

Not all bits have equal value: Investigating users' network QoS requirements

A. Bouch, and M.A. Sasse
Department of Computer Science,
University College London,
London WC1E 6BT

ABSTRACT

The number of Internet users is expected to triple between 1998 and 2002^[1] largely because of new applications (such as videoconferencing) and new services (such as e-commerce). This shift in usage imposes higher Quality of Service (QoS) requirements at different levels of granularity. It also means that the traditional Internet way of managing quality (best-effort) has to be replaced by a more service-oriented approach. The aim of this paper is to investigate end-users' cognitive and perceptive QoS requirements. We present empirical results on user QoS preferences and QoS graduations. Guidelines for translating these results into metrics that can be used to guide resource allocation mechanisms are discussed.

Keywords: Quality of Service, pricing, User-centered approach, Human-Computer Interaction

1. INTRODUCTION

Shared networks can now support many applications, ranging from data-driven applications, such as email, to real-time applications, such as multimedia conferencing. While a vision of future networks offers the potential to break traditional barriers in communications and commerce, the current service quality to users is often unacceptable^[2]. In the networking community, different technical Quality of Service (QoS) requirements have been ascribed to applications, however, traditional QoS metrics such as response time and delay no longer suffice to fully describe quality of service as perceived by users. The network community often assumes there is a straightforward relationship between objective quality, measurable at the network level (e.g. delay, jitter) and users' subjective perceptions of quality. This means that conditions such as level or fluctuation of users' demand are assumed prior to testing. Recent work with users suggests that the relationship is much more complex. For example, users may judge a relatively fast service to be unacceptable unless it is also predictable, visually appealing and reliable^[3].

The best-effort service of the Internet can no longer deliver acceptable levels of QoS to many users^[2]. Traffic produced from different applications, or service classes, can be characterized through an associated payment^{[4][5][6]}. Resource allocation schemes allocate service resources according to the assumed objective QoS requirements of applications. However, it is currently not known to what extent changes in objective QoS metrics are perceived by, and impact the behavior of users. The salience of QoS dimensions is determined, not by the fact that they are technically implementable, but by their semantic value. In networked multimedia applications in particular, variations in quality at the network level are not directly linked to the subjective assessment of quality received by the user, the latter depending on the nature of that user's task^[7].

The aim of this research is to define users' requirements for network QoS, and investigate the relationship between users' assessment of QoS in different contexts, and the conceptual models that motivate that assessment. The ultimate aim of the research was to provide models that can be used to predict users' demand for QoS in priced and unpriced situations. These models should aid the integration of users' requirements for QoS into systems design, and therefore have immediate practical benefits for systems designers and users themselves. The model is supported with results from a set of experiments. Section 2 of this paper describes the empirical work. Section 3 looks at the generation of conceptual models. In Section 4 we discuss guidelines and recommendations for resource allocation mechanisms, and how users' requirements might be translated into lower-level metrics. Section 5 presents the conclusions drawn from the work.

2. EMPIRICAL WORK

2.1 Aims

Our research was structured in 3 key stages:

1. Capture of users' evaluations of objective QoS in different contexts.
2. Construction of models that reflect users' requirements for network QoS and charging mechanisms.
3. Integration of findings from previous key stages with technical and economic perspectives.

To meet these aims we used both qualitative and quantitative research methods. The predictive validity of factors identified in qualitative studies is limited if they cannot be shown to affect behaviour during empirical work. Empirical work is also needed to establish the ability of factors to be generalised across larger numbers of users, thus providing scalable solutions problems involving user requirements capture. This paper reports a set of studies that test users' responses to different levels of QoS, and their preferences for different charging mechanisms (see Table 1).

Table 1. Summary of studies

Study	Participants	Set-up
1	25	Quality evaluation
2	24	Quality control against budget
3	30	Quality control against budget/network information
4	30	Informal discussion groups

2.2 Studies 1 & 2: What is valuable quality?

Most published proposals for partitioned networks assume users' assessments of the quality they receive mirrors the objective quality delivered at the network level (measurable through characteristics such as packet loss and delay). In Studies 1 and 2 users' QoS requirements for interactive audio were investigated to demonstrate that these assumptions may not be correct.

