Handling multicast video over non multicast networks

Michail Tolkas
University of Western Macedonia
Kozani, Greece
e-mail: mtolkas@media.mit.edu

Michail Bletsas
MIT Media Lab
Cambridge, MA, USA
e-mail: mbletsas@media.mit.edu

Pantelis Aggelidis
University of Western Macedonia
Kozani, Greece
e-mail: paggelidis@uowm.gr

***Abstract*— In this paper we evaluate the basic methods of video transmission, focusing on multicast, the popularity of which has increased over the last years. Although the specific technique is able of reducing significantly the consumption of network resources, such as the bandwidth, the specifications of IEEE 802.11 family networks do not support the mechanisms for reliable multicast transmission. Therefore, multicast video transmission over those wireless networks is considered to be almost inapplicable. In this paper we propose the use of VPN technologies for deploying an efficient and low cost infrastructure of multicast video transmission over wireless networks.**

Keywords; Multicast over wireless networks, VPN, packet encapsulation

I. INTRODUCTION

Multicast is a form of communication by which information is to be delivered from one or more sources to multiple receivers. This differs from unicast, which is one-to-one, and broadcast, which is one-to-all. Multicast is one-to-some (dynamically) communication [1]. *Unicast*, the most common, is simply point-to-point communication between two devices on the network, such as a PC and a file server [2]. Unicast is appropriate for the majority of applications, but it falls short in several collaboration software areas.

A network *broadcast* allows one station on the network to simultaneously talk to all devices contained in the same broadcast domain, or subnet. But broadcasts have their failings as well.

IP multicasting is the answer for most of the issues. It is a communication method by which a single station can transmit to multiple receivers simultaneously, but unlike the "one or everyone" possibilities of unicast and broadcast, the transmitting machine can specify a specific group of machines to receive the information. This is accomplished by transmitting to a multicast IP address, which can be conceptualized as a TV channel. Machines interested in receiving the information simply "tune in", using IGMP protocol that we'll discuss in depth later, to the particular multicast address that contains the data stream of interest. IP multicasting, when implemented properly, lets shared data streams to be transmitted over the network once and solely to those recipients who want to receive the information.

II. MULTICAST ADDRESSING

In multicast communication, we are immediately faced with two problems: 1) how to identify the receivers of a multicast packet and 2) how to address a packet sent to these receivers [3]. In the case of unicast communication, the IP address of the receiver (destination) is carried in each IP unicast datagram and identifies the single recipient. In the case of broadcast communication, all nodes need to receive the broadcast packet, so no destination addresses are needed. In the case of multicast, however, we now have multiple receivers.

Does it make sense for each multicast packet to carry the IP address of all of the multiple recipients? While this approach might be workable with a small number of recipients, it would not scale well to the case of hundreds or thousands of receivers. The amount of addressing information in the datagram would swamp the amount of data actually carried in the packet's payload field.

In the Internet architecture, a multicast packet is addressed using address indirection. That is, a single identifier is used for the group of receivers, and a copy of the packet that is addressed to the group using the single identifier, is delivered to all of the multicast receivers associated with that group. In the Internet, the single identifier that represents a group of receivers is a Class D multicast IP address. The group of receivers associated with a Class D address is referred to as a multicast group.

III. MULTICAST ROUTING

A. Internet Group Multicast Protocol (IGMP)

Part of the appeal of IP multicasting is that the multicast traffic is present only on those subnets where one or more hosts are actively requesting it [2]. Before transmitting a given multicast stream onto a subnet, a router needs to know if any machines on that subnet want to receive that multicast. For IP networks [4], the Internet Group Multicast Protocol (IGMP) is an IP datagram protocol between routers and hosts that allows group membership lists to be dynamically maintained.

B. Making Switches Multicast-Aware

Because IGMP control messages are transmitted as multicast packets, they are indistinguishable from multicast

data at Layer 2 [5]. The following approaches were developed to resolve this mentioned issue [6].

