Measurements of an Internet-Enabled Camera

Todd Hansen[*] NLANR/MNA

Hans-Werner Braun HPWREN and NLANR/MNA

October 16, 2001

1 Introduction

Since the introduction of Internet-enabled devices such as video cameras, phones and PDAs, the bandwidth requirements of each device continue to become more varied. Even devices that provide the same types of functionality tend to use different amounts of bandwidth. For example, cameras that use motion-JPEG compression tend to consume more bandwidth than cameras which use MPEG compression. These are implementation choices made by vendors, but ultimately they have an impact on the network traffic created as well as the overall user experience. An example of this dichotomy is that motion-JPEG may consume more bandwidth, but it is easily viewed through a modern Web browser. MPEG encoding, while lower bandwidth, generally requires the installation of a separate software/hardware package, sometimes at a cost to the user.

In this paper, we will cover the network characteristics of one Internet-enabled device in an attempt to better understand its impact on the network. We will not be looking at how the user experience is impacted by this particular device or how this device compares to others. Those are topics for future work or in some cases, topics better covered by general audience publications.

2 The Camera

As part of the High Performance Wireless Research and Education Network (HPWREN) [1] objectives, we instrumented our network with devices to conduct measurements [3]; one of these devices is an Internet-enabled camera (see Figure 1). The camera is a Spectra II from Pelco [4] with pan-tilt-zoom capabilities; and we used an Axis 2401 Video Server [5]. The video server takes the video from the camera and converts it to motion-JPEG which is viewable in a Web browser. It then takes user commands from the Web browser and converts them to RS-485 commands which are then sent to the camera. For these tests we set the camera at one specific location and locked out all other users. One important aspect of the motion-JPEG format is that it sends the entire frame each time, whereas MPEG sends only changes from the previous frame with an occasional full frame. A more thorough description of the camera construction is available on the HPWREN Web site [6].

After leaving the video server, the video traf-

[*]National Laboratory for Applied Network Research/Measurement and Network Analysis Group, University of California, San Diego, San Diego Supercomputer Center, 9500 Gilman Dr., La Jolla, CA 92093-0505, E-mail: tshansen@nlanr.net

Figure 1: The camera located at Mt. Woodson

fic is sent over a small network as shown in Figure 2. The video server is connected to a Cisco switch (running VLAN switching), then connected to a Cisco 3660 router on Mt. Woodson. This router sends the traffic over a 45 Mbps full-duplex Tsunami radio link to the San Diego Supercomputer Center (SDSC), where the data is sent through a set of VLAN switches to a Cisco 3640 router. At that point, it goes through another router called 'medusa' which then sends it over another in-building VLAN segment to the measurement computer. Each set of tests was conducted from the measurement computer as `tcpdump` was run. No packets were lost by tcpdump and the measurement computer was never overloaded by the data. The slowest link was the 45 Mbps full-duplex link from Mt. Woodson to SDSC.

The maximum bandwidth consumed by the camera was 9.972 Mbps (sd $\pm$81.4 kbps, sample size 3) with 10 simultaneous users using 'huge size' images with minimal compression. During this test each user saw 0.997 Mbps (sd $\pm$59.9 kbps, sample size 30). Each additional user further reduced the per user bandwidth while increasing the aggregate bandwidth at a slower rate. There is a strong indication that by having more than 10 simultaneous users will result in higher aggregate bandwidth consumption; however we did not test larger numbers of users. We expect that the user experience would degrade significantly with each additional user as individual bandwidth use decreases.

On a per user basis, the highest individual user throughput was 2.30 Mbps which occurred only once and was most likely an error because it was not repeatable. This occurred during a test with two other simultaneous users.

The idle round trip time during the tests was relatively low at 2.257 ms (sd $\pm$0.189, sample size 60). Each of the tests involved one or more users transferring streaming motion-JPEG images for 30 seconds. The output was written to `/dev/null` to avoid disk latency issues. HTTP over TCP was used as the transport protocol for the video. Therefore the motion-JPEG was sent over TCP as one large transfer for 30 seconds per user. All of the tests were run on June 18th 2001 from 12 pm to 7 pm. No known anomalies occurred during this time.