2.2.1 Study 1

Different loss patterns were generated on audio files, each particular level of loss being stable for 17 seconds. This time interval was chosen so that the experiment could address a range of loss rates whilst avoiding participants becoming over-familiarized with any particular rate. To reflect users' opinions of dynamically changeable media quality accurately, it is necessary to gather quality ratings in a continuous fashion. A recording slider, the QQuality ASSESSment Slider, *QUASS*, (Figure 1) was developed for this purpose^[8]. An adapted version of the *SERVQUAL* analysis^[9] was administered to see if there was any difference between participants' expectancies of the audio quality they were to receive. This form of analysis, applied successfully in many market research contexts, requires respondents to rate their expectancies of future quality, using an interval scale. Participants were told that the experiment was designed to establish when audio quality becomes unacceptable. All participants heard 4 passages of speech of 2 minutes each, in 4 corresponding conditions. Participants were asked to use *QUASS* to continuously assess the variable objective quality they received. The audio packet loss generated on these narratives reached a maximum of 20% and followed a random pattern.




Figure 1. QUASS (budget)

Results indicate that, when participants expect a low level of audio quality, subsequent ratings are comparatively high; similarly, when participants expect a high level of audio quality, subsequent ratings are comparatively low. This trend is illustrated in Figure 2 and Figure 3, which compare the scores obtained for individual participants whose expectancy ratings differed in the *SERVQUAL* analysis.

These results show that participants base their dynamic assessment of audio quality on an assessment of future quality. The evolving nature of this assessment is shown in the fact that the correlation between initial expectancies and dynamic ratings becoming statistically insignificant as the experiment progresses. This indicates that participants alter their expectancies based on the pattern of loss received during their interaction with the system. How accurate users' expectancies may turn out to be depends on the extent to which the audio quality received is predictable at the time when such expectancies were formed.




Figure 2. Dynamic QoS ratings by Expectancy: A




Figure 3. Dynamic QoS ratings by Expectancy: B

To investigate the impact of predictability upon users' assessment of media quality, it is necessary to investigate the dynamic variation of ratings for subjective audio quality in more detail. This requires analysing the pattern - in addition to the magnitude - of ratings given during the experiment. To examine the role of predictability in users' audio quality assessment, two figures were extracted from the subjective ratings, for each interval of audio. An initial response, just after a change in objective quality, and a second response, just before the next change in objective quality. The objective quality is therefore stable between these two measurements. It was hypothesised that, if predictability is an important determinant of users' subjective assessment, such ratings should rise during the time when the objective loss rate remains stable. Indeed, such a result would suggest that predictability is a more salient determinant of perceived quality than the absolute objective loss rate received.

Within each condition, the average subjective ratings given by participants correlate with the objective loss received at any one time. However, perhaps the most important finding is that the overall pattern of loss throughout a passage has more influence on participants' subjective quality ratings than the absolute amount of objective loss. This is shown by the fact that a subjective quality rating given for a high objective loss rate can be higher than a rating given for a comparatively low objective loss rate. Table 2 illustrates average subjective quality ratings, for various objective loss rates, amongst conditions. The figures in Table 2 illustrate that widely different ratings are given for identical loss rates.

Table 2. Subjective QoS

Loss rate	Subjective QoS rating		
	0	10	15
Condition			
1	63, 65	68, 39	35
2	69, 89	65, 40	41
3	50	45, 25	39
4	-	38	38




Figure 4. The Predictability Effect: Condition 1

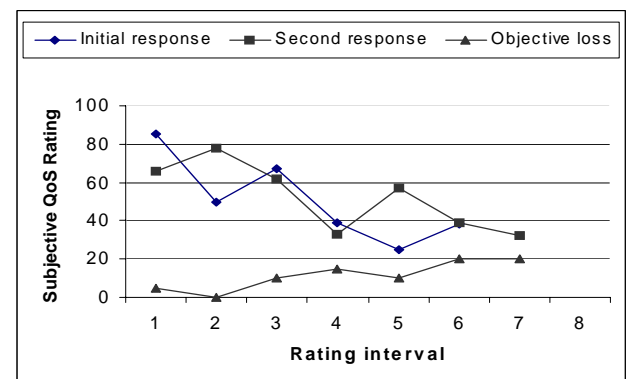


Figure 5. The Predictability Effect: Condition 3

The presence of the *predictability* effect, as indicated by these results, means that an identical objective loss rate was rated as being of a higher quality after an interval of time. For example, ratings given for a loss rate of 10%, whilst being initially lower than those given for 0% objective loss, rose within the 17 second interval to a higher level than those given for 0% objective loss. This suggests the importance of stability in the quality received, as opposed to its absolute magnitude. Figures 4 and 5 illustrate the *predictability* effect.

