

Network Game Traffic: A Broadband Access Perspective^{*}

Amit Sinha, Kenneth Mitchell^{*}, Deep Medhi

School of Computing and Engineering, University of Missouri-Kansas City

Abstract

Playing network games is fast becoming a significant application on the Internet. Most research related to traffic behavior due to users playing network games has relied on data acquired either by collecting traces at the game server, by polling the game server through queries, or by collecting traces in a LAN environment. In this paper, we study traces collected over a period of one week on the access networks of four different markets serviced by a large ISP that provides broadband fixed wireless (BFW) access service to users. Based on the amount of traffic generated, its popularity among the subscribers in all the markets, and its sensitivity to Quality of Service (QoS) requirements due to its interactive nature, we identify Counter-Strike, a *fast-action* game as the application for our study. We discuss player behavior and game traffic behavior in various markets to identify any effects of the migration to broadband services for the *last-mile* access. We also investigate some of the factors that would influence the provisioning of QoS, namely latency and bandwidth usage. We observe that the distribution of the duration of users' game playing is heavy-tailed. On the data sets we studied, we found that the latency between the users and the game servers falls below the upper bound suggested by researchers.

Key words: Game Traffic Behavior; User Behavior; Broadband Access; Traffic Modeling; Quality of Service; Heavy-tailed Distribution

1 Introduction

The advent of technologies such as Digital Subscriber Line (DSL), Broadband Fixed Wireless (BFW), and Cable modems in “last mile” access networks

^{*} This work has been supported in part by a grant from the National Science Foundation NSF ANI-0106640

^{*} Corresponding Author

Email address: mitchellke@umkc.edu (Kenneth Mitchell).

have enabled the development of user applications that require large bandwidth along with low delay and jitter. These applications include streaming audio/video, online games, Internet Protocol (IP) telephony, peer-to-peer (P2P) file sharing, and other applications.

Network game traffic has generated considerable research interest due to the fact that it poses a great challenge to the existing network infrastructure in order to satiate its requirements. Borella [1] presents source models for Quake I (a *fast-action* game) using data collected in an Ethernet LAN environment. His findings are that the client packet size can be modeled using a deterministic distribution while that of the server follows an extreme distribution. The packet interarrival times follow extreme distributions for both the client and the server. It is assumed that the impact of the upper heavy-tail is minor with respect to the aggregate traffic dynamics while constructing these models. Farber [2], verifies the models presented by Borella. In [3], Henderson and Bhatti present models for user behavior in network games from a game server perspective. They poll game servers to find the interarrival time between players arriving at the game server and the length of a player's session and find that the distribution of the session length is exponential, while the interarrival times follow heavy-tailed behavior. Feng *et.al.* [4] studied data collected at a single game server for Counter-Strike, emphasizing the impact of this traffic on networking equipment.

While most work is focused on finding and specifying the kind of service guarantees required by network games, not much has been said about the perspective of the access provider. Questions like (a) what percentage of subscribers are actually involved in playing network games; (b) what is the distribution of the game traffic being generated upstream on the access; (c) what is the relationship between the user arrival process with the time-of-the-day and day-of-the-week with intent of playing games; (d) what is the distribution of the duration of a subscriber's active game session in a single market; and (e) what QoS provisioning is required for the broadband access network to support network games? still remain unanswered.

Previous work observed the activity on the game servers, where the user activity presented an aggregated view of multiple markets. These results might not reflect the behavior of a single market for the purpose of provisioning and designing an access network. In [5], we have shown that traffic behavior in different access networks may depend on the underlying Medium Access Control (MAC) protocol. We further specify in [6] that the session length and volume behavior of various TCP applications for different access schemes are very similar, but the TCP session interarrival time process depends on the underlying MAC protocol. Understanding the behavior of game traffic on BFW access networks, therefore, is desired in order to analyze performance and facilitate network design.

In this paper, we analyze traces collected from four different markets of a large ISP that provides BFW access service. We try to answer the questions listed above with the hope that it would be of interest to other service providers. Further, we make an attempt to answer the question of whether the current best-effort broadband access service is capable of providing the services required by these games. Games are generally categorized as (a) *fast-action*, (b) *slow-action*, and (c) *strategy* type games, and their requirements vary. It is known that the most stringent network delay requirements are associated with *fast-action* games. We chose a game based on the following factors:

- it is the most popular interactive game for the markets under study,
- statistically significant data are available for analysis,
- it is highly sensitive to QoS requirements,
- the game is of interest to other researchers (so that we can make some comparisons)

We found that “Counter-strike”, a *fast-action* game satisfies the criteria above. Our analysis is specific to “Counter-strike”, with the hope that this may be of benefit to other network service providers and designers.

This paper is organized as follows: Section 2 describes the broadband access protocol and the data set collected. In section 3 we describe the method we use to identify the data related to Counter-Strike from the trace. In section 4 we present the game traffic behavior highlighting the weekly trend in game traffic, bandwidth usage, packet load, and latency observed. Section 5 discusses the user behavior such as the interarrival process and duration with respect to the network games. Section 6 covers the approach to model those behaviors of network games which would greatly influence the access network. Finally, we conclude this paper in Section 7.

