established between the CH and the MH. When a flow handoff is user-triggered as shown in Fig. 2, the MH sends a BU message with the new flow binding policy enclosed to the proxy (or the HA if the proxy is unavailable, written as proxy/HA hereafter for brevity), as depicted in Step 3. Note that the MH may prefer to use the targeted (secondary) interface for the signaling especially when the source interface is going down. We assume that the new flow policy indicates that one of the flows (flow 2) needs to be handed over from the current interface to the secondary one. The initial flow policy should have been installed at the proxy/HA during the initial registration stage at the foreign (or the home) domain. Regarding the format of a flow policy record for an MH and related action indication, the mentioned Flow ID option defined in [4] is exploited.

Upon receiving the BU, the proxy/HA authenticates and verifies the BU. If the authentication and verification are successful, the proxy/HA updates the pre-installed flow binding policy of the MH and replies with a BA indicating the acknowledgement. Afterwards, the proxy/HA switches flow 2 to its interface connected to the access network corresponding to the MH's secondary interface. The switching is achieved by tunneling (IP-in-IP encapsulation). Consequently, flows are distributed between the interfaces as desired.

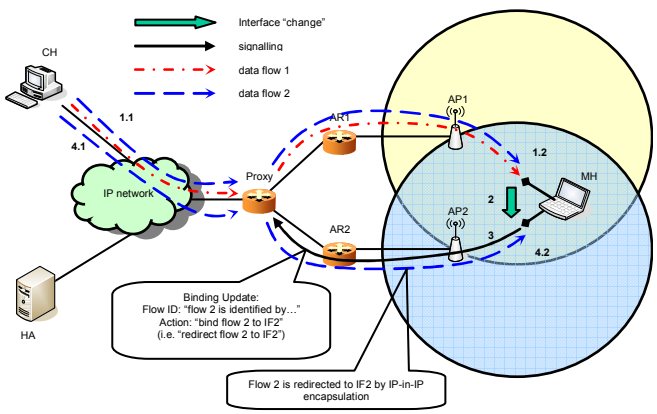

Fig. 2. User-initiated handoff support: overview

2) Support for Network-Initiated Handoffs: When a flow handoff is initiated by the network as illustrated in Fig. 3, a three-way negotiation process between the network (the proxy/HA triggered by the NIS) and the user (the MH) is proposed as follows.

Firstly, the proxy/HA formulates the new flow binding policy and signals the change(s) to the MH (Step 2.2) upon receiving the network trigger from the NIS (Step 2.1). To fully exploit standardized work, we propose to extend the MIPv6 Binding Refresh Request (BRR) message by introducing a new mobility option to accommodate the information on the flow binding change(s). This extension is based on the Flow ID option defined in [4]. A new flag is also defined in the "Reserved" field of the option to indicate if the new policy is negotiable. Note that such a new extension is supported by the extensibility of the BRR message as defined in the MIPv6 specification [2]. By extending the BRR message rather than defining a new message, the BU and BA messages enhanced for user-initiated handoffs [3][4] can be largely reused. Moreover, the built-in security specified in the original MIPv6 messages can be naturally utilized and thus no additional efforts on security are provoked here. The MH's current interface in use (IF1) should be targeted for the BRR since the proxy/HA may not be able to guarantee that the other interface (IF2) is ready unless IF2 is already in use for simultaneous parallel flows via the proxy/HA.

Secondly, on receiving the BRR, the MH may accept or reject the new flow binding policy by sending back a BU with an explicit reply, e.g. by repeating the new policy if accepted or repeating the existing policy if rejected, or simply by setting a flag. IF2 may be preferred for the signaling to show (by the MH) and to double check (by the proxy/HA) its availability.

Thirdly, the proxy (or the HA) acknowledges with a BA. In this BA, the proxy/HA may indicate the final decision. This indication can be optional if the change is mandatory, the network has the final say, and it is known that the targeted interface is ready for the handoff. A flag can be set in the BRR or the BA message to indicate if this change is negotiable. If it is not negotiable, the BRR serves as a decision notification

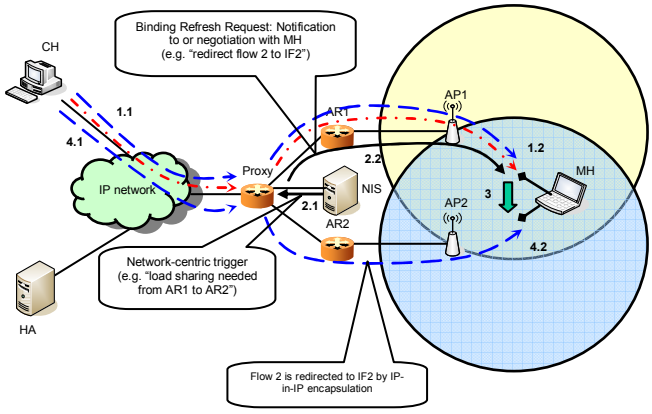

Fig. 3. Network-initiated handoff support: overview

only and the MH can only accept the change(s). Finally, if desired selected flows are handed off to the targeted interface.

