The Development and eStadium Testbeds for Research and Development of Wireless Services for Large-scale Sports Venues

Xuan Zhong, Hoi-Ho Chan, Timothy J. Rogers, Catherine P. Rosenberg and Edward J. Coyle
The Center for Wireless Systems and Applications (CWSA)
Purdue University, West Lafayette IN 47907-2035
{zhongx, hchan, tjrogers, cath, coyle}@purdue.edu

***Abstract*—We provide an overview of two closely related testbeds, a development testbed located in the CWSA wireless lab and the unique *eStadium* testbed located in Purdue's Ross-Ade Football Stadium. In the development testbed we have studied such high-bandwidth, delay-sensitive applications as on-demand streaming of video clips to a large number of wireless clients. This includes optimizing video streaming performance in wireless settings and characterizing the relationship between wireless LAN channel conditions and user-perceived quality. The lessons learned in this development testbed have been scaled up into the fully operational *eStadium* testbed in Purdue's Ross-Ade Stadium. This large-scale "Living Lab" enables measurements of football fans' use of and experiences with wireless devices to access infotainment content during Purdue football games. These measurements, some of which are summarized in this paper, have led to further research issues that are being addressed in the development lab. These include selective broadcast and multicasting of video clips to enable services to a very large number of wireless clients.**

***Index Terms*—Wireless LAN, Network Analysis, Video Delivery, Mobile Users, Location-dependent Admission Control**

I. INTRODUCTION

The 802.11b wireless networking technology known as Wi-Fi has gained great popularity and is now deployed in many public places, including football stadiums, basketball arenas, hotels, airports, and many private homes. A number of cities, including Philadelphia and San Francisco, are proposing to make 802.11 or other wireless technologies available to their citizens by placing wireless access points on light poles or other structures throughout their cities.

The increasing popularity of Wireless LANs makes 802.11 infrastructures good candidates for the development of multimedia applications that users can access via their PDAs, Wi-Fi enabled smart phones, and other wireless devices. Successful development of these applications is, however, a significant challenge because of 802.11's characteristics and the limitations of these small mobile devices. These include high bit-error rates, dynamic channel conditions, power constraints, limited storage capabilities, etc.

Many measurement studies have been performed on wireless LAN infrastructures to explore their channel characteristics and MAC layer performance. Wireless multimedia applications have also been an increasingly common focus of research as advances in hardware and software enable more capable handheld wireless devices.

Despite this increased attention, the effect of the wireless channel's characteristics on the performance of multimedia applications need further investigation in order to develop the best approach to the distribution of multimedia content over a time-varying channel. These further studies can facilitate the design of an advanced wireless video delivery system with technologies such as data rate adaptation, adaptive bandwidth allocation, and intelligent admission control.

To carry out these further studies, we have built two experimental testbeds. One enables us to study the relationship between the performance of video streaming applications and the characteristics of the wireless channel in a well-controlled environment. The other enables us to test the lessons learned in the controlled environment in a real venue with real users. This "real" testbed consists of the wireless infrastructure that has been installed throughout Ross-Ade Stadium at Purdue. The experiments performed in this testbed are measurements of the behavior of fans and the infotainment applications they access via wireless devices during games in the stadium. This testbed with real users in a real sports venue is known as the *eStadium Living Lab* [1].

The relationship between these two testbeds is bi-directional. We use the development testbed to solve problems encountered in the live testbed and deploy new applications perfected in the development testbed in the live testbed.

This paper is organized as follows. Section II provides background information and related work on wireless LAN characteristics and wireless video delivery; Section III describes our development testbed setup and experimental studies; Section IV presents the *eStadium Living Lab* application. Measurements from the living lab are analyzed and the wireless network's performance is evaluated. The paper concludes in Section V with a discussion of future work.

1-4244-0106-2/06/$20.00 ©2006 IEEE

II. BACKGROUND AND RELATED WORK

Characteristics of wireless local-area networks are studied via large-scale traffic measurements in [2-5]. Network usage patterns, mobility and application behavior are characterized, which is in turn important for designing and deploying wireless network systems and applications. A specific measurement study on wireless PDAs in a campus wireless LAN infrastructure is carried out in [6] and access and mobility patterns are characterized.

A review of the latest technologies for wireless video delivery is available in [7]. Technologies such as error concealment, packet scheduling, and joint source-channel coding are pointed out to be essential to achieve adaptability to varying bandwidth and other dynamic wireless channel conditions.

There are two main techniques for video distribution – unicast (one-to-one) and multicast (one-to-many). The access patterns of most information systems follow the rule that most requests are for a small portion of the data. A video server can take advantage of this property by waiting for requests and serving them together in one multicast stream instead of responding to each request with a unicast stream. However, this will cause service delays at the receiver's side. A technique called "patching" is introduced in [8] to leverage multicast technology for Video-on-Demand applications. Each client is able to patch into multiple streams on the same video that started at different times in order to obtain a complete playback of the video. A design methodology to ensure the optimality of the technique is also presented in [8].

