Distributed Mobility Management for Efficient Video Delivery over All-IP Mobile Networks: Competing Approaches

Dong-Hoon Shin, Arizona State University
Danny Moses and Muthaiah Venkatachalam, Intel Corporation
Saurabh Bagchi, Purdue University

Abstract

The recent proliferation of multimedia mobile devices and a variety of mobile applications are generating an enormous amount of data traffic over mobile networks. The key driver of the mobile traffic growth is mobile video. Currently, mobile networks are evolving to the 4G system, which has a flatter architecture and provides all-IP-based mobile broadband service. In all-IP mobile networks, IP mobility management is a key function that allows mobile nodes to continue their communications even when their point of attachment to the IP network changes. Existing mobile networks employ a centralized mobility management scheme where all intelligence is concentrated in one end-point system, rather than being distributed through the internet. However, this cannot satisfactorily support mobile videos, which demand a large volume of data and often require QoS such as session continuity and low delay. This motivates distributed mobility management (DMM) solutions that can efficiently handle mobile video traffic. In this article, we survey different approaches for DMM in standards development organizations such as IETF and 3GPP, and also in research organizations. We focus on three different DMM approaches that are currently being considered by the IETF: PMIPv6-based, MIPv6-based, and routing-based DMMs. We provide a qualitative analysis to compare the three DMM approaches and discuss which DMM approaches are more suitable for efficient mobile video delivery.

Over the past few years, multimedia mobile devices such as smart phones and tablet PCs have rapidly proliferated. Along with this, mobile data traffic has increased dramatically due to a huge number of mobile users enjoying a variety of applications, such as web surfing, instant messaging, mobile gaming, social networking services (e.g., Facebook and Twitter), and video streaming services (e.g., YouTube and Netflix). The key driver of the explosive mobile traffic growth is mobile video, and this trend is expected to intensify in the near future; it has been reported that mobile video traffic was 52 percent of total mobile data traffic by the end of 2011, and it is forecasted that mobile video traffic will account for over 70 percent of total mobile data traffic by 2016 [1].

Despite the high demand for mobile data, telecom operators have been observing that their average revenue per user (ARPU) in mobile data is rapidly decreasing. To revert this, telecom operators are eagerly seeking to improve their network performance and efficiency, as well as to reduce the costs expended on network operation and maintenance. A major focus of such efforts is on solutions to efficiently handle a large volume of mobile video traffic.

Mobile networks are currently evolving from the third generation (3G) to the fourth generation (4G) to meet the growing demand for the mobile broadband traffic with satisfactory user experience. The Third Generation Partnership Project[1] (3GPP) Evolved Packet System (EPS), commonly referred to as the 4G Long Term Evolution (LTE), has a flatter architectural design and provides all-IP mobile broadband service, where user data are packetized into IP packets to be delivered over underlying IP networks. In all-IP mobile networks, as mobile nodes may frequently change their point of attachment to the IP network (i.e., the access router to which they are connected), IP mobility management is a key function to allow mobile nodes to continue their communications despite changing their access router.

The 3GPP EPS has adopted the Proxy Mobile IPv6 (PMIPv6) [2] and Dual Stack Mobile IPv6 (DSMIPv6) [3], which were standardized by the Internet Engineering Task Force[2] (IETF), for network-based and host-based mobility management, respectively. They both employ a centralized mobility anchor (i.e., the packet gateway) in the core network.

This material is mostly based on the work that was done when Dong-Hoon Shin was an intern in Intel Corporation as a Ph.D. student.

[1] *http://www.3gpp.org*

[2] *http://www.ietf.org*

0890-8044/13/$25.00 © 2013 IEEE

That is, in EPS, all mobility-related processing is performed in the core network with a hierarchy of the serving gateway and the packet gateway, which requires all IP traffic to be routed via the centralized gateways. We refer to this as the *centralized hierarchical architecture* to support IP mobility. The centralized hierarchical architecture is a legacy from the General Packet Radio Service (GPRS) on the (2G) Global System for Mobile Communications (GSM), but with the improvement of fewer hierarchical levels. This centralized architecture was motivated largely by two fundamental requirements:

- To keep the IP address fixed for the life of the packet data session
- To reuse the legacy Authentication and Key Agreement (AKA) protocol used for circuit-switched voice communication

These requirements are demanded by mobile operators to maintain control over services in the home network and also to maintain their existing system of distributing user credentials in secure subscriber identity modules (SIMs) [4].

However, the centralized mobility management is not efficient for handling a large volume of mobile data traffic. The limitations include poor scalability, inefficient use of network resources, and packet delay (the reasons for which are elaborated in the next section). These limitations stand out when the centralized mobility management needs to support mobile videos, which demand a large volume of data and often require quality of service (QoS) such as session continuity and low delay. This motivates distributed mobility management (DMM) solutions to efficiently handle the ever increasing mobile traffic, the major portion of which carries video traffic.

