# QoSPlan: A Process for QoS-Aware IP Network Planning using Accounting Data and Effective Bandwidth Estimation

Alan Davy, Dmitri Botvich, Brendan Jennings
Telecommunications Software & Systems Group,
Waterford Institute of Technology, Waterford, Ireland
{adavy, dbotvich, bjennings}@tssg.org

*Abstract* - **In this paper we describe QoSPlan – a generalized process for preparing information relevant to QoS-aware IP network planning. QoSPlan is designed to reduce the cost of deployment and maintenance of traditional network monitoring systems for service providers, while maintaining a sufficient degree of precision. The process involves analysis of pre-existing accounting data to estimate a network-wide demand matrix. Part of this estimation process relates to the generalization of QoS-related effective bandwidth measurements taken from a set of collected packet traces. Based on experience, we offer recommendations on how to appropriately realise QoSPlan to maximise its accuracy and effectiveness when applied to different network traffic scenarios.**

## I. Introduction

Currently, establishing input for the network planning process depends on the use of dedicated hardware devices collecting large volumes of network traffic data that is then analyzed to identify a network configuration design reflecting estimated demand and specified Quality-of-Service (QoS) requirements. The use of dedicated hardware means the approach is expensive, incurring costs in hardware procurement and maintenance, in addition to significant training and operational costs.

We argue that the network planning process can be made more cost-effective, whilst maintaining a sufficiently high degree of accuracy, by reusing alternative sources of information residing within the network accounting system. In previous work [1] we proposed an initial QoS-aware network planning architecture based on network accounting data as an alternative source of estimating network demands. We also proposed a method of capturing a relationship between estimated network demand and required effective bandwidth levels specific to outlined QoS targets through packet trace analysis. Furthermore, we performed an economic analysis comparing deployment and operational cost of our proposed architecture compared to a traditional network planning system based on dedicated metering equipment [2], which showed that relative cost savings can be as high as 80%.

In this paper we apply our experience from previous work and propose QoSPlan – a generalized process service providers can follow to prepare input for QoS-aware IP network planning. We decompose the process into three phases, namely: (1) acquisition of data, (2) analysis and mediation of data, and (3) preparation of input for QoS-aware IP network planning. Whilst describing these phases we provide recommended settings and guidelines to improve accuracy in estimation of effective bandwidth levels of monitored traffic. We believe QoSPlan will enable service providers or network operators reduce costs in network planning scenarios where the level of accuracy can be somewhat flexible, for example, the planning of QoS requirements on traffic over periods of several hours or days.

This paper is organized as follows: §2 discusses related work in the areas of QoS-aware network planning and effective bandwidth estimation techniques. §3 outlines our proposed network accounting and planning architecture. §4 presents a particular scenario related to an IPTV service deployment to provide an example of a typical usage domain for our proposed process. We describe QoSPlan in §5, providing salient recommendations based on experience, illustrated by relevant results. §6 summarizes the paper and outlines areas for future work.

## II. Related Work

The network planning process generally requires three sources of input [3]: (1) attributes associated with the current traffic demands on the network which collectively specify their behavioral characteristics; (2) attributes associated with resource constraints on the network topology; and (3) a constraint based routing definition framework which plans routing of traffic subject to (1) and (2). This paper focuses on outlining a generic process for the establishment of QoS requirements associated with the current traffic demand, thus addressing requirement (1) of input for network planning.

Traditional planning approaches analyze vast amounts of network traffic to generate an in-depth model of network behavior. From this, various network configurations are proposed to optimize network performance metrics, such as maximum link utilisation, or minimum packet delay. These network planning approaches are effective in what they are

This work has received support from Science Foundation Ireland via the Autonomic Management of Communications Networks and Services programme (grant no. 04/IN3/I4040C) and the 2005 Research Frontiers project Accounting for Dynamically Composed Services (grant no. CMS006).

designed to do, however this level of accuracy comes at significant cost. Research into alternative sources of input to the planning process with the goal of reducing the cost of the planning process has been undertaken, see for example [4, 5, 6]. These works propose deriving a network demand matrix via various analysis techniques, including packet sampling to determine trajectory flow of traffic through the network [4] and analysis of IP flow records to estimate network demand [5, 6]. A vital consideration with these approaches is how to integrate QoS-related information on traffic.

