NETWORK QUALITY OF SERVICE:

WHAT DO *USERS* NEED?

Anna Bouch
University College London
UK

M. Angela Sasse
University College London
UK

ABSTRACT

The number of heterogeneous networked applications is constantly increasing. It is likely that network resources will have to be partitioned according to the different Quality of Service (QoS) demands made by the users of these applications. One way of implementing a partitioned network – favored in technical literature – is in terms of quality-based pricing (e.g.[1],[2],[3]). 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). It is also assumed that users are prepared to pay more for higher levels of objective QoS when they need it. In this paper, we demonstrate that these assumptions may not be correct. We report an experiment in which users' QoS requirements for interactive audio were investigated. During the experiment the QoS received was linked to an expendable resource. We also established participant's attitudes to dynamic pricing during semi-structured interviews. Our results suggest that users' assessments of the value of QoS received is influenced by a number of different factors – hence, the same level of quality will receive different assessments in different circumstances. We also found that a *Predictable* level of QoS is a key requirement for users of networked applications. The results from the interviews suggest that, when users accept a pricing scheme they develop strategies to optimize their use of an expendable resource. We conclude by outlining the implications of our findings for the design of resource allocation and pricing schemes for the future Internet.

KEYWORDS:

QUALITY OF SERVICE, USER REQUIREMENTS, PRICING, PREDICTABILITY

1. INTRODUCTION

1.1 The challenge for network designers

The Internet offers the potential to break traditional barriers in communications and commerce, and change the way in which we work [4]. There may be several reasons for the increasing growth in the use of data networks [5]. Popular argument suggests that an increase in the efficiency of data transport is responsible [6]. The increase in network use can, however, not solely be explained by improvements within the network infrastructure. Whilst technological considerations may have helped realize the potential uses of networks, it is arguably the *value* of what can be done through technology that drives users' demand for network services. Applications such as the World Wide Web and real-time multimedia conferencing have attracted new users to the Internet. In order to keep these users, a sufficient level of Quality of Service (QoS) to sustain those applications, has to be delivered.

Therefore, one of the major challenges for network designers today is how to manage the increasing demand for network resources. Shared networks are able to support a wide range of applications with varying performance demands. The increasing support for multimedia applications has added a further complex dimension to the problem of managing network resources in a dynamic manner. The introduction of dynamic or usage-based pricing is often seen as the only fair way of partitioning a finite resource. A number of different pricing mechanisms have been proposed, from bidding for access [7] to quota systems [8]. Most of these suggestions are based on existing technical mechanisms for partitioning and policing network resources in a fair manner. Perhaps as a consequence of this focus, it is largely assumed that users' willingness to pay for the quality they receive depends solely on the objective levels of QoS measurable at the network level, such as the levels of packet loss.

Previous research has found that purely technical specifications of pricing schemes fail to adequately measure user satisfaction. A study into responsive pricing [9], showed that including a measure of user requirements for QoS increases user satisfaction *and* network performance. Thus, understanding what the user requirements are is a prerequisite for designing a system that will perform efficiently as a whole, in the real world.

The subjective QoS dimensions that describe user requirements have not been fully established. Previous research suggests that a number of interconnected factors influence users' requirements for QoS. Demands for a certain type and level of network performance have been shown to vary widely, depending on the task users undertake [10]. For example, requirements for high video performance are more prominent in interactive tele-teaching tasks than in listening to lectures [11]. A consistent finding is that QoS received by users should concur with their expectancies [11], [12]. Furthermore, it has been established that large quality variations should be avoided for audio transmission [13].

1.2 Developing conceptual models

Results from experiments that included 'user objectives' in the network control loop allowed [9] to argue that providing users with pricing feedback reveals the value that users place on QoS. This measure can then be used by network designers to predict the magnitude of user-demand. However, it is currently not known how factors such as different types and levels of pricing feedback affect users' perceptions of the value of a resource, and the requests for quality they make in response to those perceptions.

As a result of our research into the factors that influence users' assessment of QoS, we have developed a set of models. These models describe how certain factors impact users' assessment of QoS, and provide a framework for predicting users' responses to the type and level of QoS they receive. Furthermore, the models can be used to predict users' reactions to different pricing schemes. Figure 1 presents a simplified version of the model that shows how one concept is influenced by another. For example, it shows that *Predictability* is a high-level QoS parameter of crucial importance in users' perceptions of the QoS they receive. *Predictability* leads to *Confidence* in the pricing scheme responsible for allocating QoS resources [12]. We established the predictive validly of some of the concepts contained in the models by conducting a series of empirical investigations [14]. However, these results were obtained with a passive listening task and therefore should not be taken as a *direct* measure of users' quality requirements for other – interactive – tasks.

