Quality is in the Eye of the Beholder: Meeting Users' Requirements for Internet Quality of Service

Anna Bouch[1], Allan Kuchinsky, Nina Bhatti
Internet Systems and Applications Laboratory
HP Laboratories Palo Alto
HPL-2000-4
January, 2000

E-mail:A.Bouch@cs.ucl.ac.uk
{kuchinsk,nina@hpl.hp.com}

Internet, Quality of Service, user perception

Growing usage and diversity of applications on the Internet makes Quality of Service (QoS) increasingly critical [15]. To date, the majority of research on QoS is systems oriented, focusing on traffic analysis, scheduling, and routing. Relatively minor attention has been paid to user-level QoS issues. It is not yet known how objective system quality relates to users' subjective perceptions of quality. This paper presents the results of quantitative experiments that establish a mapping between objective and perceived QoS in the context of Internet commerce. We also conducted focus groups to determine how contextual factors influence users' perceptions of QoS. We show that, while users' perceptions of World Wide Web QoS are influenced by a number of contextual factors, it is possible to correlate objective measures of QoS with subjective judgements made by users, and therefore influence system design. We argue that only by integrating users' requirements for QoS into system design can the utility of the future Internet be maximized.

[1] current address: Department of Computer Science, University College London, Gower Street, London, WC1E6BT, A. Bouch@cs.ucl.ac.uk

Quality is in the Eye of the Beholder: Meeting Users' Requirements for Internet Quality of Service

Anna Bouch
Computer Science Department
University College London
Gower Street, London WC1E 6BT UK
+44 171 504 4436
A.Bouch@cs.ucl.ac.uk

Allan Kuchinsky, Nina Bhatti
Hewlett Packard Labs
1501 Page Mill Road
Palo Alto, CA 94304 USA
+1 650 857 7423
+1 650 857 5610
{kuchinsk,nina@hpl.hp.com}

ABSTRACT
Growing usage and diversity of applications on the Internet makes Quality of Service (QoS) increasingly critical [15]. To date, the majority of research on QoS is systems oriented, focusing on traffic analysis, scheduling, and routing. Relatively minor attention has been paid to user-level QoS issues. It is not yet known how objective system quality relates to users' subjective perceptions of quality. This paper presents the results of quantitative experiments that establish a mapping between objective and perceived QoS in the context of Internet commerce. We also conducted focus groups to determine how contextual factors influence users' perceptions of QoS. We show that, while users' perceptions of World Wide Web QoS are influenced by a number of contextual factors, it is possible to correlate objective measures of QoS with subjective judgements made by users, and therefore influence system design. We argue that only by integrating users' requirements for QoS into system design can the utility of the future Internet be maximized.

Keywords
Internet, Quality of Service, user perception.

INTRODUCTION
The population of users of the World Wide Web is expected to grow from 100 million in 1998 to 320 million by 2002 [24]. While a vision of the future Internet offers the potential to break traditional barriers in communications and commerce, the current service quality to users is often unacceptable [7] [19]. With increasing usage of Internet services, the topic of providing adequate Quality of Service (QoS) for the Internet has become a focus of research. Traditional QoS metrics such as response time and delay no longer suffice to fully describe quality of service as perceived by users. The success of any scheme that attempts to deliver desirable levels of QoS for the future Internet must be based, not only on the progress of technology, but on users' requirements [10].

An increase in demand for access to network bandwidth has led to suggestions that the Internet should implement classes of service according to the QoS needs of applications ([11],[12]). Although resource allocation schemes differ in details, the premise from which they are constructed is fundamentally the same. This premise is to allocate service resources according to the assumed objective QoS requirements of applications. Network service providers are indeed now able to offer a wide range of facilities and services designed to address variability in the QoS requirements of applications. An example is the deployment of adaptive content delivery, where content is altered to take into account application QoS requirements and varying network conditions[2][3]. However, it is currently not known to what extent changes in objective QoS metrics are perceived by, and impact the behavior of users.