Typically, network planning processes perform QoS-related planning based on effective bandwidth requirements per traffic class. We propose a method of relating required effective bandwidth to associated network demand per traffic class, specific to outlined QoS targets, where the effective bandwidth is defined, in line with [7], as the minimum amount of bandwidth required by a traffic stream to maintain specified QoS-related targets. There are many methods of theoretically estimating the effective bandwidth value for different traffic models, see for example [8, 9]. We adopt an empirical approach based on packet trace analysis; this approach is similar to those outlined in [10, 11]. We use the concept of an effective bandwidth coefficient to capture a relationship between estimated network demands from accounting records and specified QoS targets per traffic class.

## III. Network Accounting and Planning Architecture

The architecture illustrated in Fig. 1 extends the traditional usage based network accounting architecture to facilitate network demand and effective bandwidth estimation processes. The architecture draws input from two sources, namely mediated data from collected accounting records, and packet traces analyzed to estimate effective bandwidth levels for particular QoS targets.

At the Mediation entity, we propose to deploy additional planning mediation rules to allow us to extract appropriately formatted information from accounting records. This raw information is then used to generate a QoS-aware network demand matrix, outlining current network demands between edge node pairs within the network, per traffic class. We specify our approach to estimating the demand matrix from accounting records in [1], and evaluate this approach in [2].

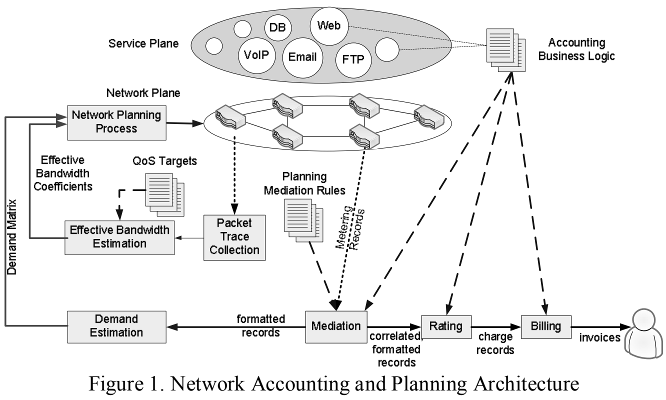


Figure 1. Network Accounting and Planning Architecture

This analysis demonstrates relative error in estimation of demand from accounting records to be around 10% for the scenarios modelled. We believe this loss in accuracy is acceptable, as network planning for future traffic demands influenced by human usage patterns can never be accurately predicted. For effective bandwidth coefficient estimation, a set of packet traces are collected from various edge nodes on the network. Each trace is collected specific to a traffic class over a relatively short time frame (on the order of minutes). The collection of traces is then analyzed offline to estimate an effective bandwidth coefficient value for each traffic class.

## IV. Demonstrative Scenario

In this section we describe an IPTV service provision scenario representing a typical application domain in which QoSPlan could be profitably deployed. Fig. 2 depicts a scenario involving a single IPTV service provider, a core network under that service provider's control, and a number of customers. Customers are connected to the network through a high speed access link such as xDSL. IPTV services typically have very strict guidelines on Quality of Service [12], such as minimum packet delay of $10^{-7}$, with high bandwidth requirements of up to 20Mbps. The IPTV service provider wishes to plan its network to ensure adequate bandwidth is being provisioned to maintain specified QoS delay targets on traffic. The service provider wishes to reduce the cost of preparing input for the QoS-aware network planning process, by avoiding installation of expensive network monitoring equipment and reusing existing accounting information collected via an existing metering infrastructure.

Within the scenario accounting data is already collected at a number of different levels. Firstly at the IP layer, edge routers collect IP flow [13] records to monitor levels of traffic passing between end node pairs over the network. These flow records are used with RADIUS [14], or DIAMETER [15], accounting records collected at the Digital Subscriber Line Access Multiplexer (DSLAM) to associate traffic monitored with customers for billing purposes. Finally service-related billable events are collected from customer premises equipment (CPE) in the format of IP Data Records (IPDRs) [16] (or similar metering record formats). For the purpose of our proposed network planning process, the packet level IP flow records are fundamental to estimating mean network demand between edge nodes of the network. QoSPlan reuses these records, already collected for network accounting purposes, and enhance them with QoS-related information to