Several wireless video applications have been developed recently. WiVision [9] uses wireless LANs to distribute real-time media to mobile users. The operational prototype is tailored for laptops as the client devices. In their setting, the video server pushes data to the wireless devices by broadcasting MPEG-1 video streams over UDP. They can only broadcast two video streams simultaneously with acceptable quality due to the bandwidth limitations of 802.11b.

Cao, et al performed empirical measurements of wireless media streaming traffic and performance in ad hoc IEEE802.11b networks in a classroom setting [10]. The results show that up to 8 clients can be supported with good media streaming quality, with each client receiving a separate 400 kbps video stream and 128 kbps audio stream. Performance degrades sharply when a ninth client is added because all nine clients must then share the same overloaded wireless medium.

There are also some ongoing commercial projects such as the NFL's wireless game plan, which delivers data, voice, and video applications over a wireless LAN infrastructure for coaches and fans in the stadium [11]. This particular wireless multimedia application focuses on videoconferencing between fans and football players as well as videoconferencing between football players, coaches, and other team staff.

III. DEVELOPMENT TESTBED

In this section we describe the development testbed used in our studies. The motivation for setting up the development testbed is to create a systematic experimental/research environment for studying the characteristics of wireless LANs and measuring the achievable performance of video streaming over 802.11b. This will in turn help us replicate and solve problems discovered in the *eStadium* environment during game days. It will also facilitate the development of new applications.

A. Testbed Hardware and Software Configuration

An experimental testbed was set up to study the interaction of video streaming and wireless LANs and understand the relationships between channel conditions and user-perceived video quality. The testbed layout is illustrated in Figure 1.

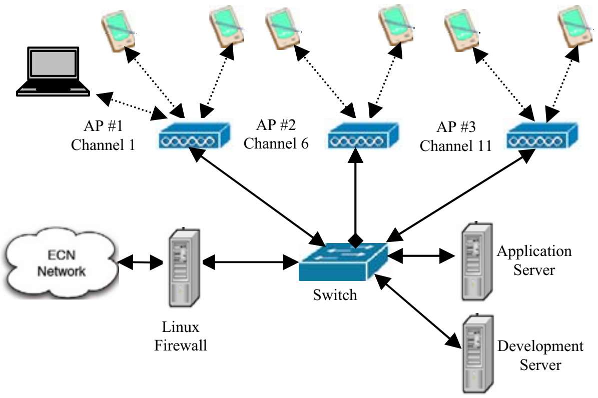

Fig. 1: Logical layout of the development testbed

There are three computers, two laptops, three access points and a number of PDAs in our testbed. We set up a Linux machine as a firewall. Everything behind the firewall (our servers and the access points) uses private IP addresses assigned by a DHCP server on this machine. We also set up a RADIUS authentication server on the Linux machine so each AP can have a centralized list of authorized devices. All DHCP messages and the APs' syslog messages are sent to the application server.

All three access points are Cisco Aironet 1200s with IOS software configured in infrastructure mode. They are assigned channel settings 1, 6, and 11, respectively.

The application server and development server are on a Gigabit Ethernet LAN. The application server is an Intel Xeon 2.8 GHz processor running Windows 2003 Server OS with 2GB RAM and two 160GB hard drives with RAID 1 configuration. It is configured with application software that includes SQL server 2000, Windows Media Service 9.1 and IIS 6.0. The development server is also running Windows 2003 Server OS with applications that include IIS 6.0, SQL Server 2000 and Visual Studio .NET 2003. The development server is used for ASP.NET application development. After the development has been completed, the application will be installed on the application server for the further tests.

The linux firewall is an Intel Pentium 2.4 GHz processor with 512MB of RAM that runs the Ubuntu OS with Linux kernel version 2.6.10. This computer is also configured to run an Apache web server, DHCP server, RADIUS Server and Syslog-ng server.

Network traffic measurements are collected via AiroPeek NX, an IEEE 802.11 network analyzer developed by WildPackets Inc. The 802.11 packet analyzer can be configured with different parameters and filters only packets of interest based on the design objective of the experiment. We also use Iperf [12], an open-source software package, to generate simulated UDP and TCP data traffic on the clients and servers.

To replicate the heterogeneous users in real football games, we use PDAs and wireless compact flash (CF) cards from a variety of venders. Our PDAs are Compaq IPAQs and Dell Axims. Wireless cards available include: DELL USI CF, D-Link Air DCF-660W, Symbol Spectrum 24, Compaq WL110, and Linksys WCF12.

All of the experiments described below were performed around midnight or during weekend hours to minimize the channel contention and noise caused by other users and other wireless networks. We also tried to carry out experiments at the same positions and maintain the same angles between the wireless devices and APs in order to ensure repeatability.