The raw measurements and initial analysis for this paper are available on the HPWREN Web site [9].

3 Parameters

In this section we will discuss the four parameters that are the most critical with respect to the camera's impact on the network. These parameters are: Image Size, Degree of Image Compression, Number of Users, and Rate Limiting.

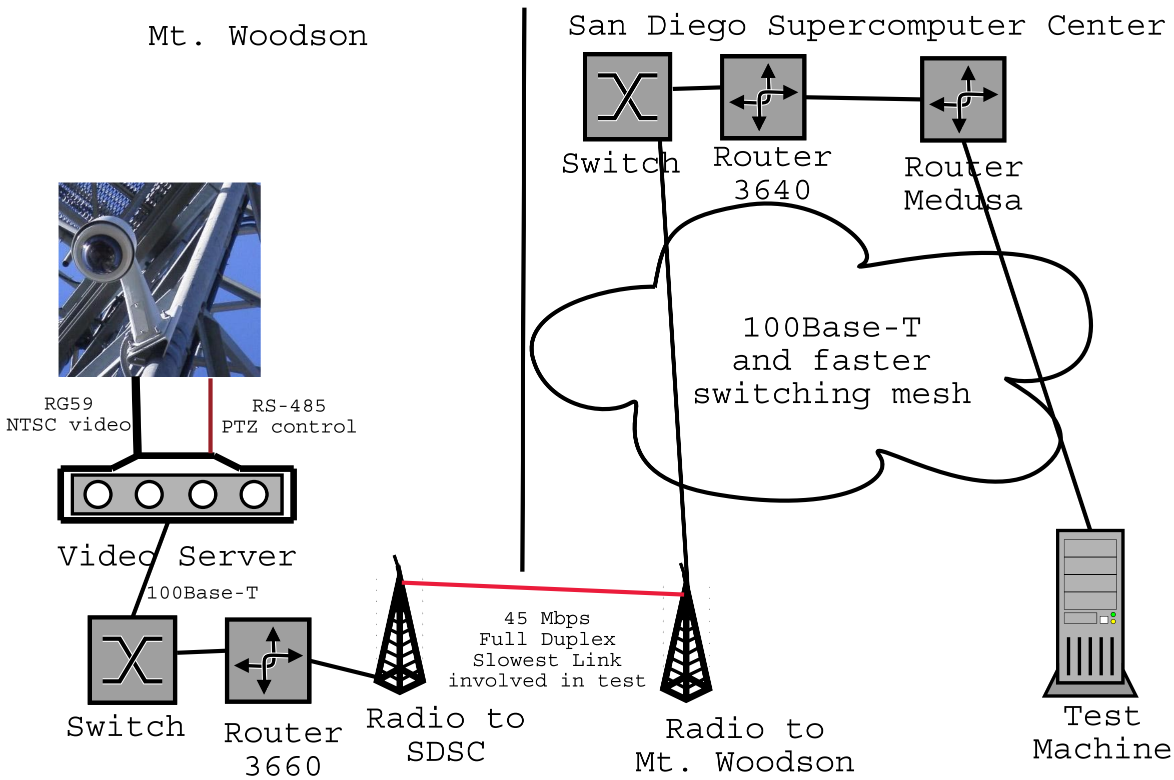

Figure 2: Map of test network

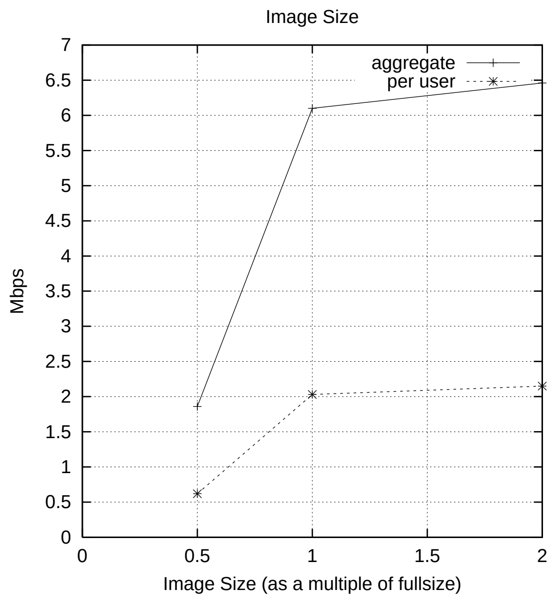

Figure 3: How image size affects network load (in Mbps)