- *IGMP Snooping*: Referred to as "IGMP Snooping," this approach requires that the switch decode the IP header (Layer-3 information) by examining IP protocol field in order to separate out IGMP messages from normal multicast traffic.
- *Cisco's Group Management Protocol (CGMP)*: The second approach, CGMP, is proprietary to Cisco and involves a router-to-switch multicast-group information exchange protocol.
- *Group Address Resolution Protocol (GARP)*: A third approach is the IEEE's GARP protocol, whose primary purpose is to maintain VLAN group information.

IV. Multicast Distribution Trees and Forwarding

A. Distribution Trees

Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network to deliver traffic to all receivers [5]. Distribution trees may be formed as either source-based trees or shared trees.

Source-based distribution trees build an optimal shortest-path tree rooted at the source. Each source/group pair requires its own state information [notated as (S,G), pronounced S comma G, in which S is the IP address of the source and G is the multicast group address], so for groups with a very large number of sources, or networks that have a very large number of groups with a large number of sources in each group, the use of source-based trees can stress the storage capability of routers.

Shared distribution trees are formed around a central router, called a *rendezvous point* or core, from which all traffic is distributed regardless of the location of the traffic sources. The advantage of shared distribution trees is that they do not create lots of (source,group) state information in the routers. The disadvantage is that the path from a particular source to the receivers may be much longer, which may be important for delay-sensitive applications. The rendezvous router may also be a traffic bottleneck if there are many high data rate sources [7].

B. Distribution of Receivers

One criterion to determine what type of tree to use, relates to whether receivers are sparsely or densely distributed throughout the network (for example, whether almost all of the routers in the network have group members on their directly attached subnets) [5]. If the network has receivers or members on every subnet or the receivers are closely spaced, they have a dense distribution. If the receivers are only in a few subnets and are widely spaced, they have a sparse distribution. The number of receivers does not matter; the determining factor is how close the receivers are to each other and the source.

C. Multicast Forwarding

In multicast, the source is sending traffic to an arbitrary group of hosts represented by a multicast group address. The multicast router must determine which direction is upstream (toward the source) and which direction (or directions) is downstream. If there are multiple downstream paths, the router replicates the packet and forwards the traffic down the appropriate downstream paths—which is not necessarily all paths. This concept of forwarding multicast traffic away from the source, rather than to the receiver, is called reverse path forwarding [RPF].

D. Reverse Path Forwarding

The idea behind RPF is simple, yet elegant [3]. When a router receives a broadcast packet with a given source address, it transmits the packet on all of its outgoing links (except the one on which it was received) only if the packet arrived on the link is on its own shortest (unicast) path back to the source. Otherwise, the router simply discards the incoming packet without forwarding it to any of its outgoing links. Such a packet can be dropped because the router knows that it either will receive, or has already received a copy of this packet on the link that is on its own shortest path back to the sender. Fig. 1 illustrates the RPF algorithm.

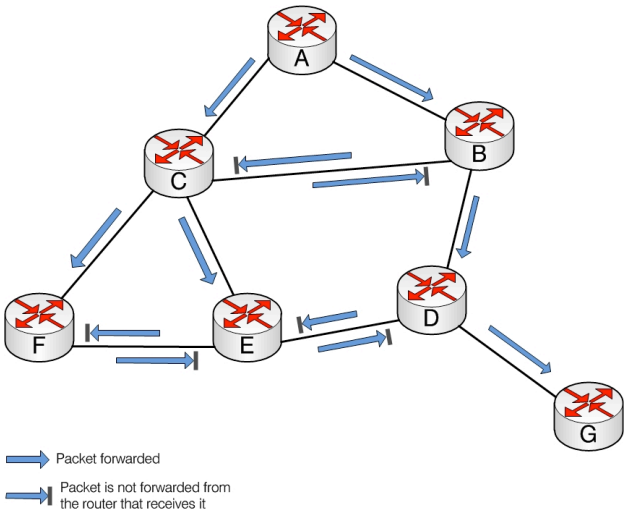

Figure 1 : Reverse Path Forwarding

V. IP Multicast Routing Protocols

Multicast routing protocols fall into two categories: Dense-mode (DM) and Sparse-mode (SM). DM protocols assume that almost all routers in the network will need to distribute multicast traffic for each multicast group (for example, almost all hosts on the network belong to each multicast group). Accordingly, DM protocols build distribution trees by initially flooding the entire network and then pruning back the small number of paths without receivers. SM protocols assume that relatively few routers in the network will be involved in each multicast. The hosts belonging to the group are

widely dispersed, as might be the case for most multicasts in the Internet. Therefore, SM protocols begin with an empty distribution tree and add branches only as the result of explicit requests to join the distribution [4]. The DM protocols are most appropriate in LAN environments with densely clustered receivers and the bandwidth to tolerate flooding, while the SM protocols are generally more appropriate in WAN environments [4]. The most widely deployed protocols are:

