

It Ain't What You Charge, It's The Way That You Do It: A User Perspective of Network QoS and Pricing

A. Bouch, M. A. Sasse
Department of Computer Science
University College London
Gower Street, London
WC1E 6BT
UK
{A.Bouch,A.Sasse}@cs.ucl.ac.uk

Abstract

Shared networks, such as the Internet, are fast becoming able to support heterogeneous applications and a diverse user community. In this climate, it becomes increasingly likely that some form of pricing mechanism will be necessary in order to manage the Quality of Service (QoS) requirements of different applications. So far, research in this area has focussed on technical mechanisms for implementing QoS and charging. This paper reports a series of studies in which users' perceptions of QoS, and their attitudes to a range of pricing mechanisms, were investigated. We found that users' knowledge and experience of networks, and the real-world *Task* they perform with applications, determine their evaluation of QoS and attitude to payment. Users' *Payment Behavior* is governed by their level of *Confidence* in the performance of salient QoS parameters. *User Confidence*, in turn, depends on a number of other factors. In conclusion, we argue that charging models that undermine *User Confidence* are not only undesirable from the users' point of view, but may also lead to user behavior that may have a negative impact on QoS.

Keywords

Quality of Service, charging, pricing, user and service profile management

1. Introduction

Shared networks – such as the Internet – can today support a wide variety of services and applications, including those that are real-time. There is, however, a persistent concern that the current 'best-effort' service model does not take account of the differing performance requirements of such applications. Real-time audio and video applications, for example, may be more sensitive to delay than asynchronous applications, such as email. Network technology is able to provide

different levels of Quality of Service (QoS) in accordance with the requirements of applications.

The aim of most current technical proposals that address differential QoS provisioning is to configure a situation in which network bandwidth can be deployed to maximum efficiency. An essential part of this efficiency is said to be the provision of incentives to the user* [1]. Drawing on traditional econometric theory, such incentives are often described in terms of objective quantities that can express users' payment behavior. An example is a measure of the magnitude of demand [2]. These variables can then be placed into equations for predicting network efficiency. In our experience, users' behavior can be predicted more accurately by analyzing their mental constructs and motivations, and the tasks they perform when using network services. To date, little attention has been paid to users' views in the debate of QoS and pricing mechanisms – probably because gathering and interpreting such data requires skills not commonly found in the technical network community. Yet, any QoS implementation and charging model that is not accepted by users is likely to fail in the real world.

This paper presents an initial set of models representing users' mental constructs relevant to this problem. The concepts highlighted by these models should be seen as *complementing* the economic and technical concepts that must be considered in ensuring the efficient function of the network system as a whole. The paper begins with a brief description of the studies through which the data were collected, and the method used to generate the models. Sections 3 and 4 provide a detailed description of the concepts and models identified, and examples of predictions that can be made about user behavior. We conclude that, even though (a) more experimental studies will be required to test their predictive validity, and (b) they need to be simplified, the models are an important first step towards incorporating a realistic notion of user behaviour into QoS and pricing.

2. Research method

Previous research on audio and video quality over shared networks has shown that there is no straightforward way of predicting perceived quality from objective measures [3]. Perceived quality depends on a number of interdependent factors, the *task* that the user is performing being the most significant. Based on these findings, we would expect users' perception of QoS and associated payment behavior to be determined by a similar set of inter-related concepts. To extract such concepts and relationships from a set of data, we utilized grounded theory methods [4]. These methods are well established in social science research and have been successfully applied to describe and predict other complex user behaviors, e.g. with computer passwords [5], [6].

Grounded theory allows the extraction and categorisation of relevant concepts from qualitative data, such as statements made by users in interviews and focus

* The term "user" in this paper always means "end-user", i.e. individuals using network services at work or home. With corporate users, additional factors need to be taken into account.

group discussions. The data used in these studies were obtained from a series of focus groups. The advantage of this method for eliciting users' views is that it encourages the development of ideas that are important to the participants themselves, and does not presuppose a specific agenda for discussion. Indeed, the method is particularly useful in discussing participants' 'meanings and ways of understanding' [7].

Two groups of users – each comprising 16 individuals – took part in the focus group sessions. All users had some experience with using network applications such as email, WWW, and videoconferencing, but the users in the “advanced” group had completed a module on computer networking as part of their degree programme (for reasons outlined below). Each user-group was divided into sub-groups to achieve a group size that was manageable for group discussion, with each session lasting for approximately an hour. In these focus groups, users discussed QoS and pricing with the help of a trained facilitator.