Previous research shows that users may judge a relatively fast service to be unacceptable unless it is also predictable, visually appealing and reliable [5]. Additionally, many QoS parameters have been found to interact in users' judgements of quality. For example, Ramsay et. al. [23] found that Web pages that were retrieved faster were judged to be significantly more interesting than their slower counterparts. Although it is often recognized that a measurement of user satisfaction must be included in the assessment of the efficiency of the network as a whole [16], research is needed to identify how the cost of resource allocation for the service provider relates to the value of that resource to users.

We present results from our study into how users define and perceive Internet QoS. We chose to study the influence of different levels of latency. Latency is defined as the delay between a request for a Web page and receiving that page. Previous research has established the salience of such parameters to users [5][8].

RESEARCH QUESTIONS

When designing technologies for electronic services, it is important to understand the interplay between different stakeholders and the extent to which providing benefits to one stakeholder may result in increased costs to another. The questions that we asked in this study examine the requirements of users while considering the need to provide system-level developers with information from which to make design decisions that meet users needs.

We asked the following main questions:

To what extent is there a mapping between objective and subjective QoS?

The definition of QoS from the perspective of users and systems-level network designers diverges. From the perspective of the Internet Service Provider (ISP), QoS approaches that focus on optimizing objective QoS may inform resource allocation mechanisms at a systems level. However, it is not clear to what degree such improvements result in perceived improvements from a user perspective. This problem is compounded by the observation that any definition of objective thresholds that might be established is subject to context-dependent behavior. Subjective QoS thresholds are therefore not fixed throughout a user's interaction. For example, users tolerance for quality is influenced by the particular task in which they are engaged at a particular time.

In addition to research into the partitioning of network resources are claims that levying a charge for premium services is the only way to reflect the value of quality to users and provide economic incentives to service providers [21]. One of the approaches for deploying different levels of QoS is the use of differentiated service mechanisms . A strength of this approach is that it allows traffic flows generated from applications to be aggregated. However, in order to allocate resources to aggregate flows in the appropriate manner, it is first necessary to identify common requirements for QoS that can be associated with those flows. It is not clear that users are willing to pay a premium for any perceived improvements to quality of service, nor have the thresholds above which users perceive the QoS they receive to be qualitatively superior been established.

This question was posed to investigate the level of objective QoS that was considered acceptable to users and to find out if general subjective thresholds could be identified below which an objective level of quality would not be tolerated, in the context of Internet commerce.

What contextual factors influence users' tolerance of Internet QoS?

Our perspective challenges the assumption that there is a strict correlation between objective levels of quality received by users and their perceptions of that quality. For example, demands for a certain type and level of network performance have been shown to vary widely depending on the amount of browser feedback provided [4]. A consistent finding is that QoS received by users should concur with their expectations but that these expectations change according to the pattern of quality received [5], [8]. We set out to investigate the factors that influence users' perception of QoS by asking if users' tolerance for QoS:

- Depends upon the type of task in which they are engaged?
- Changes as the time they spend interacting with the Web site increases?
- Changes if new Web pages are brought up incrementally, or all at once?

What underlying conceptual models influence users' judgements of QoS?

Previous research has established users' models of network operations and how those models influence the levels of quality users are willing to tolerate [5]. Several factors intervene in users' judgements of network quality and play a large part in forming their expectations of future QoS. In addition to investigating contextual factors, we wanted to understand the reasons behind the influence of such factors on users' judgements of QoS. This question, therefore, was designed to investigate how users' conceptual models related to their perceptions of QoS.

METHOD

Research approach

Our research approach combined the gathering of quantitative and qualitative data. This approach was adopted to determine if and where thresholds exist below which users will not tolerate levels of QoS. Objective metrics such as these are the most direct way of informing resource allocation mechanisms. We conducted experimental work to provide information on tolerance thresholds. Additionally, participants were asked to speak aloud about the QoS they received as they performed the task. Through the capture of these verbal protocols, we hoped to record participants' dynamic opinions of QoS. Focus group studies were conducted to gain qualitative data in order to address how contextual factors that influence the definition of thresholds relate to users' conceptual models.

The next section describes the methods we used to answer these research questions.