2.2.2 Study 2

In this experiment users' QoS requirements for interactive audio were investigated. In the experiment the QoS received was linked to an expendable resource. We also established participant's attitudes to dynamic pricing during semi-structured interviews.

Participants used QUASS to control the quality they receive in a dynamic fashion. To encourage participants to interact with the slider, QUASS was set to decrease its slider position - and the corresponding quality received by the user - by 0.1% per second. This budget diminishes in proportion to the magnitude of quality requested. No automatic decrease in quality was configured for conditions that contained a budget.

Participants were able to control the volume of the audio via the Robust Audio Tool (RAT)^[10]. The video tool *vic*^[11] was run to set the task in a representative multimedia environment. The quality of speech was degraded via the *Forwarder* software, which drops the required number of audio packets before forwarding them. 25 participants took part in all conditions of the study and 2 different configurations of loss were used in the experiment:

- *Stable* loss: In this condition the position of the slider corresponded to a certain level of packet loss – e.g. a slider position of 0 would configure a loss rate of 50%.
- *Variable* loss: In this condition, the position of the slider corresponded to a range of loss values. For example, a slider position of 0 would result in any value between the values of 0% and 50% loss, chosen by a random function.

All participants took part in 4 conditions:

1. *No-budget, stable* loss rate.
2. *No-budget, variable* loss rate.
3. *Budget, stable* loss rate.
4. *Budget, variable* loss rate.

In each condition users were required to play a word guessing-game in collaboration with the co-experimenter, where the participant would take turns in describing/guessing a word. During 5 minutes participants were asked to use QUASS to control the quality that they were receiving.

In both the *budget* and *no-budget* conditions, far less quality was required for the *stable* loss condition than for the *variable* loss condition. A comparison of maximum loss levels, averaged amongst participants in Conditions 1 and 2 is shown in Figure 6. These results show that users are willing to tolerate greater objective loss rates when those loss rates are relatively predictable. Most participants who noticed a difference in quality between conditions, said that this related to the type of quality configured in the budget conditions only.




Figure 6. Mean maximum loss rates (budget)

These findings suggest that participants used UI feedback in post-interaction assessments of quality, rather than a direct assessment of the objective levels of loss received. Since users' QoS expectations are developed with reference to their assessment of preceding QoS received, this finding may have important implications for users' future quality judgements. Results from Studies 1 and 2 suggest that, contrary to commonly made assumptions, there is no direct association between objective loss levels and subjective QoS judgements. Indeed, it is apparent from our findings that it may be possible to relieve the network infrastructure of some of the responsibility for providing network performance feedback to the user.

Manipulation of the relatively limited amount of information provided by the network can be performed at the application and UI level.

2.3 Study 3: Linking value to different pricing mechanisms

In Study 3 we asked: *what type and frequency of network feedback to users require when interacting with a priced network?*

The experiment involved listening to a recording of a 10-minute interview in which an actor played the role of a candidate who was interviewed for a university place. During the experiment users were asked to manipulate audio quality (packet loss), using the QUASS slider. To situate the task in a realistic context, video was streamed at maximum quality. Participants were either told that tasks involve a measure of task completion (*measured* tasks), or that no measure is required (*unmeasured* tasks). This distinction is made to manipulate the importance of the task. In the measured scenario, participants were asked to answer specific questions concerning the candidate’s responses in the interview. The experiment uses software that allows users to choose the levels audio quality they receive. The software allocates a budget to users prior to interaction. Participants are given panel of three menus. The menus are labeled *Current Values*, *Set Preferences* and *Future Predictions*. Under *Current Values* users can view statistics of current network conditions, price and their remaining budget. Under *Set Preferences* users can opt to keeps the audio quality at the current level or, to keep current price the maximum: selecting this means that the participant won’t pay more than is currently being paid. The *Future Predictions* option allows users to view predictions of future network quality and price. An integral part of the UI is the Your Budget interface. What this shows depends on whether it is selected from the *Future Predictions* (Figure 7) or the *Current Values* (Figure 8) menu. On selecting from the *Future Predictions* menu the display will show the Battery Life. This is a prediction of the number of minutes the participant can carry on before their money runs out, if the current level of quality continues to be requested.