2 Access Technology and Data Collection

The Broadband Fixed Wireless (BFW) service we consider uses Multi-Channel Multi-Point Distribution Service (MMDS) as the wireless transport from customer premises back to the head end. MMDS operates in the 2.1 GHz to 2.7 GHz band and can support distances up to 30 miles between sites. The access allows a downstream speed of up to 10 Mbps and an upstream speed of up to 256 Kbps. Note that in BFW access, the upstream bandwidth is shared (like cable modems). It has its own *proprietary* MAC. Fig. 1 presents the typical architecture of the BFW service structure.

This paper considers measurements from four different sites located in four different states in the United States, where broadband data services are deployed




Fig. 1. Typical BFW service structure

consisting primarily of best-effort Internet applications. The flow level UDP and TCP measurements are extracted from *tcpdump* [7] using *tcptrace* [8]. These measurements represent a 50% sampling, i.e., data were collected for 1 minute in every 2 minute interval. *tcptrace* can produce several different types of output containing information on each connection seen, such as *start-time*, *end-time*, *source and destination IP addresses and port numbers*, *total bytes and packets upstream and downstream*, *upstream and downstream throughput*, *round trip times*, *window advertisements*, etc.

The trace collection started on a Thursday 00:00:00 GMT in the markets referred to as A, B, C and D throughout this paper. The trace collection in market C is for a duration of five days, while that from other sites last for a week. The customer bases for these markets include both residential and business customers. The ratios of the subscriber population in markets A through D with respect to A is 1, 1.65, 1.39, and 1.64, i.e., market B and D have the highest number of subscribers while market A contains the smallest subscriber population.

3 Identifying Counter-Strike Traffic

Communication between the network game client and the game server takes place on non-standard port numbers (port numbers higher than 1023 [9]). These port numbers could be used by other unknown applications too, creating confusion between game traffic and other applications when only port numbers are used to make the classification. Therefore, this method of identification is not suitable for modeling the traffic and user behavior for network games. It is important that the data extracted for the purpose of analyzing traffic and user behavior reflect the true game traffic. Hence, we were interested in finding an alternative approach to identify the traffic for Counter-Strike with minimal error. The first step in this direction would be to understand the communication process and architecture of the game we are interested in.

3.1 Counter-Strike Architecture

Counter-Strike [10] is a modification to the Half-Life game and is one of the most popular network games played over the Internet. It belongs to the class of *fast-action* multi-player, online, first-person shooters games that have been a significant source of network game traffic [11]. It is built on a client-server architecture with multiple clients communicating with a central server (game server) that keeps track of the global game state.

The communication process that takes place before a client can start playing the game on the network is as follows:

- (1) The communication process uses UDP as its transport layer protocol
- (2) The client connects to the master server on port number 27010 and queries it for the list of game servers. There are two master servers for Counter-Strike, *half-life.east.won.net* and *half-life.west.won.net*.
- (3) The master server sends the list of available game servers along with the possible delays between the client and various game servers.
- (4) The player then selects one of these servers and sends a connection request. The most common client port used is 27005 and the game server port is 27015 or 27016 for the game connection. In case there are firewall restrictions, a client could set up a client port number other than 27005.
- (5) Once the connection between the client and the game server is set, the client can start playing.

3.2 Filtering Process

Now that we know the communication process between the client and servers before the client can start playing, we track the path of this communication process in the data set that we have. The steps involved are as follows:

- (1) First, we filter the flows from the UDP trace where the destination IP address belongs to one of the master servers with the destination port number 27010. We record the IP address of the sources in those flows that are requesting the list of game servers.
- (2) Next, we scan the data set for the activity of each of these sources immediately after querying the master server. All the flows that either have the client port number 27005 or destination port number 27015 or 27016 within an interval of a minute since the query are assumed to be generated by the game. Every consecutive flow with the same IP address and port number pair that starts within a one minute interval since the previous one ended is considered to be part of the game traffic.

- (3) We then verify the destination addresses from the queries made on the master servers using QStat [12], a program designed to display the status of the game servers along with their addresses.

While step-1 above does not have any caveats, there might be questions about the effectiveness of step-2 and step-3. We believe step-2 will work with the data set that we have, since the customer base consists of residential and business users. We assume that business customers are not actively involved in playing network games, while in the case of residential customers, the user has full control of the firewall, so reconfiguring the client port number would be very rare. On the other hand, it is rare but possible that game servers would use a different port number than the commonly used port number. Therefore, step-2 just ensures that all the accepted flow records either have the commonly used client port number as source port or the commonly used server port number for the game as its destination port. We then verify all the destination addresses with the game server addresses using step-3. However, the time of data collection and the time we use QStat queries are not synchronized. Therefore, there arises the possibility of rejecting valid game traffic records for which the game server address is not listed by the QStat. Among the rejected records at step-3, we accept the ones with source port number 27005 and destination port number of 27015 or 27016 as valid game records, since the possibility of another application using these port numbers at a time immediately after the source queries the master server is rare. We follow the above procedure for each market to minimize the error and to ensure the validity of the traffic data generated by Counter-Strike¹.

4 Game Traffic and Network Activity

The first step toward identifying the issues faced by access service providers with respect to network game traffic would be to understand its behavior. To do that, a proper characterization of the game traffic is required. In this section, we focus on identifying the traffic characteristics which are typical of a network game. We observe the dynamics of traffic contributed upstream by the game with respect to the day of the week. We also look at the network behavior, load, and usage.