Participants

There were 30 male participants, aged between 18 and 68, in the study. Previous research indicated that no control was required for age [6]. It was essential that a appropriately homogenous group of users was selected. This was essential because users with different amounts of knowledge and experience of Web QoS have different expectations of QoS [5]. The following criteria were applied when selecting participants. All participants must:

- Use the Internet for at least 2 hours per week
- Have made at least 2 purchases on the Internet in the last year.
- Have at least an intermediate level of self-assessed skill with using computers.

Male participants were selected for the study as it has been shown that there are gender differences in visual perception and learning [17]. Males were identified as the most frequent users of Internet commerce services [22].

Task

The task involved purchasing a home computer system using the HP Shopping Village Web site [14]. This is a frequently used site ranked first for retail revenue generated by e-commerce [13]. This grounded the task in a real-world context. During the task participants were asked to purchase each component of the computer system separately. Participants accessed 22 Web pages during the task. To study whether users' requirements for QoS were similar for similar sub-tasks, a set pattern of actions was repeated through the task. For each component purchased participants were required to:

1. View a class of similar products.
2. Select a specific product from a class of products.
3. Add the chosen product to their shopping cart.
4. View the contents of their shopping cart.

Having a set pattern was important in determining if users' tolerance changed over time. If the tasks had been widely variable then any change in tolerance could be ascribed to the variation in what participants were asked to do, and not to a genuine accumulation of frustration.

The task was designed so that all participants followed the same path through the Web site.

Experimental conditions

During the experiments, users were presented with Web pages that had predetermined delays ranging from 2 to 73 seconds. The choice of this range of speeds was guided by speeds that users had perceived to be qualitatively different in previous research [8][23]. Each user took part in all 3 conditions. The same task was used for each condition.

Investigating Latency: Non-incremental loading

This part of the study investigated whether the latency between requesting a page and receiving it influenced user perceptions of the delay of page delivery. Varying the latency in this condition had the effect that the page where the link had been clicked remained displayed in the browser until the next page had been loaded. This next page was then brought up in its entirety. Predetermined delays were injected into the page loading process. There were two patterns of variation applied to each of two conditions where latency was investigated.

Pattern 1 (Figure 1) mimicked a random pattern of delay. In pattern 2 (Figure 2) the delay generated on the Web pages formed a more regular, relatively smooth pattern.

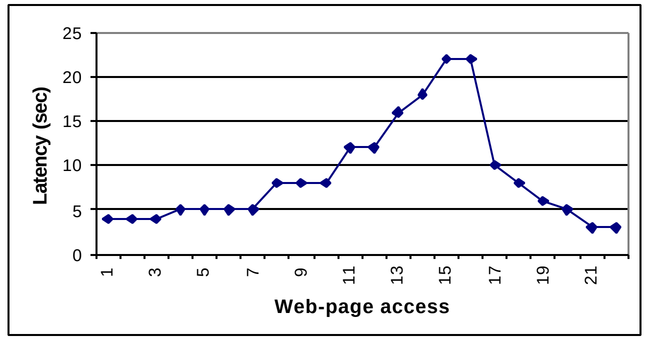

Figure 1: Latency pattern A: Random quality

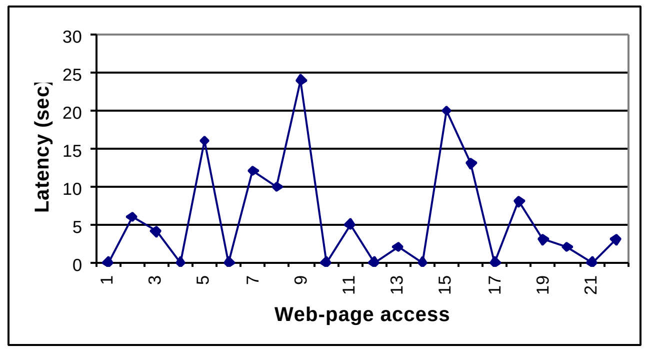

Figure 2: Latency pattern B: Regular quality

1. Classification of latency (condition 1)

During condition 1, participants were asked to perform the task and rate the latency received for each Web page access. An interface (Figure 3) was developed to register ratings. Participants were directed to click one of the buttons in this interface for each Web page accessed. Participants were told that the black button should be used to indicate that the quality was totally unacceptable.

