By Kuan-Ta Chen,
Polly Huang, *and*
Chin-Laung Lei

HOW SENSITIVE ARE NETWORK QUALITY?

Game-playing time is strongly related to network QoS, helping determine user satisfaction and deliver better service quality to online gamers.

Network quality is critical to real-time, interactive, online game play. Yet even though today's best-effort Internet does not provide quality of service (QoS) guarantees, the popularity of online gaming is only increasing. This popularity may be because QoS is not really a must-have characteristic or because players are accustomed to struggling with and playing through unfavorable network conditions.

Ironically, despite the popularity of online games, complaints about high "ping times" and "lags" continue to surge in game-player forums (such as www.gamedev.net/community/forums/). It would appear that most players view network latency and loss as a major hindrance to an enjoyable gaming experience. Here, we investigate whether players are really as sensitive to network quality as they claim. If they are, we would like to answer whether they quit games earlier than they might otherwise due to unsatisfactory network conditions.

A number of studies have sought to evaluate the effect of network quality on online gamers [1–3, 5, 7–9, 11]. In most of them, a series of games is played in a controlled network environment in which the gaming experience of subjects is graded either subjectively or objectively. Subjective tests require participants to report their feelings about playing games with different levels of network QoS. While this approach might capture human perceptions in a controlled situation, it cannot measure how players react to poor game quality in real life, as other factors may outweigh their QoS-intolerance and affect their decisions. Besides, it is very costly to scale up such studies by including more human subjects. Meanwhile, objective evaluations are usually based on user performance in a specific context (such as the number of kills in a shooting game, the time needed to complete each lap in a racing game, and the capital accumulated in a strategy game). However, game scores are highly dependent on player skills, game design, and content, so results are not comparable and generalizable across different games.

ONLINE GAMERS TO

Psychologically, the pleasing sensation players experience in online games is analogous to being in the flow state after taking a mood-changing substance like marijuana. One study [6] found that the most significant factors related to the flow experience in online gaming are skill, challenge, and involvement. The game-playing experience can be described as so pleasurable that gamers are unaware of the time that passes while they're playing the game. According to the theory laid out in [6], if the feeling of involvement in the virtual world is diminished by network lags, users become more conscious of the real world, mitigating their sense of time distortion. Some players may simply decide to quit a game as soon as they detect unacceptable lags. For these reasons, we conjecture that game-playing time is affected (to some extent) by the network quality players experience.

There are certainly exogenous reasons other than the quality of network conditions that might affect gamers' decisions to continue with (or leave) a game. For example, they might leave a game due to their schedule constraints or mental state. On the other hand, they may be tied by social bonds, no matter how unstable the game play might be. Although the behavior of individual players is highly variable and unpredictable, we show here that, overall, there is a consistent and strong link between player departure decisions and network QoS.

NETWORK QUALITY VS. GAME-PLAYING TIME

Our conjecture is validated by real-life traces taken from the commercial massively multiplayer online role-playing game Shen Zhou Online developed by UserJoy Technology in Taiwan (www.ewsoft.com. tw). MMORPGs (such as EverQuest and World of Warcraft) are networked computer role-playing games in which a large number (perhaps 10,000) of players interact with one another in a virtual world. While many game genres are round- or stage-based, a design that allows players to regain consciousness of the real world, the adventure in role-playing games is continuous; no explicit mechanism forces players to take a break.

An uninterrupted sense of immersion in a virtual environment is one of the main attractions of MMORPGs because it entices players into a flow experience; other attractions are elaborate reward cycles and evolving social networks [10]. The importance of the sense of immersion implies that network quality, a key factor in maintaining a perfect sense of immersion, strongly influences the flow experience of MMORPG players; therefore, game-playing time should be shorter if the network conditions that users experience are unsatisfactory.

Another reason for investigating MMORPGs is that they are relatively slow-paced compared to other popular genres (such as first-person shooter) that usually require players to make split-second decisions. In addition to the pace of a game, there is a great deal of variety in how players control their virtual characters. In fast-action genres (such as shooters) players must instruct characters "what" actions to take and "how" to perform them; for example, to move a character to a new location, a player must control each step it needs to take (such as three steps west followed by five steps north). On the other hand, in slow-action gen-

res (such as MMORPGs and war strategy games) players instruct characters only "what" to do; that is, they point out only the location the character should move to; it then automatically heads toward it via a route that is either predetermined by a game designer or computed on-the-fly.

By any standard, MMORPGs are slow-action games with less-stringent service requirements than fast-action games. Therefore, they could be viewed as a baseline of real-time interactive online games such that if network QoS frustrates MMORPG players, it should also frustrate players of other game genres.