Reference [8] suggests that many concepts applied to the networking domain are designed by highly trained experts and reflect their understanding ‘*of underlying hardware*’. It is a basic tenet of human-computer interaction that most users, however, reason about technology in terms of the *task* they want to complete, rather than building a mental analogue of the technology itself. Reference [9] has shown that this applies to networking: he found that the amount and type of knowledge possessed by Internet users varies widely, and that many had to re-think their ideas of network operations in order to use that network efficiently. We can, therefore, expect that users with different levels of knowledge and experience will reason about QoS and pricing differently. We therefore recruited focus group participants with different levels of knowledge:

- *Advanced Users* with syntactic and semantic knowledge of network operations, and experience with both real-time and data-driven networking tasks.
- *Novice Users* with a limited amount of syntactic and semantic knowledge of network operations, and limited experience in real-time network tasks, but familiar with data-driven network tasks such as email and the WWW.

3. Users' perception of QoS parameters

Network research has focused on defining QoS parameters from a technical point of view. However, it is the users' perception of the performance of salient QoS parameters that, in the end, determines network usage [10]. One of main aims of our work was to study the assessment that users make of specific dimensions of quality, and identify the salient QoS parameters that can be applied to different applications.

The qualitative data obtained from focus groups were analyzed using the grounded theory coding techniques. OSI (Open Systems Interconnection) standards state that QoS parameters can be defined at different levels of the network. According to the OSI framework, QoS is dependent on the lower layers responsible

for network operations. Users, however, reason about QoS in terms of the application, and therefore QoS parameters located at the application level are considered by users to be the most salient. Therefore, if QoS is acceptable at the level of the application, the QoS that is attributable to the lower OSI layers becomes irrelevant from the users' point of view. This suggests that charges for both data-driven and real-time tasks should be made for QoS parameters occurring at the level that is closest to the user in terms of conceptual importance, and the location of physical interaction. At the application level, however, there are clear differences between users' opinions of the relative importance of QoS parameters, depending to the type of *Task* being performed.

3.1 Task variables

Technical considerations of QoS often argue that traffic can be characterised into two basic categories: inelastic and elastic. Traffic that is associated with real-time applications is characterised as delay-intolerant, or *inelastic*. Traffic produced from data-driven tasks, on the other hand, is more delay-tolerant and therefore classed as *elastic*. This distinction between traffic types overlaps partly with users' views, but there are some subtle differences. The QoS dimensions associated with the timely delivery of data (e.g., *Speed*, *Delay*) may be connected to real-time applications, whilst *Throughput* is associated with data-driven tasks. However, perceived QoS parameters depend on the real-world tasks that users perform, and the salience of such parameters lies in their relevance to this Task, rather than in the characterisation of traffic produced as a result of users' operations performed in order to achieve the task. For example, *Manipulation*, associated with real-time video tasks, refers to the ability to manage the video image in terms of operations such as resizing. Clearly, users' need for this QoS parameter is dependent upon the *Value* placed on what is seen in an image.

Users' perception of the number of real-world *Task* operations that an application is required to perform also influences their assessment of QoS. For example, users conceptualize email delivery as a single operation: deliver the message. The novice users in these studies did not realize that the searching of several network sub-routes was required of the network in order to locate the recipient of the message. The *Work* for this operation is therefore judged to be a small, fixed amount. In contrast, WWW-searching applications were seen to require a number of operations. The *Work* for these applications is therefore greater than zero. We found that, when *Work* is assessed as equal to or close to zero, the user requires little or no *Feedback*. If *Feedback* is supplied for such applications, users' judgements of QoS are likely to be negative. This phenomenon is arguably caused by the real-world task metaphor possessed by novice users. These users form QoS expectations based on, for email, the high-level task of *Delivery*. Indeed, they possess no model of the detailed sub-tasks that may be involved in this operation; consequently, *Feedback* about the sub-tasks involved in an email interaction is meaningless and annoying.

