Applications and Performance Analysis of Bridging with Layer-3 Forwarding on Wireless LANs

James T. Yu and Chibiao Liu
School of Computer Science, Telecommunications, and Information Systems
DePaul University, Chicago, IL, USA
{jyu, cliu}@cs.depaul.edu

Abastract

This paper presents an in-depth study of applying the bridging technology with layer-3 forwarding (L3F) in Wireless Local Area Networks (WLAN). L3F addresses a limitation of wireless communications at the Medium Access Control (MAC) layer, and uses the information at the network layer (IP address) to forward packets. It has the flexibility of IP routing without the complexity of routing and subnet configurations. The detailed procedure of L3F is presented in this paper, along with thorough performance analysis using a high capability traffic generator and analyzer. The performance results, as measured by throughput and latency, show that the L3F performance is comparable to traditional layer-2 bridging and significantly better than IP routing. This paper also presents two practical applications of using L3F: one is in the Small Office and Home Office (SOHO) to interconnect multiple LAN segments, and the other is in the multi-hop ad hoc wireless networks.

Key words: Wireless LAN, bridging, layer-3 forwarding,, multi-hop

1. Introduction

Wireless Local Area Network (WLAN), as specified in IEEE 802.11 [1], is a fast growing area where people use it at many different environments with various applications. The current standard supports two operation modes:

- An **ad hoc** wireless network has a collection of wireless stations and these stations have a direct wireless connection with each other. However, because of the characteristic of Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), only one active connection is supported at a time.
- An **infrastructure** wireless network has a wireless access point (WAP) as the central control device. WAP is a layer-2 bridge between wired (802.3) [2] and wireless (802.11) networks. Wireless stations cannot directly communicate with each other; instead, they have to use WAP as the bridge for communications.

In addition to the above two operation modes, there is another operation called **wireless bridging**. The purpose of wireless bridge is to interconnect two LAN segments via wireless communications. Each LAN segment has a wireless bridge which communicates with the wireless bridge at the other LAN segment. The network diagrams for the above three operation modes are illustrated in Figure 1.

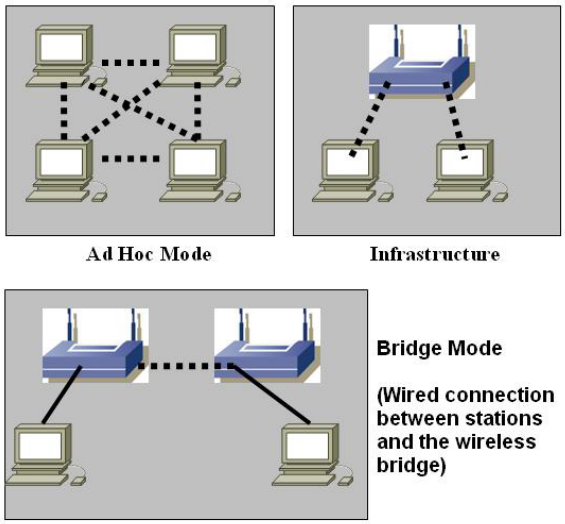

Figure 1. Wireless Operation Modes

To support these different operation modes, the 802.11 standard specifies four address schemes as shown in Table 1 [3].

Table 1. 802.11 Address scheme

Mode/Case	Addr1	Addr2	Addr3	Addr4
Ad Hoc	DA	SA	B-ID	N/A
Infrastructure WAP=>STA	DA	Send WAP	SA	N/A
Infrastructure STA =>WAP	Recv WAP	SA	DA	N/A
Bridge	Recv WAP	Send WAP	DA	SA

STA: Wireless Station **B-ID:** Basic Service Set ID
DA: Destination Address **SA:** Source Address

A WAP (or a wireless bridge) accepts only those frames with its own MAC address. When a station

sends a frame over the wireless media, the station must know the MAC address of the WAP and put this address in the wireless frame. A wireless adapter (also known as NIC) of workstations supports only ad hoc or infrastructure modes, and does not support the bridge mode.

We identified two limitations with the current implementation of WLAN. The first limitation is usually observed in the home or Small Office and Home Office (SOHO) environment [4]. The primary need in these environments is to share the Internet connection and computing resources (such as printers, FAX, and storage devices) as illustrated in Figure 2a. Note that the Wireless Access Point (WAP) in the figure is also an IP router and multi-port Ethernet switch. One traditional solution is to use an Ethernet switch (see Figure 2b) to connect separate LAN segments into one IP subnet. However, this *wired* configuration needs cabling inside the wall and requires licensed professional to do the work.