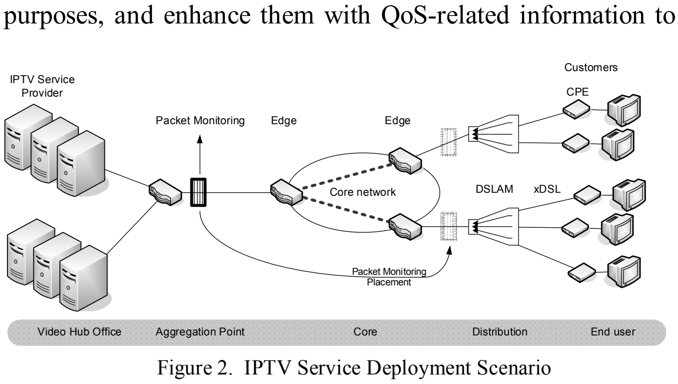


Figure 2. IPTV Service Deployment Scenario

be usable as input to an appropriate QoS-aware network planning process.

## V. QoSPlan: A Process for QoS-aware IP Network Planning using Accounting Data and Effective Bandwidth Estimation

The QoSPlan process outlined in this section formalizes the procedures and decisions a service provider must adopt to generate relevant and accurate QoS-related input for QoS-aware IP network planning. The following knowledge and capabilities are necessary prerequisites for its adoption:

1. The availability of accounting data capturing network demand between end nodes per traffic class. For example, in the IPTV scenario IP flow records will be collected at the ingress routers of the network;
2. Information about accounting data collection parameters, such as sampling intervals. We demonstrate later that configuration of these parameters can affect accuracy in estimating network demand from accounting data;
3. The ability to map end nodes to corresponding entry points of sourced and sinked traffic across the core network, i.e. a mapping between end nodes and edge nodes. This is vital in preparing a network wide demand matrix;
4. The ability to collect short packet traces from ingress edge node interfaces, for further analysis. From our IPTV scenario, we highlight appropriate packet trace collection points, where a monitoring device can be physically positioned;
5. Knowledge of specified QoS targets being imposed on QoS traffic classes. For example, for IPTV services these targets are very strict, as outlined by the DSL-Forum [12]. On other types of traffic such as web and email, QoS targets can be more relaxed, leading to different resource requirements.

We are in a position to make recommendations on how the process should be realised, based on the results of our previous evaluation work [1, 2]. This will aid others applying QoSPlan in making vital decisions and improving efficiency in network planning and accuracy in estimating effective bandwidth. We break down QoSPlan into three phases (illustrated in Fig. 3): (1) acquisition, (2) analysis and mediation, and (2) preparation. In the following we outline a number of steps to be taken within each phase, highlight important observations from experience, and make some specific implementation recommendations.

### A. Phase 1: Acquisition

QoSPlan depends on the acquisition of two independent sources of data: packet traces and network accounting data.

#### 1) Phase 1a: Acquiring Packet Traces

Packet traces are analyzed to establish a relationship between network traffic behavior and specified QoS targets. A relatively large number of packet traces must be collected per traffic class from various edge node locations within the network to ensure a general view of network traffic behavior is captured. Trace collection points should be at the ingress edge of the network, since we require that data collected is for traffic unshaped by the network itself. We wish to characterize behavior of traffic before it enters the network, as the main objective of network planning is to configure the core network in such a way as to manage incoming traffic demands. Packet traces also need to be collected over an appropriate duration. If packet trace durationss are too small or large, behavioral analysis of the traffic can lead to misleading results. We recommend packet trace durations of around 4-5 minutes.

A single monitoring device can be used to directly collect packet traces from the network. This device can be moved between ingress points within the network, and is only operational during collection. This differs from dedicated network monitoring devices that are required to be operational over long time periods and to be deployed network wide.