Image Size	Throughput	Std. Dev.
0.5	1.86 Mbps	±53 kbps
1.0	6.10 Mbps	±68 kbps
2.0	6.46 Mbps	±197 kbps

Table 1: Aggregate throughput vs. image size (sample size 5)

Image Size	Throughput	Std. Dev.
0.5	0.62 Mbps	±97 kbps
1.0	2.03 Mbps	±21 kbps
2.0	2.15 Mbps	±86 kbps

Table 2: Per user throughput vs. image size (sample size 15)

3.1 Image Size

Image size controls the number of pixels in the image sent from the video server to the Web browser. The camera always sends the same NTSC format video to the video server. This parameter mainly effects the encoder and the data path to the Ethernet interface. Figure 3 shows the effect of image size on the bandwidth consumed. There are three image sizes in the camera which we approximate to 0.5, 1.0, and 2.0 times full size. These image sizes are: ‘half size’ (176x112), ‘full size’ (352x240), and ‘huge size’ (704x480) respectively. It is possible to set an image scaling factor which takes a smaller image and multiplies it by 2 or 4. This factor was not adjusted. These tests were run with three users at minimal compression.

As can be seen in Figure 3 and Tables 1 and 2 the ’half size’ (0.5) image drastically reduces the network traffic relative to ’full size.’ This may be due to the fact that the encoder is only capable of interpreting a fixed number of images per second from the analog video signal. Therefore, the video server may have been operating at its maximum frame rate (we were not pushing any other limits). From ’full size’ to ’huge size’ images a far less significant load increase was observed. However, when in ’huge size’ we saw a much lower frame rate than in ’full size’. Additionally, the video server’s manual states that the maximum frame rate in ’huge size’ is 10 frames per second (fps) whereas in both ’full size’ and ’half size’ modes it is 30 fps [7, page 45]. We suspect that there may be an internal limit to the bandwidth of the images from the encoder to a theorized on-server image cache. In theory, this cache could be used to serve up images to multiple users in order to reduce the number of images that the encoder must produce. Further

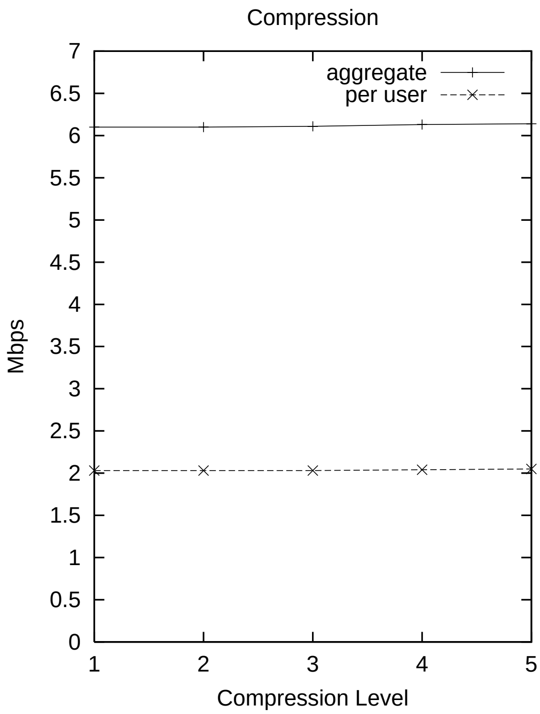

Figure 4: How compression affects network load

Compression	Throughput	Std. Dev.
1	6.10 Mbps	±68 kbps
2	6.10 Mbps	±19 kbps
3	6.11 Mbps	±5.6 kbps
4	6.13 Mbps	±12 kbps
5	6.14 Mbps	±18 kbps

Table 3: Aggregate throughput vs. level of compression (sample size 5)

Compression	Throughput	Std. Dev.
1	2.03 Mbps	±21 kbps
2	2.03 Mbps	±7.7 kbps
3	2.03 Mbps	±3.1 kbps
4	2.04 Mbps	±5.1 kbps
5	2.05 Mbps	±6.8 kbps

Table 4: Per user throughput vs. level of compression (sample size 15)

supporting this theory is the fact that in rare instances an image from half an hour ago will be given in response to a single image request. Figure **??** shows a theoretical model of the video server.