To play Shen Zhou Online, thousands of players pay a monthly subscription fee at a convenience store or online. As is typical of MMORPGs, players can engage in fights with random creatures, train themselves to acquire particular skills, partake in commerce, or take on a quest. With the help of the Shen Zhou Online staff, we recorded all inbound/outbound game traffic of the game servers located in Taipai, Taiwan, a total of 15 140-game sessions over two days.

We extracted the network performance for each session based on the sequence number and flags in the TCP packet header. We computed network latency based on the difference between the acknowledgement and the delivery time of nonretransmitted game packets and inferred the network loss rate based on the receipt (or lack of receipt) of TCP acknowledgements. The observed players were spread over 13 countries, including China, India, Hong Kong, and Malaysia, and hundreds of autonomous systems, manifesting the heterogeneity of network-path characteristics—therefore the generality of this trace. On average, players stayed for 100 minutes after joining a game. This statistic is consistent with the statistics of

Figure 1. Relationship between game session time and network QoS factors.

Korean MMORPGs played in Japan [4], showing that the average session time ranges from 80 to 120 minutes. However, the difference in individual game-playing times is quite large; for example, the shortest 20% of sessions span less than 40 minutes, but the top 20% of players spend more than eight continuous hours in a game.

We use graphical plots to illustrate the difference in the game-playing time of sessions experiencing different levels of network quality (see Figure 1). The figure depicts the association of game-playing time with network latency, network delay variation, and network loss rate, respectively. All three plots indicate that the more serious the network impairment players experienced, the sooner they were likely to leave the game. The changes in game-playing time are surprisingly significant. For instance, gamers who experienced 150msec latency played four hours on average, but those experiencing 250msec latency played for only one hour on average—a ratio of 4:1. Moreover, variations in network delay and network loss induce more variation in game-playing time (note the span of the y-axis of the graphs).

Having demonstrated that players are not only sensitive but also reactive to network quality, we still must define "good quality" and "bad quality." For example, given two QoS configurations—"low latency with moderate variation" and "moderate latency with low variation"—which one is better? To answer, we propose a model that describes the changes in game-playing time due to network QoS. This allows us to grade the overall quality of game playing based on a set of network performance metrics (such as latency and loss) in terms of user satisfaction.

MODELING GAME-PLAYING TIME
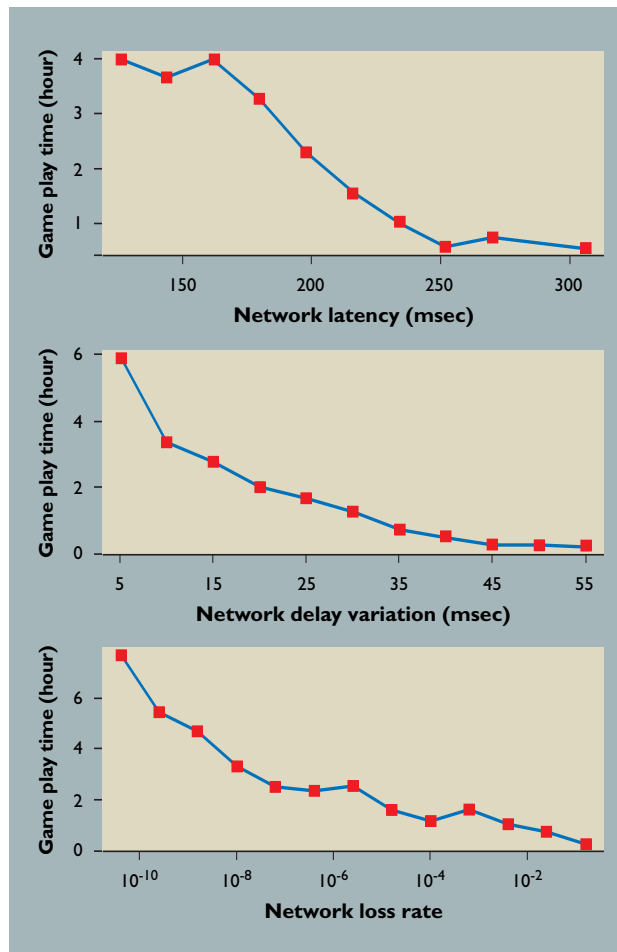We find that the survival model, which is often used in the medical field to describe the relationship

FORECASTING WHEN A PLAYER WILL LEAVE A GAME
could provide useful hints for system performance optimization and resource allocation.