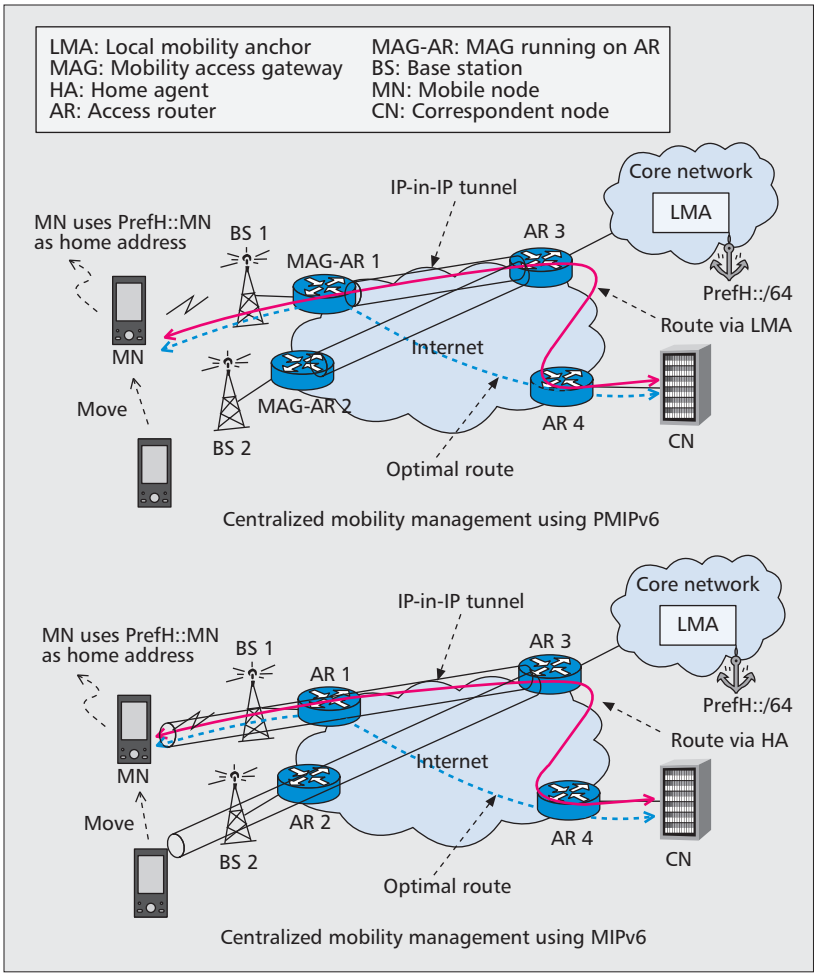

Figure 1. *Overview of centralized mobility management employed in current mobile networks.*

In this article, first we present a brief overview of the centralized mobility management employed in the current mobile networks and discuss its limitations. We then present three main approaches to DMM that are currently being considered by the IETF DMM Working Group[3] — Proxy Mobile IPv6 (PMIPv6)-based, Mobile IPv6 (MIPv6)-based, and routing-based DMM. Then we give a brief overview of the existing approaches in 3GPP standards and also some DMM solutions proposed in the research domain. We provide a qualitative analysis to compare the three DMM approaches in IETF, and discuss which DMM approaches are more suitable for efficient mobile video delivery. Finally, we present concluding remarks.

Why Distributed Mobility Management?

Current mobile networks, specifically, the 3GPP's 3G Universal Mobile Telecommunications System (UMTS) and 4G EPS, employ centralized mobility management, notably PMIPv6 and DSMIPv6, to handle network-based and host-based mobility schemes, respectively, as depicted in Fig. 1. PMIPv6 is a network-based mobility management scheme, based on the well-known host-based MIPv6 [5], where mobility-related signaling is performed by a network entity called a mobility access gateway (MAG) on behalf of the mobile node (MN). DSMIPv6, on the other hand, is an extension of MIPv6 to support the dual IP stack (i.e., both IPv4 and IPv6). These approaches employ a centralized mobility anchor, which is a local mobility anchor (LMA) and a home agent (HA) for (DS)MIPv6, in the core network to support IP mobility. Specifically, when an MN connects to the network, it is allocated a home IP address (or prefix) anchored at the centralized mobility anchor, which keeps track of the changes of the MN's point of attachment, that is, the serving access router (AR), by maintaining a mobility context of the MN. With knowledge of the MN's location, an IP tunnel is established between the centralized mobility anchor and the MN's serving AR (for PMIPv6), or the MN directly (for (DS)MIPv6) so that the MN's traffic can be redirected. The IP tunnel can be established through an IP-in-IP encapsulation, which allows an IP packet destined for an IP address to be redirected to another IP address.