Figure 7. Your Budget showing Battery Life




Figure 8. Your Budget interface showing Current Budget

Results show that a number of intervening factors influence whether users consider a certain price appropriate for a particular amount of audio quality. Table 3 presents a summary of results for each experimental hypothesis.

Table 3. Results for each experimental hypothesis

Hypothesis	Result
H1	High task importance = more Quality profile requests
H2	Low task importance = more Price profile requests
H3	Profile selected = lower quality tolerated
H4	Low budget = Price profile prioritized over Quality profile
H5	Future statistics selected more frequently than current statistics
H6	Battery life selected most frequently

Hypothesis *H1* was confirmed by the results of the study. Participants were more likely to select a profile where QoS has priority over price when they were involved in an important task. The frequency with which users selected a QoS profile over a price profile was compared between a measured and unmeasured task scenario.

- *Users are more likely to select a profile where QoS has priority over price when they are involved in a relatively important task.*

Hypothesis *H2* was not confirmed by the results of the study. This means that:

- Users are not more likely to select a profile where price has priority over QoS when they are involved in a relatively unimportant task.

Figure 9 shows the feedback required by participants for both measured and unmeasured tasks.




Figure 9. Number of participants selecting each feedback option

Hypotheses *H3* and *H4* concerned the levels of quality users would accept depending on the amount of control they were given over that quality. For each participant who selected a QoS profile, the lowest level of packet loss was compared to the lowest level of packet loss accepted under manual control. Statistical tests reveal that *H3* and *H4* can be confirmed by the results of the study. This means that:

- Users will tolerate lower levels of quality if they select a profile that controls for them. This result is statistically significant.
- Users are more likely to prioritize price over quality when their budget is relatively low.

Hypotheses *H5* and *H6* were confirmed by the results of the study. This means that:

- Users request feedback concerning future statistics more frequently than feedback concerning current statistics.
- Users find information concerning Battery Life the most important type of feedback.

These results can be seen in Figure 9. Statistical analysis shows that Battery Life is selected a significantly higher number of times than other options.

3. MODEL GENERATION

3.1 Study 4: Qualitative methods

Qualitative data is needed to explore users' motivations and conceptual models behind metrics derived from experiments. We used focus group discussions and grounded theory methods^[12] to generate qualitative data. Users were asked to consider different charging scenarios. Figure 10 is a simplified version of the model generated from qualitative data.

The introduction of an explicit payment for a network service adds, from the users' point of view, a dimension of value to perceived QoS. Confidence is gained through users' assessment of a situation as low Risk, with Risk defined as the chance of paying too much for the QoS received.

To assess Risk, users consider several sub-concepts; the relevance of the different sub-concepts depends on users' level of knowledge and experience. Users' Expectations of quality received, for instance, is influenced by the levels of congestion advanced users believes to be present. The risk of receiving poor performance from the network is dependent on the user's goal, and therefore:

$$\text{Risk assessment} = \text{Task importance} + \text{Feedback concerning predicted performance}$$

Predictability is one of the most fundamental QoS drivers from a user's perspective. The concept of Risk is directly linked with that of Predictability - a low Risk situation is one that is predictable. Data suggests that users are prepared to accept charging mechanisms that are predictable. It is important for users to be able to predict different technical QoS drivers depending on the task that they are performing. For example, for email applications it makes less sense to talk of Predictability in terms of time; email is usually an elastic application. For time-dependent WWW tasks, however, it makes far more sense to attribute Predictability as a measure of QoS that has to be reassessed. A measure of traffic type may be therefore necessary in deriving service charges.

Predictability is also implied in the level of control required by users over their payments for QoS. Findings from our studies suggest that users prefer to be able to dynamically change the levels of QoS they receive in line with the value given to the Task being performed. Therefore, dynamic pricing needs to provide Feedback on network congestion, which would enable users to predict the Risk involved in making certain payments.