B. *Experimental Study on Video Streaming*

Windows Media Service 9.0 can stream media using the following application protocols: Real-Time Streaming Protocol (RTSP), Microsoft Media Server Protocol (MMS) and HTTP protocol. RTSP enables end users to fast-forward, rewind, and pause the playback of on-demand streams. RTSP works with the Real-Time Protocol (RTP) and Real-Time Control Protocol (RTCP) and uses either TCP or UDP for delivering streams. MMS, developed by Microsoft, can, like RTSP, support VCR-like functionalities at the client side for unicast video-on-demand streams. Protocol rollover [13] is the process that Windows Media Player uses to decide which protocol to use for receiving a stream based on the client's settings and network firewall conditions. When a user opens a URL specified with "*mms://*", the following protocols are used for data delivery in the following order:

1. RTSPU (RTSP using UDP)
2. RTSPT (RTSP using TCP)
3. MMSU (MMS using UDP)
4. MMST (MMS using TCP)
5. HTTP

The media player (WM Player 8.9.1.0) on many of our PDAs only supports the last three of these protocols for video streaming. We designed the experiments in our testbed to evaluate the effect of background traffic on user-perceived video quality.

A wireless sniffer is set up close to the PDA client to capture all packet traces to get a direct observation of the protocol being used for delivering video clips. We found that if there is a firewall between the PDA and the video server, then the network option on the PDA must be modified by unchecking the UDP option or we would experience video freezing and jerky playback. We also found that MMS streaming over TCP outperforms MMS streaming over UDP under poor network conditions, but it experiences longer delays during session initialization and results in more control overhead.

By varying the number of wireless clients, their mobility pattern, their distances to the access points, the level of background traffic load, we observed the user-level video playback performance. This enabled us to determine the maximum number of concurrent video streaming users that the wireless LAN can support. We found that when the network is overloaded or the link quality is poor, the clients experienced long delays to set up their initial connections to the stream server, and sometimes even dropped the wireless connection.

C. *Measurements and Analysis of Wireless LANs*

Our experimental studies of wireless LAN characteristics are helping us to better deploy complex applications and improve the quality of video distribution by taking advantage of channel characteristics.

Due to the 802.11 MAC's CSMA/CA scheme, the existence of multiple senders will markedly reduce the achievable throughput. There is no built-in quality of service (QoS) support in 802.11; all the clients have equal access to the medium. As the number of clients increases, there is a greater chance of frame collisions, backoffs, and retransmissions. This kind of contention results in increased latency at the client's end. All clients spend more time trying to gain access to the medium instead of transmitting and receiving data frames. This further induces transport layer timeouts and possible paused transmissions, or even dropped application sessions. For web browsing, email downloading, or some other client-server applications, a good rule of thumb is to set a limit of 25 users per AP [14]. But video streaming has a minimum required data arrival rate for continuous playback or the video stream will experience freezing or discontinuities when it is being starved. After extensive tests, we determined that the number of concurrent video streaming users per AP is limited to 10, accounting for some bursty data traffic from HTTP or email applications.

The power level settings for the APs in the stadium are additional parameters that should be tuned to maximize the available channel capacity. The power level at which an access point transmits and its receiver sensitivity determine its coverage area. When many APs are operating in the same area, there is a tradeoff between coverage and interference.

Determining the power settings that yield the best combination of coverage and quality of service is a significant challenge. In small areas with high user densities – which occurs in the stadium project when there are a lot of visitors within a small area during the football game trying to access the WLAN – we need to lower the power of the AP they are associated with and increase the minimum association rate to 5.5 Mbps. This will ensure that these users close to the AP will receive good service and keep the "bad apple" users from getting associated with the AP at a lower data rate.

IV. THE ESTADIUM "LIVING LAB" – ROSS-ADE STADIUM

In this section, we review the *eStadium* project and explain the system's design and implementation. Wireless traffic measurements and findings are also presented. The locations of the 802.11 access points in the stadium are shown in the figure at the end of the paper.

A. Brief Overview

Purdue's *eStadium* project is part of a campus-wide living laboratory for the study of wireless communications. Traditionally, scientific experimentation is carried out in a laboratory. However, to facilitate real-world, hands-on learning and experimentation by students, Purdue University has created the concept of a "living laboratory." The living laboratory uses the "city" of Purdue University as a unique space for experimentation while serving in its traditional role. *eStadium* is a collaborative partnership consisting of the Center for Wireless Systems & Applications (CWSA), Information Technology at Purdue (ITaP), and Purdue's Intercollegiate Athletics department. The focus of *eStadium* is the creation of a "living" laboratory at Purdue's Ross-Ade Football Stadium by (1) equipping it with an 802.11b based wireless network, and (2) designing a set of exciting applications for game day use by football fans and wireless researchers. A fully functional *eStadium* application system is available for trial at http://estadium.purdue.edu/estadium.