The centralized approach might have been reasonable at the time the existing mobile networks were designed. However, it has obvious limitations in handling a large volume of mobile data traffic, due to the involvement of the centralized mobility anchor in handling mobility signaling and routing for all registered MNs. This leads to the following drawbacks (we highlight a subset of the drawbacks mentioned in [6]):

- Lack of dynamic mobility support
- Suboptimal routing
- Low scalability
- A single point of failure and vulnerability to attacks

First, the centralized approach lacks dynamic mobility sup-

[3] *http://datatracker.ietf.org/wg/dmm*

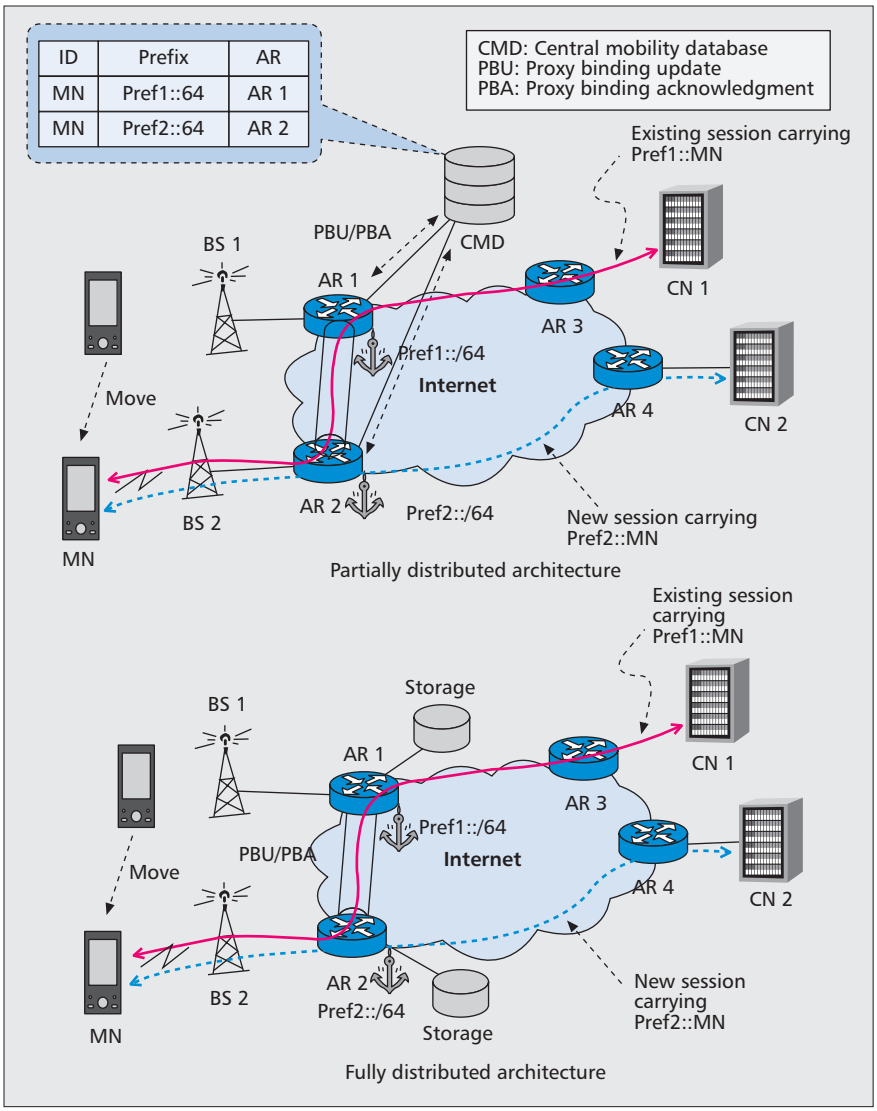

Figure 2. *Overview of PMIPv6-based DMM.*

port. That is, it has to always support the IP mobility even when it is not needed. For example, even if MNs are stationary and stay attached to their initial serving AR, the centralized mobility anchor has to still maintain their mobility contexts and also process their packets for tunneling. Due to this unnecessary IP mobility support, the centralized approach not only wastes the network resources (i.e., CPU, memory, routers, etc.), but also incurs an additional delay in packet delivery. Second, the centralized approach yields suboptimal routes via the centralized mobility anchor, which may significantly increase the length of routes, especially when an MN and a correspondent node (CN) are close to each other but both of the communicating nodes are far from the centralized mobility anchor. This also leads to inefficient use of network resources (i.e., backhaul routers) and packet delay. Third, the centralized approach is not scalable with a large number of MNs. Since the centralized approach requires all traffic to be routed via the centralized mobility anchor, the centralized mobility anchor has to maintain and manage the mobility contexts of a large number of MNs, and also to handle a large volume of data traffic. As a result, the centralized mobility anchor and the backhaul routers around it will be overloaded, leading to a bottleneck. Lastly, the centralized approach makes the system vulnerable to a single point of failure and further makes it an easy target for attacks by an adversary, due to the system design relying on a single entity for the entire system to function.