An idea that balances the need for Predictability with the need to attribute value to a particular transaction is the use of quotas. Quotas for certain important QoS drivers can be bought prior to network use. Users would then have a (dynamic) choice whether to use their quota for any particular transaction. It is evident that this mechanism involves a trade-off between the user's need to be able to predict the network's performance and the need for the dynamic evaluation of network conditions. The need for dynamic decision making is shown in the fact that decisions of whether to use quotas are dependent, not only on the type of task undertaken, but also on the context in which the task is performed. For example, the urgency associated with file download is a critical factor in determining whether the quota is used.

4. DESIGN GUIDELINES

4.1 Recommendations

Based on the results reported in this paper it is possible to make certain recommendations for the design of a network resource allocation system:

1. **Predictable – i.e. consistent – QoS is essential.** Given that lower – but consistent – QoS is rated higher than higher – but variable – QoS, it may be appropriate to intentionally lower the quality delivered to the user to a level that can be maintained. We are not suggesting that the network performance should be degraded; rather, quality should be regulated at the application level – e.g. by buffering media quality. Our findings have consequences for the configuration of dynamically adaptive applications (e.g.^[13]). The quality delivered by these applications changes dynamically in line with network conditions. Whilst this mechanism is representative of the fluctuating nature of network quality it is essential that those applications employ traffic shaping at the application level in order to provide a consistent service.
2. **Feedback is essential for predicting QoS.** Network feedback should enable users prediction of future QoS. We have shown the extent to which users depend on the feedback they receive from the UI. Indeed, it is arguably this feedback that ensures that the QoS remains predictable. In situations of high *Risk*, users should be provided with the option of configuring feedback dynamically.
3. **Consider differentiated service schemes.** Much current debate in the networking community has centred around the costs and benefits of providing differentiated versus integrated services^{[14][15]}. We have shown that, potentially, users – or the applications that represent their preferences – require feedback concerning future quality in order to make accurate assessments of that quality. This entails that a feasible service scheme must abstract much of the complexity

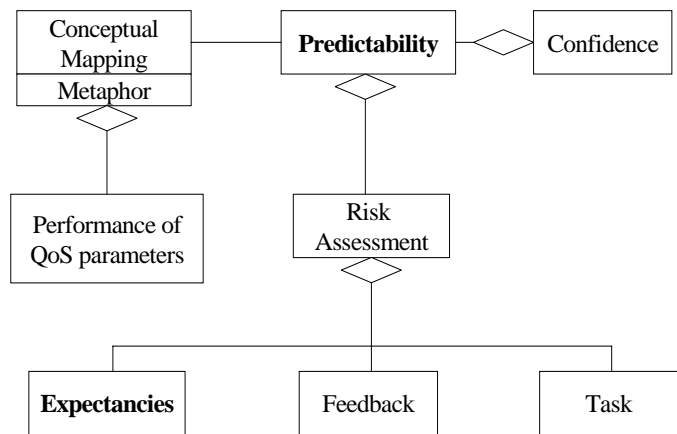


Figure 10. Simplified conceptual model for priced QoS

involved in maintaining information about network congestion away from the core of the network. The provision of differentiated service mechanisms is therefore preferable to the provision of integrated services as the latter involves the maintenance of state within the routers themselves, and may not be sufficiently scalable.