During the games, users are provided with wireless PDAs (users may also bring their own PDAs or smart phones) to enjoy applications such as video clips of instant replays, up-to-the-minute game statistics, player and coach biographies, and some other wireless "infotainment" services. In addition to providing interactive content to the fans, *eStadium* provides practical, real-world learning experiences for students involved in the project. Within the unique scope offered by the university environment, students can develop, deploy, and test next generation wireless systems, applications, and technologies.

B. Wireless Network Infrastructure

The Purdue Air Link (PAL) is an 802.11b-based virtual private network (VPN) that is supported by ITaP. PAL provides wireless access spanning more than 110 buildings on campus. The coverage is provided by over 1100 Cisco Aironet 1200 series APs. Ross-Ade Stadium is covered by 19 PAL APs. Users with a valid Purdue career account can logon to PAL to access the Internet. For security reasons, PAL uses IPSec and VPN technologies to encrypt the traffic between the wireless devices and the VPN gateways. Users wishing to use PAL must install special software on their device and logon with their Purdue Career Account before being granted access to the Internet.

The users who bring their personal PDAs to the stadium on game day may not have a PAL account, and may have customized hardware and software settings on their PDAs. Hence the objective is to make the *eStadium* applications accessible to all users regardless of the settings of their devices. We do not assume any control over the device settings of users who bring their own PDAs. It is also undesirable to require users to install proprietary software to access *eStadium* services. In order to support coexisting heterogeneous users without compromising any aspects of network security, it is necessary to modify both the wireless access network infrastructure, and the application settings at the server side. We opted for a two-part solution consisting of an infrastructure that provides wireless access to PAL clients in the stadium and a separate infrastructure that provides open access for *eStadium* patrons without a PAL account.

Wireless coverage in Ross-Ade Stadium is provided by Cisco Aironet 1200 AP based on IOS software which is capable of configuring multiple Service Set Identifiers (SSID). Each SSID runs on a different Virtual LAN (VLAN) on the same AP, allowing multiple logical networks to coexist in the same physical infrastructure.

Users with PAL accounts can get public IP addresses and can access the Internet. Users without PAL accounts can only get associations with the "estadium" SSID. They get class C private IP addresses from the eStadium DHCP server, and can only access the *eStadium* wireless network. The "estadium" network is an open network without any authentication or encryption. Thus, we are able to support both types of users without requiring any modifications to the client devices.

A logical layout of the *eStadium*'s wireless network is shown in Figure 2. Complete details of the set-up within the stadium can be found in the figure at the end of the paper. All of the APs are connected to different switches located on each floor in the stadium, which are then connected to the main router in the stadium. This router routes the traffic differently depending on which VLAN the packet is from. The packets from the PAL VLAN are routed directly to the VPN gateway. The packets from the *eStadium* VLAN are routed to the *eStadium* servers. The *eStadium* servers consist of a DHCP server, two identical content servers, a video server, and a set of load balancing computers to isolate the content servers from possible Denial-of-Service attacks as well as increasing the scalability of the content servers. The packets coming from the PAL VLAN are routed directly to the PAL VPN gateway.

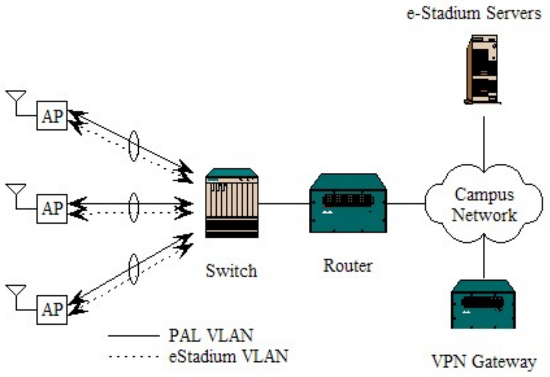

Fig. 2: *eStadium* network architecture

C. Implementation of the Video Distribution System

Windows Media Encoder 9.0 is used to convert the available analog composite video signals to WMV format video clips using the following parameters: data rate 176 kps, frame rate 15 frames/sec, spatial resolution 320×240, and maximum packet size 1400 bytes -- based on the tradeoff between bandwidth usage and user-perceived quality. The encoder saves the encoded files directly onto the video server by mapping the video server from the encoder machine. The association of the video files and the play stats are made by a human using an ASP.NET administrative web page, it will then automatically update the video links that point to the *eStadium* video streaming publishing point set up on the video server's Windows Media Service 9.1.

The architecture of the *eStadium* video distribution system is shown in Figure 3. The approach currently used for video distribution is unicast Video-on-Demand (VOD). This is not scalable because an access point can become overloaded with traffic as the number of associated users increases. When the channel is approaching its maximum capacity, video quality will degrade or the network connection will be dropped. Since all the clients associated with one AP share the same transmission medium, this kind of performance degradation is seen by all of the clients, not just one. Besides this, the wireless users can roam from one AP to another AP, which requires a location-dependent admission control mechanism in order to limit video traffic on a per-AP basis.