- Distance Vector Multicast Routing Protocol (DVMRP) [RFC1075] was the first protocol designed for multicasting. DVMRP implements source-based trees with reverse path forwarding (RPF) and pruning [3].
- Protocol-Independent Multicast (PIM) routing protocol consists of two separate protocols, PIM Dense Mode and PIM Sparse Mode (PIM-SM). The PIM-SM [RFC4601] protocol includes both Any Source Multicast (ASM) and Source-Specific Multicast (SSM) functionality. Whereas PIM-SM has been designed to avoid unnecessary flooding of multicast data, PIM-DM [RFC3973] assumes that almost every subnet at a site had at least one receiver for a group.

VI. Multicast over Wireless Networks

The robustness of the current Internet is due, in large part, to the End-to-End congestion control mechanism of the Transmission Control Protocol (TCP) [3]. With the emergence of multimedia applications, large varieties of these applications are based on UDP and are not responsive to network congestion. The best option for applications that require collaborative communication is multicasting because it can simultaneously distribute multimedia data to multiple users efficiently. However, IP multicast applications are based on best effort delivery, which means that there is no guarantee that provides for reliable data delivery.

Additionally, the IEEE 802.11 standard supports multicast transmissions by simple broadcasting without any feedback (ie. acknowledgement). This means that the multicast sender only performs Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) [8] before transmitting a data frame. There is no MAC-layer recovery on multicast frame as in unicast. As a result [9], the reliability of multicast is reduced due to the increased probability of lost frames resulting from collisions, interference or errors.

According to the IEEE 802.11 specifications [10] (section 9.2.7, page 317) for the multicast MPDU (MAC protocol data unit) transfer procedure:

"[...] only the basic access procedure shall be used. Regardless of the length of the frame, no RTS/CTS (Request to Send / Clear to Send) exchange shall be used. In addition, no ACK shall be transmitted by any of the recipients of the frame. [...] The broadcast/multicast message shall be distributed into the BSS (Basic Service Set). The STA (Station) originating the message shall receive the message as a broadcast/multicast message. Therefore, all STAs shall filter out broadcast/multicast messages that contain their address as the source address. Broadcast and multicast MSDUs (MAC service data unit) shall be propagated throughout the ESS (Extended service set).

There is no MAC-level recovery on broadcast or multicast frames [...]. As a result, the reliability of this traffic is reduced, relative to the reliability of individually addressed traffic, due to the increased probability of lost frames from interference, collisions, or time-varying channel properties. "

"[...]. There are no guarantees that the submitted MSDU will be delivered successfully. Broadcast and multicast transport is part of the data service provided by the MAC. Due to the characteristics of the WM (wireless medium), broadcast and multicast MSDUs may experience a lower QoS, compared to that of unicast MSDUs".

The lack of feedback results in three problems [11]: 1) no contention window (CW - a packet transmission delay) adaptation, 2) no retransmission, and 3) no rate adaptation (an additional issue that we'll discuss below). These problems cause fairness, reliability and efficiency issues, respectively.

First, the collided multicast frame is just dropped without any retransmission at the MAC layer because there is no way to know about the collision of a multicast frame due to the absence of an acknowledgement. The current IEEE 802.11 standard supports only an unreliable service. As the number of other flows increases, the loss rate of multicast frames also increases. Second, no feedback signaling means that an access point (AP) cannot easily collect the state information of STAs that are participating in a multicast group. Hence, most commercial APs use a fixed and low transmission rate (typically, one of basic rates) for multicasting in order to guarantee that as many multicast packets as possible can be received successfully.