between patients' survival time and the treatment they receive, provides a good fit for describing the reactions of game players to network quality [3]. The derived model takes network QoS factors as the input and computes the departure rate of online players as the output. The regression equation is defined as follows:

$$\textit{departure rate} \, \alpha \, \exp(\, 1.6 \times \textit{network latency} \, + \\ 9.2 \times \textit{network delay variation} \, + \\ 0.2 \times \log(\textit{network loss rate})).$$

It illustrates that the player departure rate is roughly proportional to the exponent of the weighted sum of certain network performance metrics, where the weights reflect the effect of each type of network impairment. The coefficients can be interpreted by the ratio of the departure rates of two sessions. For example, suppose two players—A and B—join a game at the same time and experience similar levels of network loss and delay variations, except that their network latency is 100msec and 200msec, respectively. The ratio of their respective departure rates can then be computed by exp((0.2 . 0.1) × 1.6) ∼ 1.2, where 1.6 is the coefficient of network latency. That is, at every moment during the game, the probability that A will leave the game is 1.2 times greater than the probability that B will leave the game.

Given the strong relationship between game playing time and network QoS, we can "predict" the former if we know the latter. Forecasting when a player will leave a game could provide useful hints for system performance optimization and resource allocation. To assess the feasibility of forecasting game playing time, we compared actual time and model-predicted time for Shen Zhou online players (see Figure 2). The

Figure 2. Actual vs. model-predicted game-playing time for session groups sorted by their risks.

graph sorts sessions, grouping them by player risks, which are defined according to the first equation, to estimate the level of player intolerance to poor network conditions. We observed that, at the macro level, the prediction is rather close to the actual time observed, suggesting that a service provider can predict how long a given individual player will stay after joining a game and optimize resource allocation accordingly.

Please note that while this methodology is applied to all kinds of online games, the exact equation about players' QoS sensitivity may depend on a specific game design (such as the transport protocol and client-prediction techniques used), varying among games.
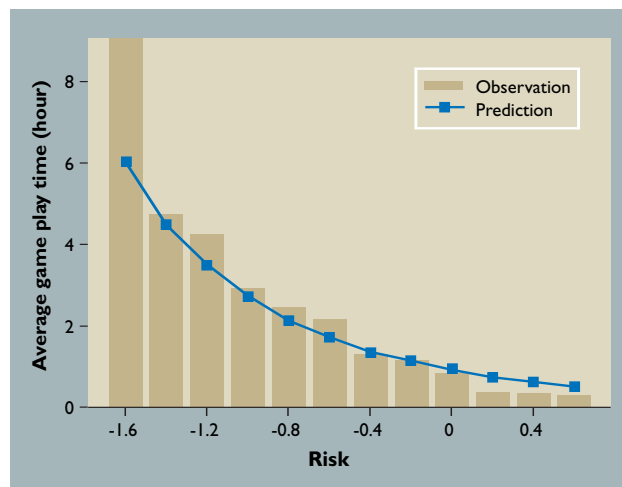
APPLYING PLAYER SENSITIVITY
Even though we are unable to change the fact that the Internet is not QoS-guaranteed, perhaps leaving players dissatisfied due to poor network conditions, we can still improve matters by exploiting their sensitivity to network QoS.

Improving user satisfaction. Given that we can quantify the risk of players leaving a game due to unsatisfactory QoS, systems can be designed to automatically adapt to network quality in real time in order to improve user satisfaction. For example, we might enhance the smoothness of game playing in high-risk sessions by increasing the packet rate (if the risks are caused by long propagation delay or random loss on a noisy link rather than transient congestion) or the degree of data redundancy, so players would be less likely to leave prematurely. Scarce resources, including processing power (for handling player requests and computing game states) and network bandwidth (for dispatching the latest game states), could be allocated more appropriately to sessions based on session risk scores. Resource allocation could be deliberately biased toward high-risk sessions characterized by poor network quality, while simultane-

ously maintaining a reasonable level of player satisfaction in low-risk sessions.

Optimization of network infrastructure. Based on the first equation, network quality can be assessed by way of a single value, making it possible to compare different levels of quality in terms of user satisfaction. For example, according to our scoring, players prefer a setting of 200msec latency and 0.1% loss rate over a setting of 100msec latency and 1% loss rate, because packet loss is less tolerable than network latency. As different performance metrics usually reflect design trade-offs, this score is useful for optimizing overall user satisfaction. For instance, as our model indicates that players are less tolerant of large delay variations than they are of high latency, providing a smoothing buffer at the client side—incurring additional latency but smoothing the pace of game play—would improve the overall game experience.