¹ It is possible that some of the filtered traffic records belonged to Half-Life (the older version of Counter-Strike). We believe that this would not affect our analysis since Half-Life is relatively outdated, would therefore have insignificant contribution, and is intrinsically similar to Counter-Strike.

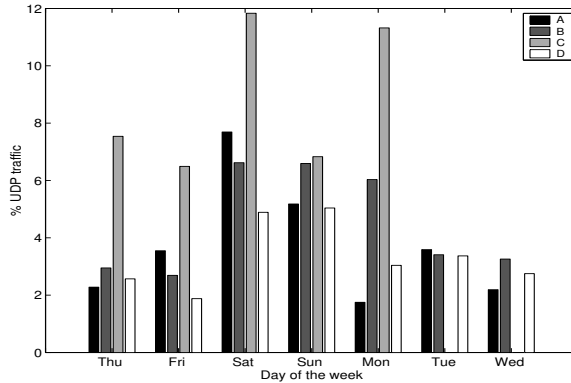


Fig. 2. Trend in % of UDP traffic contribution in upstream direction

4.1 Traffic Contribution Upstream

McCreary and Claffy [11] observed that network games are contributing an increasingly large proportion of traffic on the backbone. One major concern of an access service provider is the upstream bandwidth, since this is much less than the downstream capacity. So, it is important to know the network game traffic contribution in the upstream direction. Fig. 2 displays the percentage of traffic contributed by clients playing Counter-Strike to the total UDP traffic in the upstream direction. The fact that the peak contribution from just one network game approaches 12% in certain markets is significant. The broad trends in various markets peak somewhere between Friday evening and Sunday night. This percentage of traffic contribution depends on the number of subscribers playing the game on a particular day and also the variation in the relative contribution of other applications to the total UDP traffic. Either way, the fact is that the network game contributes significantly to the total upstream UDP traffic. Therefore, further investigation to understand its behavior and its impact on the access network is desired.

4.2 Bandwidth Consumption and Packet Load

Since there is a significant upstream traffic contribution made by Counter-Strike, it is important to investigate the bandwidth consumption by this application. Table 1 indicates the total bandwidth usage and packet load generated per user. Feng *et.al.* [4] show that the bandwidth usage per user averages around 40 kbps, owing to the fact that current game designs target to saturate the narrowest last-mile link, which at the time they were designed, happened to be 56 kbps modems. We observe that while the mean bandwidth usage for various markets lie below this average, the maximum bandwidth usage exceeds this value. As an access provider, the concern is more about the bandwidth usage in the upstream direction. Moreover, for BFW and cable access,

Table 1
 Statistics of bandwidth usage and packet load

Markets	Bandwidth (kbits/s)				Packet Load (packets/s)			
	Total		Upstream		Total		Upstream	
	Mean	Max	Mean	Max	Mean	Max	Mean	Max
A	12.38	110.52	3.49	10.37	16.48	65.08	9.43	26.17
B	20.16	68.70	6.93	44.42	30.63	159.21	19.59	156.81
C	15.14	60.01	5.58	50.51	23.27	156.66	13.30	155.49
D	14.09	52.28	4.33	30.67	18.93	68.03	11.41	62.56

the upstream bandwidth is shared between the subscribers. The average per user upstream bandwidth usage for game traffic varies from 3 kbps to 7 kbps (maximum from 10 kbps to 50 kbps) in different markets. As of now, there are only a very small fraction of subscribers (approximately 1% or less) in various markets who are actively playing network games. As network games gain more popularity, this number is going to increase. Unfortunately, the bandwidth usage would definitely put an upper limit to the number of users that can be playing simultaneously and experience the service quality required, especially in a shared access environment. The number of players who can play simultaneously will also be affected by the activity of other subscribers who are consuming bandwidth at that time. This would be an important factor to be considered while designing the network for QoS provisioning for such traffic.

Table 1 also presents the packet load generated by the network game in various markets. Feng *et.al.* [4] have observed that the rate of incoming packets at the game server is very high, though the size of the packets is small. We observe average packet load generated per user varying from 16 to 30 pps, with a maximum packet load from 65 pps to 160 pps. This reiterates the point put forward in [4] that the routers would have to be carefully provisioned to minimize the packet loss and delay caused by routing highly periodic bursts of small packets generated by *fast-action* network game applications like Counter-Strike.

4.3 Latency

Several works discuss the sensitivity of the *fast-action* real-time multi-player games to network latency [13–15]. According to these works, the amount of delay that can be tolerated by a player varies somewhere between 100 milliseconds to 150 milliseconds. Henderson looked at the network game Half-Life and found that latencies above 225-250 milliseconds are intolerable for players [3]. There are two different approaches researchers have suggested to overcome this issue. One school of thought believes in implementing QoS in the network

Table 2
 Statistics of RTT

Markets	RTT (millisecond)			Delay (milliseconds)
	Mean	Stand. Dev.	C^2	
A	112	89	0.62	11.5–100.5
B	84	74	0.78	5.0–79
C	82	35	0.17	23.5–58.5
D	89	60	0.45	14.5–74.5

and providing the requisite service to the users who are interested in network games for some price. The second school of thought [16] believes in designing the game protocol in such a way so that it adapts and compensates for the latency and other network related issues. On the other hand, Henderson and Bhatti show with their experiment [17] that the players adapt themselves to the latencies and probably would rather play games without QoS than pay for the service.