This paper reports results from a set of experiments and semi-structured interviews. The aim of the studies was to investigate the concepts important to users when

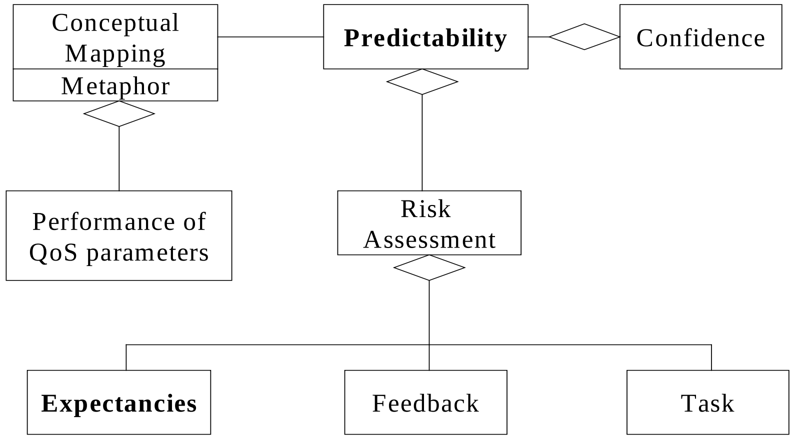

Figure 1: Conceptual model (simplified)

deciding to pay for the QoS they receive. The studies explore the assumption that there is a *direct* correlation between objective QoS and subjective evaluations of that QoS. The experiment was also designed to validate the central role of the concept of *Predictability* in users' requests for QoS in both a priced and non-priced situation and how users' expectancies of future levels of QoS influence their subsequent judgements of that QoS.

A secondary aim of the studies was to explore the influence of feedback from a simple pricing scheme on the level of quality requested by participants, and their attitudes to dynamic pricing after the event. Since audio quality is the most important determinant of subjective quality in real-time tasks such as multi-media conferencing [15], assessment was focused on audio quality.

The following section describes the experimental method, and section 3 provides a detailed discussion of the results. In section 4 we conclude that users' assessments of the QoS received, and their attitudes to pricing mechanisms are influenced by a number of potentially complex factors. Although future work is needed to determine user requirements for QoS and pricing schemes under different tasks, our results clearly demonstrate the importance of maintaining a predictable level of QoS under any pricing scheme. The results also show that making this provision need not place additional overhead on the internal operations of the network but is, in part, achievable by the configuration of pricing feedback at the UI.

2. METHOD

2.1 Tools

To gain reliable information about required levels of QoS, it is necessary to allow participants to control the quality they receive in a dynamic fashion. The QUASS QUality ASsessment Slider (QUASS) [16] has been adapted to work in conjunction with network-level software to allow users to control the levels of 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. The slider can also be configured to display a budget (Figure 2). 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) [17]. The video tool *vic* [18] 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.

2.2 Experimental set-up

The configuration of the systems set up for the experiment is shown in Figure 3. Computer A ran RAT and QUASS, and was connected to Computer B on the same port number. On Computer B, RAT was started and connected to Computer C, which ran the Forwarder software. During the experiment audio packets generated by users on Computer A affected the audio packets that were forwarded from the co-experimenter (on Computer B). Since speech from Computer A to B did not pass through the Forwarder, the co-experimenter did not receive degraded audio at any time.

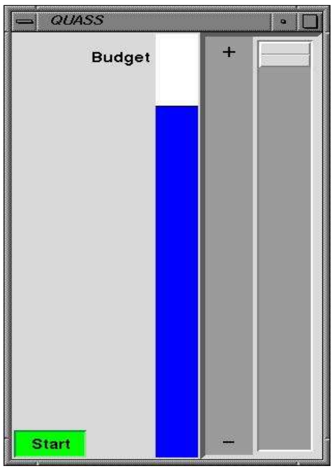

Figure 2: QUASS (priced)

2.3 Participants

25 participants took part in all conditions of the study. Participants were selected according to the following criteria:

- They should be from both technical and non-technical academic backgrounds.
- They should possess experience with using networks (e.g. Internet usage).
- They should have limited experience with real-time audio applications.

2.4 Procedure

2.4.1 Experimental conditions

There were 2 different configurations of loss used in the experiment:

1. *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%.

2. *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. In Condition 2, a new level of loss within the appropriate range was chosen every 5 seconds. In Condition 4 a new level of loss within the appropriate range was chosen every second. This difference was due to the auto-decrement function in Condition 2.