To give a full picture of the signaling and operations for the proposed network-initiated flow handoff management, Fig. 4 demonstrates a sample QoS-triggered flow handoff scenario. Step 1 shows two established flows from the CH being tunneled by the proxy (or the HA if the proxy is unavailable) to Interface 1 (IF1) of MH via AR1. In Step 2, the NIS collects (periodically or informed by other entities including MHs) and inputs selected QoS measurements to predefined intelligent network selection algorithms. Assuming a flow handoff is to be triggered as shown in Step 3, the NIS sends the new flow handoff policy to the proxy/HA, which would authenticate and verify the request and acknowledges it (Step 4). In Step 5, the network-triggered flow handoff based on HMIPv6 (or MIPv6) is initiated as mentioned. Consequently, as an example, one of the two flows is redirected to another interface by the proxy/HA tunneling as shown in Step 6. The protocols used in Steps 2 and 4 can be non-MIP based and are currently being specified. One candidate approach is to extend ICMPv6 messages to realize the signaling.

Note that the illustrations only demonstrate the downlink flows, and selected downlink flows are handed over to another access network by the proxy/HA on a flow handoff. Regarding an uplink flow handoff, the MH itself switches selected flows to another interface by tunneling in the similar manner.

IV. CONCLUSION

In this paper, we have proposed a hybrid framework to support flow handoffs for multihomed mobile hosts. The proposal is comprehensive as the framework can handle both user- and network-triggered handoffs in a unified platform.

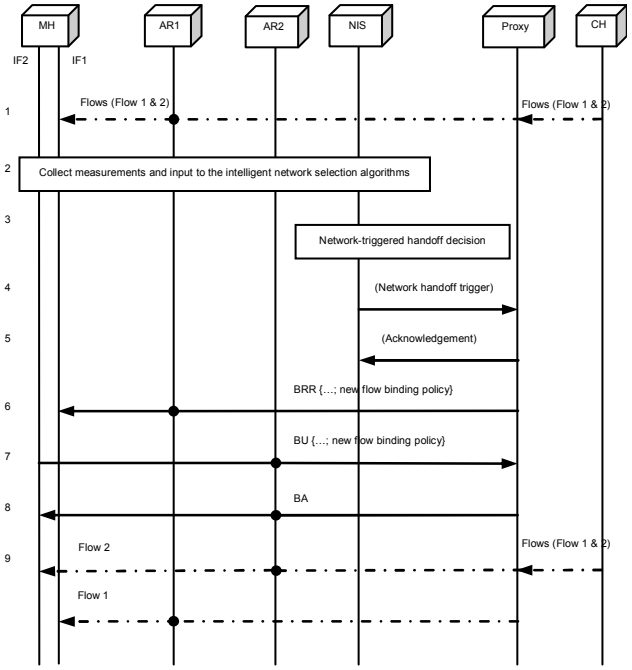

Fig. 4. Network-initiated handoff support

Great flexibility is provisioned in the proposed architecture. Firstly, both the interests of the user and the network are taken into account and well balanced through mechanisms especially the three-way negotiation process. Secondly, incremental deployment scenarios can be handled depending on the availability of a local proxy in a foreign domain. When the visited a foreign domain has no local proxy, signaling and operations are based on MIPv6. Otherwise, hierarchical address management based on HMIPv6 and downlink flow handoff support can be implemented in the proxy, which is an evolved HMIPv6 MAP. Consequently, route optimization and localized handoff signaling can be achieved so that the communications and signaling are more efficient, and handoff responding delays are reduced compared with the MIPv6-based case.

Furthermore, the proposed architecture aims at a standardized solution. The proposed handoff protocol is mainly based on MIPv6 and HMIPv6, and thus specification efforts could be minimized. Operations on multiple address registration and flow binding in MIPv6/HMIPv6 can be based on the IETF proposals [3][4] being standardized. In addition, we propose extensions to MIPv6 BRR message to signal the network-initiated flow handoff trigger from a proxy/HA to an MH. This approach gracefully combines the standard work into a uniform architecture whilst fully exploiting the merits of standardized messages (e.g., the built-in security).

To sum up, the proposed unified architecture can support both user- and network-initiated flow handoffs in a flexible and standardized way. Future work will define the intelligent network selection algorithms, and further evaluate the advantages of this design especially in the aspect of the balanced benefits for both the mobile users and the network, and compare different design choices in terms of signaling and data delivery costs, flow handoff response delay, and etc.

ACKNOWLEDGMENT

This work is sponsored by the EU IST MULTINET project (www.ist-multinet.org). The authors would like to thank for our project partners involved in the discussions on this topic.

REFERENCES

[1] E. Gustafsson and A. Jonsson, "Always Best Connected," *IEEE Wireless Communications*, vol. 10, no. 1, pp. 49-55, Feb 2003.

[2] D. B. Johnson, C. Perkins, and J. Arkko, "Mobility support in IPv6", *IETF RFC 3775*, Jun 2004.

[3] R. Wakikawa, T. Ernst, and K. Nagami, "Multiple care-of addresses registration," *IETF Internet Draft*, <draft-wakikawa-mobileip-multiplecoa-05.txt>, work in progress, Feb 2006.

[4] H. Soliman, N. Montavont, N. Fikouras, and K. Kuladinithi, "Flow bindings in mobile IPv6," *IETF Internet Draft*, <draft-soliman-monami6-flow-binding-00.txt>, work in progress, Feb 2006.

[5] R. Stewart, Q. Xie, K. Morneault, and et al., "Stream control transmission protocol," *IETF RFC 2960*, Oct 2000.

[6] R. Moskowitz and P. Nikander, "Host identity protocol architecture", *IETF RFC 4423*, Aug 2005.

[7] H. Soliman, C. Catelluccia,K. E. Malki, and Ludovic Bellier, "Hierarchical mobile IPv6 mobility management (HMIPv6)," *IETF RFC 4140*, Aug 2005.