4. **Feedback requirements are task-specific.** *Predictability* may be a concept that applies across many tasks – predictable quality is a requirement that can be applied to virtually any networking application. *Strategy Formation*, however, is clearly task dependent: only within the context of the particular speaking and listening task was it possible for participants to formulate the strategy described. This suggests that there is unlikely to be a generic set of required performance characteristics, or strategies, for all types of networking task. The research reported in this paper therefore needs to be extended to investigate how user requirements for QoS and UI feedback relate to other tasks, and how these requirements vary according to the value placed by users on task performance.
5. **Dynamic re-evaluations of quality are not required for all interactions.** Users should be able to configure default preferences with regard to payment. Our results show that, optimally, users should be able to determine the influence future QoS requests may have on a finite resource. As with pricing feedback, the actual price should be set within a specified range. Only when situations of high cost, or where there is high risk that delivered QoS will not meet the specified range, should the network ask for adjustments in payment. In accordance with user and task requirements, an application might request *Feedback* from the network only if a specified amount of change occurred to the internal network state. This implies the need to configure agent software at the application level. Currently suggested software that fits the flexible needs of users may act as a ‘QoS-Broker’, thus encapsulating much of the complexity of dealing with variable network congestion within an automated process^[16].
6. **QoS requirements can be classified according to task characteristics.** This suggests that profile-based pricing schemes should be used. On a technical level, the encapsulation of default preferences in the application is perhaps best served through the use of profile-based pricing schemes that have the ability to interact with adaptive software at the application level. Not only does this approach alleviate the computational load on the network, but also affords heterogeneous task requirements to be represented at an appropriate level of flexibility. Users after all, regard the network as a tool that enables them to perform particular real-world tasks. The task is therefore users’ desired focus, rather than the price of interaction.
7. **Charge for an entire interaction.** Users’ evaluations of expenditure suggest that they make dynamic assessments concerning the value of an interaction, as a whole, against their assessment of its price, meaning that users’ think about QoS as attached to an entire interaction. Charging schemes might therefore be designed to attach or estimate the cost of an interaction before that interaction starts, in order to provide accurate data to fit users’ expectations. Users’ assessments of value are therefore reflected by the Commitment they make to reach the goal of their task. This task is represented by the interaction as a whole.

4.2 Translating user requirements into metrics

To bridge the gap between user requirements that are more easily expressed in higher level conceptual terms, and the language in which system requirements are expressed is not trivial. The expression of subjective experience in terms of magnitudes means that much of the semantic character that defined those requirements in the first instance is lost. However, from a design point of view, user requirements that cannot provide guideline metrics for resource allocation algorithms are of limited benefit. We can express subjective opinions in a series of relationships between variables. This involves producing econometric equations. How is it possible to derive econometric equations from data captured in this research when such data represents disaggregated conditional variables and when parameters to traditional econometric equations use aggregate variables? Clearly no direct substitution can be made. However, it is possible to show:

- That top-level aggregate conditional variables discovered in this research are essential components to econometric equations.
- How sub-concepts such as Expectations interact with other concepts to produce aggregate effects: For example, how the magnitude of one variable changes in relation to another.

Subjective concepts can be translated into objective metrics through experimental data. Dependencies and causal relationships can therefore be defined. For example, from a study into users’ tolerance for Web-page latency^[17] we discovered that users become less tolerant as the duration of the session increases. The phenomenon of duration causing

users to be more critical of performance can be expressed as a function that takes into account the duration of the session. Here we express the *utility* of a session of requests as a number between 1 and -1, where *utility* >0 indicates that acceptable performance thresholds have been exceeded, *utility* = 0 indicates that the service is exactly acceptable, *utility* <0 is unacceptable performance. Total utility of a session of length *N* can be expressed as:

$$Utility = \frac{\sum_{i=1}^N \frac{threshold(i) - latency(i)}{threshold(i)}}{N}$$

threshold(i) = threshold of acceptability of access *i*
latency(i) = delay in completion of page *i*
N = length of session

The definition of equations allows us to describe users behaviour in response to their requirements. Experimental work is then needed to pinpoint the key magnitudes that describe the percentage of users lost at what threshold of low quality. Likewise, the work with audio packet loss, reported in this paper, shows that there are thresholds where most users find the quality intolerable, and where increasing the quality yields diminishing returns.

5. CONCLUSIONS

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 work reported in this paper has shown that a vision of future network service must be based on an old principle, that economies are ultimately service-driven. As Peter Drucker^[18] points out:

“Quality in a product or service is not what the supplier puts in. It is what the customer gets out and is willing to pay for. [...] Customers pay only for what is of use to them and gives them value. Nothing else constitutes quality.”

This paper has addressed the question of what users actually need from network QoS. Our results show that this is not simply the ability of the network to provide a higher, level of QoS. Rather users need a predictable level of QoS that allows them to make accurate value judgements about the quality that they receive. Our findings show that it is not safe to assume a correlation between objective levels of QoS and users’ subjective evaluations of that QoS. Rather, pricing mechanisms should must address connections between the *Value* ascribed to certain levels of quality, and the *representation* of that quality to users. Our results demonstrate that objective levels of quality are less important to users than receiving the expected amount of quality that enables them to perform the chosen task. Furthermore, the ability of the network to provide users with an appropriate amount and type of *Feedback* about the media quality they are likely to receive leads to acceptable level of *Predictability*.

