Hierarchical Broadcasting in the Future Mobile Internet

Cristian Hesselman and Henk Eertink
Telematica Instituut, The Netherlands
{hesselman, eertink}@telin.nl

Abstract

We describe an architecture for the hierarchical distribution of multimedia broadcasts in the future mobile Internet. The architecture supports network as well as application-layer mobility solutions, and uses stream control functions that are influenced by available network resources, user-defined preferences and application-level streaming policies. We are implementing the architecture in two environments, one with and one without Mobile IP support.

1. Introduction

The mobile Internet of the future will be characterized by extreme heterogeneity in terms of administrative domains (public and private, stationary and moving, with or without mobility support), network technologies (ranging from fixed, to Bluetooth-like personal area networks, to WLAN hot spots and 3G cellular networks), and mobile hosts (varying from laptops down to PDAs) [1, 3, 4]. One of the challenges in such an environment is to provide seamless services that allow users to move across network technologies and administrative boundaries [4].

In this paper, we present ongoing work on the use of application-level policies [17] that support this challenge for multimedia broadcast services (e.g., TV, radio, e-cinema, e-learning, etc. [7]). These policies for instance control the availability of TV channels based on the type of subscription of the end-user, or on the availability of server or network resources. To enforce the policies, we designed (and are currently implementing) a mobility and quality-aware content distribution platform that runs entirely at the application-level and does not require any QoS support or mobility functions from the IP substrate.

In this paper, we will look into the design of our platform's functions. In Section 2, we will first discuss our model of multimedia broadcasts that take place in the the future mobile Internet. We then present the design of our platform in Section 3. In Section 4, we consider the status and future directions of our work. We conclude with a brief summary in Section 5.

2. Model

We model multimedia broadcasts in terms of channels and streams (Section 2.1), and the Internet's infrastructure in terms of domains (Section 2.2).

2.1 Channels and Streams

Our platform allows so-called *channels* (e.g., 'world news') to be broadcast at various perceptual *quality levels* [10] (e.g., 'TV' or 'video phone' quality [7, 9]). Each quality level ultimately maps to an encoded and packetized digital *stream*[1]. In the rest of this paper, we use the term channel to refer to the logical content, and the term stream to refer to a digital stream that carries the channel at a particular quality level using a well-defined packet format (e.g., based on MPEG-4 [18] and RTP [12]).

2.2 Domains

Motivated from a business perspective, we decompose the Internet's infrastructure into network domains and streaming domains (cf. [6, 7, 8])

Network Domains

In our model, the network infrastructure of the Internet consists of *network domains* that offer IP-level connectivity to end-users. A network domain consists of

[1] The system should also support layered encoding (i.e., quality levels that are obtained via a combination of streams), but this adds considerable complexity to the description of the model. We refer to [5] for a more elaborate overview that includes both perceptual quality and layering.

one or more subnets, which may be based on different network technologies (e.g., UMTS, Ethernet, Bluetooth, etc.). The coverage areas of subnets may form an overlay [1, 11]. For example, subnet S3 (see Figure 1) forms an overlay for subnets S1 and S2. A network domain *may* furthermore support QoS (e.g., based on IntServ) or mobility (e.g., using Mobile IP), but our platform does not require such functions.

Multi-mode terminals (or: *mobile hosts*) are able to connect to several subnets, possibly simultaneously. For instance, a laptop equipped with an 802.11 wireless LAN interface, an 802.3 Ethernet interface and a UMTS connection (e.g., via a Bluetooth interface to a mobile phone) may be able to (simultaneously) connect to subnets of the corresponding technologies. The subnets may belong to network domains of different administrative authorities.

We assume that mobile hosts automatically receive IP addresses for their interfaces. The application may or may not be aware of these address changes [19], depending on the availability of a network-layer mobility solution such as Mobile IP [16].

Streaming Domains

In the lack of widespread deployment of IP multicast, we expect that current content infrastructure trends will continue to achieve scalability. This means that content distribution can be hierarchical and that adaptation will take place as close to the user as possible [2, 6, 8, 13, 14].

In our view, the application-level infrastructure of the Internet will consist of streaming domains that interact with one another at the level of channels. Our work concentrates on *mobility-aware streaming domains* (MSDs), which are streaming domains that deliver channels to mobile hosts at mobile-specifc quality levels. For example, MSD2 (see Figure 1) delivers channel ‘world news’ to a mobile host at quality level ‘mobile TV’.