Due to these limitations, the centralized approach is not efficient to support mobile videos, which demand a large volume of data traffic and often require QoS such as session continuity and low delay. This motivates current mobile systems to further evolve toward a flatter architecture employing DMM, thereby providing efficient means to handle mobile video traffic, such as distributed and dynamic anchoring, optimized routing, and signaling with low latency.

IETF DMM Approaches

In this section, we present three main approaches that are currently being considered by the IETF DMM Working Group: PMIPv6-based, MIPv6-based, and routing-based DMM.

PMIPv6-Based DMM

PMIPv6-based DMM decouples the role of the centralized LMA into three basic functions:

- Allocation of home prefixes/addresses
- Tunneling of data traffic
- Management of bindings of an MN's home address (HoA) and MAG's proxy care-of address (CoA)

PMIPv6-based DMM redeploys these three functions in the mobile network, as shown in Fig. 2. The first two functions are deployed on the edge of the mobile network (i.e., collocated with ARs). In other words, ARs also function as a mobility anchor to MNs. The data plane thus consists purely of ARs. On the other hand, there are two alternative approaches to deploy the third function as the control plane [7]. It can be deployed in either a centralized or distributed manner, which leads to the overall PMIPv6-based DMM architecture being either partially or fully distributed.

In the data plane, each AR has a unique set of global prefixes for which it functions as a mobility anchor, so mobility anchors are distributed over ARs. The set of IP prefixes for each AR is assigned statically when deployed, or is managed dynamically through a standard protocol, such as the Dynamic Host Configuration Protocol (DHCP). An MN can have multiple concurrent sessions that are anchored at different ARs. Figure 2 illustrates an example of this. When AR1 detects the attachment of an MN, it allocates the MN an IP prefix, Pref1, from its prefix set. The MN then configures an IP address, Pref1::MN, with the given prefix Pref1, and starts a communication session using Pref1::MN. When the MN moves and is attached to a new AR, AR2, it is given another IP prefix, Pref2, by AR2 and configures another IP address, Pref2::MN. The MN can start a new session using the new IP address Pref2::MN, while continuing the ongoing sessions anchored at AR1 (which still carries the original IP address, Pref1::MN). To continue the ongoing session, an IP tunnel is established between AR1 and AR2.

For the control plane, two alternative approaches to store, update, and retrieve the bindings of an MN's HoA and a MAG's Proxy CoA are presented in [7]. The first approach, shown in the top half of Fig. 2, employs a centralized mobility

database to maintain a global view of MNs' bindings, that is, a complete knowledge of the ARs to which each MN is currently attached and the ARs at which each MN's active sessions are anchored. The centralized mobility database registers and updates MNs' bindings by exchanging control messages with ARs. Through the exchange of control messages, ARs obtain the bindings of the MNs' active sessions anchored at other ARs for tunnel establishment. The second approach, shown in the bottom half of Fig. 2, has a mobility database distributed among ARs. That is, each AR is only given a local view of the bindings of only the MNs anchored at itself. ARs exchange control messages with each other to register, update, and retrieve bindings in their cache. Thus, in the second approach, both the data plane and the control plane are handled entirely by ARs.

In addition to the basic operations above, PMIPv6-based DMM needs a network-based mechanism to effectively support simultaneous and dynamic anchoring. That is, it has to allow an MN to simultaneously send/receive traffic anchored at different ARs and determine the right prefix (which is anchored at the serving AR), without requiring any special support in the protocol stack on MNs. For this, one solution [8] currently being considered by the IETF DMM Working Group introduces a logical interface at the IP stack of ARs, called the distributed logical interface (DLIF). The basic idea of DLIF is that for each MN, the serving AR creates a DLIF associated with the mobility anchor of each active prefix, so it can expose itself to the MN as multiple (logical) routers. Then the MN can see multiple logical routers with a different IP prefix. However, the MN does not yet know which IP prefix is the one that is anchored at its serving AR. To resolve this issue, when the serving AR sends router advertisements over the DLIFs, it includes a zero prefix lifetime for the prefixes that are not anchored at the serving AR, so they will be deprecated for the new session. In this way, the MN will start a new session using the right prefix and also be able to continue the ongoing sessions without changing their IP addresses.