Table 2 presents the average round-trip time (RTT) from the user to the game server and back, along with its standard deviation and squared coefficient of variation (C^2). The RTT values have been averaged over all the game servers that the user connects and plays on from each market. Interestingly, we find that for all the markets, the latency between the user and the game server falls below the upper bound suggested by researchers, so that its impact on the user performance is negligible. Therefore, in the markets that we are investigating, latency does not seem to be an issue with the network game players. This delay can increase though, for example, in case of a network outage. But an adaptive game design would be able to compensate for these occasional blurs.

5 User Behavior

The interest of the service provider lies in providing acceptable QoS to its subscribers. Therefore, it is important to identify the user² behavior with respect to network games. The percentage of subscribers playing Counter-Strike in markets A, B, C and D were 0.66%, 1.08%, 0.83% and 0.38% respectively for the duration of data collected. The total number of active players was more in market B and C than in market A and D, though market D has one of the largest subscriber population among the four. This suggests that the

² We use the term “user” and “player” interchangeably to refer to the subscriber who plays the network game “Counter-Strike” at least once during the entire observation period.




Fig. 3. Proportion of users playing Counter-Strike

fraction of subscribers involved in the network game may not necessarily be proportional to its total population. Our interest is to see if we can identify the common behavior of these users in different markets and also highlight the differences in their attributes.

5.1 User Activity

First, we would like to identify trends in user behavior versus the day of the week. For this, we find the proportion of unique users playing Counter-Strike on different days to the total number of unique users that played Counter-Strike at least once during the entire observation. Fig. 3 gives the proportion of players active on different days with respect to the total number of players in their market for the duration of the trace. The common trend for all the markets is that this proportion increases as the weekend approaches, peaking on either Friday, Saturday or Sunday. This is in agreement with the notion that more users would have relatively free time to play on these days than the weekdays.

5.2 User Session Duration and User Interarrival Times

Next, we present both the duration of the user's game playing session and the arrival process of requests to play the game. In Fig. 4, we plot the play duration (on log-scale) versus the time of arrival of users for various markets. This gives a pictorial view of how these two factors vary with the time of the day and the day of the week. In almost all the markets, there is a period of time in a day when there are very few to no users playing games. This period falls between late night and early morning. The exact hours vary depending on the market, indicating varying user behavior. The arrivals are clustered between Friday evenings and Sunday nights, indicating that a lot more users play the

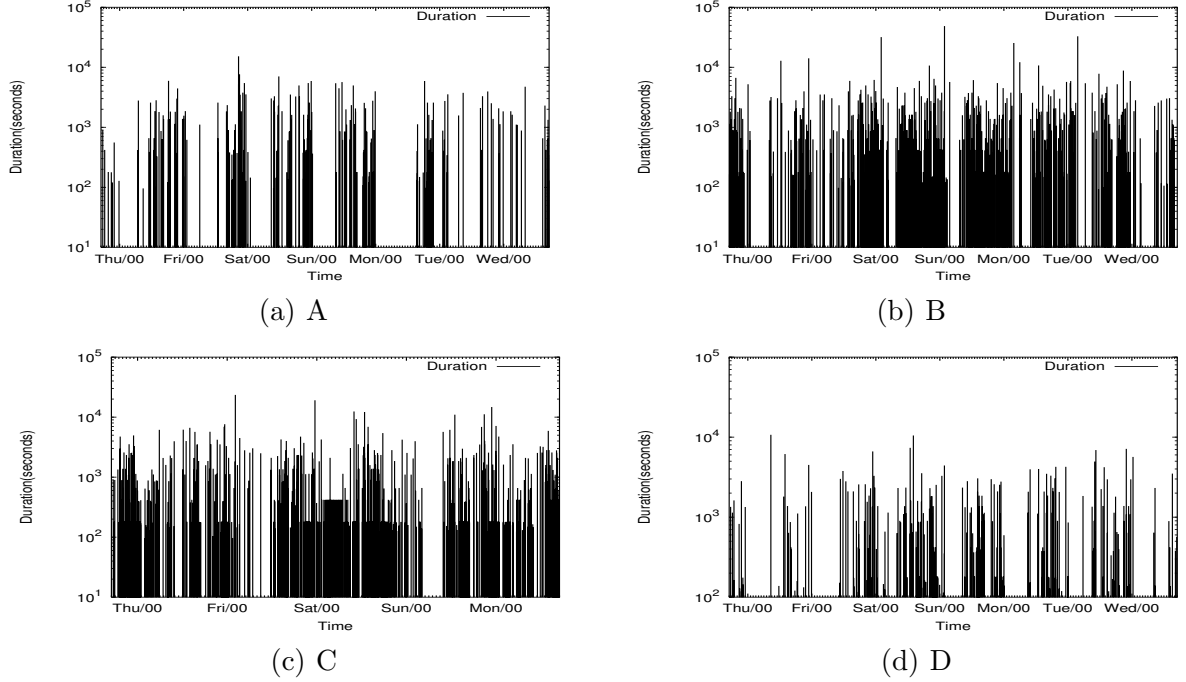


Fig. 4. User arrival processes and game durations in different markets

Table 3

Statistics of user durations and interarrival process

Markets	Game Duration (seconds)					Interarrival Time (seconds)		
	Min	Mean	Max	Variance	C^2	Mean	Variance	C^2
A	95	1253	15242	2830861	1.80	2343	30860579	5.62
B	91	724	48606	4784400	9.14	414	1641233	9.59
C	93	683	23460	2344160	5.03	353	794559	6.41
D	101	948	10704	2023697	2.25	1654	17669072	6.46
Aggregate	91	775	48606	3450999	5.76	–	–	–

game on the weekends than on weekdays. This is in agreement with traffic trends observed in Section 4. Also, there appears to be more game traffic in markets B and C as compared to A and D, which is just a natural outcome of more active player population in the former markets.