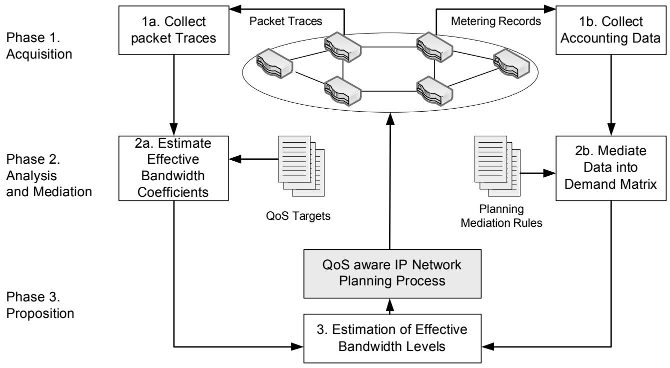

Figure 3. Process for developing input for QoS aware network planning

#### 2) Phase 1b: Acquiring Accounting Data

We estimate network demand per traffic class from available accounting data. To achieve this we must ensure the required information is available from the accounting data. As QoSPlan estimates demand per traffic class, each accounting record must contain QoS-related information that can identify a volume of traffic to a traffic class. To identify where on the network demand flows, accounting data needs to identify source and destination end points to enable QoSPlan map demand appropriately to the network. A common method of accounting for network demand is though the use of IP flow record collection. This approach captures traffic in the form of flows between two end nodes where packets are mapped to flows based on matching IP header information. Table 1 below depicts a sample flow record where 1032 packets matching the source and destination IP address and type of service passed a particular observer interface.

A common approach taken to further reduce the overhead acquisition of accounting data has on network performance, is to use sampling in the observation of packets [13]. This can have a knock-on effect of reducing accuracy in estimating network demand, as we have shown in [1, 2]. For the case of bursty elastic traffic such as HTTP, email and FTP, sampling at high intervals can dramatically reduce the accuracy in estimation of network demand, as this type of traffic is not well represented by flow-based accounting. The process of

TABLE I. Sample Accounting Data Flow Record

| Src IP | Dest IP | DSCP | Size | Packets | Start | End |
|---|---|---|---|---|---|---|
| 10.37.2.22 | 10.34.1.118 | EF | 332049 | 1032 | 12:35.31 | 12:48.22 |

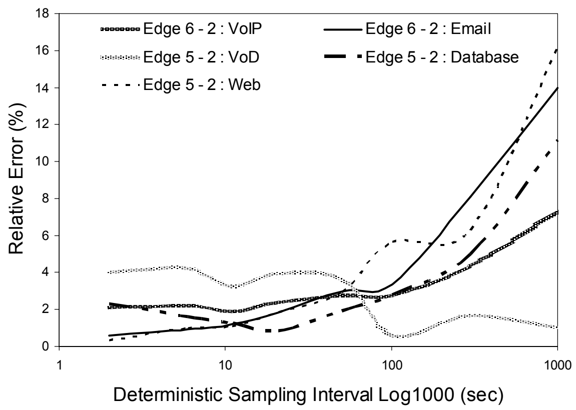


Figure 4. Demand estimate relative error for deterministic packet sampling.

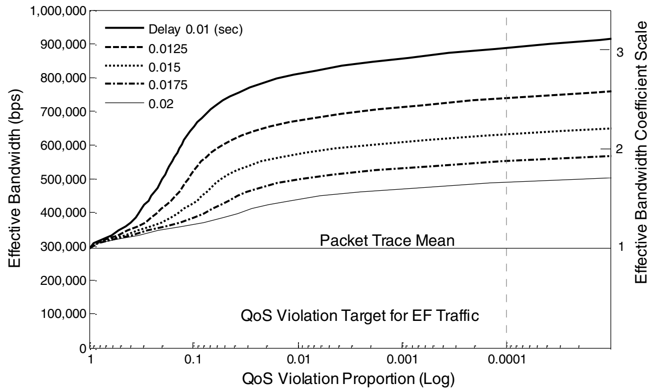


Figure 5. Varying QoS delay targets on effective bandwidth estimates.

sampling can miss large quantities of data passing a point. On the other hand, stream based traffic such as Voice-over-IP (VoIP) and Video-on-Demand (VoD) translates well to flow record representation. Sampling at high intervals has little impact in estimating demand from associated records as traffic streams between end points are relatively constant in comparison to elastic traffic.

Fig 4 demonstrates the variation in accuracy as sampling intervals increase for different application traffic. In this case the relative error in estimating network demand flow records increases up to 18% for elastic traffic such as email, HTTP and database, when a packet sampling interval of 1 in 1000 packets is used. However, relative error remains lower than >6% for streaming traffic such as VoD and VoIP. This suggests that streaming traffic throughput remains more stable over longer periods, justifying a larger sampling interval. We recommend that if configuration of sampling is possible, low sampling interval for elastic traffic, say 1 in 200 packets, should be used. For stream-based traffic a higher sampling interval can be used, such as 1 in 2000 packets.