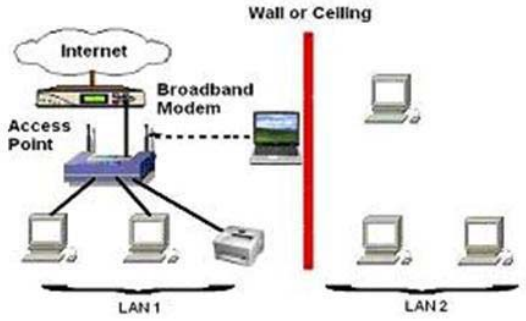

(a) Network Connection

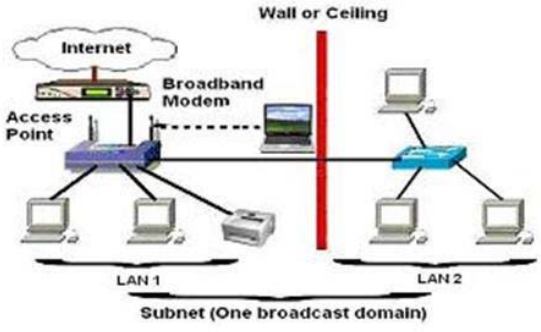

(b) Legacy Solution (wired)

Figure 2. Wired SOHO Network Configuration

The solution to the cabling problem is to apply wireless technology as illustrated in Figure 3a. However, one issue with this wireless configuration is the need of a wireless adapter for each workstation. Another issue is performance impact where all wireless workstations are competing for the shared Radio Frequency (RF) channel and degrade the performance [5]. The third issue is workstation configuration where changes made to the WAP (such as the WEP security key) require manual reconfiguration on every wireless workstation.

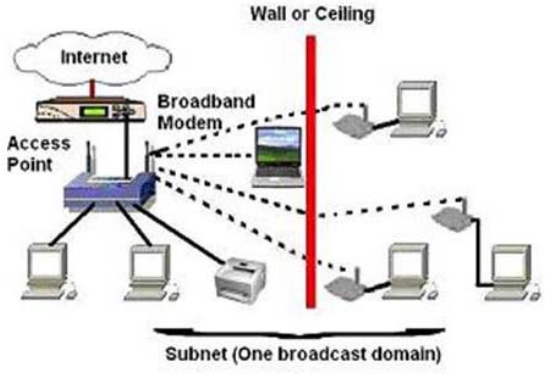

(a) Current Wireless Solution

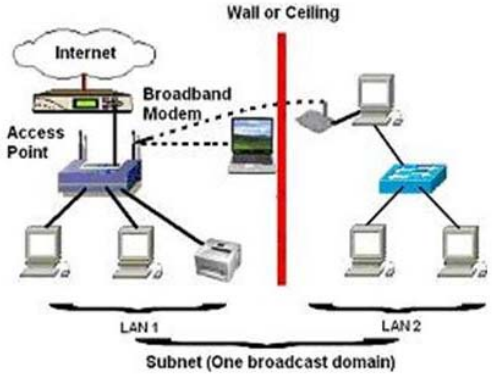

(b) Ideal Wireless Solution

Figure 3. Wireless SOHO Network Configuration

A better solution to the wireless configuration of Figure 3a is to implement the **bridging** function in one wireless workstation and this wireless bridge forwards the traffic to other workstations using wired connection as illustrated in Figure 3b. Because this is a bridging configuration, all workstations are on the same IP subnet (one single broadcast domain). The workstations connecting to the wireless bridge does not know (and does not need to know) whether it is connected to the router via wired or wireless media. Unfortunately, wireless adapters do not support the bridge configuration illustrated in Figure 3b. WAP that supports the bridge configuration is more expensive and more complex than a typical WAP.

The second limitation of wireless communication is the relatively short distance. Testing results show that a typical 802.11b/g communication can reach only 300 ft (with boosting antenna) for 802.11b/g, and the performance goes down significantly as the distance increases. One purpose of the wireless bridge is to

extend the distance range of wireless communication, and this application is known as the wireless multi-hop environment as illustrated in Figure 4. Unfortunately, the current standard of wireless ad hoc network does not support this bridging communication as described in Table 1.