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-

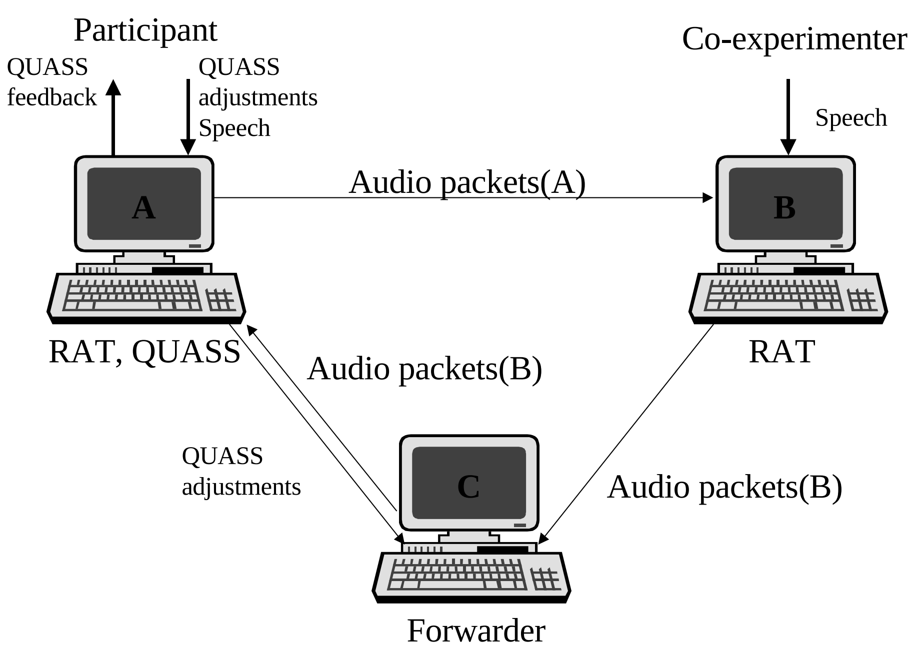

Figure 3: Experimental set-up

experimenter, where the participant would take turns in describing/guessing a word. Thus, the task was relatively equally partitioned into speaking and listening. Each condition was timed and participants asked to stop after 5 minutes. During the 5 minutes participants were asked to use QUASS to control the quality that they were receiving. For all conditions participants were instructed to move the slide-bar so that the quality they received was at the minimum level where the task could still be performed. In Conditions 3 and 4, a budget indicator was displayed, and participants were asked to perform the task maintaining as much of this finite resource as possible.

2.4.2 Semi-structured interviews

All participants were asked a series of 4 questions during a semi-structured interview. These questions related to the experiment in which they had taken part, and QoS pricing in general. Each interview lasted approximately 5 minutes. The questions asked of each participant were:

1. *Did you notice a difference in quality between the conditions where no budget was applied?*

The purpose of this question was to explore participants' perceptions of the quality received without the influence of any potentially distracting elements, such as a budget. The question was also asked to elicit the manner in which participants described the quality and the references made to the representation of this quality in the UI.

2. *Did you find it difficult to operate the slider and play the game at the same time?*

In a previous study of a distance learning application, users found it difficult to operate the QUASS slider to rate audio and video quality, and participate in interactive tasks at the same time. Even though here users were required to rate audio quality only, it was felt that a possible conflict between dynamic QoS requests and task performance should be investigated. Such conflicts may have a potential impact upon the applicability of requests for dynamic QoS to many Internet applications.

3. *Did you feel that the budget made a difference to the way you operated the slider?*

4. *How would you feel about being charged in this manner for Internet services?*

Our previous research suggested that the introduction of a resource that is linked to QoS requests may encourage users to re-evaluate the quality they receive in terms of the *Risk* that that quality will not be of appropriate *Value* [12]. These questions were therefore asked to elicit participants' subjective opinions about the presence and use of the budget, both within the context that they had experienced in the experiment, and for the Internet in general.