In the current implementation, a Location Discovery system (LODS) is used to find the AP for a given device. It uses syslog messages to locate a device - finding which AP it belongs to based on its IP address. The location detection process does not need additional devices like GPS or additional measurements like SNR to determine a wireless device's location. Instead it only uses the system log messages sent by the DHCP server and access points to locate a device. Syslog messages from APs provide the MAC-to-AP mapping. Syslog messages from DHCP server provide the MAC-to-IP mapping. These messages are sent to the video server via some UDP port and stored in two different tables in an SQL server database. We then join these two tables to get the IP-to-AP mapping in real time.

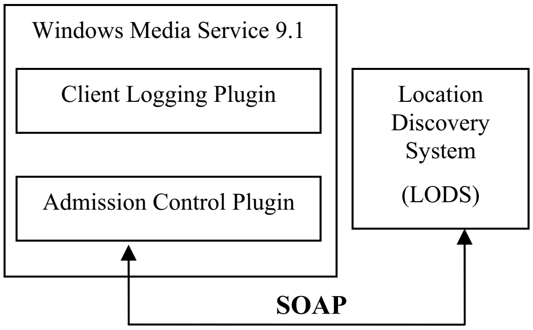

Fig. 3: *eStadium* video distribution system architecture

The admission control plugin is used to limit the video traffic on each AP. It is implemented as a custom plugin under Window Media Services 9.0 SDK. It runs as a Windows service on the video server and listens on a particular port and communicates with LODS using the Simple Object Access Protocol (SOAP).

We also enable the client logging plugin when we set up the publishing point for *eStadium* video clips on the video server. The client-receive logs indicate how the client received the content. We will show some results parsed from the client logs during the football games in the next section.

D. Data Collection and Results

All the applications developed for *eStadium* are written in ASP.NET using C#. Data is collected from applications running as Windows services and stored in the SQL server. This data is then further analyzed by programs written in Python and graphed in RRDtool [15].

We focus on data collected from the 2004 and 2005 home games. We used four techniques to trace client's WLAN usage and the video streaming performance: Syslog messages, SNMP polls, wireless sniffers, and client receive logs.

Syslog Messages: The *eStadium* APs are configured to send syslog messages to the LODS database on the video server whenever clients are authenticated, associated, re-associated, disassociated or de-authenticated. Each message consists of a timestamp, a message type and the MAC address of the station. Since the *eStadium* network uses open authentication, we do not have messages related to authentication. Association messages are generated when a client first associates with an AP. A client actively monitors the SNR from periodic beacons so it may choose a different AP to re-associate with. When a client no longer needs the network connection, it disassociates with the current AP. Actual syslog message types are dependent on wireless card venders, since some vendors never use re-association, always using association instead. We have not seen disassociation messages from Cisco APs; instead, we have seen de-authentication messages with the reason listed as "disassociated because sending station is leaving (or has left) BSS". Therefore, we treat these de-authentication messages as disassociation messages.

Session duration is defined to be the time between when a wireless device joins the wireless network and when it leaves the network. In order to calculate a session's duration, we need to parse the syslog messages from APs that contain the information about the association and disassociation of wireless devices. The APs send the messages with a timestamp accuracy of 1 second whenever an association or disassociation event happens, which means we can capture mobility information in a real-time fashion. Any associated messages that arrive less than 30 seconds after the previous association or re-association message are treated as re-associated messages instead of the creation of a new session. They thus indicate a roam. A roam happens when a station switches to a different access point from a previously associated one during a session.

The number of roams for the Wisconsin game, Michigan game, Ohio State game, and Indiana game in the 2004 football season were 1417, 4729, 3985 and 3335, respectively. Assuming that each game is 4 hours long, then at least one roam happened every 15 seconds on average. We later found out that the large number of roams that occurred in the Michigan game was probably due to devices jumping back and forth between APs and failing to latch onto a single AP.

Table 1 gives a snapshot of a "weird" session in which the device jumps back and forth between different APs in a short time. These jumps could be caused by overlaps of the coverage areas of different APs. After the Michigan game, we modified the channel settings for the APs so neighboring APs occupied different, non-overlapping frequency bands – channels 1, 6, and 11. We didn't see any "weird" sessions after that.

Time	Access Point
15:16:44.37	172.19.127.188
15:16:44.40	172.19.127.189
15:16:44.98	192.168.64.7
15:16:46.77	172.19.127.188
15:16:46.80	192.168.64.7
15:16:48.68	192.168.64.9

Table 1: A "weird" session showing the times at which a specific client device associates with different APs.

Figure 4 and Figure 5 show the empirical CDFs for the durations of sessions and the number of roams for each 2005 home football game, respectively. As can be seen from figures 4 and 5, user patterns are similar across games.