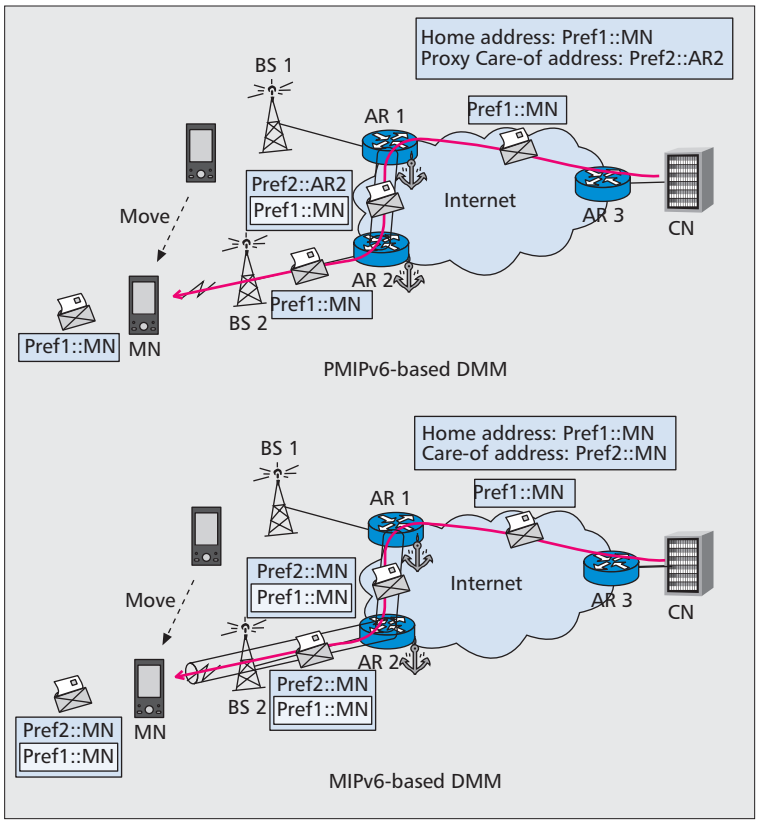

Figure 3. *Change of destination address in packet header in PMIPv6-based and MIPv6-based DMM.*

MIPv6-Based DMM

Due to the similarity in the base protocols (i.e., PMIPv6 and MIPv6), MIPv6-based DMM [9] can employ the fully and partially distributed architectures used for the PMIPv6-based DMM, but with a few differences in the protocol operations. In the data plane, MNs' traffic is encapsulated or decapsulated by MNs on behalf of their serving ARs. Figure 3 depicts this difference by showing how the destination address changes in the PMIPv6- and MIPv6-based DMM, respectively, as a packet is delivered from the CN to the MN. In the control plane, instead of ARs, MNs initiate their mobility-related signaling with the anchoring ARs (in the fully distributed architecture) or with the central mobility database (in the partially distributed architecture) in order to register and update the bindings associated with mobility anchors. Thus, MIPv6-based DMM is a host-based mobility management scheme, which requires modification on the IP stack of hosts (i.e., MNs), while PMIPv6-based DMM is a network-based mobility management scheme, which requires no such modification and thus also works for legacy IP hosts.

Unlike PMIPv6-based DMM, MIPv6-based DMM supports route optimization (RO) due to its base protocol. This allows an MN to directly communicate with the CN by informing the CN of its CoA. Specifically, when an MN moves and is allocated a new CoA from a new AR, the MN sends a message, called binding update (BU), to the CN to inform the CN of its new CoA. Upon receiving the BU, the CN sends an acknowledgment, called binding acknowledgment (BA), to the MN to confirm receipt of the BU. Since the CN now knows the MN's new CoA, it can send packets directly to the MN using the MN's new CoA, thereby leading to the packets being routed through an optimal path. On the other hand, RO brings up an issue of location privacy [6]. Also, RO significantly increases the signaling overhead as it causes an additional return routability (RR) procedure for security. Specifically, RO requires the additional RR procedure to give the CN assurance that the MN is indeed reachable at the MN's claimed CoA as well as at the MN's HoA. Furthermore, the signaling for the RR and BU/BA procedures should be performed with all of the MN's CNs.

In addition, there are other proposals for dynamic mobility management employing the two (fully and partially) distributed architectures based on PMIPv6 and MIPv6. Bibliographic references to these proposals along with brief descriptions are provided in [10].