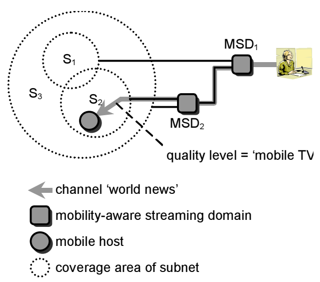

Figure 1. Overview.

An MSD is either the original source of a channel or a surrogate that acts on behalf of the source [6]. For example, MSD2 (Figure 1) is a surrogate for channel ‘world news’, while MSD1 is its source. MSD1 could be a news network such as CNN; MSD2 could be an Akamai-like organization.

A surrogate MSD may process an incoming channel to construct (additional) quality levels, for instance using on-the-fly frame dropping and transcoding [9, 13].

MSDs may connect to subnets directly [2] or indirectly through backbone connections [6, 8]. A subnet and an MSD may belong to the same administrative authority (e.g., a walled-garden telco operator) or to different administrative authorities (e.g., a ‘pure’ ISP operating the subnets of a network domain and an Akamai-like organization operating the MSD).

Observe that there need not be a relation between subnets, geographical location and MSDs (see Figure 1).

3. Platform

Our platform provides the control logic that allows a mobile host to select a proper stream for a channel from one or more MSDs (cf. Figure 1). The platform also allows the mobile host to remain connected to the channel (not the stream!) while it moves. This may require the host to switch between quality levels (adaptive broadcasting), possibly of different MSDs.

In this section, we discuss the services that the platform provides and illustrate their behavior by means of a scenario in which a user gains access to channels through an electronic program guide (EPG). When the user has selected a channel (e.g., ‘world news’), the EPG starts a viewer for that channel. The viewer allows the user to control the reception of the channel (e.g., start and stop) and displays a quality label (e.g., ‘mobile TV’ or ‘mobile video phone’) that reflects the current perceptual quality of the channel.

The services of the platform are: a registration service (Section 3.1), a directory service (Section 3.2), and a channel control service (Section 3.3).

3.1 Registration Service

The registration service registers a user with one or more MSDs that can be reached from the mobile host’s current location (cf. Figure 1). The service is available on mobile hosts through a local software component called the virtual registrar. It is invoked by the EPG at start-up.

Each MSD operates a registrar [22] that allows the virtual registrar to register the user with the MSD. Registration is based on the user’s policies (e.g., number of MSDs to simultaneously register with) and allows for the exchange of user preferences, explicit authentication (e.g., to enable billing), and the exchange of terminal capabilities. The MSD can use this information as an input for a policy-driven procedure that decides which quality levels it supports under the current conditions. For example, the registration information allows for MSD-policies that enable quality level ‘mobile TV’ only when there are at least 10 interested mobile hosts that can handle this quality level.

Registrars use soft state for registrations. This means that the virtual registrar must regularly reregister the user with an MSD. If it does not, the MSD's registrar will eventually consider the registration stale and will remove its state.

The local virtual registrar also registers the end-user when it detects a new MSD, typically as a result of the mobile host moving into a new subnet. The virtual registrar can for instance learn about the availability of a new MSD through a DHCP option [20] at the time it receives an IP address.

Figure 2 shows the structure of the registration service for the example of Figure 1.

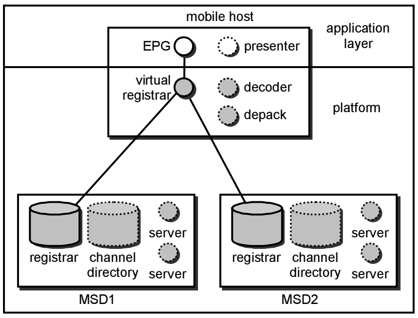

Figure 2. Registration service.

3.2 Directory Service

The directory service retrieves the names of channels that are supported in the MSDs with which the virtual registrar has registered the user. The service is available on mobile hosts through a local component called the virtual directory. It is invoked by the EPG after the virtual registrar has successfully registered the user with at least one MSD.

The virtual directory contacts individual, reachable, MSDs to retrieve the names of the channels they support. Each MSD operates a channel directory for this purpose. A channel directory contains a controller object (see Section 3.3) for each channel that the corresponding MSD supports. The virtual registrar retrieves the names of these channels from the controllers and passes them to the EPG. The virtual registrar can use a request-response mechanism to interact with directories, or it can rely on repeated announcements from the channel directory (cf. SAP [21]).

The directory service is also invoked by the virtual registrar if a new MSD appears or when one disappears as a result of mobility. This will update the set of list of channel names. The virtual directory informs the EPG of new or obsolete channel names.