3.2 Image Compression

The image compression tests were very interesting. Originally, we expected the image compression to reduce the bandwidth consumed. While running the tests it became clear that the bandwidth was not reduced. Upon finishing the analysis, it was apparent that the bandwidth consumption increased just slightly with each higher level of compression (see Figure 4). It can be summarized that: compression does not affect the bandwidth utilization to any large degree. However, subjective analysis suggests that it increases the frame rate. This is a parameter that should be investigated further on the basis of user experience. This test was run with a full-size image and three simultaneous users. Figure 4 and Tables 3 and 4 show compression as values 1-5, which corresponds to the following levels of compression: minimal, low, medium, high, and very high respectively. The manual gives no numerical relation between the levels of compression, therefore the graph is arbitrary in that respect.

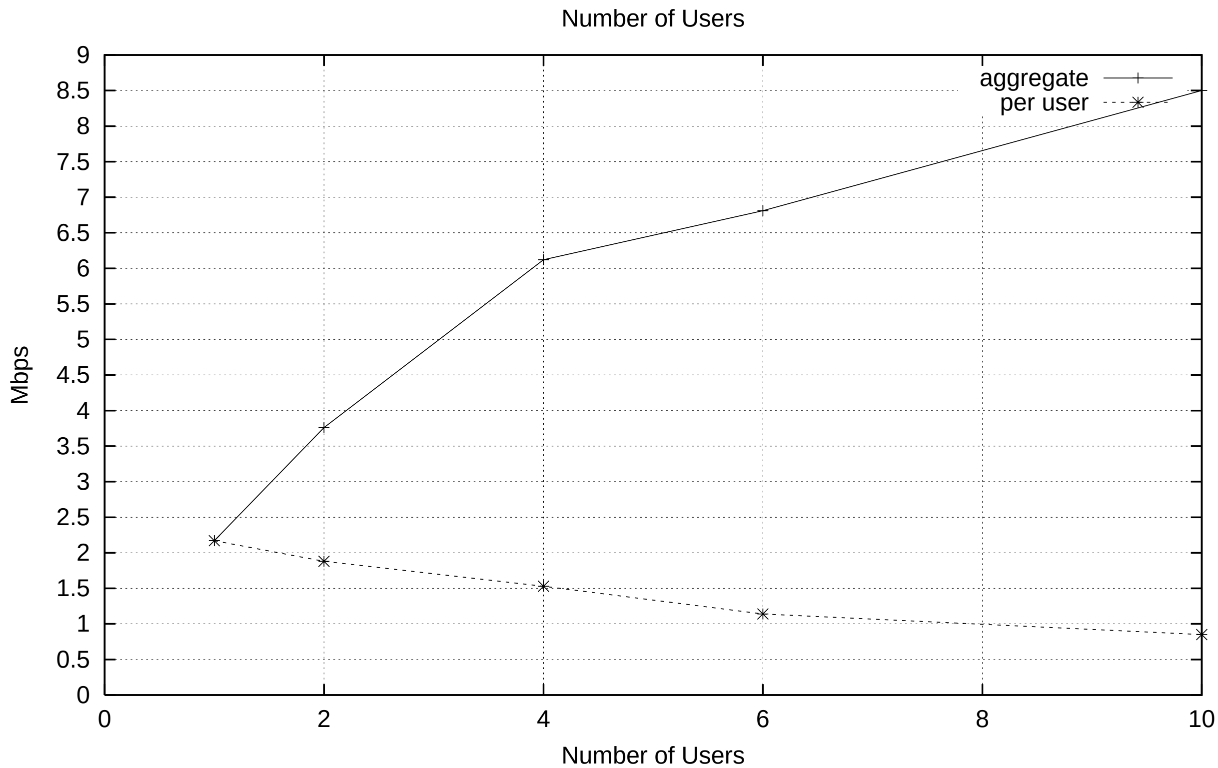

Figure 5: How the number of users affects network load

Number of Users	Throughput	Std. Dev.
1	2.17 Mbps	±6.9 kbps
2	3.76 Mbps	±113 kbps
4	6.12 Mbps	±68 kbps
6	6.81 Mbps	±50 kbps
10	8.50 Mbps	±202 kbps