Routing-Based DMM

Routing-based DMM [4] makes use of a routing protocol to support mobility, instead of a tunnel setup protocol as in PMIPv6- and MIPv6-based DMM. Thus, an optimal route can be used between two MNs. In this DMM approach, the network is structured in a hierarchy of three layers (core, aggregation, and access), as shown in Fig. 4. Note that each router can be connected to more than one router in the upper layer, and can even be connected directly to its peer routers in the same layer. The connectivity between the routers is thus more like a mesh structure and less hierarchical than that in current mobile networks (shown in Fig. 1). The ARs, each of which has its own set of IP prefixes (or addresses), are collocated with base stations, and act as the first-hop routers for MNs. The routers in the aggregation layer are configured as route reflectors for the ARs connected to them. That is, they aggre-

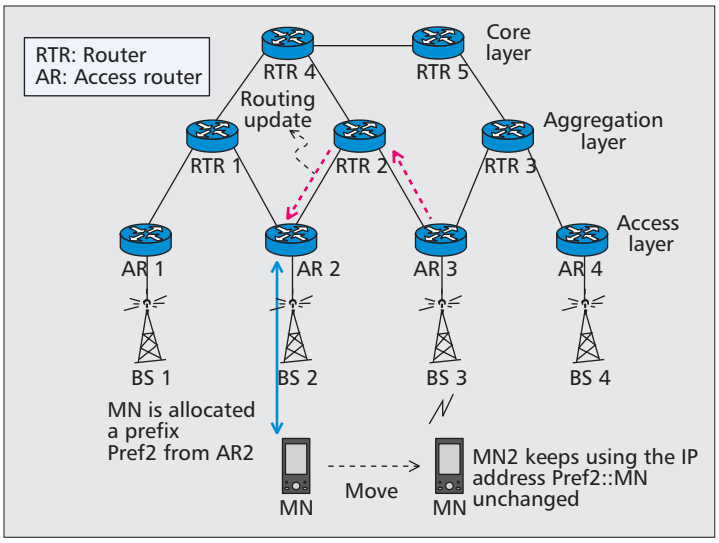

Figure 4. *Overview of routing-based DMM.*

gate the assigned prefixes advertised by ARs for the core routers in the upper layer and also reflect all subprefixes advertised by any AR to all the other ARs in the cluster of which they are in charge.

Upon initial attachment, an MN is allocated an IP address (or prefix) from the AR to which it is attached through a standard protocol (e.g., DHCP). Then the MN updates its Domain Name System (DNS) record to point its hostname to the IP address assigned, while the serving AR updates the reverse pointer in the in-addr.arpa (for IPv4) or ip6.arpa (for IPv6) space to point to the MN's hostname. That is, the MN and the serving AR control the forward and reverse mappings separately.

When a handover occurs, the new AR to which the MN moves first looks up an IP address using the MN's hostname obtained during the authentication. If found, the new AR performs a reverse lookup to confirm that some AR has actually assigned the IP address to the MN's hostname. If this is confirmed, a routing update is performed. The new AR creates a Border Gateway Protocol (BGP) update message containing the MN's IP address and sends the message to its peers to announce the new route. If the MN moves to an AR in the same cluster where the MN's IP prefix was originally assigned by an AR, the BGP update message will be sent to the parent routers in the aggregation layer, by which the update will be reflected down to all the other routers in the same cluster. Otherwise, the aggregation router propagates the update up to the core layer, which reflects the update down to all the other aggregation routers and then down to all the ARs in the access layer. With the route update, the packets destined to the MN can be routed to the new AR through an optimal path. In addition, to prevent the disruption of the MN's ongoing sessions during the handover process, after the MN moves, the previous AR should forward to the new AR the packets destined for the MN's IP addresses that it has received.

Other Approaches in 3GPP and the Research Domain

In this section, we briefly introduce the existing approaches in 3GPP to handle the ever increasing mobile data traffic and also some DMM solutions in the research domain to evolve the architecture of the current mobile networks.

The 3GPP continues to make ongoing efforts to alleviate the traffic load on the mobile core network. The prevalent schemes are Local IP Address (LIPA) and Selected IP Traffic Offload (SIPTO) [11]. LIPA enables an IP-capable MN connected via a femtocell (i.e., a small base station) to access other IP-capable entities in the same residential or enterprise IP network, without the user plane traversing the mobile operator's core network. LIPA is achieved by collocating a local gateway with the femtocell and enabling a direct user plane between the local gateway and the femtocell. On the other hand, SIPTO enables a mobile operator to offload certain types of traffic (e.g., best effort Internet traffic) at a network node close to the MN's point of attachment to the access network. This is achieved by selecting a set of gateways (i.e., serving gateway and packet gateway) geographically or topologically close to the MN's point of attachment. For 3GPP Release 11, there is currently a work item called LIPA Mobility and SIPTO at the Local Network (LIMONET) [12], which aims to provide mobility support for LIPA between femtocells within the same local gateway and also enable SIPTO at the local network. However, LIPA and SIPTO with LIMONET provide only localized mobility support (i.e., support for devices moving within a small geographical region).