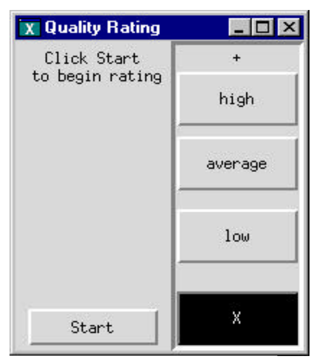

Figure 3: UI

2. Control of latency (condition 2)

During condition 2, participants were again asked to assess the latency of each Web page. However, in this condition they were told that if they found the delay of the Web page unacceptable, they could click a button labeled 'Increase Quality'. The effect of this button was to immediately bring up the requested page. Previous research suggested that this would be a valid measure of users' requirements for speed [4]. This experimental set-up contrasted users' opinions about tolerance of QoS, captured in their direct classifications, with what can be inferred about users' tolerance from their behavior when they controlled the quality.

Participants were split into two groups for the investigation of latency in conditions 1 and 2. 15 participants received pattern 1 for both classification and control of latency, while the remaining 15 participants received pattern 2 for both classification and control of latency. Table 1 shows the experimental conditions applied.

Table 1: Experimental conditions

Condition	Page loading	Pattern	Participants
1: Classify	All at once	Random	15 (Grp 1)
2: Control	All at once	Random	15 (Grp 1)
1: Classify	All at once	Regular	15 (Grp 2)
2: Control	All at once	Regular	15 (Grp 2)
3: Classify	Incremental	Random	30

Investigating Incremental Loading (condition 3)

This part of the study (condition 3) investigated whether users would be more tolerant of delay when Web pages loaded incrementally instead of all at once. Previous research suggests that providing continuous feedback reassures users that the system is working and gives them something to look at while waiting [18]. However, Nielson [20] points out that standard browser feedback, provided in the form of progress bars, fails to communicate the amount of the page that has been completed. Loading Web pages incrementally can address this shortcoming, while providing users with visually useful feedback.

The flow of information between Web server and client was manipulated to cause the Web pages to load in parts. In this condition, participants would receive the banner of the next page as soon as they clicked a link. This was followed by text, and, later, graphics.

Participants were asked to evaluate the time it took for the whole Web page to complete. All participants received pattern 1 in this condition.

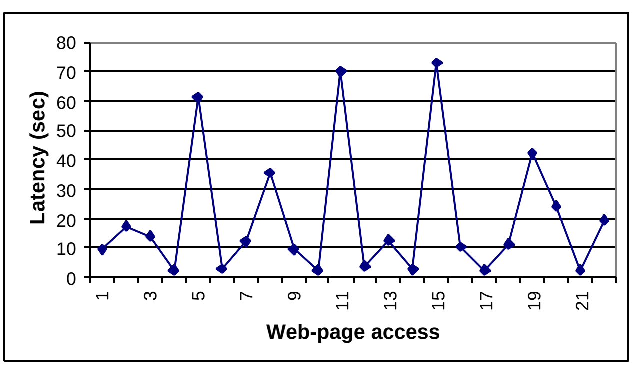

Figure 4: Latency pattern (incremental loading)

Figure 4 shows the mean delay taken by each web-page to complete in condition 3. These measurements were taken using client-based software that captures latency received by users with 100% accuracy [9].

FINDINGS

There were 3 key findings to the research:

- A mapping between objective QoS and users' subjective perceptions of that QoS can be identified and quantified.
- This mapping is influenced by a number of contextual factors including the type of task in which the user is engaged, the method of page loading, and cumulative time of interaction.
- Users' conceptual models underlie the influence of contextual factors on subjective perceptions of QoS.

To what extent is there a mapping between objective and subjective QoS?

Verbal protocols indicated that participants used the 'Low' button when they found the QoS was unacceptable; very few participants used the black button. We aggregated these responses when conducting a set of Chi-squared tests for statistical significance.

Classification of Latency (condition 1)

A significant number of participants assigned certain levels of latency to particular levels of service. Table 2 shows this classification. Figure 5 shows the number of users that classified particular levels of delay in certain categories. For example, it shows that 14 participants classified an 8 second delay as 'High' while 4 participants classified it as 'Low'. Not every participant registered a rating for every level of delay.