3. RESULTS

3.1 Quantitative data

Results from 25 participants were included in the analysis. Some participants experienced a conflict between attempting to operate the slider, and playing the game at the same time. This is likely to be the case in the assessment of any dynamic QoS parameter, but

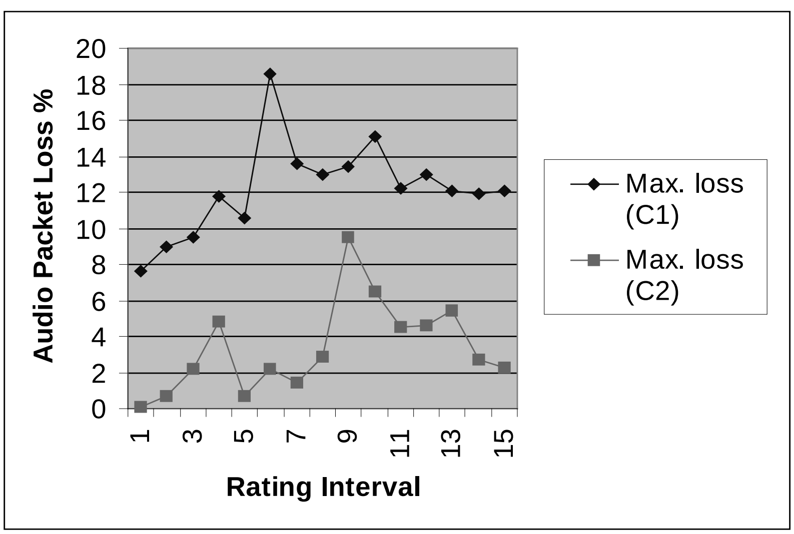

Figure 4: Mean maximum loss rates (no-budget)

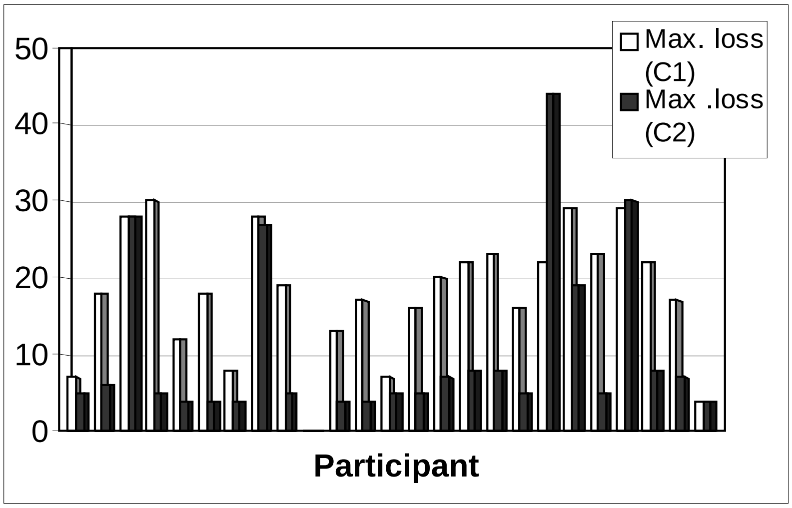

Figure 5: Maximum loss rates per participant (Conditions 1&2)

is important when considering potential applications for quality-based pricing, as it suggests the need for encapsulated functionality in order to minimize the number of real-time operations required by the user (see section 4).

3.1.1 The Effect of Predictability (No-Budget Conditions)

In the *no-budget* conditions, there is a marked difference between the two conditions in terms of the levels of quality requested by participants: 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 and configured for Conditions 1 and 2 is shown in Figure 4 ($t = 12.15$ $p<0.001$). These results show that users are willing to tolerate greater objective loss rates when those loss rates are relatively predictable. Results for each participant are illustrated in Figure 5 ($t = 7.21$, $p<0.001$).

3.1.2 The Effect of Predictability (Budget Conditions)

In the *budget* conditions, the quality requested for the *stable* loss condition also differed significantly from the *variable* loss condition: far less quality was required for the *stable* loss condition compared to the *variable* loss condition. A comparison of maximum loss levels, averaged amongst participants, configured for both conditions is shown in Figure 6 ($t = 11.38$, $p<0.001$). Maximum levels of loss received by each participant are shown in Figure 7 ($t = 10.25$, $p<0.001$).