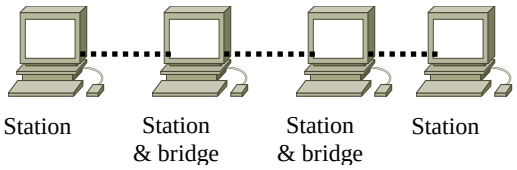

Figure 4. Multi-hop in Wireless Networks

This paper describes a new bridging technology, called bridging with layer-3 forwarding (L3F). It solves the above two limitations of the current wireless communications, and has the potential for other new applications. Technically, we can use IP routing to solve the problems addressed by layer-3 forwarding. However, routing requires multiple IP subnet configurations, and the network becomes more difficult to manage. We also performed extensive testing of the layer-3 forwarding using Windows XP workstations and a high capability traffic generator. The results show that L3F has significantly better performance than IP routing as measured by *delay* and *throughput.*

2. Bridging with Layer-3 Forwarding

2.1 Layer-2 Bridging and IP Routing

Bridging is typically considered a layer-2 (data link layer) technology. A bridging device maintains a **forwarding** table which shows the mapping relation of the L2 address[1] and the physical port [6]. When the device receives an incoming packet, [2] the source address is *learned* to create an entry in the forwarding table, and the destination address is used to determine the *outgoing* port. In the case of unknown destination address, the packet is treated as a broadcast packet.

Routing, on the other hand, is a layer-3 function. Given the popularity of the Internet Protocol (IP), almost all the routing is based on the IP address, and the routing device is called router. A router maintains a routing table for the mapping of IP addresses and outgoing ports. In the SOHO or small networks, the routing table is typically manually configured (static routing) without a dynamic routing protocol. A router accepts only those packets with the destination MAC address to itself, while a bridge accepts all incoming packets regardless of the destination MAC address. In general, bridging configuration is preferred on Local Area Network (LAN) for easier network configuration and better performance, while routing configuration offer more flexibility and better security control.

2.2 Layer-3 Forwarding (L3F) Process

The Layer-3 forwarding (L3F) scheme presented in this paper is based on the Layer-2 forwarding but uses IP addresses to determine the outgoing port. The concept of Layer-3 forwarding is similar to proxy ARP [7] which constructs the forwarding table using the same algorithm as the layer-2 bridging. However, the content of the forwarding table is similar to the IP routing table. The L3F bridge is functioning as a proxy ARP server to proxy ARP requests between the LAN segments.

For a station that sends packets to a L3F bridge, it is not aware of the existence of the L3F bridge, just like it is not aware of the existence of a layer-2 bridge. However, in order for the L3F bridge to process the packet, the station must use the MAC address of the L3F bridge as the destination MAC address. As a result, the station's ARP table should have the MAC address of the L3F bridge, instead of the real MAC address of the destination node. The construction and use of the L3F table and the ARP table are illustrated in the following example (Figure 5).

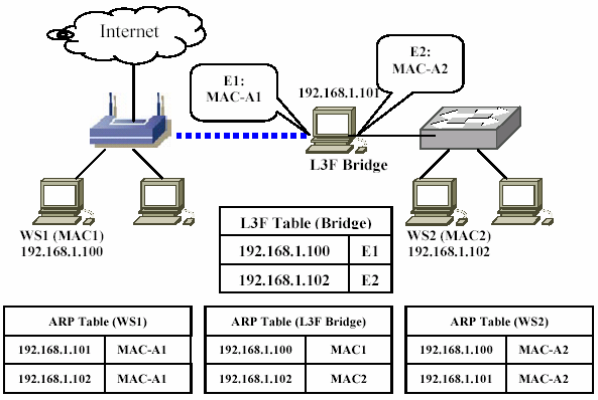

Figure 5. Wireless Network using L3 Forwarding

In the above example, when WS1 sends a packet to WS2, WS1 looks up its ARP table for the MAC address of WS2 (192.168.1.102) which shows the MAC address of the L3F bridge (MAC-A1). As a result, the packet uses MAC-A1 as the destination MAC address. When this packet is received by the L3F bridge, it looks up its L3F table and finds an entry for the destination IP address (192.168.1.102). The L3F bridge uses its ARP

[1] The layer-2 address of Ethernet is Medium Access Control (MAC) address.