Table 2: Classification of latency

Rating	Range of latency Condition 1	Range of latency Condition 3
High	0 – 5 sec	0 – 39 sec
Average	> 5 sec	> 39 sec
Low	> 11 sec	> 56 sec

Previous Web usage studies have found that delays greater than 10 seconds encourage users to believe that an error has occurred in the processing of their request [20]. Our finding that the threshold where QoS is judged as 'Low' is around 11 seconds is therefore consistent with previous work.

Control of Latency (condition 2)

We observed, in condition 2, that there was a large standard deviation among participants in terms of their tolerance of latency. Although the average tolerance was 8.57 seconds in this condition, the standard deviation was 5.85 seconds.

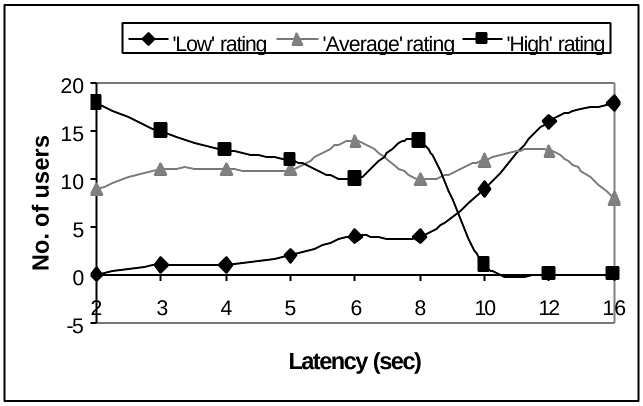

Figure 5: Ratings of latency (condition 1)

It is not possible for us to conclude from this condition that users will tolerate a specific amount of latency before finding that QoS unacceptable. Multiple regression analysis revealed that the number of hours participants used the Web significantly influenced their tolerance for latency in this condition. Higher levels of Web usage were associated with less tolerance for delay during interaction ($p<0.01$).

The large standard deviation observed when participants were asked to control latency may have been due to differences amongst participants in terms of their risk-taking behavior. Participants differed in terms of whether they took advantage of the fact that there was no penalty for pushing the button to increase quality. This difference is also suggested by the fact that there was no correlation between participants' tolerance when classifying latency and their tolerance when controlling latency.

Classification of Latency (condition 3)

The upper bound of latency assigned by participants to each classification in condition 3 are almost 6 times higher in each case compared to the classifications made in condition 1 (see Table 2).This indicates that users are more tolerant of latency when Web pages load incrementally than when there is a delay followed by the display of the page in its entirety.

What contextual factors influence users' tolerance of Internet QoS?

The data we gained from verbal protocols and focus groups indicates that participants were strongly influenced by their expectations of delay when responding to the QoS they received in the experiment. We found that there was almost unanimous agreement amongst participants concerning the factors that help form these expectations.

The amount of time users allocate to the task

Our findings suggest that users anticipate the time it will take them to perform on-line tasks. This anticipation helps form their expectations of the time it should take them to complete the task. Our results suggest that when the process of completing a task is disrupted by unanticipated delays a conflict arises between users' expectations and the QoS received. That QoS is likely to be rejected:

- *'So I'll be sitting there for half an hour so I'm set for that…so a lot of it depends on the time I anticipated I had when I set out'.*
- *'If I'm going to buy something that I need to do research on, mentally I'll allocate more time'.*

Understanding systems-level operations

Since our participants were selected on the basis of their experience in Internet shopping, they were likely to possess an understanding of the manner in which networks store and route data. This understanding was shown to influence the QoS participants expect from the network. Specifically, users are likely to tolerate less delay in receiving a Web page if they believe that it should be easily accessible from memory:

- '(It should be fast because) *I've already been there, it should be cached'*

Users are more likely to be tolerant of delay at certain times of the day, or for 'busy' Web sites:

- *'Understanding that there's a lot of people coming together on the process makes us more tolerant'.*

The company whose products are advertised

Participants in our study believed that companies that are more commercially successful should possess the financial means to supply at least adequate levels of QoS, 100% of the time. This expectation means that users are less likely to accept delays incurred while interacting with a site that promotes the products of such companies:

- *'Because the companies are so huge they should pour money into their web-sites, should have fast sites. If I try to get on those sites and they're slow then I'm not as patient'.*
- *'This is the way the consumer sees the company…it should look good, it should be fast'.*

Duration of interaction

In all conditions we found that users' tolerance for delay decreased as the length of time they spent interacting with the system increased. In all cases this finding is statistically significant ($p<0.01$). The effect is more powerfully significant for condition 3 ($p<0.001$). Maximum tolerance was indicated by the point at which the participant clicked the 'Increase Quality' button so that the Web page would appear.