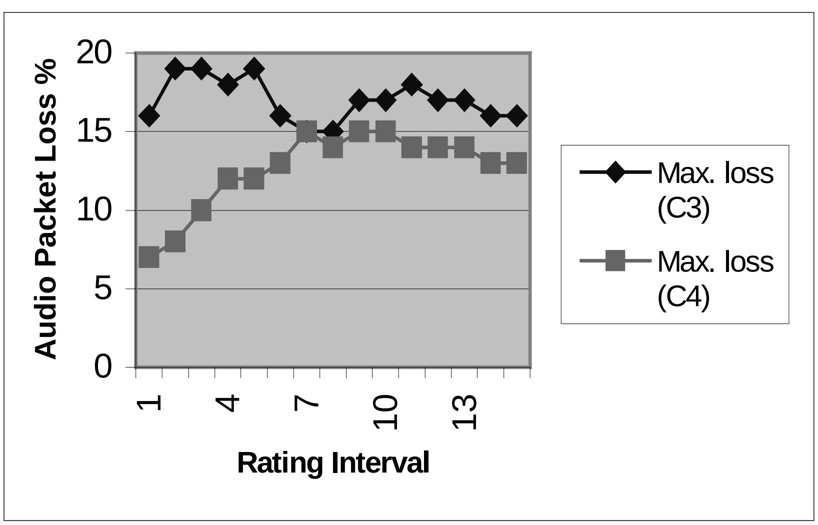

Figure 6: Mean maximum loss rates (budget)

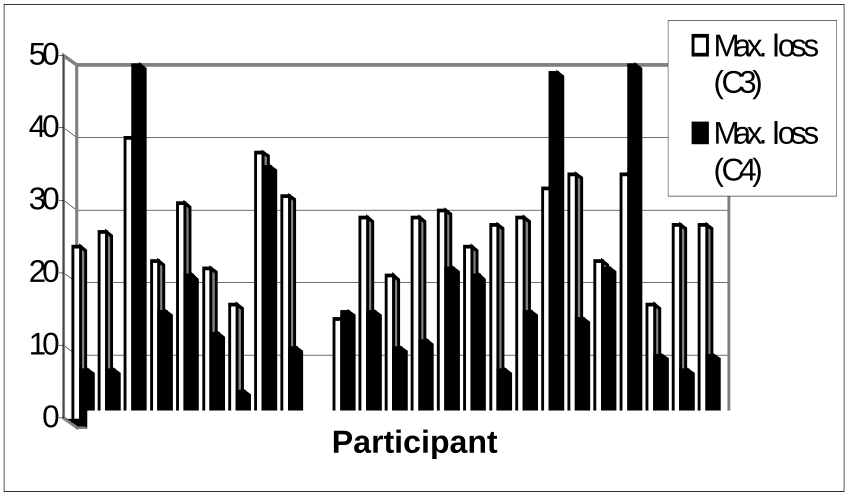

Figure 7: Maximum loss rates per participant (Conditions 3&4)

These results show that users will request higher levels of quality when that quality is relatively unstable.

3.2 Qualitative data

As we have previously stated, the use of objective measures alone is not sufficient to establish a framework from which users' requirements for QoS can be predicted. Interviews were conducted to establish the reasons behind significant differences in users' behavior during the experiments.

3.2.1 The Importance of Feedback

Question 1: Did you notice a difference in quality between the conditions where no budget was applied?

Table 1: Responses to question 1

Question	"No" (C1,C2)	"Yes" (C1,C2)	"Yes" (C3,C4)
1	10	11	11

A highly significant number of participants placed the slide bar at positions where a relatively low level of quality would be received in conditions where the loss was stable, compared to when the loss was variable. This result shows that users prefer lower, but stable QoS to levels that are higher, but variable – and therefore unpredictable. In response to the first interview question (see Table 1), the majority of participants who said that they could notice a difference between conditions attributed their answer to the fact that they could place the slider at different positions in the interface. For example:

'…I thought the quality in the second one was a lot better and I remember it that way because I put the slider down a lot more.'

Most participants who noticed a difference in quality between conditions, said that this related to the type of quality configured in Conditions 3 and 4 only. This finding may be due to the fact that the presence of pricing feedback at the UI required that participants re-evaluate the quality they received. A higher amount of concentration resulted in participants noticing the differences between these two – latter – conditions. For example:

'…I noticed it in the second lot…the budget one was a lot worse than the first lot…It's sort of like concentrating on that thing (budget) *going down and the dial* (slider)*, connecting those two'.*

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. We have already established (in sections 3.1.1 and 3.1.2) that users value a predictable pattern of quality. Arguably, it is the ability of the system to meet users' expectancies that determines that system's predictability. Results from interview data show that, for a large number of participants, an important tool in accomplishing this would be the feedback given to the user via the UI.

These results 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.

3.2.2 Strategy Formation

Question 3: Did you feel that the budget made a difference to the way you operated the slider?