3.2 Differences between users

Perceived QoS parameters depend not only on the *Task* being performed, but also on users' knowledge and experience. Inexperienced users tend to reason about applications and quality in terms of a real-world metaphor. We already mentioned that novice users in this study had no conception that email requires routing of information through a network infrastructure. Instead, these users applied the metaphor of the postal service to email applications. They therefore expected these applications to be wholly elastic. With increasing knowledge about networking, the influence of such real-world *Conceptual Metaphors* on perception of QoS decreases, enabling advanced users to conceptualize tasks according to the type of traffic produced.

Reliance on a *Conceptual Metaphor* leads novice users to attribute *Boundaries* to tasks in terms of the amount and type of QoS that it is appropriate for that user to receive. For example, a novice user will associate *Throughput* with Internet searching applications regardless of whether these applications contain real-time traffic. This may be due to the consideration that the *Conceptual Metaphor* applied to a network application is relatively static, compared to the potential changes in the traffic characterization associated with that application.

3.3 Virtual distance

The term *Virtual Distance* refers to the physical distance between local and remote users. This distance is virtual as what is represented is not network connectivity and routing paths taken by data as it travels from sender to recipient. Instead, novice users are prepared to accept a lower general level of QoS if the physical distance between local and remote users is high. In videoconferencing, for example, users will tolerate a higher level of audio packet loss from a participant who is on a different continent than from somebody in the same building. They also rate sessions with places that may be physically close, but take a long time to reach (because of lack of public transport, or congested roads) more favorably than those places they could get to more easily. The concept of *Virtual Distance* shows that users make a real-world cost-benefit assessment which affects their assessment of the QoS.

The current focus on the management of QoS within the network infrastructure has yielded a wealth of routing and queuing schemes that support mechanisms such as differentiated service [11]. The implications of our findings, however, is that management of QoS is not only required at the network level, but at the level of the application. Considerations of both *Feedback* requirements, and the concept of *Virtual Distance*, indicate that the comparison novice users' make with the real world influences their assessment of QoS. To support QoS from the users' point of view, QoS management has to be more flexible. Indeed, the users' view is influenced by the *Value* users place on the *Task* being performed. To be effective, QoS management must be able to translate higher-level QoS requirements into traffic characterizations whose profiles are understandable by the network.

3.4 QoS and pricing

One benefit of identifying salient QoS parameters could be that it would provide the basis for defining a charging mechanism based on subjective QoS. Basing a charging mechanism on salient QoS parameters does, however, not mean that such a charging mechanism would be automatically acceptable to users. Reference [12] for example, found that users prefer flat-fee pricing schemes, where they have limited control over the QoS they receive, to usage-based pricing schemes. Users also tend to overestimate their network usage, which indicates the need for *Feedback* on such usage. It has been suggested that it is important to determine the amount and type of *Feedback* required by users in order to promote desired behavior [13]. Based on results of experiments into responsive pricing, it has been argued that user requirements are at the key to an efficient network. It was found that extracting users' QoS requirements through a feedback loop allows price-sensitive users to reflect on the *Value* of quality received [13]. Results from these simulations show that technical network efficiency (measured by packet loss) and user satisfaction increased when *Feedback* concerning congestion and pricing is supplied to the user. These findings suggest that engineering measures of QoS alone, whilst addressing network efficiency from a technical point of view, are an imperfect measure of the economic efficiency of that network. The latter dimension must include a measure of the *Value* that the user places on the QoS. Indeed, it is the concept of *Value* that provides the link between QoS and pricing.

4. Pricing schemes

The second major aim of our research was to explore the acceptability of various pricing mechanisms that are currently discussed for shared networks. In the focus groups, we asked users to consider different pricing scenarios, using the previously outlined definitions of QoS. One priority was to establish the influence of *Feedback* requirements on the acceptability of pricing mechanisms. The schemes discussed by users included*:

- *Smart Market* [14]: Users submit a bid price with each packet they wish to submit to the network. The price to send a packet therefore varies as the degree of congestion on the network varies. Users do not pay the price actually bid, but the bid price of the first packet not admitted to the network.
- *Quota Pricing* [15]: In this scheme, weights can be assigned to packets and pre-paid quotas are reduced in proportion to the sum of the weighted packets admitted to the network. In this fashion the user may, for example, send elastic traffic such as email with a delay priority/weight of zero.
- *Paris Metro Pricing* [16]: The user is required to pay more to use a particular queue, although the architecture of this queue is identical to those that are

* If any of these schemes were not raised by users during the focus group, the moderator introduced and explained them before asking users to discuss how acceptable these mechanisms were.

cheaper. The idea is that the queue that is more highly priced would attract less traffic and therefore be able to deliver a higher level of QoS.