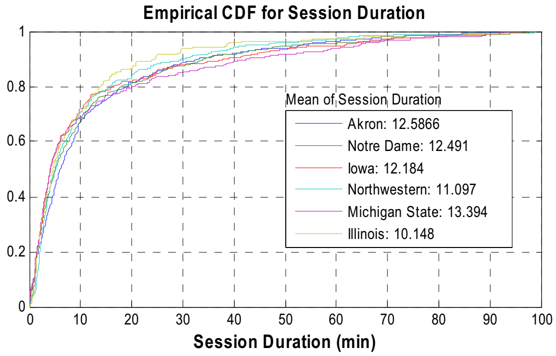

Fig. 4: Empirical CDF of session durations

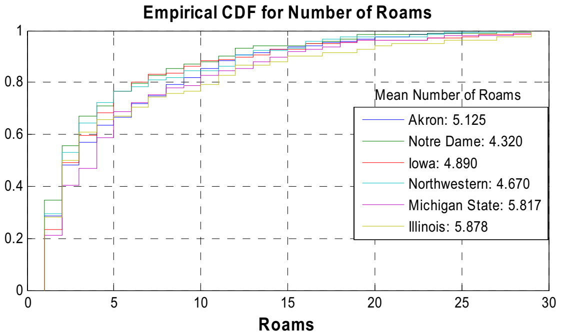

Fig. 5: Empirical CDF of roams

Figures 6 and 7 show the video request patterns from the year 2004. The request pattern in the Ohio State game is fairly consistent. In the Indiana game, the request rate dropped considerably after the second quarter, probably because Purdue was already ahead by so many points. It is interesting that users still wanted to look at video clips from previous quarters even in the 4th (last) quarter.

The number of users in two games stayed at a certain level and then dropped near the end of the game, although there were different video request patterns. This indicates that users use the devices for other applications when they are not viewing video instead of turning off the devices.

The study of usage patterns can guide us in refining video distribution solutions in a more efficient manner; e.g. multicast steaming with advanced video scheduling based on the popularity of the video clip. This will also help with scalability issues in the video distribution system if the number of users increases sharply in the future.

SNMP Polls: We use the Simple Network Management Protocol (SNMP) to poll each AP every minute, querying the AP's inbound/outbound traffic and errors. We use cacti [16], a front-end to RRDtool [17] that stores all necessary traffic information in a round-robin manner and creates and populate graphs with data from a MySQL database.

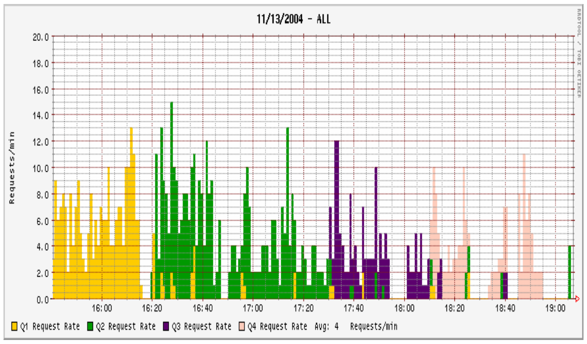

(a) Video Request Rate

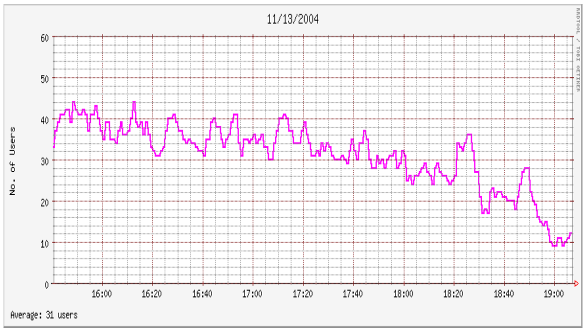

(b) Number of Wireless Users

Fig. 6: Video Request Rate; Number of Wireless Users (Ohio State Game, Nov. 13, 2004)

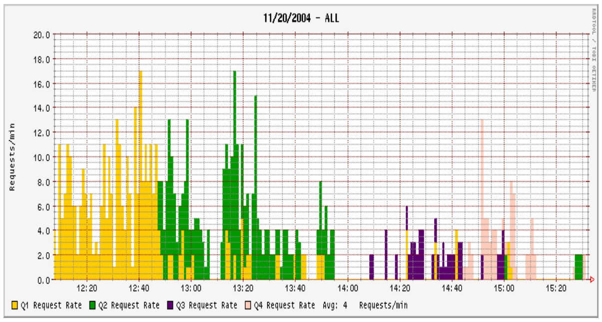

(a) Video Request Rate

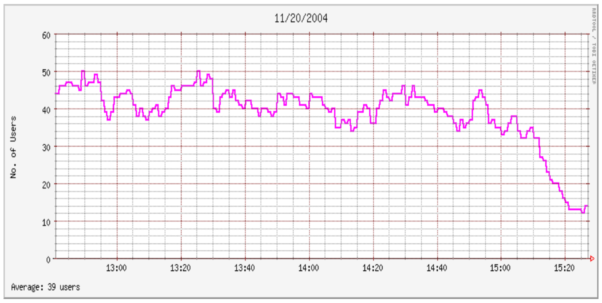

(b) Number of Wireless Users

Fig. 7: Video Request Rate and Number of Wireless Users (Indiana Game, Nov. 20, 2004)