Table 2: Responses to question 3

Question	"No"	"Yes"	Use of strategy
3	2	20	12

As expected, the introduction of a finite resource into Conditions 3 and 4 had the effect of increasing participants' tolerance to loss. However, the combination of quantitative and qualitative data may reveal further, more interesting and complex results (see Table 2). Having established that UI feedback is crucial to some users in forming a post-interaction assessment of quality, it is now necessary to ask *why* it is so important. Analysis of participants' responses to post-experiment questions about the effect of the budget upon their requests for QoS shows that a number of participants use the way in which the UI presents QoS feedback to formulate a *strategy* for pre-meditating future QoS, and conserving their budget:

'I'd got it into my mind that I couldn't make it go below the budget…and what you end up doing is making trade-offs between does it matter if the quality is not good when I'm speaking…so you work out little strategies.'

This strategy was developed due to the configuration of the UI:

'It's because it had a finite amount and you'd think, what happens when it gets to the bottom, do you get cut-off?'

The results in this study illustrate the behavioral consequences of strategies that users form according to UI feedback. Figure 4, for example, shows that quality requests fluctuate widely throughout the experiments, but follow similar patterns between conditions. Interview data suggests that this effect may be the result of participants decreasing their quality requests when they were speaking in order to prevent their budget from decreasing, and increasing such requests when they were required to listen.

The discovery that users formulate strategies based on UI *Feedback* and *Task* requirements can be placed into context. Our previous research has suggested that, when a finite resource is involved, users assess the risk of sacrificing part of that resource against the expected value of future QoS [12]. They use this *Risk Assessment* when making requests for a certain amount of QoS (see Figure 1). It is therefore likely that participants in these studies formulated strategies based

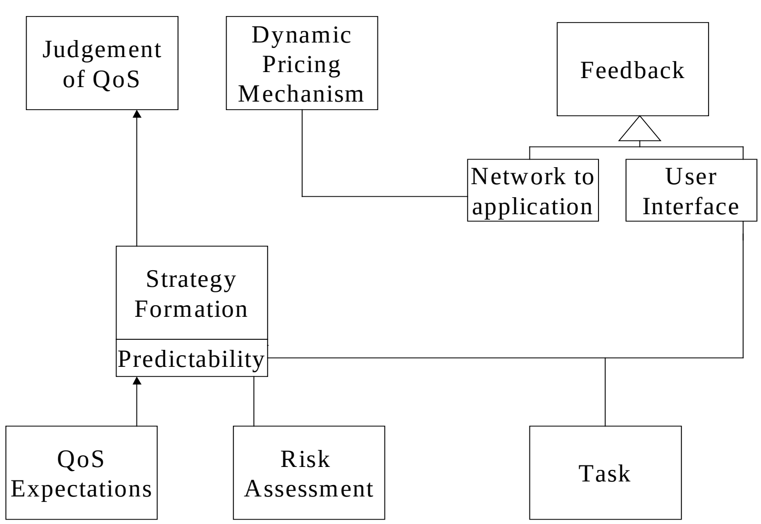

Figure 8: Strategy Formation in context

on the amount of risk associated with their perception that the QoS received will be in accordance with what is expected. If that quality is predictable, then any strategy that users form will be reliable, and the particular pricing mechanism is more likely to be accepted. These influences on *Strategy Formation* are shown in Figure 8. Figure 8 illustrates that the strategy users form has implications for the judgements of the QoS received. For example, in Conditions 1 and 2 users followed the simplest strategy of positioning the slider at a position where the task could easily be completed – there was no need to formulate a more complex strategy.

During the conditions where a finite resource was represented by the UI (Conditions 3 and 4), participants had to consider, more directly, the *Value* of the quality they received according to the *Task* being performed. Quality received during the sub-task of listening was considered of greater value than that received when the participant was speaking. The formation of an appropriate strategy affected participants' assessment of the objective QoS received: some participants who did not notice a difference between the quality in Conditions 1 and 2, *did* notice this difference in Conditions 3 and 4:

'I noticed it in the second lot…it really unnerved me in the budget one…I noticed it because I'd just got used to doing it (in Conditions 1 and 2)*, and working out where it would go.'*

'It got fuzzier quicker, only in the last one (Condition 4)*'*

Thus, the feedback at the application level that indicates the presence of a finite resource has a crucial influence on subjective assessments of quality. However, it is according to the particular strategy users form that value is ascribed to (parts of) the interaction.

3.2.3 Default Strategies

Question 4: How would you feel about being charged in this manner for Internet services?