VII. Packet encapsulation using VPNs

To avoid the downsides described in the previous sections, our proposed solution consists of the "conversion" (encapsulation) of the multicast packets to unicast. Since the unicast transmission includes feedback mechanisms in the IEEE 802.11 specification, all the above-mentioned problems are eliminated. Our workaround for the packet encapsulation makes use of Virtual Private Networks (VPN). In a VPN, the server creates a secure communication session with the clients using unicast transmission. When a (wireless) client requests a multicast data stream, initially sends the IGMP request, but this time encapsulated as a VPN (unicast) packet. After the server receives the request will forward the multicast traffic to the client. Fig. 2 illustrates the general procedure:

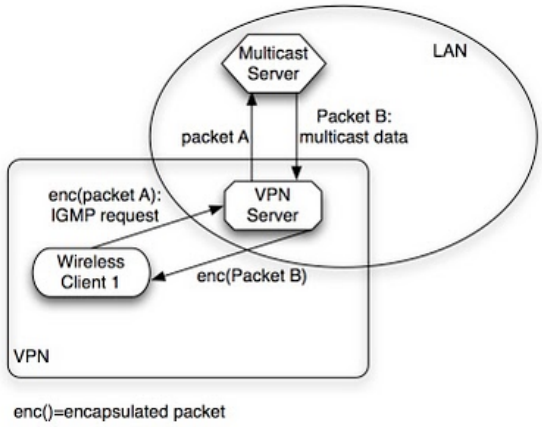

Figure 2: Packet encapsulation with VPN

VIII. VPN IMPLEMENTATIONS

A. *OpenVPN implementation*

Our first attempt will be using the OpenVPN open source VPN solution (version 2.1.3). Our configuration consists of a VPN server connected on the same LAN with the multicast server. OpenVPN operates in two modes, Bridging and Routing. Although bridging mode by default isn't appropriate for the multicast transmission, after a lot of experiments we concluded that OpenVPN server (at least until version 2.1.3) uses only one "pipe" to send data to the clients, so even in routing mode, the same stream arrives to all the clients. That's an undesirable behavior for our goal, because the stream turns out being transmitted throughout the VPN as Broadcast instead of Multicast.

So, if for example we have 10 clients connected to the VPN, and each client requests to join a different multicast group, the result for the server would be to send all 10 streams to all clients. And in a case like this, it's most likely that the routers of the clients will not stand the traffic. A worst-case scenario for this setup would be, each client requesting an HD channel (with almost 11-12Mbps in demanding channels like ESPN HD), so 10x11Mbps = 110Mbps for each router.

Fig. 3 is a Wireshark screenshot from one VPN client and illustrates the downside of OpenVPN in routing mode. In this setup, the two VPN clients subscribed to two different multicast groups, ".18" and ".6". The network interface from one of the two clients, instead of receiving only the stream that requested (channel ".18" or ".6") ends up receiving both.

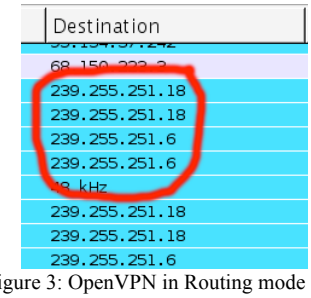

Figure 3: OpenVPN in Routing mode

B. *PPTP implementation*

Because of the failure of the testing version of OpenVPN, we also tried the option of PPTP VPN that comes by default with Windows Server. The results of this test were successful and we managed to have an efficient wireless multicast video transmission over Wi-Fi.

Windows Server comes by default with a PPTP VPN service, which we'll use along with an IGMP proxy to set up our properly working multicast network.

On the Windows Server 2003 we will add the "Remote Access/VPN Server" Role and also use the "Domain Controller" Role for Active Directory authentication if needed.

We need to make sure that the physical interface (usually something like Local Area Connection) is set to "IGMP proxy" and the "Internal" interface (virtual VPN interface) is set as an "IGMP router". Also we might need to change the "IGMP protocol version" to get it properly working. For example, "Version 3" will not work for some versions of MacOS X.

Figure 4 shows how the encapsulated packets arrive to the physical network interface of the wireless client.