Wireless Sniffers: Syslogs and SNMP provide basic statistics on traffic, wireless users, and mobility. Wireless sniffers allow us to get packet-level statistics such as protocol usage distribution, packet retransmissions, etc. By setting appropriate packet filters, we are able to record the activities of each PDA during an entire football game.

Client Receive Logs: Windows Media Services 9 series supports client receive logs [18]. The client receive log indicates how the client received the video files. In our database on the video server, we record every client's streaming entries with detailed client receive logs such as the average bandwidth at which the client is connected to the server, number of packets from the server that are received correctly on the first try, number of client requests to receive new packets using UDP resend, number of packets recovered after being resent through UDP, etc.

Figure 8 shows results on video streaming performance per AP for the 2005 Northwestern vs. Purdue game. The APs' physical locations on each floor shown at the end of the paper are mapped to this 3-D plot. The yellow bar shows the actual value for one game; the green bar shows the average value for all games. Note that the number of video requests in each game only reflects the requests with complete playback. We didn't count requests that were aborted during the video playback by the user or because of loss of the network connection.

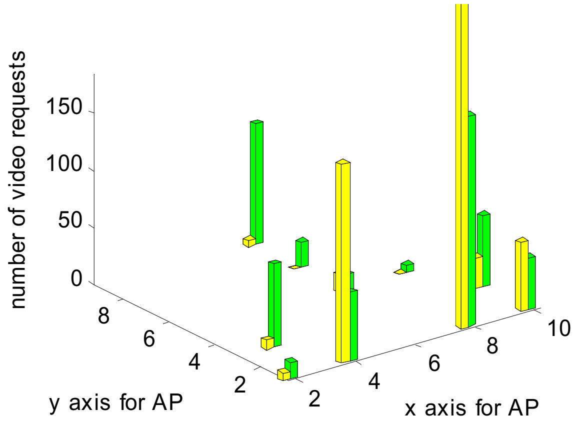

(a) Number of video requests

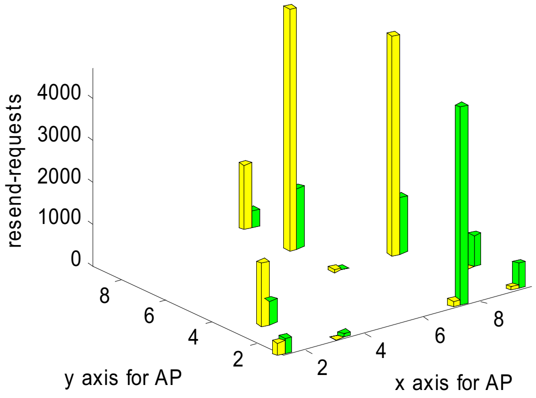

(b) Number of UDP resend-requests

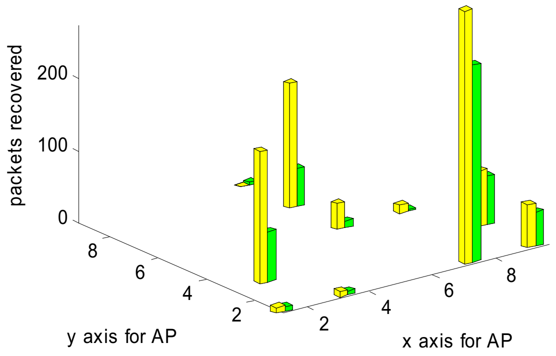

(c) Number of packets recovered via UDP resend-requests

Fig. 8: Video streaming results for the Northwestern game

V. CONCLUSIONS AND FUTURE WORK

Based on our experimental data and observations, it seems that the value of "UDP Resend Requests" is a good measure of client-perceived wireless network conditions. "UDP Resend Requests" is the number of times that clients ask the Windows Media Services server to resend data packets that were not received. This value is high when the server cannot reliably send packets via UDP. Thus, it is a good indicator of wireless network overload.

We are developing a more advanced admission control that is based on feedback about client-perceived network conditions. Information regarding these conditions, such as the packet loss rate and the actual video playback rate at the client, can be obtained from Windows media streaming client logs. This information can be further processed to implement an adaptive admission control scheme that reflects the dynamic channel conditions. This will enable us to make better use of the channel's capacity.