Table 3: Responses to question 4

Question	For	Against	Depends
4	8	11	8

The majority of participants in this study would not wish to be presented with a budget mechanism such as that employed in these experiments, for all Internet interactions. Many felt that the mechanism conveyed too much feedback. This encouraged participants to concentrate on the relationship between the budget display and slider objects, and the consequent formation of a strategy that may have removed participants' concentration from the task at hand. These findings agree with the ideas of Wakeman (Ian Wakeman, University of Sussex, personal communication, 1998), in that the budget mechanism used in these experiments may have required the user to perform unnecessary sub-tasks. By bringing the user into the control loop at a higher level of granularity, we may encourage the acceptance and use of QoS pricing in what is, after all, a traditionally free environment.

From one point of view, therefore, our findings suggest that users should be able to state how much they are willing to pay for a certain amount of quality during the conduct of the task. This should be done in one operation, prior to interaction. In a dynamic environment, however, it is questionable whether users will be able to state the value of an object within the task, prior to interaction. Indeed, what is assessed in the former scheme is the *perceived*, rather than the actual, value of that object. Furthermore, in some circumstances users do not wish to be presented with *any* pricing mechanism. The key point here, then, is that any pricing mechanism would need to provide the appropriate levels of encapsulation required by users in order to be representative of the *Value* ascribed to the task, at any particular time.

In order to ensure that the system is as predictable as possible, participants felt that they would prefer a base-rate QoS charged according to a flat-fee coupled with incremental costs for extra quality:

'…most of the time you might not be that desperate to get something really quickly'.

This would enable users to configure higher quality, dynamically, in situations where the value of the task is relatively high. In situations where dynamic pricing was applied the majority of participants in our study felt it would be important to be provided with a chance to configure the pricing feedback in two main situations:

- Where the cost of the interaction was likely to be expensive:

'Yes, you want the information to be able to make those decisions yourself…if it's a period of day that's going to be expensive I'd like to be able to turn it on'.

- Where the risk of poor quality is relative high and the subjective value of task performance is high:

In this situation the necessity to formulate a strategy to minimize this risk is dependent upon users' ability to predict the level of quality received. As previously discussed (see section 4.1) system predictability can be made explicit through the use of UI feedback. The provision of dynamic pricing minimizes *Risk* by providing the user with a sense of control:

'..if you can control the quality, or if the quality goes you're able to control the quality and get charged accordingly.'

The optimal solution to the trade-off between an ultimately predictable flat-fee mechanism and dynamic quality based pricing that reflects the value of the interaction may involve users providing a set of predefined preferences. With regard to the UI used in this particular experiment, comments from some participants suggest the use of encapsulated functions that would be responsible for monitoring quality levels:

'..there could be something that could swap the quality, put the quality down when you're speaking ..because you could save half your budget that way'.

'I think I would say "give me a range" and I would just pick a range instead of constantly changing it'.

'I think perhaps you could have a trial period where you got used to the level you liked and then you'd fix it at that afterwards. That level would be fairly predictable afterwards.'

This system would enable feedback to be configured to the UI only when the network could not meet such preferences. Although users' react to a predictable, priced system by formulating a *Strategy* that enables them to minimize the *Risk* inherent in any transaction, as in these experiments a single strategy might be applicable for an entire *Task*. These findings suggest the use of task profiles, in which specific performance requirements could be encapsulated and translated into objective criteria.

4. CONCLUSIONS AND FURTHER WORK

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. We have also shown that, for users to formulate appropriate strategies, priced network quality must be presented to them according to the task that is their goal. Although focused on audio, the conclusions of these findings have implications for the design and implementation of networked multimedia applications and the underlying network.

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. [19]). 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 centered around the costs and benefits of providing differentiated versus integrated services [20],[21],[1]. 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 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. ***Users require a specific level of feedback to formulate appropriate strategies.*** Dynamic pricing feedback should, therefore, not be displayed at all times. Our findings suggest that it is possible to provide too much feedback to the user. Users will apply a *Strategy*, which they have previously formulated, to a new situation – until UI feedback informs them that this strategy is no longer viable. A pricing mechanism that provides an inappropriately large amount of *Feedback* distracts users from their primary *Task*, and is just as likely to be rejected as one that provides none at all. Instead, *Feedback* should be provided when there is a relatively high cost in using the network, and the user must re-assess the value placed on the *Task*.

6. ***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 (e.g. [22]).

7. ***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.

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. QoS mechanisms that focus on the most technically efficient way of capturing users' preferences for the deployment of QoS may miss the point – i.e. that such preferences are influenced by the configuration of *Feedback* at the user's point of interaction, where the *Task* is primarily supported.

Much of the technical literature on network QoS suggests the efficiency of an application can be enhanced by providing functionality exclusively at the network layer. 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*.