Task variation

For users who possess an understanding of the way that networks store and route data, the type of real-world task in which they are engaged is likely to have an influence on the amount of delay tolerated:

- *'Like when you're comparing I expect that to take a little longer because it's going to have to go out and get information'.*

Qualitative data suggests that participants expect different tasks to take longer than others. From this information, we were able to classify tasks according to participants' expectations of the latency each task should incur. High tolerance tasks were:

- Comparing several items.
- Viewing the shopping cart.

By comparison, low tolerance tasks were:

- Returning to a previously accessed page.
- Viewing a class of products.
- Adding to the shopping cart.

During the experiments we found that users tolerated different levels of latency depending on what they were doing. A closer look at Figure 5, for example, shows that more participants classified an 8 second delay, corresponding to comparing different printers, as 'High' quality, compared to an objectively lower 6 second delay, corresponding to viewing a class of monitors. Statistical tests show that users will accept more delay when they are comparing products or viewing the contents of their shopping cart than when they are viewing a class of products or adding to the shopping cart, ($p<0.01$).

What underlying conceptual models influence users' judgements of QoS?

Qualitative data showed that participants possess a conceptual model of the way that networks store and access information. This conception influenced their tolerance for delay. For example, our results showed that tolerance for delay associated with specific tasks was dependent on *a)* if the task required accessing a database, e.g. from which to compare products, or *b)* if the task involved a calculation to be made, e.g. in calculating the total spent from the items placed in the shopping cart. Although no pages were cached in the experiment, it was participants' conception of this technology that made them relatively intolerant of delay when re-visiting previous pages.

Additionally, we found that participants' expectations of the commercial setting in which the task was placed influenced their tolerance of delay. For example, viewing the shopping cart was a high tolerance sub-task:

'when I brought up my shopping cart I figured it would have to compile a bit longer so I was more willing to wait a little bit for it to come up'.

Our results show that the longer users interact with a site, the less latency is tolerated. This effect is more pronounced when the latency experienced is more unacceptable. In condition 3, users tolerance of delay declined more rapidly than in condition 1, where the latency received was lower.

Qualitative data suggests that participants in this study felt that any cumulative slowness on Web pages suggested that the security of their purchase was compromised:

'If it's slow I won't give my credit card number'.

'I'd say, you haven't got your resources figured out, you're a poorly managed outfit, I don't trust you any longer'.

Once users perceive that their security has been compromised, no purchase will be made and the main purpose of any commercial Web site becomes critically damaged. It is therefore crucial for systems designers to understand the effect of cumulative frustration, especially as it is typically in the later stages of interaction that users are likely to commit to a purchase.

Qualitative data may explain the reasons behind increased tolerance for delay when Web pages load incrementally. These reasons center around the utility of the feedback provided by incremental loading. Participants were less frustrated when feedback in the browser showed them that the network was processing information:

'As long as you see things coming up it's not nearly as bad as just sitting there waiting and again you don't know whether you're stuck'.

Some participants in our study used the standard browser feedback to indicate activity in the network:

'I think if I had a way to know that it loading I wouldn't mind the old page intact until it was ready to switch'.

Typically, these participants did not prefer incremental loading. This finding suggests that users require feedback to be assured that the network is continuing to process their request. Either browser feedback or incremental loading can provide this feedback.

Our results show that users' conceptual models of the way in which networks operate can significantly influence their tolerance of QoS in predictable ways. Consequently, an understanding of users' conceptual models, and, perhaps more importantly, the behavior which is driven by them, is a crucial step in accommodating user demand.