Table 5: Aggregate throughput vs. number of users (sample size 3)

Number of Users	Throughput	Std. Dev.
1	2.17 Mbps	±6.9 kbps
2	1.88 Mbps	±131 kbps
4	1.53 Mbps	±172 kbps
6	1.14 Mbps	±148 kbps
10	0.85 Mbps	±92 kbps

Table 6: Per user throughput vs. number of users

3.3 Number of Users

The number of clients was varied to show how multiple users affect the overall network bandwidth consumed by this device. As a direct result, we were also able to determine the per user bandwidth. Figure 5 shows the results of these measurements. On inspection it can be seen that the number of users increases the aggregate bandwidth consumed by the device. At the same time, the bandwidth of each individual user decreases, resembling a tug of war over video frames. This effect is positive because it means that as more users are added the utilized bandwidth becomes more efficient. We anticipate additional users will impact the user experience through a slower frame rate and higher latency. Note that even at 10 users the overall bandwidth is still increasing, however, at a slower rate. Leaving the question: at what point will the aggregate bandwidth reach maximum? These tests were run with 'full size' images at minimal compression.

3.4 Rate Limiting

Rate limiting tests had rather disappointing results. The rate limiting control provided by the video server seems to have a limited effect on network utilization. Given the number of tests needed for each of these measurements, we ran it only on a small subset of the possible rate limit settings. The only setting that appeared to cause a change in the traffic was the 0.5 Mbps setting. In addition, this setting did not lower either the aggregate throughput or the per user throughput to the manufacturers specified amounts during our tests. The results are shown in Figure 6. Knowledge regarding how the rate limiting was implemented by the manufacturer would have given a better understanding of these results. Limiting the amount of bandwidth consumed by this device is critical to HPWREN, therefore it is likely that some rate limiting based on Cisco Express Forwarding (CEF) [8] will be implemented.

A second possibility regarding the manufacturer's implementation of this rate limiting is that it could have been based on adjusting the TCP window size. Upon further inspection of the traces, we found nothing that would point to the use of this method. We are unable to say conclusively that the rate limiting functionality does not work, just that it did not work properly in our test configuration.