In the research domain, there are some proposals [13, 14] to create an evolved architecture of the current mobile networks based on the PMIPv6- and MIPv6-based DMM approaches. In [13], a DMM solution is developed in the context of the European Union's Multimedia Transport for Mobile Video Application (MEDIEVAL) project,[4] which aims to evolve today's mobile Internet architecture for efficient video transport. The MEDIEVAL DMM solution employs a hybrid approach of the PMIPv6- and MIPv6-based DMM, where PMIPv6-based DMM is used when an MN is in a localized mobility domain that consists of a set of access routers, while MIPv6-based DMM is used when the MN moves to another localized mobility domain. Also, it provides multicast mobility support for efficient video traffic delivery. On the other hand, [14] presents an evolved 3GPP architecture supporting the PMIPv6- and MIPv6-based DMM approaches, and describes in detail how to implement them in the existing 3GPP architecture.

Comparative Analysis of IETF DMM Approaches

In this section, we first provide a qualitative analysis to compare the three IETF DMM approaches, in terms of three main characteristics used for the evaluation of mobility protocols:

- Data and signaling overhead
- Handover latency
- Packet delay

We then discuss which DMM approaches are more suitable for efficient mobile video delivery.

Data and Signaling Overhead — We first consider packet overhead. Routing-based DMM requires no IP tunnel. On the other hand, both PMIPv6-based DMM and MIPv6-based DMM without RO both, when a handover occurs, need to establish an IP tunnel for ongoing sessions (but not for new sessions initiated after the handover). Hence, PMIPv6-based DMM and MIPv6-based DMM without RO incur 40-byte overhead (for an IPv6 packet header) due to packet encapsulation. For MIPv6-based DMM without RO, the overhead is incurred even over a wireless link (which connects the MN to the network). On the other

[4] *http://www.ict-medieval.eu/*

hand, when RO is used for MIPv6-based DMM, no IP tunnel is required, and thus the overhead can be removed.

We next consider signaling overhead. For session setup, routing-based DMM requires both the MN and the AR to exchange control messages with the DNS to register the forward and reverse mappings between the MN's hostname and IP address, while PMIPv6- and MIPv6-based DMM do not need such a registration. For handover, routing-based DMM requires the new AR (to which an MN moves) to perform a reverse lookup for each of the MN's IP addresses to confirm the mapping from each IP address to the MN's hostname, followed by the routing update. The signaling overhead due to the routing update can be very high if an MN moves to an AR in a different cluster, since in such a case the routing update might be propagated through a large part of the network. On the other hand, in PMIPv6- and MIPv6-based DMM, the control messages are exchanged between two entities for each of the MN's bindings; thus, the signaling overhead increases in proportion to the number of the MN's IP prefixes/addresses anchored at ARs other than the serving AR.

Handover Latency — It is known that, compared to PMIPv6, MIPv6 has inferior performance in handover latency [15]. The heaviest contribution to the handover latency of MIPv6 is made mainly by the binding signaling delay. On the other hand, the handover latency of MIPv6-based DMM is improved from MIPv6 due to the localized and distributed mobility management being introduced, and depends strictly on the underlying network topology. Unlike PMIPv6- and MIPv6-based DMM, routing-based DMM requires, in addition to the confirmation of the MNs' IP-address-to-hostname mappings, the routing updates to be reflected in all the ARs in a cluster, the number of which can be substantially large when an MN moves to an AR in a different cluster. Therefore, the routing update procedure may cause a large delay. This results in routing-based DMM incurring higher handover latency compared to PMIPv6- and MIPv6-based DMM.

Packet Delay — In routing-based DMM, packets are routed through an optimal path. On the other hand, in PMIPv6-based DMM and MIPv6-based DMM without RO, packets are routed through a suboptimal path when a handover occurs. If RO is used for MIPv6-based DMM, the routing path can be optimized at the cost of compromising MNs' location privacy and also increasing the handover latency. Thus, in terms of packet delay, routing-based DMM outperforms PMIPv6- and MIPv6-based DMM, and MIPv6-based DMM, when RO is used, outperforms PMIPv6-based DMM.

Based on the qualitative analysis made above, the PMIPv6-based DMM seems most suitable for real-time interactive video applications (e.g., videoconferencing and interactive gaming) due to its low handover latency. MIPv6-based DMM seems suitable for supporting real-time video applications that can tolerate some amount of delay (e.g., video on demand with a large buffer) and also non-real-time applications (e.g., video downloading) due to its good performance in handover latency and packet delay. Routing-based DMM would be the most efficient way to support non-real-time video applications due to its superior performance in packet routing; however, the scope of its usage seems limited to less mobile users due to its high signaling overhead.