IMPLICATIONS FOR SYSTEMS DESIGN

The perspective of this work is not only to understand user behavior relating to QoS but to interpret those findings into solutions for real-world problems.

Servers implement mechanisms that dynamically control the processing and delivery of information in response to users' requests. These functions are performed by using scheduling algorithms that allow requests to be processed in a specific order, for example, by Earliest Deadline First scheduling. By prioritizing requests in this way, the server can process the highest number of requests whose delivery still has utility to the user. Our research has defined objective thresholds below which delivering information has less utility. Service providers should design with these kinds of thresholds in mind.

Knowing the latency users will tolerate also allows servers to provide the appropriate feedback to the browser. Our results have shown that providing information concerning the processing of a request can significantly increase users' tolerance of poor QoS. For example, if a server receives a computationally intensive request it can calculate if it will be possible to process the request within the specified deadline. The server is then able to calculate the trade-offs concerning the number of requests that can be delivered to users within the required time-frame. Informing the user if their request will be delayed above the established threshold for tolerance implicates that the QoS of the task as a whole is seen as better:

'I think it's great…saying we are unusually busy, there may be some delays, you might want to visit later. You've told me now. If I decide to go ahead, that's my choice'.

Our results have shown that there are contextual factors that influence users' tolerance for latency These factors can be used in prioritizing requests in the server. Users will accept more delay when performing tasks that they believe to be computationally intense. Technology exists that can tag packets of information that are generated from a certain task [2]. The server could then queue and service tagged packets accordingly. Tasks perceived as computationally intense, for example, could be tagged according to a certain profile and placed near the back of the server's queue. We have been able to show that a significant number of participants classify certain levels of QoS into different classes of service. This information can be used by the designers of priority service schemes in installing processing deadlines for each class of service.

A central finding in our study was that users' tolerance for latency decreases over the duration of interaction with a Web site. This phenomenon can be considered when performing server scheduling optimizations.

IMPLICATIONS FOR FURTHER HCI RESEARCH

Traffic smoothing techniques exist that can shape the pattern of latency received by users [3]. Participants who were given latency pattern 2 (Figure 2) made classifications that were lower than participants given pattern 1 ($p < 0.05$). This effect was found for both short, e.g. 3 seconds, and long, e.g. 16 second delays. Our comparisons were limited in that we investigated this affect by comparing randomly generated QoS with a single arbitrarily selected pattern. With the goal of gauging the effect of a predictable service, further work is needed to establish this trend for other patterns of latency where the magnitude of delay is more precisely controlled.

All participants in our study were given the conditions in the same order. Further work is needed to establish the effect of the order of exposure on acceptance of delay. Exposure to delay in former conditions could introduce memory effects and effect participants' subsequent expectations of delay:

'You get a bit spoiled I guess once you're used to the quickness, then you want it all the time'.

Our results have shown that users are more tolerant of delay when the Web pages load incrementally. However, the range of objective delays given in this condition was wider than those given for the condition used as a comparison, that where the page was displayed all-at-once. Although our results possess a high level of statistical significance one has to consider that the classifications given by participants in our study were relative. Further work is therefore needed to establish the effect of incremental loading between conditions with an identical range of latency.

Our study was specific to a Web shopping task. Further studies of users' perceptions of QoS should investigate the validity of our findings in different domains, such as an entertainment Web site.

It has been shown that users judge the acceptability of real-time audio quality based on visual interface feedback, as opposed to making direct evaluations of objective levels of quality [4]. Further work is needed to investigate the influence of the other contextual factors reported in this paper on users' perceptions of real-time QoS.

The combination of results from different domains and applications, would make it possible to create generalized conceptual models for predicting changes in tolerance. Our research represents an important step in identifying such relationships, and indicating the need for technology to meet user requirements.

CONCLUSIONS

This study was designed to investigate users' requirements for Internet QoS. We specified a set of objective thresholds that reflect users' subjective assessments of quality. We showed that:

- The task in which users are engaged, the length of time they have been interacting with a site, and the method of page loading affects the acceptability of QoS.
- Tolerance of delay is influenced by users' conceptual models of how the system works.
- Poor Web site performance leads to poor company image and often compromises users' conceptions of the security of the site.

