

Tech. Topic 5: Network Support for Distributed Gaming

I. INTRODUCTION

In this week's technical topic on Network Support for Distributed Gaming, we have three papers[1], [2], [3] as part of the reading list, one of which is optional [3]. In this report, we choose to present a review of "A Traffic Characterization of Popular On-Line Games" [2], which is an extension to their previous work [4]. In our current days, interactive on-line computer games have been gaining more and more popularity and are attracting larger portions of the population. According to reports referenced in this paper, Internet games are accounting for a growing fraction of the aggregate Internet load. From a networking point of view, it would be useful to analyze the traffic behavior of these popular on-line, multi-player games to understand the implications of these traffic sources on the existing network infrastructure.

II. PAPER 1 REVIEW: "A TRAFFIC CHARACTERIZATION OF POPULAR ON-LINE GAMES"

This paper extends over the authors' previous paper [4] by comparing the results of Counter Strike (CS) traffic to other existing on-line games. When a paper extends over previous work, it is important to indicate what exactly the extension over previous work is and to motivate and present an extension which presents a contribution over what has already been done. This is clearly presented in the last paragraph of the Introduction. Furthermore, after understanding what the contribution is, we can clearly see how it is clearly and accurately reflected in the abstract of this paper. The contribution presented is very interesting since, more importantly than analyzing the characteristics of one computer game which might lose popularity, it finds or compares findings between other existing computer games.

The first thing that came to mind is whether World of Warcraft (WoW) had been released at the time this study had been performed, seeing as WoW is the most widely played massively multi-player online role-playing game (MMORPG). However, it turns out that WoW was released on November 23, 2004, whereas this paper was published in 2002, and indicates that the games they study were the most widely used at the time.

This paper is properly motivated: online games are gaining popularity, they are accounting for a growing fraction of the aggregate Internet load, therefore it is important to analyze their characteristics to understand the possible network support that can be provided. Here, we are not only looking at the characteristics of one game, but multiple games and finding common trends.

In the beginning of the Introduction, the authors reference [5], a work from year 2000 which studies the trends in wide area IP traffic and shows that the fraction of online game traffic seems to be on the rise. The authors mention that the results from their paper confirm the results from the year 2000 paper; this provides good support to the motivation for this work.

It is clear that the authors have properly referenced and properly place their related work to avoid unjustified or unfounded assumptions/statements. They gain greater credibility by doing so.

The authors, at the end of the Introduction, indicate how they characterize the games players themselves in terms of their session-time and geographic distributions. These characteristics or results are a very interesting contribution, particularly that pertaining to the geographic location as gamers generally aggregate in certain locations. Integrating geographic factors into the analysis process has useful advantages, such as providing substance to deployment considerations.

The authors appear to understand the world of gaming pretty well. The results in this paper do not lose any of their realism. For example, they understand the time relevance of game popularity, although current games, such as WoW, seem to be gaining more and more popularity as the population grows.

An interesting question that always comes to mind is: is the population of gamers going to cap? The gamer population is increasing, at a rate that is not presented in this paper. However, it might be a social phenomenon that that rate is capped, as is the number of gamers.

The authors of this paper provide a background study of the characteristics of Internet gaming for readers that are unaware of games, their behavior and characteristics. A background study is also very useful to all readers at the same time since it clarifies the points which the authors are trying to address, and particularly assures the gamers that the authors know what they are talking about.

Then, the authors move on to provide a summary of the CS traces. The trace setup sounds reasonable, and the information that they include in the summary includes the relevant information necessary to get a higher level understanding of the traffic passing through the server. They do a good job in explaining the results that they have. They also mention a few of the assumptions that they made, particularly that, although they were able to effectively analyze this single server, the results

in this study may not necessarily apply to overall aggregate load behavior of the entire collection of CS servers. This is an important point to mention, however, they do not provide much of an argument with or against what they are doing, and how this point may or may not affect the importance of their results, or how this point is exactly relevant to what they are trying to do.

To further understand the dynamics of the trace summary, they move on to provide more accurate graphs, which show interesting and motivating results.