Since Cisco Aironet 1200 series APs support 802.11e, we would also like the admission control module to use client feedback to dynamically manipulate the 802.11e QoS parameters to improve the QoS for the clients associated with each AP.

One major problem in the current *eStadium* video distribution system is that it is not scalable. We propose two solutions. The first uses multicast video streaming. One challenge is that multicasting needs router support. Since the video server for *eStadium* is located across campus from the stadium, the video traffic must pass through Purdue's core router to get to the stadium. Enabling the core router to support IP multicast will cause many security concerns. Our solution is a "last-hop" scenario – to set up a proxy server in the stadium that supports unicast from the video server to the proxy server; the proxy server then multicasts the video to the wireless clients. Video multicast also requires the client's support, so an application needs to be installed at the client side.

Our second solution is to use a hybrid operation mode that allows clients to switch to the ad-hoc operation mode of 802.11 when the AP is approaching its channel capacity or a few mobile clients in physical proximity want to exchange video files with each other. In this case, it is favorable to let the clients form a transient ad-hoc network, perform the video file transfer, then switch back to infrastructure mode. Another benefit of hybrid operation is energy conservation because the power required to make a local connection is much less than that required to connect to the AP. To enable the hybrid mode, we need to implement a low level API that allows switching between infrastructure mode and ad-hoc mode.

Our next goal is to prepare the current *eStadium* system for IEEE 802.16 (WiMax) technology. WiMax can provide WiFi-like functionality with greatly improved range, reduced network latency, and greater throughput. This makes it an excellent candidate for video streaming applications. There are no WiMax products available yet, but they will be soon. For now, we plan to use OPNET (a network simulator) to do some simulations based on *eStadium* scenarios. We will use real *eStadium* network measurements to define the traffic profiles for the simulation. We'd also like to simulate new bandwidth allocation schemes in different traffic situations and do some performance analysis and optimization.

ACKNOWLEDGMENT

We acknowledge the support of all members of the eStadium team. Computer Graphics Technology: Prof. Ron Glotzbach, Mei-Qi Ren, Dorina Mordkovich. ITaP: Jim Bottum, Julie Kercher-Updike, Doug Magers, Mark Aiman, Brandon Case, Nathan Bohlmann, Blair Gillam. Purdue Intercollegiate Athletics: Morgan Burke, Jay Cooperider, Barb Kapp.

REFERENCES

[1] H. Chan, "eStadium: Wireless Applications and traffic measurement in a wireless network supporting a set of heterogeneous users," Master Thesis, School of Elect. and Comp. Eng., Purdue Univ., West lafayette, IN, May 2005.

[2] D. Kotz and K. Essien, "Analysis of a campus-wide wireless Network," *In Proceedings of ACM MobiCom*, 2002.

[3] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," *In Proceedings of ACM MobiCom*, 2004.

[4] M. Balazinska and P. Castro, "Chacterizing mobility and network usage in a corporate wireless local-area network," *In Proceedings of ACM MobiSys*, 2003.

[5] F. Chinchilla, M. Lindsey and M. Papadopouli, "Analysis of wireless information locality and association patterns in a campus," *In Proceedings of IEEE INFOCOM*, 2004.

[6] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol 9, no. 2, pp. 40-55, Apr. 2005.

[7] M. Etoh and T. Yoshimura, "Advances in wireless video delivery," *In Proceedings of the IEEE*, vol. 93, no. 1, pp. 111-122, Jan. 2005.

[8] Y. Cai and K. A. Hua, "Sharing multicast videos using patching streams," *Journal of Multimedia Tools and Applications*, Vol. 21. No. 2, pp. 125-146, 2003.

[9] Pradipta De, Srikant Sharma, Andrew Shuvalov and Tzi cker Chiueh, "WiVision: A Wireless Video System for Real-Time Distribution and On-Demand Playback," *In Proceedings of the First IEEE International Conference on Consumer*

Communications and networking Conference, pp. 575–580, January 2004.

[10] X. Cao, G. Bai and C. Williamson, "Media streaming performance in a portable wireless classroom network," *In Proceedings of IASTED European Workshop on Internet Multimedia Systems and Applications (EuroIMSA)*, Grindelwald, Switzerland, pp. 246-252, Feb. 2005

[11] C. Metz. [on line]. Available: http://www.pcmag.com/article2/0,1895,1843232,00.asp

[12] [on line]. Available: http://dast.nlanr.net/Projects/Iperf/

[13] [on line]. Available: http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wmform95/htm/protocolrollover.asp

[14] P. Roshan and J. Leary, 802.11 wireless LAN fundamentals, Cisco Press, 2004.

[15] [on line]. Available: http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/

[16] [on line]. Available: http://cacti.net/.

[17] [on line]. Available: http://www.rrdtool.org/.

[18] [on line]. Available: http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wmsrvsdk/htm/clientreceivelog.asp