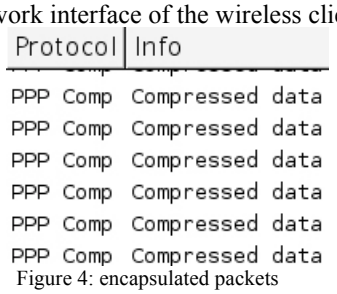

Figure 4: encapsulated packets

Figure 5 shows the actual UDP multicast frames at the virtual PPP interface. We see that the client receives only the packets from the multicast group that previously joined.

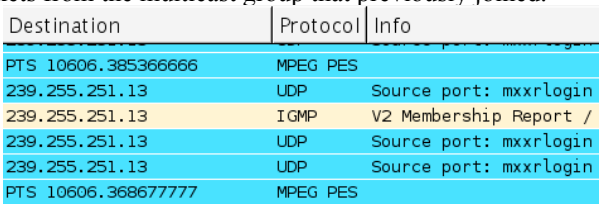

Figure 5: Multicast frames

Finally, Figure 6 shows another example in which we see both the packets and the actual video, using the open source VLC media player.

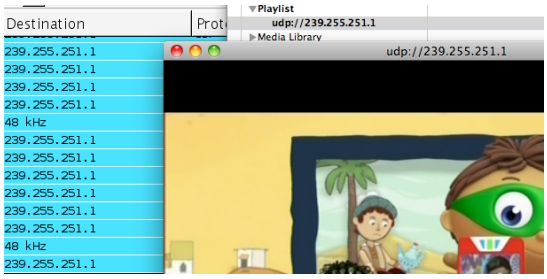

Figure 6: Multicast video over Wi-Fi

IX. Conclusion and future work

In this paper we presented the basic methods of video transmission, focusing on multicast. We analyzed the restrictions caused from the specifications of IEEE 802.1. We proposed and implemented an infrastructure for efficient wireless multicast video transmission using VPNs in networks where previously was not applicable.

The lack of documentation on this subject proves that although multicast is in the research field for more than 15 years, there are only a few applicable solutions, with the majority of those being expensive.

At the time of this writing, the proposed implementation is deployed at the MIT Media Lab, supporting the existing DirecTV system and partially used in the Glass Infrastructure project of the institution.

Future work of this project consists of a detailed design and deployment of a cross-platform architecture for supporting the wireless transmission of multicast video.

References

[1] Matrawy, A.; Lambadaris, I.; , "A survey of congestion control schemes for multicast video applications," Communications Surveys & Tutorials, IEEE , vol.5, no.2, pp.22-31, Fourth Quarter 2003.

[2] T. Tannenbaum; "IP Multicasting: Diving Through the Layers", Network Computing, November 15, 1996.

[3] Keith W. Ross, James F. Kurose; "Computer Networking: A Top-Down Approach (5th Edition)", Addison Wesley, 2009, ISBN-10: 0136079679.

[4] Electronic Publication: Cisco Systems: “Overview of IP Multicast.”

[5] Electronic Publication: Cisco Systems: “Internet Protocol (IP) Multicast”

[6] Electronic Publication: Intelligraphics: “IP-Multicasting Technology: History and Overview”

[7] Electronic Publication: Cisco Systems: “IP Multicast Deployment Fundamentals”

[8] Min-Te Sun; Lifei Huang; Arora, A.; Ten-Hwang Lai; , "Reliable MAC layer multicast in IEEE 802.11 wireless networks," Parallel Processing, 2002. Proceedings. International Conference on , vol., no., pp. 527- 536, 18-21 Aug. 2002

[9] Xiaoli Wang; Lan Wang; Wang, Y.; Zhang, Y.; Yamada, A.; , "Supporting MAC Layer Multicast in IEEE 802.11n: Issues and Solutions," Wireless Communications and Networking Conference, 2009. WCNC 2009. IEEE , vol., no., pp.1-6, 5-8 April 2009

[10] IEEE Standard for Information technology - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications

[11] Nakjung Choi; Yongho Seok; Taekyoung Kwon; Yanghee Choi; , "Leader-Based Multicast Service in IEEE 802.11v Networks," Consumer Communications and Networking Conference (CCNC), 2010 7th IEEE , vol., no., pp.1-5, 9-12 Jan. 2010