Table 3 presents the statistics for the game durations and the user interarrival times for different markets. It is interesting to note that the mean interarrival time increases with the increase in the mean game duration. Since the player population is relatively small in the markets we are studying, this may indicate that there is a finite set of users playing the game in each market – the longer they play, the longer they take to come back again. But the mean by itself does not say much about the behavior. The variance for this process is high

for all markets. The squared coefficient of variation (C^2) is very high for all the markets indicating intense burstiness in the arrival process.

We also observe very high variability in the users' game durations for all markets, and the very high C^2 indicates that the game duration process is very bursty, especially for market B and C. We have game durations as small as 90 seconds and as large as around 13 hours. The small game durations are likely due to dropped connections, users not finding the game to their liking and quitting early, or due to novice players losing out early and quitting. The large game durations are probably due to expert players who like to play for long periods of time, or due to users who forget to exit from the game by mistake. Since we are investigating the same application in different markets, we observe the aggregate statistics for the game duration as well. The overall mean of the game duration is 775 seconds.

6 Approach for Modeling Behavior

So far in this paper we have discussed the game traffic behavior and the user behavior. It is known that analytical models are a basic requirement for planning and designing of a network in order to provide desired service to the subscribers. In this paper, we do not attempt to come up with analytical models for the behaviors, but instead look at the possible approach for modeling. We assume that with time every market would have a similar high number of subscribers interested in playing game on the Internet and the behavior in all the markets would converge to that of the markets with maximum proportion of its subscribers involved in network games. Therefore the behavior in markets B and C are important, since they have the two highest percentage (number) of players in their customer base.

6.1 User arrivals

The variation in the number of users wanting to play network games would provide an idea to the access provider about the QoS requirements not just from the design, but also from the pricing perspective. The users in each market form a fraction of the population on different game servers and we are looking at the aggregate in each market. Therefore, this can be compared to the observations made by Henderson and Bhatti in [3] for the session membership at multiple game servers. Indeed, in Fig. 5, we observe a near sinusoidal pattern for the number of active users. In fact, this behavior remains strong for all markets throughout the duration of the trace. This is evident from Fig. 4 (Section 5) where the arrival process shows a strong time-of-the-day ef-

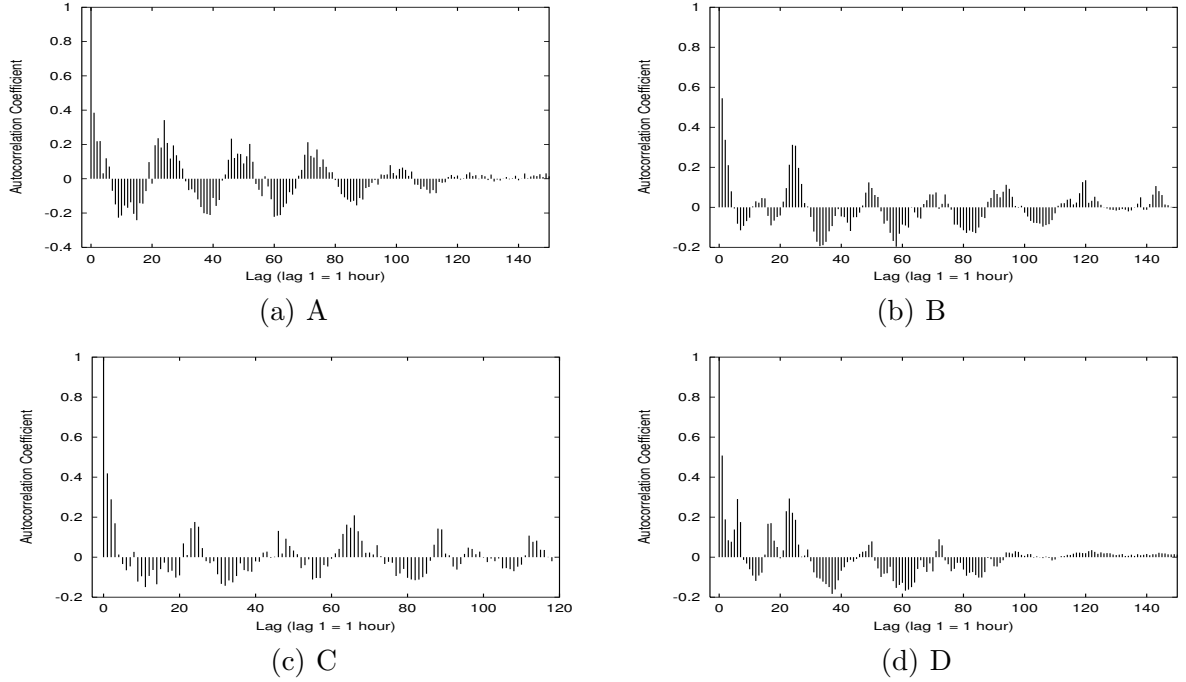


Fig. 5. Number of user

fect. Therefore, a time-series model, for example, an ARIMA (Autoregressive Integrated Moving Average) model, introduced by Box and Jenkins [18] and used in [3] would be applicable.