The process by which users judge the acceptability of a pricing scheme is best represented as a structured storyline. Table 1 is an example of a high-level storyline that represents components that describe advanced users' *Payment Behavior*.

4.1 Risk and confidence

The introduction of an explicit payment for a network service adds, from the users' point of view, a dimension of *Value* to perceived QoS. Users' level of *Confidence* that charging procedures represent *Value* was found to be the overall high-level determinant of their *Payment Behavior*. *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. Advanced users draw on *Network Concepts*, novice users do not. The *Expected Utility* of quality received, for instance, is influenced by the levels of congestion the advanced user believes to be present.

Less knowledgeable users, however, attribute the control of received quality to the operations they perform at the application level. As a consequence, their actions are not conceptually linked to the operations of the network, and they can attribute unexpected QoS to *Involuntary* actions. The ability for novice users to arrive at a positive assessment of *Risk*, in order to achieve *Confidence*, is therefore inhibited by the view that an action was *Involuntary*. This latter situation is one of *Danger*. If, for example, the user pays a relatively large amount of currency and is required to perform two acts of confirmation of his desire to send an email message, then that user may associate the request to perform such actions as *Involuntary*. *Confidence* is never associated with a situation of *Danger*. Reference [17] defines *Danger* as being the denial of '*the attributability of loss to a decision*'. This means that loss is viewed as externally generated. The loss to the user in the case of our research is that of the control over the QoS received. Control over QoS is attributed to *Network-Centric* processes, as distinct from *User-Centric* processes that are expected by novice users. In contrast, for advanced users, an understanding of the multiplexed nature of the network leads to a concept of *Collective Quality*. The presence of this concept suggests that an assessment of *Risk* will be partly based on the users' perception of the QoS received by other users of the network.

Table 1: Payment Behaviour: High-level Story-line

The advanced user may gain *Confidence* leading to *Payment Behavior*:-
By perceiving the coupling of user and network processes leading to:-

- A) A *Risk Assessment* of a transmission producing:-
 - 1. A high-risk situation characterized by:-
 - a) *Network Procedures* that are:-
 - I. *Unpredictable* according to *Feedback* which is:-
 - i. Not qualitative and quantitative.
 - ii. Slow
 - iii. Untrustworthy.
 - iv. Of inadequate amount.
 - II. Degrade *Trust (Contextual)* by opposing expected *Collective Quality*.
 - III. Have high *Probability* of unacceptable *Expected Utility* according to:-
Traffic elasticity so that traffic is:-
 - i. Elastic and the network has unacceptable reliability.
 - ii. Inelastic and the network has unacceptable capacity.
 - b) *User Conceptions* of internal network operations that:-
 - I. Degrade *Trust (Social)*
 - 2. A low-risk situation characterized by:-
 - a) *Network Procedures* that are:-
 - I. *Predictable* according to *Feedback* which is:-
 - i. Quantitative as well as qualitative.
 - ii. Fast.
 - iii. Trustworthy.
 - iv. Of adequate amount.
 - II. Promote *Trust (Contextual)* by confirming expected *Collective Quality*.
 - III. Have low *Probability* of unacceptable *Expected Utility* according to:-
Traffic elasticity so that traffic is:-
 - i. Elastic where the network has acceptable reliability.
 - ii. Inelastic where the network has acceptable capacity.
 - b) *User conceptions* that:-
 - I. Promote *Trust (Social)*.

B) An assessment of the performance of network QoS drivers leading to a situation that produces:-

1. *Concurrence with Expected Utility* which results in:-
 - a) A low performance and low *Expectancy* situation leading to:-
 - I. An acceptable level of QoS. Acceptance of pricing mechanism.
 - b) A high performance and high *Expectancy* situation leading to:-
 - I. Acceptance of good QoS. Acceptance of pricing mechanism.
2. *Inconcurrence with expectancies* which results in:-
 - a) A low performance and high *Expectancy* situation leading to:-
 - I. An unacceptable level of QoS. Rejection of pricing mechanism.
 - b) A high performance and low *Expectancy* situation leading to:-
 - I. Acceptance of good QoS. Rejection of pricing mechanism.