Concluding Remarks

In this article, we have surveyed different approaches for distributed mobility management in IETF, 3GPP, and the research domain, which can efficiently handle mobile video traffic. We have mainly focused on the three IETF DMM approaches: PMIPv6-based, MIPv6-based, and routing-based DMM. We have presented a qualitative analysis to compare them in terms of data and signaling overhead, handover latency, and packet delay. The qualitative analysis suggests that PMIPv6- and MIPv6-based DMM are more suitable for efficient mobile video delivery than routing-based DMM, and can better support delay-sensitive and delay-tolerant video traffic, respectively.

References

[1] Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016," http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html, Feb. 2012.

[2] S. Gundavelli *et al.*, "Proxy Mobile IPv6," IETF RFC 5123, Aug. 2008.

[3] H. Soliman, "Mobile IPv6 Support for Dual Stack Hosts and Routers," IETF RFC *5555*, June 2009.

[4] P. McCann, "Authentication and Mobility Management in A Flat Architecture," IETF Internet-Draft, draft-mccann-dmm-flatarch-00, work in progress, Mar. 2012.

[5] C. Perkins, D. Johnson, and J. Arkko, "Mobility Support in IPv6," IETF RFC *6275*, July 2011.

[6] D. Liu, "Distributed Deployment of Mobile IPv6," IETF Internet-Draft, draft-liu-mext-distributed-mobile-ip-00, work in progress, Sept. 2011.

[7] C. Bernardos *et al.*, "A PMIPv6-Based Solution for Distributed Mobility Management," IETF Internet-Draft, draft-bernardos-dmm-pmip-01, work in progress, Mar. 2012.

[8] C. Bernardos and J. Zuniga, "PMIPv6-based Distributed Anchoring," IETF Internet-Draft, draft-bernardos-dmm-distirubted-anchoring-00, work in progress, Mar. 2012.

[9] B. Sarikaya, "Distributed Mobility IPv6," IETF Internet-Draft, draft-sarikaya-dmm-dmipv6-00, work in progress, Feb. 2012.

[10] H. Chan *et al.*, "Framework for Mobility Management Protocol Analysis," IETF Internet-Draft, draft-chan-dmm-framework-gap-analysis-06, Nov. 2012.

[11] 3GPP, "General Packet Radio Service (GPRS) Enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) Access," TR 23.401, Sept. 2011.

[12] 3GPP, "LIPA Mobility and SIPTO at the Local Network," TS 23.859, July 2011

[13] F. Giust *et al.*, "A Hybrid MIPv6 and PMIPv6 Distributed Mobility Management: The MEDIEVAL Approach," *Proc. 16th IEEE Symp. Computers and Commun.*, June 2011, pp. 25–30.

[14] C. Bernardos, J. Zúñiga, and A. Reznik, "Towards Flat and Distributed Mobility Management: A 3GPP Evolved Network Design," *Proc. IEEE ICC '12*, June 2012, pp. 6855–61.

[15] K.-S. Kong *et al.*, "Mobility Management for All-IP Mobile Networks: Mobile IPv6 vs. Proxy Mobile IPv6," *IEEE Wireless Commun.*, vol. 15, no. 2, Apr. 2008, pp. 36-45.

Biographies

DONG-HOON SHIN (donghoon.shin.2@asu.edu) is a postdoctoral scholar at Arizona State University, Tempe. He received his B.E. degree from Korea University, Seoul, in 2003, his M.S. degree from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, in 2006, and his Ph.D. degree from Purdue University, West Lafayette, Indiana, in 2012. He was an intern in Intel's Wireless Standards and Advanced Technologies Group from August 2011 to June 2012. His research interests are in the areas of wireless communication networks and cyber-physical systems.

DANNY MOSES' (danny.moses@intel.com) biography was unavailable at the time this issue went to press.

MUTHAIAH (MUTHU) VENKATACHALAM (muthaiah.venkatachalam@intel.com) is the director of System Architecture in Intel's Mobile Computing Group. He has served on the Editorial Board of Elsevier's *Journal of Computer Networks*. He has 38 issued patents with several pending. Previously at Intel, he was the lead software/system architect for the IXP2300 network processor and has driven Intel's efforts on developing network-processor-based IP and ATM traffic management solutions; system architectures for broadband access, wireless access, and metropolitan optical networking systems. As the network processor architect at Intel, he has helped secure numerous design wins for Intel IXPs.

SAURABH BAGCHI [SM] (sbagchi@purdue.edu) is an associate professor in the School of Electrical and Computer Engineering and the Department of Computer Science (by courtesy) at Purdue University. He is a Senior Member of ACM. At Purdue, he is the assistant director of CERIAS, the security center, an IMPACT Faculty Fellow, and leads the Dependable Computing Systems Laboratory (DCSL). He is a visiting scientist with IBM's Austin Research Laboratory since 2011. Best of all, he gets to work with the smartest group of students in DCSL, who are not afraid to make and break real computer systems.