

Topic Review: Multimedia over Wireless

Review: An Experimental Study of Multimedia Traffic Performance in Mesh Networks

This paper was published in 2005 in the proceedings of the International Workshop on Wireless Traffic Measurements and Modeling (WitMeMo), collocated with MobiSys. This workshop is “intended to serve as a forum for scientists and engineers in academia and industry to exchange and discuss their experiences and research results ... that report on experiences obtained from operational wireless experiments in testbeds or the field.” This paper is well aligned with the goals of the WitMeMo workshop, and though it does not offer new methodologies or solutions to address the limitations of wireless multimedia networking, it does present experimental data that reveal new insights into the issues encountered transmitting multimedia streams over wireless mesh networks in various configurations.

Specifically, this paper focused examining the performance of bursty video streams and constant bit rate voice streams on an ad hoc wireless network. The paper primarily looked at the effect of network path length, packet send and data rate, and specific protocol settings on the performance of voice and video streams. Performance is evaluated across the following metrics:

- Packet latency
- Packet loss rate
- Inter-flow fairness
- Packet jitter

These metrics are useful in understanding the raw performance of network itself; however, they are not as directly useful in determining the quality of media stream and its interaction with the particular coding scheme being used. Considering that the evaluation was focused on the performance of the network when transmitting multimedia data, it would have been useful if the study had included other metrics (such as SNR) that more directly indicated the quality of video and audio as it would be perceived by the user.

That being said, the evaluation did yield interesting and useful results. I found several of the results particularly surprising. First, I was surprised to see the significant performance degradation as hop count increased. This indicates the distribution area for multimedia data over mesh networks is fairly limited. The paper, in general, is silent regarding suggested solutions to the problems it uncovers, and here is an area where I believe some guidance could have increased the paper’s contribution. For example, what protocol enhancements could be made knowing that hop count is such a significantly limiting factor?

I also did not expect the packet sending rate to have such a large impact on the number of supported data flows. In particular I found it surprising that it a much larger effect than data rate, and that nearly

the same number of video flows can be supported as voice flows. The authors offer some explanation of why this might be the case, essentially that the higher packet send rate leads to more network congestion; however, I found this reasoning to be somewhat unconvincing, or at least unsupported. It is not clear why a higher packet transmission rate would more adversely affect network congestion than a higher amount of total bytes being transmitted. A more detailed explanation of this phenomenon would have been helpful. Despite this weakness, this observation certainly is a useful one for protocol developers and can certainly lead to optimizations that minimize packet sending rate. This is a significant contribution of this work, I believe.

Another valuable observation made in the paper is that data flows are not always treated fairly. In particular, latency and packet loss rate increases, up to 300%, for video and voice streams that are started later. Again, the paper offers little explanation for why this might be the case. Nevertheless, it is a worthwhile observation that points to some obvious deficiencies in the networking protocols being used.

In general, the paper does provide useful and insightful observations. However, it misses some opportunities and does not accomplish all it sets out to do. For example, the introduction states that the real-time data delivery requirements of multimedia applications “necessitate that both the application and the network be able to adapt to the highly variable nature of wireless channels.” However, the evaluation occurs on relatively static network topologies and configurations. I expected from this statement to see an evaluation of the effects of changing network conditions on the performance of multimedia data transmission, but none was given.

Also, one of the motivations given for conducting this experimental study is that simulations often fall short and “may not reflect the performance obtained in real networks.” To help support this claim, the authors could have included a comparison of the actual observed results on the physical mesh network with results obtained from running a simulation. This could have served two purposes: 1) it could have further justified the need for an empirical study; and 2) it could have revealed the specific areas in which the simulations fall short, potentially leading to improvements in the simulation tools or commonly used models.

Finally, the authors claim that this study is “beneficial in both wireless network capacity planning and protocol design.” On the first point regarding network capacity, the paper does a fairly good job of clearly demonstrating capacity limitations and limiting factors. However, while several of the observations allude to potential areas to investigate to improve protocol performance, the authors offer little in the way of concrete suggestions, other than stating that RTS/CTS should not be enabled. I felt more could have been done to offer guidance for protocol optimizations or further research in protocol design.

In summary, I believe this paper was appropriate for the venue in which it was published. It offers non-trivial contributions that have the potential to improve the performance of mesh networks and influence the future design of protocols for these networks. However, the paper could have done more in explaining and expanding upon the findings it presents.