B. *Phase 2: Analysis and Mediation:*

The analysis and mediation phase extracts information from the acquired data relevant to the process of network planning. The process attempts to establish a relationship between network traffic behavior and required effective bandwidths for associated QoS targets through packet trace analysis. QoSPlan involves the mediation of collected accounting data into a demand matrix, mapping network demand to edge node pairs per traffic class.

1) *Phase 2a: Estimation of Effective Bandwidth from Collected Packet Traces.*

QoSPlan estimates the effective bandwidth of each collected packet trace using the empirical method outlined in [2]. In brief, this method replays the collected packet trace through a simulated FIFO queue at various queue service rates, until a service rate is found at which specified QoS targets on packet delay are maintained – this rate is taken as the effective bandwidth estimate, $R_{eff}$. Note that $R_{eff}$ is specific to the packet trace in question and to the particular QoS target.

$R_{eff}$ is then used to calculate effective bandwidth coefficient, which defines the relationship between the mean throughput of the packet trace and the estimated effective bandwidth for the trace.

$$k_i = \frac{R_{eff,i}}{mean_i}$$

The coefficient $k_i$ is calculated for all packet traces per traffic class. These ranges of coefficients allow the process to generalize a relationship between estimated network mean demands and effective bandwidth requirements to suit specified QoS targets per traffic class. A representative coefficient is calculated by taking the 95$^{th}$ percentile of these collected coefficients. The significance of this value is that the network planning process generally plans for near peak traffic. The 95$^{th}$ percentile of this range identifies the relationship between near peak traffic and its associated QoS target. If we have the mean value of this range, any traffic which exceeds this effective bandwidth / mean throughput relationship may not be adequately provisioned for within the planning process, thus increasing the level of QoS violations within the network.

An important decision to be made here is the choice of appropriate QoS targets for the traffic being analyzed. If the targets are too strict, effective bandwidth requirements may be unrealistic to provision within the network. If QoS targets are too relaxed, service performance can degrade, impacting on the customer's Quality of Experience. Fig. 5 demonstrates the effect of varying the QoS target on a single video packet trace has of the estimated effective bandwidth coefficient. The mean of the packet trace analyzed here is around 300 Kbps. To ensure QoS targets are maintained, a provisioning of 900Kbps is required for this stream. The graph shows a corresponding effective bandwidth coefficient of approximately 3. As the delay target and traffic proportions change, so does the required effective bandwidth.

Based on experience we recommend appropriate QoS delay targets for the following for three DiffServ traffic classes. Best Effort generally aggregates bursty, low priority traffic such as HTTP, email and FTP. As this type of traffic has high delay tolerance, QoS delay targets can be relaxed to about (100ms, 0.01). This target states that only 0.01

proportion of total traffic within this traffic class can experience packet delay of greater than 100ms. Assured Forwarding traffic tends to have a lower tolerance to packet delay, such as applications making database requests, require speedy responses. An appropriate QoS target for this traffic class could be (40ms, 0.001). The traffic class with the lowest tolerance to packet delay is Expedited Forwarding traffic. Traffic aggregated to this traffic class includes Video on Demand, IPTV, and VoIP. A common QoS target for this traffic class would be (0.01, 0.0001).

*2) Phase 2b: Mediate accounting data into network demand matrix.*

From the available accounting data we need to populate a matrix of network demand between edge node pairs per traffic class. The demand matrix is vital input to any network planning process. To populate the network demand captured by accounting data into the associated dimension, source and destination nodes need to be mapped to corresponding entry points to the core network, in other words access routers (edge nodes). Table 2 outlines an example static mapping between edge nodes and end nodes. This can be queried to map an accounting data record to an appropriate dimension within the demand matrix. Based on this table, the record in Table 1 would map to the dimension within the demand matrix for source edge node A and destination edge node D. We do not discuss the method of creating this mapping, although there are approaches such as interrogation of routing tables within access routers (as outlined, for example, in [6]).

Network demand is calculated and populated into the demand matrix over specified measurement intervals. If a network accounting record relates to a flow that started and ended within a measurement interval, the total traffic volume captured by the record is added to the appropriate demand matrix entry. If, on the other hand, the accounting record spans multiple time intervals, the volume must be divided proportionally between the relevant intervals. As the distribution of packet inter-arrival times and packet sizes throughout a packet stream are not generally uniform, dividing the accounting record volume proportionally between measurement intervals can lead to inaccuracies in estimation of network demand over an interval.