There are several stakeholders in the design of Internet services: server designers, network providers, advertisers, companies whose products are sold on-line, and consumers themselves. A failure to understand users' on-line QoS requirements may effect users' conception of a company's stature and commercial viability which, in turn, affects the business interests of service providers and advertisers. The future Internet will have more users and support a greater diversity of Internet applications. It has the potential to change the way in which consumers interact with companies. Our research has shown that it is possible to integrate users' requirements into systems design. Only through such integration will it be possible to achieve the customer satisfaction that leads to the success of any commercial system.

ACKNOWLEDGEMENTS

We would like to thank management at HP Labs and members of HP labs for taking part in pilot trials. We thank Angela Sasse (UCL Computer Science).

REFERENCES

1. Abdelzaher, T., and Bhatti, N. Adaptive content delivery for web server QoS. *Proceedings of IWQoS'99* (London, May 1999).
2. Bernet, Y. A Framework for End-to-End QoS Combining RSVP/Intserv and Differentiated Services. *IETF* (March 1998).
3. Bhatti, N., and Friedrich, R. Web Server Support for Tiered Services. *IEEE Network* (September/October 1999).
4. Bouch, A., and Sasse, M.A. Network QoS: What do users need. *Proceedings of IDC'99* (Madrid, September 1999).
5. Bouch, A., and Sasse, M.A. It ain't what you charge it's the way that you do it: A user perspective of network QoS and pricing. *Proceedings of IM'99* (Boston MA, May 1999).
6. Boyer, D.L., Pollack, J.G., and Eggemeier, T.F. Effects of Aging on Subjective Workload and Performance: Determinants of Age Differences in Cognitive Performance. *Proceedings of Human Factors Society 36th Annual Meeting 1*, (1992).
7. Cullinane, P. Ready, set, crash. *Telephony 3*, (Nov 1998).
8. Dunlop, M.D., and Johnson, C. Subjectivity and notions of time and value in interactive information retrieval. *Interacting with Computers 10*, 1, (1998).
9. ETEWatch for Web Browsers. http://www.candle.com/etewatch.
10. Fox, R. News Track. *Communications of the ACM* (May 1999), 9-10.
11. Fishburn, P.C., and Odlyzko, O.M. Dynamic behavior of differential pricing and Quality of Service options for the Internet. *Proceedings of ICE-98*, ACM Press.
12. Gupta, A., Stahl, D.O., and Whinston, A.B. Pricing of services on the Internet. http://cism.bus.utexas.edu/alok/pricing.html.
13. Hogan, M. The first ever report on the top 100 e-commerce businesses and the secrets of their success. *PC Computing Magazine*, June 8, 1999.
14. HP Shopping Village. http://www.shopping.hp.com.
15. Keller, J.J. Ex-MFs managers plan to build global network based on Internet. *Wall Street Journal*. January 20, 1998.
16. Mackie-Mason, J.K., and Varian, H. Economic FAQs about the Internet. In McKnight, L.W., and Bailey, J.P. (eds.). *Internet Economics*. MIT Press, 1997.
17. Morgan, K., Morris, R.L., Macleod, H., and Gibbs, S. Gender Differences and Cognitive Style in Human-Computer Interaction. *Proceedings of EWHCI'92* (Moscow, Russia, August1992).
18. Myers, B. A. The Importance of Percent-Done Progress Indicators for Computer-Human Interfaces. *Proceedings of CHI'85* (San Francisco CA, April 1985).
19. Network Reliability Steering Committee. *Annual Report 1998*. http://www.nric.org.
20. Nielson, J. *Usability Engineering*. AP Professional Press, Boston MA, 1994.
21. Odlyzko, O.M. The Internet and other networks: Utilization rates and their implications. http://www.research.att.com/~amo.
22. Perry, C. Travelers on the Internet: A survey of Internet users. *Online 19*, 2, (1995).
23. Ramsay, J., Barbesi, A., and Preece, J. Psychological Investigation of Long Retrieval Times on the World Wide Web. *Interacting with Computers 10*, 1, (1998).
24. The Internet, Technology 1999, Analysis and Forecast, *IEEE Spectrum* (January 1999).