



Modelling multicast pricing using networked games

TRISTAN N.H. HENDERSON

Department of Computer Science, University College London,
Gower St, London WC1E 6BT, UK

T.Henderson@cs.ucl.ac.uk


ABSTRACT

Current analysis of multicast user behaviour uses MBone sessions as a source of data. The limited size of the MBone, however, means that many of these sessions are sparsely populated. Existing unicast multipoint applications such as games and chat are widely-used, and from a user's viewpoint, are largely indistinguishable from multicast equivalents. As such, they may provide better data for analysing how users behave in multicast scenarios. We have recorded usage data for popular multiplayer networked games, and are using this data to develop multicast pricing schemes. Initial observations, presented here, include strong time-of-day effects and high variability in the duration of users' sessions and their inter-arrival times.

1 Motivation

- Previous analysis of multicast user behaviour, e.g. [1], looks at MBone sessions. The limited deployment of multicast, however, means that these sessions tend to be small and single-source.
- We are interested in behaviour in large-scale, multiple-source, dynamic membership, multicast sessions for modelling multicast pricing schemes.
- Use popular unicast multiplayer games such as *Half-Life* [2], since from a user's viewpoint, these are similar to their multicast equivalents.

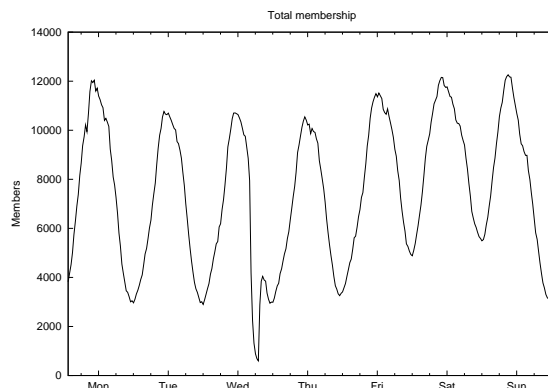
2 Methodology



- *Half-Life* is a multiplayer game, where users connect to a central server and then share and interact in a common "world"; players explicitly choose a server (by IP address) upon starting a session, and actively disconnect from a session.
- A list of 1936 servers running the *Half-Life* daemon was obtained from the master *Half-Life* server at half-life.east.won.net.
- Servers were polled at regular intervals, using a game-specific query mechanism, and the number of players, player nicknames and length of each player's game were retrieved.
- If a server failed to respond, group membership was assumed to be the same as at the previous successful poll. If a player is no longer found at a subsequent poll, they are determined to have left at some point in time between the previous poll and the current.
- These 1936 servers were polled every half an hour for one week, and from this, the 35 most popular servers were selected, and polled every five minutes for a three-day period (Friday to Sunday).
- The master list of servers does not appear to be updated very frequently; many servers were obsolete or defunct. Hence more reliable data was required. Microsoft Research have set up servers running three variants of *Half-Life* and *Quake III: Arena* [3], and have given us permission to monitor them. This was done at five minute intervals for a month.
- Additional caveats: user's nicknames may change or be duplicated; users joining and leaving between polls are neglected; the calculation of user leave times is dependent on the intervals between polls.


3 Number of users

We have observed strong time of day effects in the number of players. As might be expected for a leisure activity such as playing games, the peak days are the weekends. The peak time of day, however, is still the middle of the day.




4 User duration

The distribution of the durations of a user's session appears to be heavy-tailed, and almost follows Zipf's law [4].




4.1 Log-log plot of distribution




5 Inter-arrival times

The users' inter-arrival times also vary widely, and time-of-day effects can also be observed.



5.1 Boxplot of distribution



6 Pricing modelling

This data collection has taken place as part of wider research into multicast pricing. We are using an expected utility model in conjunction with this usage data to simulate various multicast pricing schemes, and to examine their stability and effectiveness.

We assume the existence of network externalities; part of a user's benefit from a multicast session is determined by the number of people in the session. Otherwise they would use a unicast transmission. We also assume that users share the cost of a multicast session. Expected utility theory [5] holds that agents choose the outcome with their highest expected utility; this means that *ceteris paribus*, users will choose to enter or leave a session depending on the ratio of U to P . Thus, we can model:

$$U = U(k + n^{2\alpha})$$

$$P = P(n)$$

where:

- U is the utility received by a participant in a session
- n is the number of members in a session
- k is the innate value of a session (i.e., the value to a user in the absence of other participants)
- α represents the diminishing marginal utility from additional members

By applying density functions to the distribution of U/P thresholds to users, we can model how users will interact given membership and price changes in multiple-source sessions with dynamic membership.

7 Conclusions

- Strong time-of-day effects in membership and inter-arrival times
- Heavy-tailed distribution in users' session durations
- High variability in users' duration and inter-arrival times
- Data is being used for simulating the stability and effectiveness of multicast pricing schemes

REFERENCES

- [1] K. C. Almeroth, "A long-term analysis of growth and usage patterns in the multicast backbone," in *Proceedings of the 19th IEEE Conference on Computer Communications (INFOCOM)*, (Tel Aviv, Israel), Mar. 2000.
- [2] "Half-Life." <http://www.sierrastudios.com/games/half-life/>.
- [3] "Quake III: Arena." <http://www.quake3arena.com>.
- [4] G. K. Zipf, *Selected Studies of the Principle of Relative Frequency in Language*. Cambridge, MA: Harvard University Press, 1932.
- [5] J. L. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1944.