Network troubleshooting. In order to provide continuous high-quality game service, providers must monitor network conditions between servers and customer networks, detecting problems in real time before customer complaints flood the customer service center. However, considering the large number of online users (hundreds of thousands is not uncommon for some popular games), it would be prohibitively expensive, even impractical, to monitor the status of network paths between all intercommunicated networks in real time. Instead, game operators can track user gaming time, which is much more cost-effective. Since online gamers are sensitive to network conditions, a series of unusual (premature) departures over a short period might indicate abnormal network conditions and thus automatically trigger appropriate remedial action.

CONCLUSION

Discussion of QoS provisioning for real-time interactive applications, including online gaming, is as lively as it is because of the popularity of the applications and the complexity of the problem. One of the main obstacles to providing a satisfactory user experience is that, unlike system-level performance metrics (such as bandwidth and latency), user satisfaction is intangible and unmeasurable. The key to addressing this problem is our ability to measure user opinions about network performance objectively and efficiently. We have shown that game-playing time is strongly related to network QoS and is thus a potential indicator of user satisfaction. Nevertheless, user perception is inevitably influenced by the design and implementation of each application.

How can the measure of a user's experience be generalized so the satisfaction measures for different applications are normalized and compared with one another? We leave it as an open question. Only a detailed analysis of the relationship between each component in a network system and user perceptions can deliver a clear answer. **C**

REFERENCES
1. Armitage, G. An experimental estimation of latency sensitivity in multiplayer Quake 3. In *Proceedings of the 11th IEEE International Conference on Networks* (Sydney, Australia, Sept. 28–Oct. 1). IEEE Press, Piscataway, NJ, 2003.
2. Beigbeder, T., Coughlan, R., Lusher, C., Plunkett, J., Agu, E., and Claypool, M. The effects of loss and latency on user performance in Unreal Tournament 2003. In *Proceedings of NetGames 2004* (Portland, OR, Aug). ACM Press, New York, 2004, 144–151.
3. Chen, K.-T., Huang, P., Wang, G.-S., Huang, C.-Y., and Lei, C.-L. On the sensitivity of online game playing time to network QoS. In *Proceedings of IEEE INFOCOM 2006* (Barcelona, Spain, Apr. 23–29). IEEE Press, Piscataway, NJ, 2006.
4. 4Gamer.net. Gametrics weekly Korea MMORPG population survey; www.4gamer.net/specials/gametrics/gametrics.shtml.
5. Henderson, T. and Bhatti, S. Networked games: A QoS-sensitive application for QoS-insensitive users? In *Proceedings of the ACM SIG-COMM Workshop on Revisiting IP QoS* (Karlsruhe, Germany, Aug. 27). ACM Press, New York, 2003, 141–147.
6. Ila, S., Mizerski, D., and Lam, D. Comparing the effect of habit in the online game play of Australian and Indonesian gamers. In *Proceedings of the Australia and New Zealand Marketing Association Conference 2003* (Adelaide, Australia, Dec. 1–3, 2003).
7. Nichols, J. and Claypool, M. The effects of latency on online Madden NFL Football. In *Proceedings of Network and Operating System Support for Digital Audio and Video 2004* (Kinsale, County Cork, Ireland, June 16–18). ACM Press, New York, 2004, 146–151.
8. Oliveira, M. and Henderson, T. What online gamers really think of the Internet? In *Proceedings of NetGames 2003* (Redwood City, CA, May 22–23). ACM Press, New York, 2003, 185–193.
9. Yasui, T., Ishibashi, Y., and Ikedo, T. Influences of network latency and packet loss on consistency in networked racing games. In *Proceedings of NetGames 2005* (Hawthorne, NY, Oct. 10–11). ACM Press, New York, 2005, 1–8.
10. Yee, N. The Daedalus Project; www.nickyee.com/daedalus.
11. Zander, S. and Armitage, G. Empirically measuring the QoS sensitivity of interactive online game players. In *Proceedings of the Australian Telecommunications Networks and Applications Conference 2004* (Bondi Beach, Sydney, Australia, Dec. 8–10). ATNAC, Australia, 2004, 511–518.

KUAN-TA CHEN (ktchen@iis.sinica.edu.tw) is an assistant research fellow in the Institute of Information Science at Academia Sinica, Taipei, Taiwan.
POLLY HUANG (phuang@cc.ee.ntu.edu.tw) is an associate professor in the Department of Electrical Engineering and director of the Network and Systems Laboratory at National Taiwan University, Taipei, Taiwan.
CHIN-LAUNG LEI (lei@cc.ee.ntu.edu.tw) is a professor in the Department of Electrical Engineering and director of the Distributed Computing and Network Security Laboratory at National Taiwan University, Taipei, Taiwan.