6.2 User Interarrival Time and Duration

Fig. 6 presents the log-log plot of the complementary cumulative distribution function (CCDF) for the user interarrival time process. The linear behavior of the plot indicates heavy-tailed nature for this process for all the markets. This observation is in agreement with the observations made in [3]. One interesting observation is that the scale of the behavior depends on the user population involved in playing game. Therefore, the tails of the distributions for market A and D nearly overlap and so do the distributions for market B and C.

Fig. 7 presents the log-log plot of complementary cumulative distribution function (CCDF) for users' game durations. The linear behavior in Fig. 7(a) indicates a heavy-tailed distribution of the game duration process. This is different than what was observed by Henderson and Bhatti [3], where they found the distribution to be exponential. The duration behavior is expected to be a property of the game rather than the market, so we aggregated the user durations for various markets and plotted the CCDF in Fig. 7(b) along with the tail of the exponential distribution with the same mean. The deviation from the exponential behavior is quite evident. It is rather difficult to pinpoint the

exact reasons for this difference, since the possibilities are galore. Therefore, we have listed below the differences that exist between our trace and the one used in [3]:

- *Collection method*: the traces we collected using tcpdump are more granular than in [3] where the collection of information was based on polling the game servers every 5/30 minutes,
- *Collection point*: our trace was collected on the access points of four different markets, where as the trace collected in [3] was at the game servers,
- *Population*: smaller population of clients involved in network games being observed by us, as compared to [3],
- *Observation time gap*: the time gap between the trace collected by us and the trace used in [3] is more than 2 years, so the user behavior could have changed, with more users playing for longer durations.

6.3 Upstream Bandwidth Usage

One of the most important concerns for access service providers is the upstream bandwidth, especially in the shared upstream environment. Therefore, modeling the dynamics of the upstream bandwidth usage by the network game users with respect to time is important to understand the variation in demand. One would expect the bandwidth consumption to depend both on the number of players and session durations. The near sinusoidal variation in the correlograms for various markets in Fig. 8 suggests seasonal variation for the bandwidth consumption. This sinusoidal variation remains strong throughout the trace for markets A and C, while for market B this variation persists until lag 48 (48 hours or 2 days). For market D, the sinusoidal pattern is extremely weak, but since this market contains a very small number of players, this anomaly is expected.