[2] At layer-2, a protocol data unit (PDU) is usually called a *frame*. At layer-3, a PDU is usually called a *packet*. The L3F is a technique between L2 layer-2 and layer-3. Without loss of generality, we simply reference a PDU as a packet in this paper, regardless of the layer of the PDU.

table to find the new destination MAC address (MAC2), changes the source MAC address to itself (MAC-A2), and forwards the packet to its physical interface (E2). This packet then arrives at WS2 as if it were directly sent from the L3F bridge. As in the case of layer-2 forwarding, the process is completely transparent to the sender and receiver. They do not know (and do not need to know) the existence of a L3F bridge. In a L3F network, it is a single IP subnet and a single broadcast domain.

3. Experiment Design

The purpose of the experiment is to measure the performance of L2F, L3F, and IP routing based on the network configuration discussed earlier. The Device Under Test (DUT) is Windows XP workstation which supports Layer-2 forwarding, Layer-3 forwarding [8], and IP routing.[3] The traffic generator and analyzer is the IXIA 1600 chassis with dual Gigabit ports for packedt transmission and reception. We follow the IETF RFC-2544 standard for the benchmark performance testing [9][10]. The network connection for the L3F testing is illustrated in Figure 6.

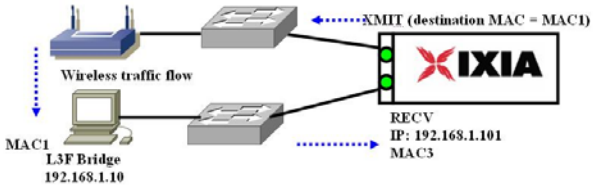

Figure 6. L3 Forwarding Configuration (single IP subnet)

Because the network is a single broadcast domain, the wireless interface (adapter) accepts only those packets with its own MAC address. To create a L3F table with the proper entries in it, we sent testing traffic from the RECV port to help the L3F bridge learn the IP and MAC addresses of the RECV port. The traffic generator then sends packets with the destination MAC address to the L3F bridge (MAC1). The L3F bridge translates the destination MAC address to MAC3 and forwards the packet to the receiving port (RECV) of the traffic analyzer. The traffic generator puts a time stamp on each outgoing packet, and the receiving port generates a traffic report every 2 seconds (the timer is adjustable.) The traffic report includes throughput, delay, bytes sent and received, packets sent and received, and lost packets. In this paper, we are interested in the *delay* and *throughput* measurements. The traffic generator allows us to send ICMP packets, raw IP packets, TCP segments, UDP segments, or any application layer traffic. For performance comparison, we present only UDP traffic in this paper. The comparison of other traffic yields the same results as the UDP traffic.

The second testing configuration is IP routing where the Windows XP station is configured as an IP router. The physical connection is the same as the L3F and is illustrated in Figure 7. In this routing configuration, we create two IP subnets (i.e., two broadcast domains). In order to create the traffic flow, the packets sent from the XMIT port must have the destination IP address of the RECV port (192.168.2.11) and the destination MAC address of the IP router (MAC1). On the IP router, we need to manually create an ARP entry for the RECV port (192.168.2.11 and MAC3) so that the IP router knows how to forward the packet to the IXIA RECV port. As in the L3F testing, we tested the routing configuration with various traffic, raw IP, ICMP, raw UDP, raw TCP, and application traffic.

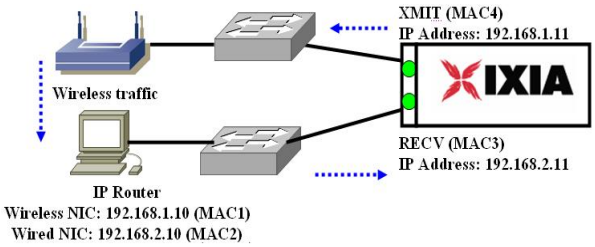

Figure 7. IP Routing Configuration (two IP subnets)

Although layer-2 forwarding (L2F) does not work in the wireless environment as discussed earlier, we are interested in its performance for benchmarking purpose. After careful studying the traffic flow, we were able to create a scenario to measure the L2F performance in a wireless environment. This is based on the broadcast nature of the 802.3 frames, the monitoring capability of the Ethernet switch, and the sniffing capability of the IXIA traffic analyzer. We configure the Windows XP station as the Layer-2 bridge with the *promiscuous* mode. In this configuration, the L2F bridge receives all frames on the shared medium, and not just the frames addressed to itself. In addition, the L2F bridge forwards the frames to all the other ports on the same broadcast domain. The physical connection is also the same as the L3F test connection, but the configuration and traffic flow are different. As illustrated in Figure 8, the traffic is generated on the XMIT port with the destination MAC address of the L2F bridge (MAC1) so that the wireless port of the L2F bridge could accept the incoming packets. Because the L2F bridge is configured in the promiscuous mode, these packets are forwarded to the wired port. Because these packets have the destination MAC address of the L2F bridge,

[3] Linux supports the standard 802.1D bridging and IP routing, but does not support the layer-3 forwarding.

the Ethernet switch would not forward them. To overcome this problem, we use the switch monitoring capability and configure one port as the *monitored* port and another port (to the IXIA RECV port) as the *mirrored* port. Traffic to and from the monitored port is copied to the mirrored port. As a result, we are able to capture the traffic at the IXIA RECV port and generate the traffic report as we did in the L3F testing.