There are further differences between advanced and novice users in the conceptualization of *Trust*. In order to form a positive *Risk Assessment*, it is necessary for users to make a positive evaluation of *Trust*. However, as novice users do not possess *Network Concepts*, *Trust* is conceptualized solely as that which is attributed to processes external to the network. (i.e. the user-ISP contract). As might be expected, having *Network Concepts* leads advanced users to explain *Trust* in terms of the behavior of the network infrastructure.

4.2 Predictability

The concept of *Risk* is directly linked with that of *Predictability* - a low *Risk* situation is one that is predictable. Clearly, the *Predictability* of network processes depends on appropriate *Feedback* being given to the user. The importance of *Feedback* has already been demonstrated in this research with regard to QoS parameters (section 3) and *Virtual Distance* (section 3.3). These considerations demonstrated that the amount of *Feedback* required by users depends on the *Task* being performed.

The relevance of *Task-dependent Feedback* in assessing the acceptability of pricing mechanisms was found to be especially apparent for novice users. Due to the use of a *Conceptual Metaphor* in the formation of *Confidence*, it is essential that *Feedback* be configured according to that metaphor. If users apply the

telephony metaphor to Web-searching tasks, for example, *Feedback* should be supplied on a ‘per-call’ basis.

Predictability is also implied in the level of control required by users over their payments for QoS. Somewhat contrary to previous findings (e.g., [12]), our findings 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, although the received QoS should be of a guaranteed level. Therefore, dynamic pricing needs to provide *Feedback* on network congestion, which would enable users to predict the *Risk* involved in making certain payments. In this case, therefore, users’ need for *Predictability* tempers their need for network processes to be encapsulated.

5. Summary: Users’ models of QoS and pricing

Figure 1 and Figure 2 present the sets of models showing the links between concepts extracted from the data. Higher-level concepts depend on an aggregate of sub-concepts, which vary with users’ level of knowledge and experience. A detailed definition of each concept is provided in a separate Glossary (see Section 8).

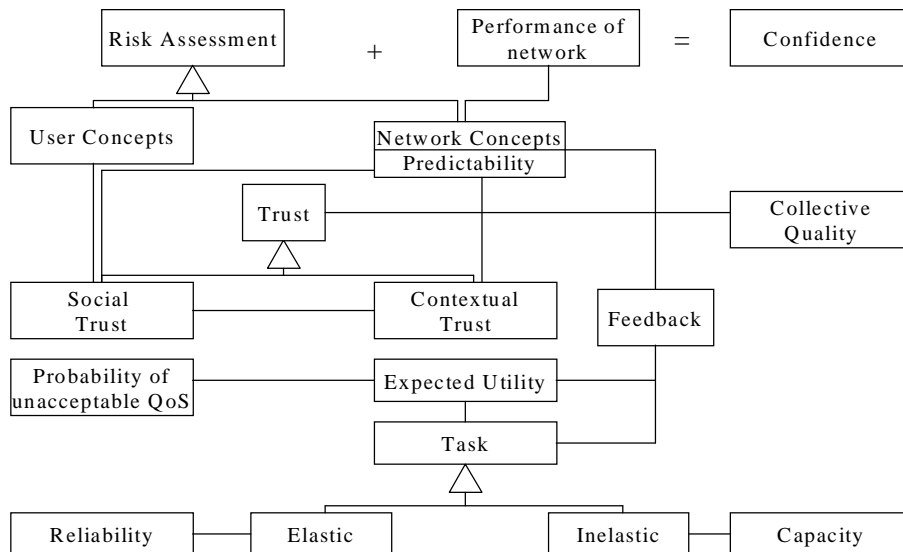


Figure 1: Concepts determining the acceptance of QoS and pricing mechanisms (Advanced Users)

