

Research Resources for Network Application Studies

Award Number: CNS-0423362

Mark Claypool, Robert Kinicki, Craig Wills

CS Department, Worcester Polytechnic Institute, Worcester, MA, 01609, USA

URL: <http://perform.wpi.edu/nas/>

Abstract

The growth of computer networks has led to increasing diversity of Internet applications, including streaming media and network games. However, without precise information on how network and system improvements benefit the networked application user, it is difficult to properly assess the benefits of new network treatments or to design the next generation networks that will effectively support the QoS of emerging applications. This research attempts to bridge this gap in understanding with three innovative projects: 1) integrating measures of network performance with user perception; 2) quality of service for network games; and 3) perceived quality of adaptive streaming media repair. With the requested research resources, we have developed an application performance studies laboratory that allows us to finely control network performance for a range of selected networked applications. Each project shares research resources in the new laboratory to measure performance for interactive applications, network games and streaming media repair, as appropriate.

1 Introduction

The growth in the popularity and capability of the Internet has led to an increasingly diverse set of Internet applications with varying network behaviors and requirements. In order to design the next-generation networks and protocols to support these applications, the Quality of Service requirements of these applications must be studied, particularly along the key metrics of latency, loss and throughput. Traditional applications such as file transfer, Usenet news and email are primarily concerned with throughput and can tolerate delays on the order of minutes. Web browsers are also concerned with throughput, but the interactive nature of browsing requires latencies on the order of seconds or at most tens of seconds. Emerging real-time applications such as IP telephony and networked games typically have the lowest throughput requirements but are even less tolerant of latency than other applications. Adaptive repair techniques promise to mitigate the effects of data loss but often at the cost of added latency. Knowing how these real-time applications react to latency and loss is the crucial first step in designing the next generation network hardware and software that will support their requirements. In addition, classifications of real-time applications according to their latency tolerance will enable designers, developers and engineers to make informed decisions on appropriate quality for classes under such architectures as DiffServ. An overview is provided in Figure 1.

We have used the research resources to conduct system and user studies for emerging applications, forming the core of three innovative projects:

1. *Prediction of Application Quality*. This project is a started bridging the gap between user perception of the application and network performance measures through cross-layer analysis (see Section 3.1).
2. *Quality of Service for Network Games*. This project studied the effect of network latency, frame rate, and frame resolution on user performance and network traffic for Internet games (see Section 3.2).
3. *Perceived Quality of Adaptive Streaming Media Repair*. This project incorporated an analytic model of adaptive repair into a streaming video system and evaluate the impact on user perceived quality and network traffic (see Section 3.3).




Figure 1: Today's applications, such as streaming video, Web browsing and network games (depicted on the left), use increasingly complex network configurations (depicted on the right). The impact of new network treatments on users of these applications is unclear, demanding new studies and methodologies for assessing network impact.

The requested resources were used to construct an environment, usable by all three projects, in which we can systematically control key system parameters while still providing the user with a realistic environment appropriate for the applications being studied. In addition, the resources allow gathering of both system and application level performance measurements so we can correlate them with user perceptual quality.

2 Resources

In order to conduct effective user studies and accurate system level measurements of application and network performance, we established dedicated high-end equipment that allowed us to focus on the effects of network treatments on Internet application performance. We also acquired network hardware to support the computer equipment and robust, realistic client-server Internet applications to run. Below, we provide a brief overview of the research resources.

Application Servers - high-performance computers for hosting streaming media servers, Web servers, and network game servers. The servers are equipped with sufficient RAM and disk space to have capacity for the most resource intensive server applications. Each server is equipped with multiprocessors to better allow multi-user tests, stressing the network without having degradation in server processing.

Multimedia Workstations - high-end multimedia workstations for effective streaming media and game play. While advances in video technology or enabling streaming over restricted environments, streaming video is still resource hungry, requiring fast processors, and significant amounts of memory. Computer games require at least as many resources as does streaming media, with the added requirement of a high-end graphics card. Both streaming media and computer games benefit from large, clear displays and effective sound systems.

Game Consoles - game consoles with high quality displays for effective game play. Game consoles, which sell more games than do PCs, have traditionally supported only multiplayer games with split-screens, but are increasingly supporting multiplayer networked games. The differences in console interfaces and end-host operating systems demands the study of console games in addition to PC games. This includes hand-held game consoles, as well.

Network Equipment - switches, hubs, cables and PC "routers" to control network treatments. Basic networking hardware (switches, hubs and cables) connects the servers and clients. Most importantly, PC routers running NIST Net¹ are used to emulate network conditions over a WAN, allowing fine-grained control tuning of latency

¹<http://snad.ncsl.nist.gov/itg/nistnet/>

and loss (and variation of each) with capacity restrictions.