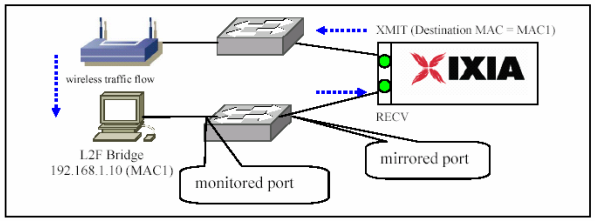

Figure 8. Layer-2 Test Configuration

4 Performance Results

The first performance test is to measure the one-way delay from the IXIA XMIT port to the RECV port. The data transmission rate is set at 1-1000 packet/sec to avoid network congestion. The wireless interface is 802.11b which has the max line rate of 11M bps. We created the data streams using raw IP, ICMP, UDP, and TCP traffic, and only the UDP data is presented in this paper. The data is collected for various packet sizes, from 64 bytes to 1500 bytes as recommended by the RFC 2544 testing standard. As illustrated in Figure 9, L2F and L3F have the same latency for various packet sizes. However, IP routing has longer latency (5-10% longer) than both L2F and L3F. This confirms our understanding that IP routing involves more overhead in packet and routing processing. It is also interesting to observe the linear relationship between the frame size and the latency which is consistent with our understanding of the *store-and-forward* scheme used in packet forwarding.

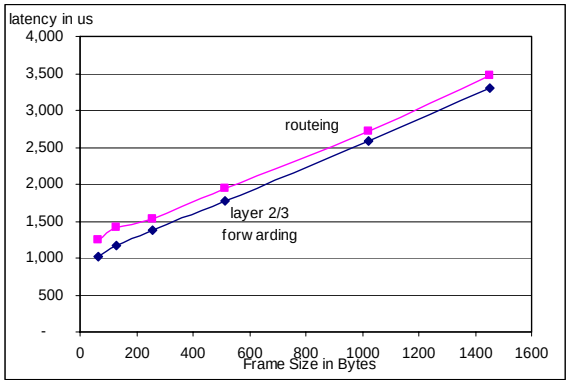

Figure 9. Delay of L2F, L3F, and Routing

The second test is to measure throughput (bps), and the data transmission rate is set at 1K-1M packets/sec. As illustrated in Figure 10, the test of L2F and L3F configurations shows a maximum throughput of 7.2 Mbps at the packet size of 1500 bytes, which is consistent with the published result in the literatures [11]. The throughput of IP routing, is 6.1M bps, which is 15% less than L2F and L3F. These performance results show that L3F has almost no additional overhead in comparison with L2F, which is the ideal case.

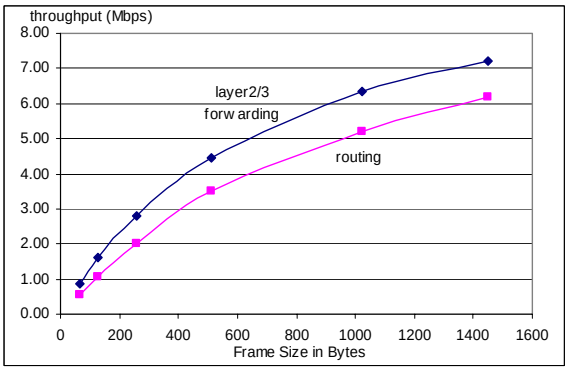

Figure 10. Throughput of L2F, L3F, and Routing

We also run a congestion test for the frame size of 1500 bytes with various input rates to find the *threshold* value of the network congestion point. When the network congestion occurs, the delay increases significantly, usually more than an order of magnitude. The results of congestion test are given in Table 2.