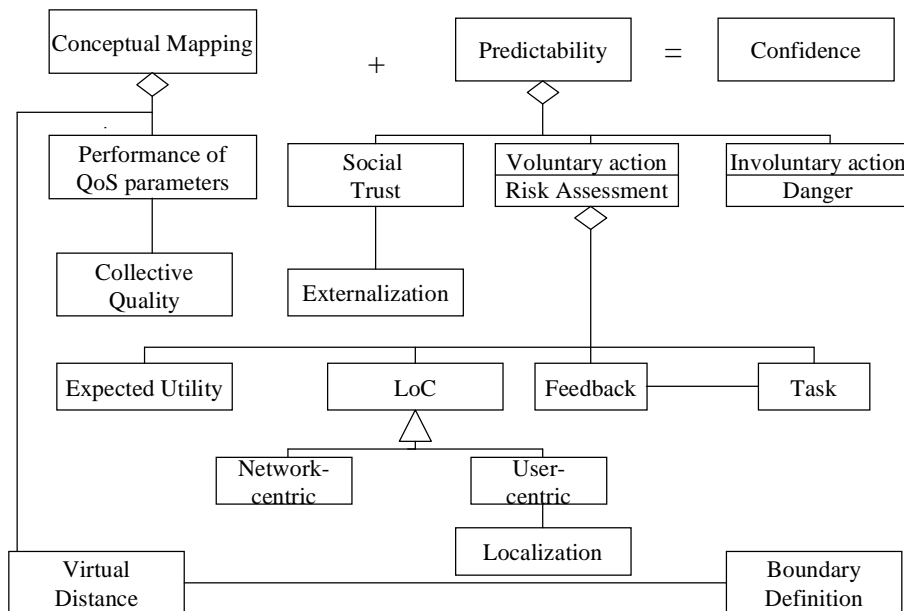


Figure 2: Concepts determining the acceptance of QoS and pricing mechanisms (Novice Users)

6. Discussion

6.1 Selecting pricing mechanisms

The findings reported in this paper have yielded a number of complex processes that determine users' perception of pricing mechanisms. Pricing mechanisms can be judged against these concepts in order to provide an indication of their acceptability. For example, *Smart Market* [14] was judged by novice participants to involve a greater *Risk* relative to other schemes discussed. This is because users felt unable to develop *Social Trust* due to the fact that the *Conceptual Metaphor* for the *Smart Market* - the auction - is a situation where bids are made face-to-face. In the latter situation bidders can barter for commodities against a known price. This metaphor would be difficult to support adequately in a multiplexed, multi-route architecture.

A pricing scheme in which quotas can be bought prior to network usage is more acceptable to users [15]. This mechanism involves a trade-off between users' need to be able to predict the performance of the network, and their need to dynamically evaluate its performance, and act on their evaluation. In addition, the quota scheme allows the *Value* of a *Task* to be considered in users' decisions of whether to use their quota.

The assessment of *Risk* is clearly an operation that users apply to future events. This is indeed why the *Expected Utility* - of QoS received in the future – influences the assessment. Users’ expectancies concerning the future state of the network are a strong determinant of their future actions [18]. The influence of *Risk* and *Expected Utility* suggests that payment could be made for the chances of poor QoS being configured from a network that has become congested. A charge would be levied, not for explicit network conditions, but for the conditions expected by the user. The theory of *Expected Capacity* is based on a similar idea [19]. The individual user’s need for predictable QoS could therefore become the very aspect for which she is prepared to pay.

The lack of *Network Concepts* in the models of novice users suggests that the management of the QoS required by an application should be partially encapsulated. Our findings suggest that the provision of the required amount of *Feedback* to such users should result in the acceptance of the configured level of QoS, and likely *Payment* for that quality. A popular suggestion for the implementation of a degree of automation at the application and network levels is the creation of user agents, constructed according to the *Feedback* requirements of the user. These user agents could interact with network agents in the network system in the manner of a ‘*QoS Broker*’ [20].

6.2 Implementing pricing profiles

A manner in which feedback encapsulation could be realized is through the implementation of a series of task profiles. The construction of each profile would be driven, not only by the need to capture users’ QoS requirements along a number of salient dimensions, but by the definition of those dimensions. Table 2 shows those factors that vary according to the *Value* placed on a particular *Task*, at a particular time.

Table 2: Example task profile for an email application: Variable Factors

<i>Concept</i>	<i>Value</i>	<i>Implementation</i>	<i>Consequence of correct implementation</i>
Purpose of Task	Inelastic	Dialogue to indicate purpose of task	Expected Utility confirmed Value of task represented
Conceptual Metaphor	Telephone	Dialogue to indicate purpose of task Feedback on ‘per-call’ basis	Expected Utility confirmed Value of task represented
Virtual Distance	High	Provide objective QoS feedback	Predictability of system enhanced
Collective Quality	High	Provide objective QoS feedback	Predictability of system enhanced

We have chosen to present the profile of a traditionally elastic application, in a situation where the purpose of the task is inelastic, in order to demonstrate the influence of the *Conceptual Metaphor* held by users. The task profile for this application requires the system to engage in a dialogue that enables the user to indicate their desire for, in this case, *Speed* of transmission, as opposed to *Throughput*. This ensures that the QoS dimension upon which users base conceptions of *Expected Utility*, is appropriate to the purpose of the *Task*.