Fig. 6. demonstrates this effect for two different traffic types, email and VoIP. The graph shows the relative error between the estimation of network demand from flow records using a number of different measurement intervals, and directly measured demand from collected packet traces over the same measurement interval. As VoIP has a more stable flow of packets estimation of demand over a short interval can be accurately captured by a flow, allowing accurate estimation of demand over small intervals. On the other hand email

TABLE II. Mapping between edge nodes and end nodes

| Edge Node | End Point Mappings | | |
|---|---|---|---|
| A | 10.37.1.*, | 10.37.2.*, | 10.37.3.* |
| B | 10.36.1.*, | 10.36.2.*, | 10.36.3.* |
| C | 10.35.1.*, | 10.35.2.*, | 10.35.3.* |
| D | 10.34.1.*, | 10.34.2.*, | 10.34.3.* |

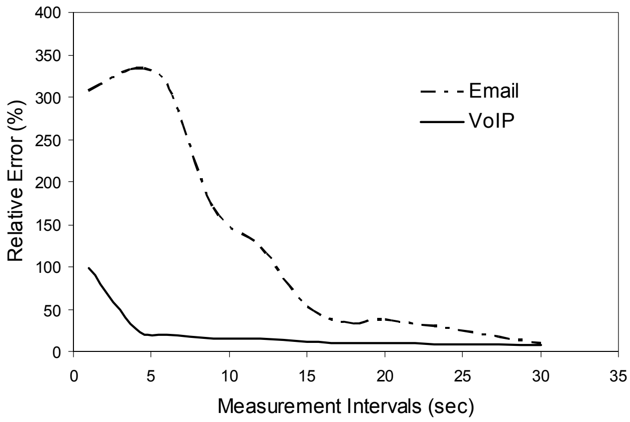


Figure 6. Demand estimate relative error with varying measurement interval.

traffic is quiet bursty, thus large amounts of traffic can group together over short periods of time, thus the choice of measurement interval has a more significant impact. By representing this traffic in flow record format, the burst intensities are levelled out over the duration of the captured flow, thus at short measurement intervals, these bursts are not represented in the calculation of network demand. This can be seen in the large degree in relative error demonstrated in Fig 6. Based on these observations we recommend a minimum measurement interval of around 1 minute, as smaller measurements can affect accuracy in demand estimation per interval.

### *C. Phase 3: Preparation*

In Phase 3 a matrix of estimated effective bandwidths per traffic class for input to the network planning process is prepared. This is achieved by multiplying the appropriate effective bandwidth coefficient by the estimated network demand between edge node pairs for the particular traffic class. A critical decision here is the choice of an appropriate representative effective bandwidth coefficient from the range collected per traffic class. As network planning is predominantly based on provisioning for near peak traffic, we recommend choosing the 95$^{th}$ percentile value of this range. The 95$^{th}$ percentile represents the relationship between near peak traffic and the associated effective bandwidth required for the specified QoS target. In our previous work, we have found a number of contributing factors that can affect the accuracy in estimating the effective bandwidth through this method. The main sources of error emerge from the inaccuracies inherent within the data collected in phase 1, such as the packet traces collected, and the accounting data collected and the inappropriate analysis and mediation of this data in phase 2. If our recommendations are followed where possible, we believe they will help reduce the degree of error resulting from the effective bandwidth estimation process.

## VI. CONCLUSIONS AND FUTURE WORK

We have presented QoSPlan – a process for preparation of input for QoS-aware IP network planning based on mediation of pre-existing accounting data and analysis of a limited number of representative packet traces collected from the network. QoSPlan outputs a matrix of estimated effective bandwidths per traffic class between node pairs, which can subsequently be used for network planning, and indeed other purposes. We contend that adoption of QoSPlan has the potential to greatly reduce the cost of network planning, as it allows the service provider replace their costly dedicated device metering architecture. To facilitate a gradual changeover it would be straightforward modify the process to utilise a hybrid architecture incorporating data collected from dedicated metering devices and data from mediated accounting data.

For future work we plan to apply QoSPlan within a specific application domain – IPTV. In particular we will apply QoSPlan's effective bandwidth estimation method as a source of input to QoS-focussed admission control strategies for Video-on-Demand traffic. We plan to augment the base information collected by QoSPlan with additional information resident in the accounting system, for example programme cost, programme duration, and file size. This information can be used in conjunction with the effective bandwidth estimates to realise admission control algorithms that maximise revenue for the service provider, whilst minimising the likelihood of QoS violations.