Figure 3 shows the structure of the directory service for the example of Figure 1.

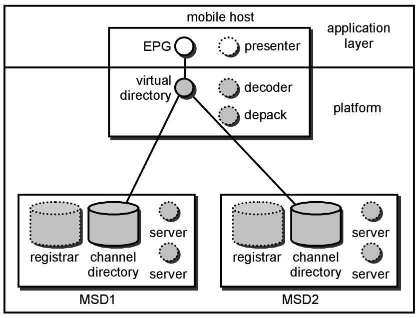

Figure 3. Directory service.

3.3 Channel Service

The channel service determines which stream from a selected channel the mobile host will receive from which MSD. The service is available on mobile hosts through a local component called the virtual controller and is invoked by a viewer. There may be several instances of the virtual controller, each of which supports one viewer.

The virtual controller allows a viewer to start and stop the reception of a channel. When the viewer starts a channel, the virtual controller automatically selects a stream from an MSD. The selection process is client specific, but will typically be based on the user's preferences (e.g., preferred quality levels, network interfaces, prices of channels from various MSDs, etc.) and on the list of streams (with quality levels) that the mobile host can receive at its current location. The establishment of the latter is based on a negotiation procedure between the virtual controller and several channel controllers (e.g., those of 'worlds news' in MSD1 and MSD2, see Figure 1), which are aware of the available resources in their MSDs. A controller for instance knows at which quality levels its channel is currently being streamed from which servers to which mobile hosts. The negotiation procedure also considers the resource availability on the mobile host and in the subnets the host can connect to. For the latter, we currently use an application-layer beaconing algorithm [15] to each registered mobile host to obtain information on network congestion[2]. This information allows the mobile host to decide which quality levels the subnets can handle.

After the virtual controller has selected a particular stream from an MSD, it will use standard streaming clients to display the stream to the end-user (e.g., it hands the URL in a stream description over to a Real Player or a Windows Media Player that contacts the appropriate server in the MSD).

[2] Alternatively, RTCP feedback mechanisms can be used; however these mechanisms are not always implemented, or are not sensitive enough.

Channel controllers use soft state for subscriptions on quality levels. This means that a virtual controller must regularly refresh its subscription to a quality level. If it does not, the controller will eventually consider the subscription stale and will remove its state. The result may be that the mobile host can no longer receive the stream.

The quality level at which a mobile host receives a channel may become unavailable when it moves into a new subnet. For example, the mobile host might move into a subnet where its current MSD is no longer available (e.g., in Figure 1 MSD2 is not available in S1), the MSD might not support the current quality level in the new subnet, and so on. In such situations, the virtual controller selects a new quality level, possibly from another MSD via another subnet. This may involve a renegotiation with one or more controllers. Renegotiation also occurs when the network resources become scarce (currently observed through the beaconing mechanism).

The change between an old and a new quality level is subject to an adaptation policy [10, 15]. For a smooth handoff, it is crucial that the virtual controller finds a new quality level as soon as possible.

Figure 4 shows the structure of the channel service for the example of Figure 1.

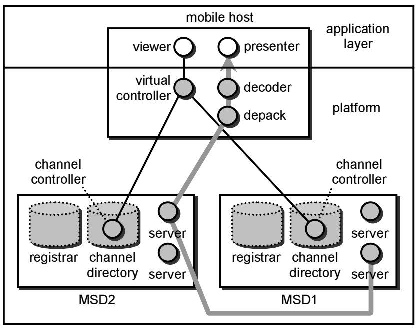

Figure 4. Channel service.

4. Status and Future Work

We are currently implementing our platform in an environment similar to the one shown in Figure 1. The implementation is based on hierarchical streaming and application-layer mobility handling, which means that the involved subnets are not mobility aware (i.e., there is no layer 3 roaming between subnets). We will use this version of the platform to experiment with application-layer policies that consider quality and mobility aspects.

We are also building an implementation of the platform that uses Mobile IPv6. The environment consists of multiple subnets and a single origin MSD (end-to-end approach). This infrastructure contains all functional elements (registration, directory lookup and channel selection), but handles all mobility aspects through Mobile IP. As a result, application-level policies cannot take mobility aspects into account since mobility is hidden from the platform by Mobile IP. This implementation also allows us to compare handoff delays with those of the non-Mobile IP solution.

In both cases we implement the software in C/C++, on an IPv6 network infrastructure, using public domain streaming servers and clients.

Our future work will be on the description of policies and the policy execution environment; and on performance implications (in terms of handoff delays and bandwidth usage) between various instantiations of the general concepts described in Section 2.