Table 2. Windows XP Bridge/Router Configuration Congestion Test (frame size =1500 bytes)

XMIT Rate (Mbps)	L2F Bridge		L3F Bridge		IP Router	
	Delay (ms)	TPT (Mbps)	Delay (ms)	TPT (Mbps)	Delay (ms)	TPT (Mbps)
0.12	3.30	0.11	3.30	0.11	3.47	0.11
1.16	3.25	1.15	3.30	1.15	3.32	1.15
5.80	3.37	5.80	3.43	5.79	**161**	**5.80**
6.96	**4.90**	**6.90**	**4.95**	**6.94**	3,000	5.82
7.19	98.8	7.15	97.6	7.14	4,000	6.0
7.25	98.7	7.17	98.2	7.17	4,300	6.12
7.31	102.6	7.22	101.7	7.17	6,000	6.12
7.54	103.9	7.15	102.9	7.15	7,000	6.1
8.12	103.0	7.22	101.3	7.20	7,340	6.11
11.60	100.9	7.19	102.0	7.21	9,345	6.14

XMIT Rate: Transmit Rate; **TPT**: Throughput;

Table 2 shows that both L2F and L3F configurations have similar performance. They reach the bottleneck at an input rate of 6.96 Mbps. At that point, their throughput is 6.90 Mbps, and the latency is 4.9ms. If input rate continues increasing, network congestion occurs and latency increases by 2-3 orders of magnitude where throughput stays the same. IP

routing, on the other hand, reaches bottleneck at a lower input rate of 5.8Mbps. If input rate continues increasing, the latency of IP routing increases by 3-4 orders of magnitude. In summary, our performance study shows that L3F configuration has comparable performance to L2F configuration, and it is significantly better than that of IP routing.

5. Conclusions

To the best of our knowledge, this is the first in-depth study of wireless bridging with layer-3 forwarding. We identified two practical applications of L3F. One application is in the home and SOHO environments to interconnect two LAN segments where there is no existing cabling infrastructure. The need in this environment is to connect a *small* number of workstations to share resources and the Internet connection. We showed that the L2F configuration does not work with the wireless access points and adapters, and wireless bridge is not cost-effective to connect a small number of workstations. The L3F configuration provides a cost-effective solution to address this issue. Another application is the multi-hop wireless ad hoc environment. Because the wireless adapter cannot work in both the bridge and ad hoc modes, L3F provides an effective solution to address the need of multi-hop configuration. We acknowledge that IP routing can perform the same function as L3F; however, we show that IP routing configuration is a lot more complex than L3F. IP routing requires manual configuration of the routing table both on the router as well as on individual workstations, which is a challenge to many users. On the other hands, L3F is transparent to end-users and requires no user configuration.

Another contribution of this paper is on the wireless experiment design using a high capability traffic generator and analyzer. The experimental framework allows us to generate various data streams and create multiple network scenarios. The results of performance analysis show that there is little overhead for L3F in comparison to L2F as measured by throughput, latency, and congestion threshold. On the other hand, L3F shows significantly better performance than IP routing as measured by throughput and latency.

Currently, L3F is available on the Windows XP environment only, and it is not available on the Linux environment yet. We are planning to develop this feature on the Linux environment and conduct more studies on it, such as implementing Spanning Tree Algorithm and Protocol (STP) to address the issue of loop topology in the network.

References

[1] "Wireless LAN MAC and PHY Specifications," IEEE 802.11-1999.

[2] "CSMA/CD Access Method and Physical Layer Specifications," IEEE 820.3-2000.

[3] B. A. Forouzan, *Local area Networks*, McGraw-Hill, 2003. pp. 337-339.

[4] Microsoft Corporation, "Home and Small Office Network Topologies" *http://www.microsoft.com/technet /prodtechnol/winxppro/plan/topology.mspx*, May, 2003.

[5] J. A. García-Macías, F. Rousseau, G. Berger-Sabbatel, L. Toumi, A. Duda "Quality of service and mobility for the wireless internet", *First ACM Wireless Mobile Internet Workshop '01 Rome, Italy*, 2001.

[6] "MAC Bridges," IEEE 802.1D-1998

[7] "Using ARP to Implement Transparent Subnet Gateways," RFC 1027, October 1987

[8] J. Davies, "How the Windows XP Network Bridge Works", http://www.microsoft.com/technet /community/columns/cableguy/cg0102.mspx, January 2002.

[9] J. T. Yu, "Performance Evaluation on Linux Bridge", *Telecommunications System Management Conference 2004*, Louisville, KT.

[10] S. Bradner, J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

[11] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," IEEE ISAC, vol. 18, no. 3, March 2000.