This paper is properly structured and outlined. The reader can clearly read through it and understand how the author's are moving from one point to the other, in a logical manner. They provide interesting graphs and results which support their analysis. Whenever a point is made in the text, there is a supporting graph. It can be said that the authors will a little too generous with the number of graphs they included. However, since the graphs served a purpose, as minor as it might be, they were still useful. I do not particularly understand the reason why the authors present graphs where variability is seemingly removed due to looking at the results at a higher level, or rougher granularity. A graph where variability is removed provides an interesting condition, but does not necessarily serve much of a purpose as a graph in itself.

The client characterization of session-time distribution is plausible within the setting of this study. However, the reader remains unsure of whether the argument holds, or remains convincing under other settings. Furthermore, there is no clear explanation of the alternatives to the Weibull distribution, although the conditions of the model can be adapted to the framework of gamers entering and leaving a game. Even though the model is most commonly used, there should be some additional reasoning involved in settling to using that distribution, other than it being more accurate than the exponential distribution that has been suggested in previous work. Furthermore, they should mention how they determined the values for the three parameters from the Weibull distribution function. Despite the lack of motivation behind using the Weibull distribution, they do provide a good analysis of its implications and the failures of using the exponential function which assumes constant failure rates.

A further useful approach the authors take is to try to reasonable explain the reasons behind a scenario that they observe. Such as pinpointing the reason behind the inexplicable high failure rate for flows of shorter duration.

When the authors begin to describe their assumptions when describing the geographic characteristic of CS clients, they also discuss the limitations of their assumptions and make a reasonable argument as to why the limitations are not unreasonable (particularly when describing the correlation between network and geographic proximity). They also discuss the limitations of their analysis, or data collection process. However, their arguments remain plausible and reasonable. Despite the limitations, the reader is not left with the impression that the results are incomplete or inaccurate, because it seems that they are considering that aspect.

They concentrate on the possible reasons behind the disparity between network and geographic proximity. This is an important point to concentrate on given the purpose of this paper. The reasons are reasonable. They further indicate that the major determinant is actually the time of day. They propose a solution, which is to globally position and re-position resources over time to match the usage patterns of players. Although the solution sounds reasonable, one can argue the practicality and the effectiveness of such methods.

Although the extension this paper provides over the previous paper is the compare results between different online games, they seem to rush through the extended section. It might be an overkill to go into details with the other online games seeing as a lot of detail was given previously for CS, which might overwhelm the reader at this point. Some observations from the new games are mentioned but are not really addressed. The authors were simply looking at similarities, not a detailed analysis of differences.

Despite the fact that the number of graphs is overwhelming, they do provide the reader with an in-depth and clear understanding of the points the authors are trying to make in an extremely clear and straight-forward manner.

The bad news and good news are accurate. The device solutions they propose are reasonable and some are currently implemented in our day. They still make sure to reference work that is relevant to the point they are making; their arguments are reliable. Their final solution is an interesting solution to the problem.

When discussing the good news, they mention that their results depend on the relatively fixed number of active players. They then mention how this might not be the case since games can exhibit variability and self-similarity. Therefore, instead, this point should have been mentioned as part of the bad news and not the good. Otherwise, if it was really meant to belong as part of the good news, they should explain how this variability and self-similarity would not largely affect the good news.

The paper appears to be very properly revised with a minimal number of English mistakes. It is an extension of a work previously done; this might have considerably contributed to the fact that the paper might have passed through multiple revisions.

Otherwise, it provides intuitive results as well as an interesting read and discussion.

REFERENCES

- [1] J. Smed, T. Kaukoranta, and H. Hakonen, "Aspects of networking in multiplayer computer games," 2001.
- [2] W.-c. Feng, F. Chang, W.-c. Feng, and J. Walpole, "A traffic characterization of popular on-line games," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 488–500, 2005.

- [3] M. Claypool and K. Claypool, "Latency can kill: Precision and deadline in online games," in *ACM Multimedia Systems Conference*, February 2010.
- [4] F. Wu-chang, C. Francis, F. Wu-chi, and W. Jonathan, "Provisioning on-line games: A traffic analysis of a busy counter-strike server," Tech. Rep., 2002.
- [5] S. McCreary and kc claffy, "Trends in wide area ip traffic patterns - a view from ames internet exchange."