5. Summary

We presented an architecture for the hierarchical distribution of multimedia broadcasts in the future mobile Internet. We believe the architecture is promising, in the sense that the surrogate MSDs can provide a scalable solution for multimedia broadcasting in a heterogeneous environment. The architecture also allows for different business models, based on open, traditional Internet-style end-to-end models, as well as the traditional 'walled garden' telecommunication provider models, and the combination of these with content distribution service providers and policy-driven edge servers.

References

[1] M. Haardt, W. Mohr, "The Complete Solution for Third-Generation Wireless Communications: Two Modes on Air, One Winning Strategy", IEEE Pers. Comm., Dec. 2000

[2] A. Dutta, W. Chen, H. Schulzrinne, O. Altintas, "Mobility Support for Wireless Streaming in MarconiNet", Proc. of IEEE Broadband Wireless Summit, Interop 2001, Las Vegas, USA, May 2001

[3] H. Yumiba, K. Imai, M. Yabusaki, "IP-Based IMT Network Platform", IEEE Pers. Comm., Oct. 2001

[4] N. Drew, M. Dillinger, "Evolution Toward Reconfigurable User Equipment", IEEE Comm. Mag., Feb. 2001

[5] C. Hesselman, H. Eertink, "Broadcasting Multimedia Channels in Future Mobile Systems", Position Paper, Proc. 6th Int'l Conference on Protocols for Multimedia Systems (PROMS'01), Enschede, The Netherlands, Oct. 2001

[6] G. Tomlinson, R. Chen, M. Hofmann, "A Model for Open Pluggable Edge Services", Internet Draft draft-tomlinson-opes-model-00.txt, July 2001

[7] M. Vernick, S. Bryden, M. Condry, G. Disher, J. Straley, W. Walkoe, "Requirements for End-To-End Delivery of Broadband Content", Broadband Content Delivery Forum, Oct. 2001

[8] DOLMEN Consortium, "Open Service Architecture for Mobile and fixed environments (OSAM)", Final Release, Version 4.0, July 1998

[9] N. Yeadon, F. Garcia, D. Hutshison, D. Shepherd, "Filters: QoS Support Mechanisms for Multipeer Communications", IEEE Journal on Selected Areas in Comm., Sept. 1996

[10] R. Liao, A Campbell, "On Programmable Universal Mobile Channels in a Cellular Internet", 4th ACM/IEEE

International Conference on Mobile Computing and Networking (MOBICOM'98), Dallas, October, 1998
[11] M. Stemm, R. Katz, "Vertical Handoffs in Wireless Overlay Networks", ACM Mobile Networking, Special Issue on Mobile Networking and Internet, Spring 1998
[12] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, Jan. 1996
[13] A. Balachandran, A. Campbell, M. Kounavis, "Active Filters: Delivering Scaled Media to Mobile Devices", 7th Int. Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV'97), St. Louis, USA, May 1997
[14] E. Amir, S. McCanne, R. Katz, "An Active Service Framework and its Application to Real-time Multimedia Transcoding", Proc. of ACM SIGCOMM'98, Vancouver, Canada, Sept. 1998
[15] C. Hesselman, H. Eertink, A. Peddemors, "Multimedia QoS Adaptation for Inter-tech Roaming", accepted for the 6th IEEE Symosium on Computers & Communication (ISCC 2001), Hammamet, Tunisia, July 2001
[16] J. Solomon, "Mobile IP — The Internet Unplugged", Prentice Hall, 1998
[17] IETF Policy WG, "Terminology for Policy-Based Management", RFC 3198, Nov. 2001
[18] R. Koenen, "MPEG-4 — Multimedia for Our Time", IEEE Spectrum, Feb. 1999
[19] A. Snoeren, H. Balakrishnan, F. Kaashoek, "Reconsidering Internet Mobility", Proc. 8th Workshop on Hot Topics in Operating Systems (HotOS-VIII), Elmau/Oberbayern, Germany, May 2001
[20] R. Droms, "Automated Configuration of TCP/IP with DHCP", IEEE Internet Computing, July-August 1999
[21] M. Handley, C. Perkins, E. Whelan, "Session Announcement Protocol", RFC 2974, Oct. 2000
[22] F. Vakil, A. Dutta, J-C. Chen, M. Tauil, S. Baba, N. Nakajima, Y. Shobatake, H. Schulzrinne, "Supporting Mobility for Multimedia with SIP", Internet Draft, draft-itsumo-sipping-mobility-multimedia-01.txt, Dec. 2001