Fig. 9 displays the varying bandwidth usage per hour with respect to the dura-




Fig. 6. User interarrival time

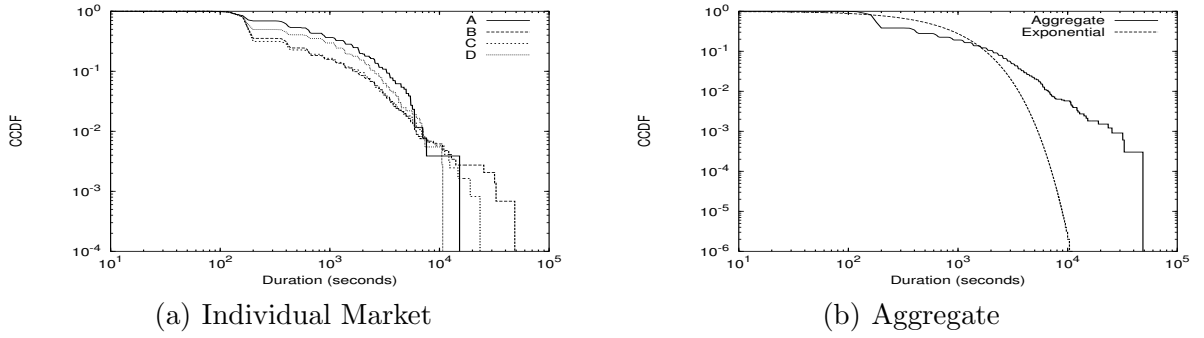


Fig. 7. User duration distributions

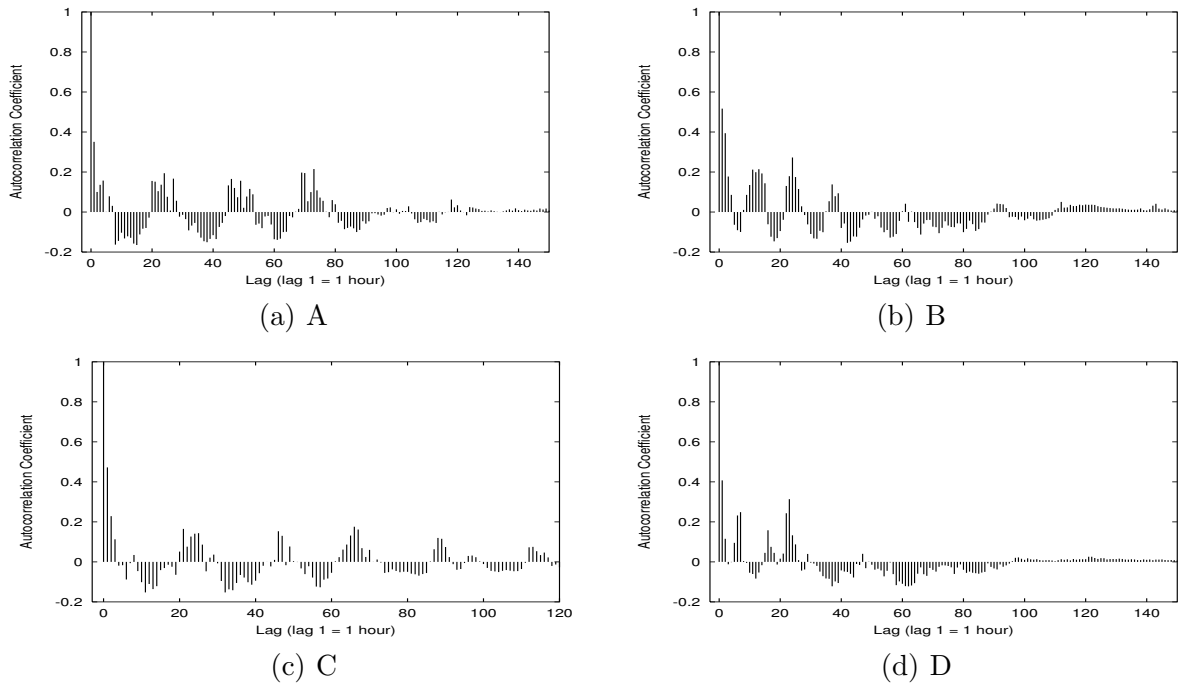


Fig. 8. Upstream bandwidth consumption

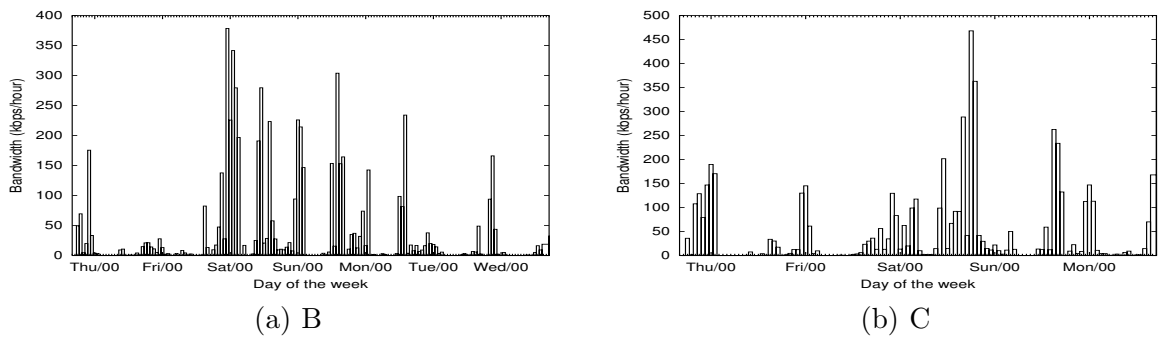


Fig. 9. Bandwidth usage by Counter-Strike players

tion of the trace for markets B and C. The bandwidth consumption upstream indicates both time-of-the-day variation and day-of-the-week variation. The

peak usage occurs around midnight and is highest between Friday evening and Sunday night in these markets. This is strongly reflected in Fig. 8 for market B where the sinusoidal pattern weakens after 48 hours, which would be the start of the weekend. Therefore, this process can not be modeled simply as a single time-series model. One approach would be to model the weekday and weekend bandwidth usage separately. We intend to do this as a part of our future work.

7 Summary and Conclusion

From the perspective of a player, the major issue facing a *fast-action* game like Counter-Strike on the Internet is the latency. A lot of work has been done emphasizing the impact of latency on the player's performance. In the trace that we collected from four markets containing customers in different geographical locations, we found that latency was well within the boundaries suggested. While the latency would increase considerably in the event of network outage (not noticed during the trace duration), we suggest an adaptive approach within the game protocol to take care of these occasional deviations.

The major issue facing the service provider is the upstream bandwidth usage, especially in the case of cable and BFW service providers where bandwidth is shared among the subscribers. We have found that the current bandwidth consumption by the users playing the network game Counter-Strike is well below the total network capacity. This would definitely change as this game and others become more and more popular. We have found that usage varies with the time-of-the-day and day-of-the-week. The seasonal variation indicates that the bandwidth consumption could be predicted using a time-series model, but the model could not be applied directly since there is a strong day-of-the-week dependency. Therefore, we propose to have two different models, one for the weekdays, and the other for the weekends. The behavior of the number of users is consistent with observations made by other researchers. Thus, we propose to model this process using ARIMA models.

We observe that the user interarrival process is still heavy-tailed but the user duration is no longer exponential but also rather heavy tailed, which reflects a change in user behavior with increasing popularity of the game over time. This study has brought up some facts regarding networking issues related to a *fast-action* game called Counter-Strike, which is the most popular game in the markets observed. These analyses form the first step toward constructing a generalized parametric model for upstream bandwidth consumption and a model for the number of users actively playing games.

Acknowledgments

We thank Sprint Corporation for making the representative data set available for this study; in particular, we thank Jim Landon, Feng Zhang and Khurram Sheikh for their help.