A recently announced scheme to allocate Web QoS combines predictable service with admission control [23]. The system defers users from a site but makes a prediction of exactly how long it will be before they can be admitted, thus encouraging them to accept a lower level of QoS in return for a predicable service. This system has yet to be tested in a fully controlled user-trial. Our research presents an important first step in confirming the influence and benefits of allowing users to predict future levels of quality and can therefore provide evidence for the potential success of this system.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the work done by Orion Hodson (UCL) in writing and enabling the Forwarder software. Special thanks are also due to Anna Watson (UCL), for playing the guessing games.

REFERENCES

[1] Gupta, A.; Stahl, D.O.; Whinston, A.B.: Priority pricing of integrated services networks. In McKnight, L.W.; Bailey, J.P. (eds).: Internet Economics. MIT Press 1997.

[2] Cocchi, R.; Shenker, S.; Estrin, D.; Zhang, L.: Pricing in computer networks: Motivation, formulation and example. IEEE/ACM Transactions on networking, 1993, 1, 6, pp. 614-627.

[3] Clark, D.: Internet cost allocation and pricing. In McKnight, L.W.; Bailey, J.P.: (eds.) Internet Economics. MIT Press 1997.

[4] Bellotti, V.: Telework: When your job is on the line. Interactions, 1997, 5, 1, pp. 44-54

[5] Odlyzko, O.M. The Internet and other networks: Utilization rates and their implications. Available at http://www.research.att.com/~amo.

[6] Keller, J.J.: Ex-MFs managers plan to build global network based on Internet. Wall Street Journal. January 20, 1998.

[7] Mackie-Mason, J.K.; Varian, H.R: Economic FAQs about the Internet. Journal of Economic Perspectives, 1994, 8, (3), pp.24-36.

[8] Bohn, R.; Braun, H.W.; Wolff, S.: Mitigating the coming Internet crunch: Multiple service levels via precedence. Applied Network Research Technical Report, 1994, GA-A21530. University of California.

[9] Mackie-Mason, J.K.; Murphy, L.; Murphy, J: Responsive pricing in the Internet. In Internet Economics. McKnight, L.W.; Bailey, J.P.: MIT press 1997.

[10] Watson, A.; Sasse, M.A.: Evaluating audio and video quality in low-cost multimedia conferencing systems. Interacting with Computers, 1996, 8, (3), pp. 255-275.

[11] Kokotopoulos, A.: Subjective assessment of multimedia systems for distance learning. In Proc. European Conference on Multimedia Applications, Services and Techniques, Milan, May 1997.

[12] Bouch, A.; 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. Proc. IFIP/IEEE International Symposium on Integrated Network Management (IM'99) Boston 24-28 May, 1999.

[13] Kawalek, J.: A user perspective for QoS management. Proc. QoS workshop, 16th September 1995, Crete, Greece.

[14] Bouch, A.; Sasse, M.A.: Network quality of service – an integrated perspective. Proc. RTAs'99, Vancouver 1-3 June, 1999.

[15] Watson, A.; Sasse, M.A.: Multimedia conferencing via multicast: Determining the Quality of Service required by the end-user. Proc. International Workshop on Audio-Visual Services over Packet Networks (AVSPN), 15-16 September 1997, Aberdeen.

[16] Bouch, A.; Watson, A.; Sasse, M.A.: QUASS – a tool for measuring the subjective quality of real-time multimedia audio and video. In May, J., Siddiqi, J., & Wilkinson, J (eds.) HCI'98 Conference Companion, pp. 94-95, 1998.

[17] Hardman, V.; Sasse, M.A.; Kouvelas, I.: Successful multi-party audio communication over the Internet. Communications of the ACM, 1998, 41, (5), pp.74-80.

[18] McCanne, S.; Jacobson, V.: vic: A flexible framework for packet video. Proc. ACM Multimedia, 1995, pp. 511-522.

[19] Gecsei, J.: Adaptation in distributed multimedia systems. IEEE Multimedia, 1997, April-June, pp.58-66.

[20] Wang, Z. (1997). User-share differentiation (USD) scalable bandwidth allocation for differentiated services. Available at: http:www/ietf.org/internet-drafts/draft-wang-diff-serv-usd-00.txt

[21] Nichols, K.; Blake, S.: Differentiated services operational model and definitions. Available at http://www.ietf.org/internet-drafts/draft-nichols-dsopdef-00.txt

[22] Nahrstedt, K.; Smith,J.M.: The QoS Broker. IEEE Multimedia, 1995, Vol.2, No.1, pp.53-67.

[23] DiCarlo, L.: HP moves to safeguard online retailers against crashes. News-Online, June 21, 1999.