# Networking for Multimedia
# Tech Topic #05: Network Support for Distributed Gaming

W.C. Feng, F. Chang, W.C. Feng, and J. Walpole, "A Traffic Characterization of Popular On-Line Games," IEEE/ACM Transactions on Networking, vol. 13, num. 3, June 2005.

Another classic example of a "take traces and discuss their meaning" work, it is a good guess that this paper was one of the first journal articles to explicitly look at online gaming, given that it appeared in Transactions on Networking rather than a special issue on entertainment systems or the like.

Five years ago WoW was not necessarily a house-hold term, but Counter-Strike was (...college house-hold, at least) and looking backwards it makes more sense why first-person shooter games were the focus at the time. But one can always make the argument that the traces were not particularly representative of online gaming traffic in the aggregate. The second section flatly states that games rise and fall in popularity in a matter of months, but still the entire paper is underlined by the assumption that focusing on first-person shooters as a genre is valid.

The analysis of the main trace can be roughly divided into discussions on predictability, packet size and characterization of clients. The periodicity of the main trace are perfect fodder for the frequency-domain analysis, and the power spectral density plots of packet rate are particularly effective. If there is a complaint regarding this analysis it is that not all phenomena are sufficiently explained, for example the concentration around 38 Hz in figure 4(a) is not addressed.

From a layout perspective the sheer number of graphs distributed on several pages make the results a bit hard to follow, particularly when flipping back and forth between pages.

The description of the session time results stand out as striking, assuming this paper is an early and novel application of the Weibull distribution as a source model for traffic on the Internet. Such an analysis, while well done, does bring to mind the question of trace relevance: what percentage of Internet traffic actually behaves like this?

In the context of the paper's FPS focus, the main trace is nicely backed up by smaller-scale traces of comparable games. These traces largely serve to reinforce previous trends and correlate multiple game sources, rather than provide additional and/or contrasting data points. The possible exception to this is the unusually small client bandwidth observed for Unreal Tournament 2003.

Section five discusses the positive and negative implications of the previous traces, but does not do a good job of describing why current networking devices will perform so poorly when faced with this type of traffic. The questionable assumption here being the fact that these devices were designed primarily with TCP traffic in mind this implies a lack of support for other traffic types.

Interestingly, n*arrowest last-mile link saturation* is a phenomenon presented in the conclusion as an important result of the paper despite the fact that its derivation is largely speculative and predicated on observed traffic rates and game designs.

M. Claypool and K. Claypool, "Latency Can Kill: Precision and Deadline in Online Games," ACM Multimedia Systems Conference, Phoenix, AZ, February 2010.

This paper is to appear in the first ACM Multimedia Systems conference at the end of the month. Not specifically related to this work, but still interesting, are two somewhat unique points about this conference: that it intends to archive its papers in a blog format to allow for (moderated) public questions and comments, and that it will host software and data archives as part of the publication process. If these services were to become regular features at major conferences, it may be interesting to track blog activity and data/software downloads as a further indicator (aside from citations) of general impact/interest.

The paper utilizes a very straightforward structure, wherein the authors motivate the field of study by stating that more users are playing games online than ever before, describe certain multimedia characteristics relevant to network and systems, and qualify their classifications based on experimentation. Overall the work is a solid representation of Doing Science.

Section 2 outlines a general game classification consisting of *Avatar* and *Omnipresent*, the former being any game played through a single character's perspective, and the latter becoming effectively (by definition) a catch-all for everything else. Specific reasoning why more detailed delineations are or are not needed is omitted. For example, a game where the user might switch first-person perspectives between a selection of virtually disjoint characters. The game classification system is used somewhat sparingly throughout the rest of the work, and is only broadly related to latency concerns to the point that its inclusion might be questioned, aside from establishing a common terminology.

Phases of an on-line game session are generically presented in section 3, and consist of *setup, synchronization, play* and *transition* periods. Again, the section serves mostly as interesting background and with a nice graphic as the meat of the paper concentrates on latency effects which are of most concern during the *play* period.

Sections 4 and 5 finally arrive at the main points the authors wish to make and back up; briefly, that player actions can be broadly categorized into a varying spectrum of *precision* and *deadline,* and that network latency has a proportionately increasing performance effect on each of these.

To validate their claims an existing open-source online game was modified to artificially induce varying amounts of network latency. Furthermore, the game software was modified to vary *precision* by adjusting the size of the in-game targets (tanks) and to vary *deadline* by adjusting the speed of the bullets. The net result is a rather clever way to make points about certain aspects of online games by choosing the parameters to manipulate wisely.

One might argue that the basis of the experiment is fundamentally flawed by virtue of ignoring human behavior given certain scenarios. For example a human player might, recognizing particular latency characteristics, start timing her actions to take advantage of consistently latent game updates. The 'bot' matches in this simulation will of course never model this, and the reason for not using human subjects was not mentioned. Although, human response to latent game sessions was not the primary goal of the paper, and from the perspective of what the authors were attempting to accomplish—proving how their model affects game play—the approach was quite effective.

While the paper certainly presents itself in a nice package, at the end of the day its conclusions were intuitive enough to question their worth. The motivation for studying these certain aspects probably should have gone more in-depth. Provided uses for the work are towards game designers, network designers, and players, but how these parties are to use the results is not very well explained. In other words, there seems to be little takeaway other than "latency should be minimized" and "real-time strategy players perhaps need not be as concerned". In short, the precision-deadline model is a nice classification of online game characteristics but its applications leave something to be desired.