## REFERENCES


[1] Davy, A., Botvich, D., Jennings, B., "An Efficient Process for Estimation of Network Demand for QoS-aware IP Network Planning", in Proc 6th IEEE Workshop in IP Operations and Management, Eds. Parr, G., Malone, D., Ó Foghlú, M., LNCS Autonomic Principles of IP Operations and Management, vol. 4268/2006, pp. 120-131, Springer Berlin/ Heidelberg, ISBN 978-3-540-47701-3, 2006.

[2] Davy, A., Botvich, D. D., Jennings, B., "On the use of Accounting Data for QoS aware IP Network Planning", accepted at 20th International Teletraffic Congress, ITC 20, Ottawa, Canada, 2007.

[3] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and McManus, J., "Requirements for Traffic Engineering Over MPLS",Informational RFC 2702, last accessed 13/03/2007, http://www.ietf.org/rfc/rfc2702.txt.

[4] Duffield, N., Grossglauser, M., "Trajectory sampling for direct traffic observation," New York, NY, USA: ACM Press, pp. 271-282, 2001.

[5] Feldman, A., Greenberg, A., Lund, C., Reingold, N., Rexford, J., and True, F., "Deriving traffic demands for operational IP networks: methodology and experience," IEEE/ACM Transactions on Networking, vol. 9, no. 3, pp. 265-280, 2001.

[6] Papagiannaki, K., Tatf, N., and Lakhina, A., "A Distributed Approach to Measure Traffic Matrices", Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement, pp 161-174, ACM Press, New York, NY, USA, ISBN 1-58113-821-0, 2004.

[7] Kelly, F., "Notes on Effective Bandwidth", in Stochastic Networks: Theory and Application, Royal Statistical Society Lecture Notes Series, F. P. Kelly, S. Zachary, and I. B. Ziedins, Eds., vol. 4, pp 141-168. Oxford University Press, ISBN 0-19-852399-8, 1996.

[8] Kesidis, G., Walrand, J., and Chang, C. S., "Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources", IEEE Transactions on Networking, vol. 1, issue 4, pp. 424-428, ISSN 1063-6692, August, 1993.

[9] Spitler, S. L. and Lee, D. C., "Integrating effective-bandwidth-based QoS routing and best effort routing", IEEE INFOCOM 2003, The Conference on Computer Communications, vol. 2, pp. 1446–1455, ISBN 0-7803-7752-4, ISSN 0743-166X, 2003.

[10] Botvich, D. D. and Duffield, N., "Large deviations, the shape of the loss curve, and economies of scale in large multiplexers", Queueing Systems, vol. 20, pp. 293-320, Springer Netherlands, ISSN 1573-9443, 1995.

[11] Baras, J. S. and Liu, N. X., "Measurement and simulation based effective bandwidth estimation", IEEE Global Telecommunications Conference, GLOBECOM, vol. 4, pp. 2108–2112, ISBN 0-7803-8794-5, 2004.

[12] Rahrer, T., Fiandra, R., Wright, S., "Triple-play Services Quality of Experience (QoE) Requirements", DSL-Forum, Technical Report TR-126, Architecture & Transport Working Group, Dec 2006.

[13] Sadasivan, G., Brownlee, N., Claise, B., and Quittek, J., "Architecture for IP Flow Information Export" IETF Internet Draft, last accessed 27/02/2007, http://www.ietf.org/internet-drafts/draft-ietf-ipfixarchitecture-12.txt, 2006.

[14] Rigney, C., "RADIUS Accounting", IETF RFC 2866 Internet Engineering Task Force, Network Working Group, www.ietf.org/rfc/rfc2866.txt, 2000.

[15] Calhoun P., Loughney J., Guttman E., Zorn G., and Arkko J., "Diameter Base Protocol", RFC 3588 Internet Engineering Task Force, Network Working Group, www.ietf.org/rfc/rfc3588.txt, 2003.

[16] Cotton, S., Kleinmann, A., Resource data Collection, Analysis & Control (RDC) For IP-Based Services, "Service Specification – IP Television (IPTV)", Version 3.5-A.0.0, IPDR.org, October 31, 2006.