References

- [1] M. Borella, Source models of network game traffic, *Computer Communications* 23 (4) (2000) 403–410.
- [2] J. Farber, Network game traffic modelling, *Proceedings of the 1st workshop on Network and system support for games* (2002) 53–57.
- [3] T. Henderson, S. Bhatti, Modelling user behavior in networked games, *ACM Multimedia* (2001) 212–220.
- [4] W. Feng, F. Chang, W. Feng, J. Walpole, Provisioning on-line games: A traffic analysis of a busy Counter-Strike server, *ACM SIGCOMM Internet Measurement Workshop* (2002) 151–156.
- [5] A. Sinha, K. Mitchell, D. Medhi, Flow-level upstream traffic behavior in broadband access networks: DSL versus Broadband Fixed Wireless, *IEEE Workshop on IP Operations and Management* (2003) 135–141.
- [6] A. Sinha, K. Mitchell, A comparative study of TCP/IP traffic behavior in broadband access networks, *Technical Report* (2004), <http://conrel.sice.umkc.edu/asinha/papers/comcom04.pdf>.
- [7] tcpdump, <http://www.tcpdump.org/>.
- [8] tcptrace, <http://www.tcptrace.org/>.
- [9] W. Stevens, *TCP/IP Illustrate, Volume1: The Protocols*, Addison-Wesley, 1994.
- [10] Counter-Strike, Counter-strike: Counter-terrorism action powered by Half-Life, <http://www.counter-strike.net/>.
- [11] S. McCreary, K. Claffy, Trends in wide area IP traffic patterns - a view from Ames Internet exchange, *ITC Specialist Seminar on IP Traffic Modeling, Measurement and Management* (2000) 1–11.
- [12] Qstat, <http://www.qstat.org>.
- [13] L. Pantel, L. Wolf, On the impact of delay on real-time multiplayer games, *Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)* (2002) 23–29.
- [14] G. Armitage, An experimental estimation of latency sensitivity in multiplayer Quake 3, *IEEE International Conference on Networks (ICON)* (2003) 137–141.

- [15] N. Sheldon, E. Girard, S. Borg, M. Claypool, E. Agu, The effect of latency on user performance in Warcraft 3, Proceedings of the 2nd workshop on Network and system support for games (2003) 3–14.
- [16] Y. Bernier, Latency compensating methods in client/server in-game protocol design and optimization, Proceedings of 15th Games Developers Conference (2001), <http://www.gdconf.com/archives/2001/bernier.doc>.
- [17] T. Henderson, S. Bhatti, Networked games - a QoS-sensitive application for QoS-insensitive users?, Proceedings of the ACM SIGCOMM Revisiting IP QoS workshop (2003) 141–147.
- [18] G. Box, G. Jenkins, G. Reinsel, Time series analysis: forecasting and control, Prentice-Hall, New Jersey, 1994.
- [19] M. Oliveira, T. Henderson, What online gamers really think of the Internet?, Proceedings of the 2nd workshop on Network and system support for games (2003) 185–193.
- [20] T. Lang, G. Armitage, A ns2 model for the Xbox system link game Halo, Australian Telecommunications Networks & Applications Conference (2003), <http://atnac2003.atcrc.com/ORALS/Lang.pdf>.

Curriculum Vitae

Amit Sinha received his Bachelor of Engineering Degree in Civil Engineering from Indian Institute of Technology-Roorkee (formerly University of Roorkee), India in 1997. He received his Master of Science Degree in Computer Science from University of Missouri-Kansas City in 2002. He is currently a Ph.D. student in the Computer Science & Electrical Engineering (CSEE) Department at the University of Missouri-Kansas City (UMKC). His research interests include traffic data analysis and modeling of telecommunications and computer networks.

Ken Mitchell received a Bachelor of Music Degree in performance from Indiana University in 1979. He received his Master of Music in performance from Indiana University in 1980, and received his M.Sc. and Ph.D. in Computer Science from the University of Missouri-Kansas City in 1994 and 1997 respectively. He is currently an Assistant Professor in the Computer Science & Electrical Engineering (CSEE) Department at the University of Missouri-Kansas City (UMKC). His research interests include analytic modeling of telecommunications and computer networks and his work has been supported by grants from the National Science Foundation and the University of Missouri Research Board.

Deep Medhi is Professor in the Computer Science & Electrical Engineering (CSEE) Department at the University of Missouri–Kansas City (UMKC). He received the Ph.D. degree in Computer Sciences from the University of Wisconsin-Madison in 1987. Prior to joining UMKC in 1989, he was a member of the technical staff in the traffic network routing and design department at the AT&T Bell Laboratories, Holmdel, New Jersey from 1987 to 1989. He was an invited visiting professor in the Institute of Telecommunications at the Technical University of Denmark during the summer of 1999, and Visiting Research Fellow at the Department of Communication Systems, Lund Institute of Technology, Lund University, Sweden during the summer of 2003. He has published over fifty five peer-reviewed papers. His research interests are in survivable network design and architecture, dynamic quality-of-service routing, next generation network architecture, IP network availability, wireless networking, network management. He is a senior technical editor of the Journal of Network and Systems Management (JNSM). He has served on the technical program committees of several IEEE conferences including IEEE INFOCOM, and IEEE NOMS. He was the Technical Program Chair of the 2003 IEEE Workshop on IP Operations & Management (IPOM'2003). Over the years, his research has been funded by Defense Advanced Research Project Agency (DARPA), National Science Foundation (NSF), Sprint Corporation and Doris Duke Charitable Foundation.