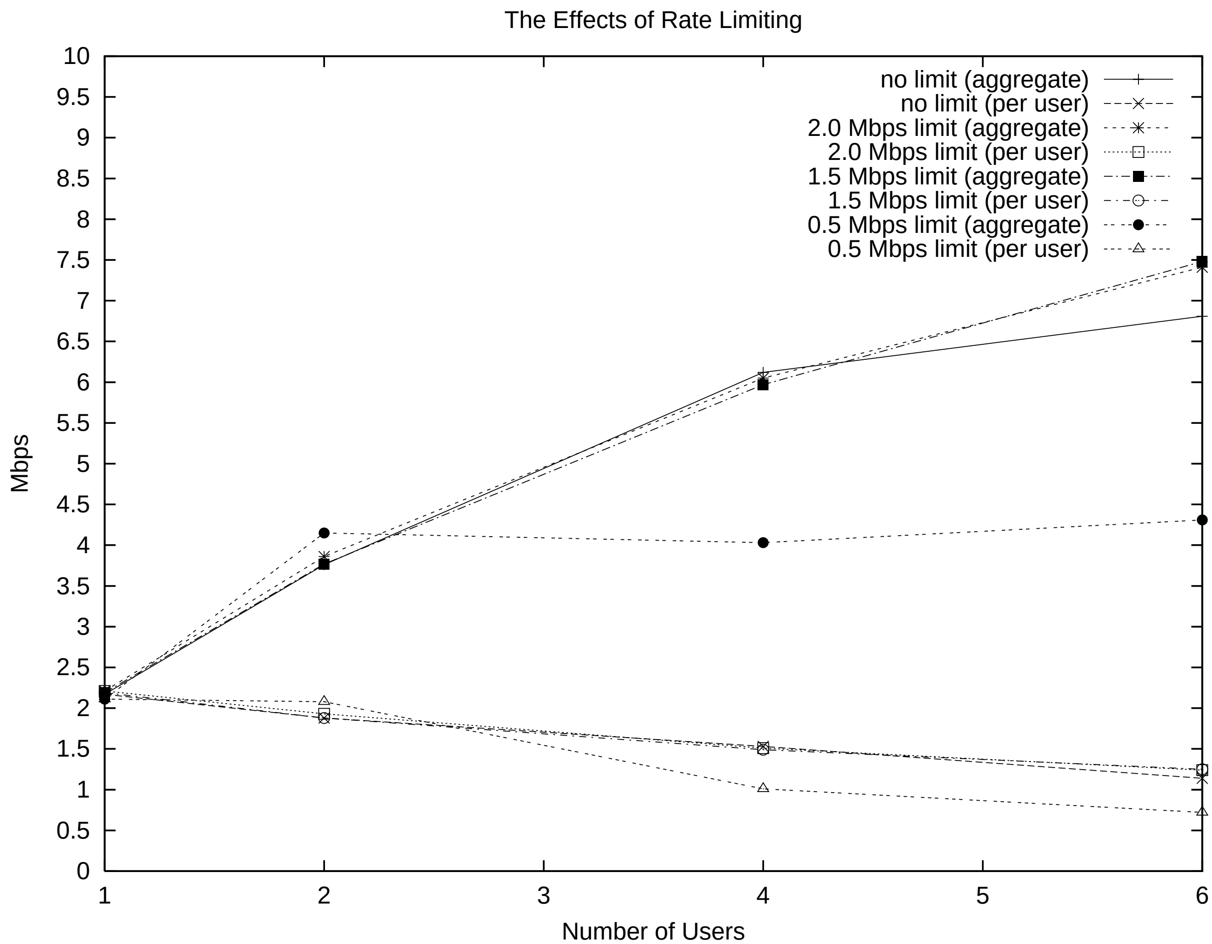

Figure 6: The effects of rate limiting

4 Conclusion

This paper presents what we feel are the best methods and variables to use to measure the impact of Internet-enabled cameras on network bandwidth. As discussed, this is important because new devices are becoming available every day and it is necessary to understand their resource requirements as well as the trade-offs that exist between various products.

In our testing, we were unable to make the rate limiting perform to the manufacturers specifications. Additionally, we are concerned about the ability of this video server to consume 9 Mbps of bandwidth.

It seems likely that an image distribution scheme based on multicast would have been a better solution. However, UDP suffers losses and multicast would require different client software.

We are interested in what others find to be the best methods and most important parameters for their devices.

5 Acknowledgments

This work was sponsored by the National Science Foundation under Cooperative Agreement No. ANI-9807479 (NLANR/MNA) and Grant No. ANI-0087344 (HPWREN) to the University of California, San Diego. We would also like to thank Maureen C. Curran for her reviews of this paper.

References

[1] High Performance Wireless Research and Education Network (HPWREN): http://hpwren.ucsd.edu/

[2] National Laboratory for Applied Network Research Measurement and Network Analysis Group (NLANR/MNA): http://moat.nlanr.net/

[3] T. Hansen, P. Yalamanchili, H-W Braun. *Wireless Measurement and Analysis on HPWREN.*

[4] Pelco: http://www.pelco.com/

[5] Axis Communications: http://www.axis.com/

[6] Mt. Woodson Camera Description: http://calorie.nlanr.net/Wireless/cameras/woodson_doc.html

[7] Axis Communications, *Axis 2400/2401 Administration Manual v1.1*, Copyright 1996-1999 Axis Communications AB

[8] Cisco Express Forwarding: http://www.cisco.com/warp/public/cc/pd/iosw/iore/tech/cef_wp.htm

[9] Raw measurements and Initial Analysis: http://stat.hpwren.ucsd.edu/tshansen/camera_tests/