As mentioned, if users perceive that *Virtual Distance* is relatively high they expect the QoS resulting from network usage to be relatively low. The same effect is seen if users perceive that the *Collective Quality* received on the network is relatively low (e.g., it is a time of day traditionally associated with congestion). Whilst this misconception may seem like a good way to reduce the demands made by users on the network, it does mean that the actual quality received will not conform to users' expectancies. Providing users with objective QoS feedback enables the accurate prediction of that QoS, and as we have shown, pricing mechanisms configured over a predictable system are more likely to be accepted in the long run.

Apart from profile dimensions that are specific to this particular application, an implemented system of profiles must consider general requirements made by users. These general requirements are based on factors arising from the real-world assessments user make concerning the operations of certain applications. Table 3, shows some of these general characteristics for the email example used above.

Table 3: Example task profile for an email application: General Factors

<i>Concept</i>	<i>Value</i>	<i>Implementation</i>	<i>Consequence of correct implementation</i>
Work	Low	Encapsulate routing processes	Positive rating of QoS
Risk	Dependent on implementation	Provide user-configurable feedback in high risk situations	Predictability of system enhanced
Confidence	Dependent on implementation	Implementation of lower level concepts	Acceptance of pricing mechanism

Including feedback showing that *Virtual Distance* is based on a misconception may not be appropriate for a user who possesses knowledge about the operations of the network. Clearly, any implemented pricing scheme must be flexible enough to represent the needs of different users. A way to do this when representing task profiles is to implement general dimensions in a core profile and allow those dimensions that are dependent on the task and user to be selectable as conditions apply.

7. Conclusions and future work

Whilst network technology is progressing towards providing performance guarantees for real-time applications in shared networks, the QoS requirements of users have so far been rarely considered. This paper has provided a description of the way in which users view QoS and pricing in a set of models that may be used for the prediction of users' acceptance of QoS and their consequent *Payment Behavior*. Additionally, results have shown that the relevance of QoS parameters and concepts associated with network pricing depend on users' level of knowledge and experience with networks, and on the *Task* being performed. While results suggest that it is the level of *Confidence* possessed by users in the performance of salient QoS drivers that is important, the sub-concepts that may form *Confidence* are variable according to *Task* and user group.

The complexity of models suggests that unfortunately, there are no simple mechanisms that will be widely applicable. Future work will help to identify the most important factors from the models, and subject them to further analysis. Arguably, the role of *Predictability* and *Feedback* should be tested whilst the user is performing a *Task* to which *Value* is ascribed. We have started a series of experiments where users dynamically adjust QoS for different tasks whilst being given feedback about the budget associated with the QoS chosen [21]. We intend to add and integrate these concepts with the current body of technical knowledge of QoS and pricing.

8. Appendix: List of concepts and their definitions

Boundary Definition: Refers to distinct QoS categorizations according to *Task*.

Collective Quality: User's perception of the QoS received by all network users.

Conceptual Mapping: The reliance on a real-world metaphor as a comparison for the QoS provided by the *Task*.

Confidence: Acceptance of a pricing mechanism and the *Value* of QoS delivered.

Contextual Trust: Trust which is placed on the operations of the network architecture.

Danger: An assessment of the user's action as involuntary.

Expected Utility: Subjective usefulness a user gains from salient QoS parameters.

Externalization: The tendency to attribute system functions to processes outside network operations.

Feedback: The amount and type of *Feedback* is a determinant of the level of *Predictability* of the system.

Involuntary Action: Where the actions of a system are perceived as unpredictable.

LoC: Refers to the attribution of control to either the network or to the user.

Localization: The attribution of received QoS to processes that are user-centric.

Metaphor: Real-world example by which novice users judge system operations.

Network-centric: The attribution of *LoC* to network-space.

Network concepts: Concepts that are associated with network actions/operations.