The complexity of the models generated from this research suggest that there is no simple experiment that can test the efficacy of all the established factors in one study. However, there are a number of areas where factors are linked to the extent that empirical work can be designed to test their influence on users’ behaviour. Distinctions between the type of QoS ascribed to different tasks have been established. It is therefore essential that the factors derived in the models be tested for validity while users are performing different tasks. The effect of users’ tasks on these concepts has been investigated during experiments by assessing users’ QoS tolerance for Email, FTP and Web interactions in addition to multimedia conferencing applications^[19]. These experiments confirmed the salience of many of the factors described in by the models. This means that the models can be used to predict users’ demands for QoS in particular contexts. The nature and level of these demands having been discovered, technical implementations can be designed to provide subjectively optimal resource allocation mechanisms.

REFERENCES

1. P. Cullinane, "Ready, set, crash", *Telephony* **3**, pp. 3-13, 1998.
2. J. Bruhns, "A Short Commentary on Usability in Northern Europe. Still a Long Way to Go". Available at <http://www.webword.com/moving/global0002.html> (2001).

3. J. Ramsay, A. Barbesei, and J. Preece, "Psychological Investigation of Long Retrieval Times on the World Wide Web", *Interacting with Computers* **10**, 1, pp. 77-86, 1998.
4. R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, "Pricing in computer networks: Motivation, formulation and example", *Transactions on Networking* **1**, pp. 614-627, 1993.
5. J.K. Mackie-Mason, and H.R. Varian, "Economic FAQs about the Internet", *Journal of Economic Perspectives* **8**, pp. 75-96, 1994.
6. D. Clark, "Internet cost allocation and pricing", in *Internet Economics*, L.W. McKnight, and J.P. Bailey, eds., MIT Press, 1997.
7. A. Watson, "Evaluating audio and video quality in low-cost multimedia conferencing systems", *Interacting with Computers* **8**, 3, pp. 255-275, 1996
8. A. Bouch, A.. Watson, and M.A. Sasse, "QUASS-A tool for measuring the subjective quality of real-time multimedia audio and video", *Proceedings HCI'98*, pp. 94-96, 1998.
9. A. Parasuraman, L.L. Berry, L.L. and V.A. Zeithaml, "SERVQUAL: A multiple-item scale for measuring consumer perceptions of service quality", *Journal of Retailing* **64**, 3, pp. 514-532, 1998.
10. V. Hardman, M.A. Sasse, and I. Kouvelas, "Successful multi-party audio communication over the Internet", *Communications of the ACM* **41**, (5), pp.74-80, 1998.
11. S. McCanne, and V. Jacobson, "vic: A flexible framework for packet video", *Proceedings of ACM Multimedia*, pp. 511-522, 1995.
12. A. Strauss, and J. Corbin, "Basics of qualitative research: Grounded theory procedures and techniques". Sage, 1990.
13. J. Gecsei, "Adaptation in distributed multimedia systems", *IEEE Multimedia*, April-June, pp.58-66, 1997.
14. Z. Wang, "User-Share Differentiation (USD): Scalable bandwidth for differentiated services," Available at: <http://www.globecom.net/ietf/draft/draft-wang-diff-serv-usd-00.html>, 1998.
15. A. Anurag, J.K. Muppala, and H.G. DeMeer, "VoIP support on Differentiated Services using Expedited Forwarding. *Proceedings of IPCCC-2000*, pp. 574-580, 2000.
16. K. Nahrstedt, and J.M. Smith, "The QoS Broker," *IEEE Multimedia* **2**, No.1, pp.53-67, 1995.
17. A. Bouch, N. Bhatti, and A.J. Kuchinsky, "Quality is in the eye of the beholder: Meeting users' requirements for Internet Quality of Service," *Proceedings of CHI'2000*, pp. 297-304, 2000.
18. P. Drucker, *Management*, Butterworth-Heinemann, 1999.
19. A. Bouch, M.A. Sasse, and H.G. DeMeer, "Of packets and people: A user-centred approach to Quality of Service," *Proceedings of IWQoS'00*, pp. 189-197, 2000.