Software - commercial software for both the clients and servers. While basic client-server software can sometimes be built for network emulation, the latest commercial software is needed to provide robust, realistic streaming media and computer game environments.

3 Highlights

The award has supported research that has thus far resulted in 1 journal publication [1] and 12 peer-reviewed conference publications [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13]. This section highlights some of the specific research contributions. More details on the research accomplishments can be found in the individual papers available at <http://perform.wpi.edu/nas/>.

3.1 Streaming Video

The research resources supported experiments that streamed commercial videos over a wireless campus network, while capturing video quality wireless LAN characteristics across network and wireless link layers [10, 8]. A forecasting analogy was used to emphasize the importance of predicting streaming video quality (the weather) given measurements of current network conditions. For a specific quality metric (the weather *prediction*), the forecasting goal is to find a measurable network parameter (a weather *predictor*) that accurately estimates streaming video quality.





Figure 2: Frame Rate Prediction by Average Wireless Capacity Figure 3: Frame Rate Prediction by IP Packet Loss Rate

Average wireless capacity is the predictor for the weather map in Figure 2. The x axis is the predictor and the y axis is the probability. The thin, horizontal illustration at the top of each figure shows a histogram of the data samples. The lack of significant vertical overlap in the map suggests average wireless capacity is an effective predictor of frame rate. In the sampling range from 0 to 54 Mbps, an average wireless capacity greater than 5 Mbps forecasts a high likelihood of good video performance, while a capacity greater than 18 Mbps always forecasts good video performance. IP packet loss rate is the predictor for the weather map in Figure 3. Unlike average wireless capacity, the IP packet loss rate is not effective for forecasting video frame rate. Only when the loss rate is under 2% or over 16% is a single forecast likely.

Overall, after examining six different predictors across network layers, the results show that the wireless received signal strength indicator and average wireless link capacity are the most accurate indicators to predict the performance of streaming video over wireless LANs. The weather forecast philosophy can be beneficial for adapting video streaming in wireless LAN environments.

3.2 Network Games

The research resources supported user studies on the effects of system performance for network games [1, 2]. The effects of latency are determined by a classification based on the player perspective (*first-person* or *third-*




Figure 4: Player Performance versus Network Latency for Several Classes of Network Games

person) and on the model (avatar or omnipresent). Practically, all games fall into one of three categories: first-person avatar (ex: first-person shooter, racing), third-person avatar (role-playing, side-scrolling) or omnipresent (real-time strategy). A meta analysis of previous work that measured the effects of latency and network games, normalized and combined, provides a way to quantify the effects of latency on network games.

Figure 4 summarizes the performance versus latency for the different classes of network games, depicted by an exponential curve fit to the previously measured data. The horizontal gray area around 0.75 in Figure 4 is a visual indicator of typical player tolerances for latency. The exact threshold depends upon the game (and to some extent, the player), but generally game performance above this threshold is acceptable while game performance below this threshold is unacceptable. Overall, games that use the avatar model of player interaction are more sensitive to latency than games that use the omnipresent model, and games that use the first-person perspective are more sensitive to latency than games that use the third-person perspective.

3.3 Media Repair

The research resources supported development of a novel video quality metric called *distorted playable frame rate* that provides estimation of the perceptual quality of video considering temporal and quality degradations. The distorted playable frame rate (R_D) provides an estimate of streaming video quality when media repair and media scaling are employed. The work included a comprehensive user study with of over 70 users that provided scores for comparing with R_D for a variety of videos.

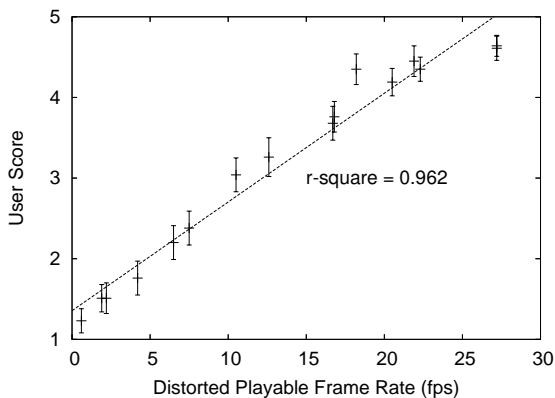


Figure 5: User Score versus Distorted Playable Frame Rate

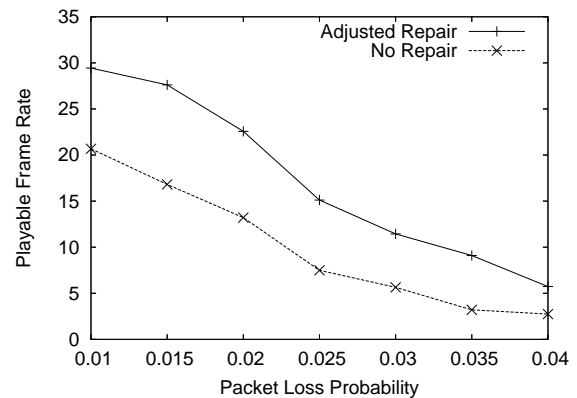


Figure 6: Effectiveness of Adaptive Media Repair

Figure 5 correlates distorted playable frame rate (R_D) and the perceptual quality expressed by the users' scores. Each data point in the figure represents the comparison of the original and degraded video clips. The