Predictability: The probability that the users' judgement of future QoS is accurate.

User-centric: The attribution of *LoC* to user-space.

User concepts: Concepts that are associated with user actions/operations.

Reliability: The degree to which the system is judged to be stable.

Risk Assessment: Assessment created by the inclusion of a value judgement on network performance. Risk is defined as the risk of not obtaining value for money from QoS drivers considered salient for the *Task*.

Social Trust: Attributed to actions of other network users, or to user-ISP contracts.

Value: The subjective usefulness of the data received.

Virtual Distance: The physical distance between the location of the user and the destination of the data unit.

Voluntary Action: Where the actions of a system are perceived as predictable.

Work: Number of real-world *Task* operations applications are required to perform.

References

[1]. Gupta, A., Stahl,D.O., and Whinston, A.B. (1995). Priority pricing of integrated services networks. In McKnight.L.W., and Bailey,J.P. (Eds). Internet economics. MIT Press.

[2]. Johnston, J. (1972). Econometric methods (2nd Ed.).M^cGraw Hill.

[3]. Watson, A., and Sasse, M.A. (1997). Multimedia conferencing via multicast: Determining the quality of service required by the end user. Proceedings of AVSPN'97 – International Workshop on Audio-Visual Services over Packet Networks, pp.189-194.

[4]. Strauss, A., & Corbin, J. (1990) Basics of qualitative research: Grounded theory procedures and techniques. Sage.

[5]. Adams, A. and Sasse, M.A. (in press). Users are not the enemy. Communications of the ACM.

[6]. Adams, A., Sasse, M.A., and Lunt, P. (1997). Making passwords secure and usable. In Thimbleby, H., O'Conaill, B., Thomas, P. (Eds.). People & Computers XII, Proceedings of HCI'97 in Bristol, August 12-15 p. 1-19. Springer.

[7]. Lunt, P.K., and Livingstone, S. (1996). Rethinking the focus group in media and communications research. Journal of Communication, 46, (2).

[8]. Shneiderman, B. (1992). Designing the user interface: Strategies for effective human-computer interaction. (3rd Edition). Addison-Wesley.

[9]. Perry, C. (1995). Travelers on the Internet: A survey of Internet users. Online, 19, (2), pp.29-32.

- [10]. Brittan, D. (1997). Spending more and enjoying it less? Technical Review 100, 5, pp.11-12. Available from:
<http://www.web.mit.edu/afs/athena/org/t/techreview/www/articles/july97/brittan.html>
- [11]. Wang, Z. (1997). User-share differentiation (USD): Scalable bandwidth allocation for differentiated services. Available from:
<ftp://ftp.ee.lbl.gov/papers/dsarch.pdf>.
- [12]. Cosgrove, J.G., and Linhart, P.B. (1979). Customer choices under local measured telephone service. *Public Utilities Fortnightly*, 30, pp.27-31.
- [13]. Mackie-Mason, J.K., and Varian, H.R. (1997). Responsive pricing in the Internet. In McKnight.L.W., & Bailey,J.P. (eds). *Internet Economics*. MIT Press.
- [14]. Mackie-Mason, J.K., and Varian, H.R. (1994). Economic FAQs about the Internet. *Journal of Economic Perspectives*, 8, (3), pp.24-36.
- [15]. Bohn, R., Braun, H.W., and Wolff, S. (1994). Mitigating the coming Internet crunch: Multiple service levels via Precedence. *Applied Network Research Technical Report, GA-A21530*. University of California.
- [16]. Odlyzko, A. (1997). A modest proposal for preventing Internet congestion. Available from <http://www.research.att.com/~amo/doc/recent.html>.
- [17]. Luhmann, N. (1991). *Risk: A sociological theory*. De Gruyter.
- [18]. Huberman, B.A., and Lukose, R.M. (1997). Social dilemmas and Internet congestion. *Science*, 277, (July), pp.535-537.
- [19]. Clark, D. (1997). Internet cost allocation and pricing. In McKnight.L.W., and Bailey, J.P. (Eds). *Internet economics*. MIT Press.
- [20]. Nahrstedt, K., and Smith, J.M. (1995). The QoS broker. *IEEE Multimedia*, 2, (1), pp.53-67.
- [21]. Bouch, A., Watson, A. and Sasse, M. A. (1998). 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.