x axis is the distorted playable frame rate R_D and the y axis is the mean rating score for all users from “Much worse” (1) to “Same” (5), shown with a 95% confidence interval. Visually, the relationship between R_D and user score is almost linear.

Using this new measure of quality, an adaptive repair system [11, 13] was developed that dynamically adjusts the amount of repair data based on the loss rate, round-trip time, and video content to achieve the best quality. Figure 6 depicts the playable frame rates for adjusted repair and no repair for a video with modest visual motion. The x axis is the packet loss probabilities, and the y axis is the playable frame rates. From the data in these figures, adjusted repair provides better video quality over the range of network loss rates.

References

- [1] K. Claypool and M. Claypool. On Frame Rate and Player Performance in First Person Shooter Games. *ACM/Springer Multimedia Systems Journal (MMSJ)*, 2007.
- [2] K. Claypool and M. Claypool. The Effects of Resolution on Users Playing First Person Shooter Games. In *Proceedings ACM/SPIE Multimedia Computing and Networking (MMCN) Conference*, San Jose, CA, USA, Jan. 2007.
- [3] M. Claypool. On the 802.11 Turbulence of Nintendo DS and Sony PSP Hand-held Network Games. In *Proceedings of the 4th ACM Network and System Support for Games (NetGames)*, Hawthorne, NY, USA, Oct. 2005.
- [4] M. Claypool, K. Claypool, and F. Damaa. The Effects of Frame Rate and Resolution on Users Playing First Person Shooter Games. In *Proceedings ACM/SPIE Multimedia Computing and Networking (MMCN) Conference*, San Jose, CA, USA, Jan. 2006.
- [5] M. Claypool, R. Kinicki, W. Lee, M. Li, and G. Ratner. Characterization by Measurement of a CDMA 1x EVDO Network. In *Proceedings of the Wireless Internet Conference (WICON)*, Boston, MA, USA, Aug. 2006.
- [6] S. Gargolinski, C. S. Pierre, and M. Claypool. Game Server Selection for Multiple Players. In *Proceedings of the 4th ACM Network and System Support for Games (NetGames)*, Hawthorne, NY, USA, Oct. 2005.
- [7] J. Gretarsson, F. Li, M. Li, A. Samant, H. Wu, M. Claypool, and R. Kinicki. Performance Analysis of the Intertwined Effects between Network Layers for 802.11g Transmissions. In *Proceedings of the 1st ACM Workshop on Wireless Multimedia Networking and Performance Modeling (WMuNeP)*, Montreal, Canada, Oct. 2005.
- [8] F. Li, J. Chung, M. Li, H. Wu, M. Claypool, and R. Kinicki. Application, Network and Link Layer Measurements of Streaming Video over a Wireless Campus Network. In *Proceedings of the 6th Passive and Active Measurement Workshop (PAM)*, Boston, Massachusetts, USA, Apr. 2005.
- [9] F. Li, M. Li, R. Lu, H. Wu, M. Claypool, and R. Kinicki. Tools and Techniques for Measurement of IEEE 802.11 Wireless Networks. In *Proceedings of the Second International Workshop on Wireless Network Measurement (WiN-Mee)*, Boston, MA, USA, Apr. 2006.
- [10] M. Li, F. Li, M. Claypool, and R. Kinicki. Weather Forecasting - Predicting Performance for Streaming Video over Wireless LANs. In *International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Skamania, Washington, USA, June 2005.
- [11] H. Wu, M. Claypool, and R. Kinicki. Adjusting Forward Error Correction with Quality Scaling for Streaming MPEG. In *Proceedings of Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, Stevenson, Washington, USA, June 2005.
- [12] H. Wu, M. Claypool, and R. Kinicki. Guidelines for Selecting Practical MPEG Group of Pictures. In *Proceedings of IASTED International Conference on Internet and Multimedia Systems and Applications (EuroIMSA)*, Innsbruck, Austria, Feb. 2006.
- [13] H. Wu, M. Claypool, and R. Kinicki. On Combining Temporal Scaling and Quality Scaling for Streaming MPEG. In *Proceedings of Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, Providence, RI, USA, June 2006.