THE DEVELOPMENT AND EVALUATION OF A MONITORING TECHNIQUE FOR M-FAC

ALI MAQOUSI and FRANK BALL

Distributed Systems Research Group
School of Computing and Mathematical Sciences
Oxford Brookes University
Wheatley Campus
OX33 1HX, UK
amaqousi, fball@brookes.ac.uk

Abstract: This paper presents the development of a measurement based flow admission control (M-FAC) mechanism, which is used in conjunction with Class Based Queuing (CBQ) to support guaranteed services in multiservice packet switched networks. M-FAC uses an analytical queuing model to decide if a new flow may be admitted into a particular class, and bases this decision on the parameters provided by the flow's source. However, once the new flow has been accepted, M-FAC continually monitors the arrival process in order to determine the true load on the class' queue. This true load is then used in subsequent admission tests. The analytical model was developed in previous work, and this paper focuses on the development of the monitoring process.

This paper introduces the monitoring process and explains how it is used by M-FAC to continually update the estimate of the instantaneous load vector, and presents the results from a simulation experiment designed to evaluate the effectiveness of this monitoring process.

Keywords: Traffic Monitoring, Multiservice Networks, M-FAC.

1. INTRODUCTION

In previous work we have been investigating a number of issues relating to the problems of providing multiservice in packet switched networks. Our work has focused mainly on the development and evaluation of bandwidth partitioning and resource control mechanisms, with particular emphasis on providing support for continuous media traffic. Although the work is aimed at packet switched networks in general, it is applicable to future IP networks that offer DIFSERVE, INTSERVE or any similar scheme that may be developed in the future.

This work was divided into three main areas:

1) A comparative evaluation of Class Based Queuing (CBQ) and Weighted Fair Queuing (WFQ), in which each mechanisms was presented with identical traffic profiles [Callinan, 2000]. An important consideration of this work being to assess the effectiveness of each system in supporting guaranteed services to continuous media sources, and not just different classes of computer data traffic.

2) The development of reactive congestion control mechanisms to support adaptive continuous media [Tater, 2000].

3) An investigation into the problems of mapping IPv6 flows onto ATM cell streams [Basu, 2000].

The results from this work, which involved both analytical modelling and simulation studies, have led to the development of a general model for multiservice support within packet switch routers.

Currently, we are using the theoretical results from our previous work in the development and

implementation of a prototype multiservice router [Li, 2000]. A major motivation for carrying out this implementation work is to further validate the control mechanisms we have developed during the theoretical phase of our research. We are also continuing with the development of a Measurement Based Flow Acceptance Control (M-FAC) mechanism for use with CBQ, which is needed to support guaranteed services to continuous media applications that require a consistent quality.

Mechanisms such as CBQ can partition bandwidth and provide isolation between different classes of traffic. However, in order to provide a guarantee to the individual flows from within a class, it is also necessary to control admission to that class. Admission control mechanisms are responsible for deciding if a new flow may be accepted, or if it should be rejected. This is a predictive process that will generally employ some form of analytical queueing model and use the traffic parameters supplied by the sources. However, it is generally difficult to obtain accurate a priori estimates for these traffic parameters, which can lead to an inaccurate control of resources. This can lead to low network utilization, or in the worst case, failure to meet QoS requirements. One solution to this problem is to employ a measurement-based admission control mechanisms that combines prediction with monitoring.

This paper presents the latest stage in the development of the M-FAC mechanism, and in particular reports the results of a simulation experiment designed to evaluate the associated traffic monitoring techniques. The remainder of the paper is organized as follows: section 2 presents a general model for multiservice that we have developed using the results of our previous work; section 3 introduces M_FAC and considers its traffic monitoring requirements; section 4 outlines a new technique for monitoring a high percentile of queue occupancy; section 5 describes the process of monitoring the mean and variance for both the packet inter-arrival times and packet lengths of a continuous media source; section 6 presents a simulation model designed to evaluate the mechanism introduced in section 5, and reports the results of validation experiments; and finally, in section 7 we conclude and outline future work.

2. A MODEL FOR MULTISERVICE

This section presents a general model for multiservice, which is based on the findings of our previous work. Throughout this work our approach has been to consider the problems of providing multiservice in a general context, with a focus on the medium to long-term future. Therefore our investigations have not been limited by legacy factors or the constraints of current commercial considerations. A general model for a multiservice router is shown below in figure 1.
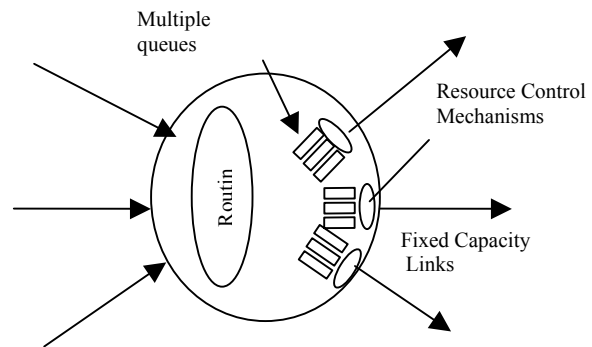
Figure 1: A General Model for a Multiservice Router

Our work was not primarily concerned with the routing stage of the model. Ideally the routing engine should be engineered to be sufficiently powerful to allow the routing process to be internally "non-blocking", thereby ensuring that resource contention takes place only at the output links. The main focus of our work has been on resource control and the link sharing model associated with the output stages of the router.

The comparative evaluation of WFQ and CBQ [Callinan, 2000] set out to compare the performance of the two systems in three main areas: the ability to provide bandwidth partitioning and isolation between different flows of traffic, particularly under conditions of heavy network loading; the response time delay characteristics, again under heavy load conditions; and the ability to be fair when distributing surplus bandwidth between competing flows, at times when network was lightly loaded. Although it was shown that both WFQ and CBQ could offer adequate isolation to continuous media sources, CBQ was shown to

offer a much lower response time delay at times of heavy loads, provided that the sources were grouped into appropriate classes. WFQ was shown to give a lower response time delay once the network load drops below a certain level. This would be expected due to WFQ being a work-conserving algorithm. However, this was not considered to be of any real advantage to continuous media since the response time delay under heavy loads would be the most important factor for determining the jitter smoothing requirements at the receiver. Another important observation was that the mathematically proven bound offered by WFQ (the Parekh Bound) [Parekh, 93] was seen to be far too conservative to be of general use.

Both systems were shown to be capable of guaranteeing a minimum throughput and providing a fair share of surplus bandwidth to competing flows of computer data. However, the work-conserving rate controller used with CBQ for the computer data classes [Floyd, 95] was seen to be sensitive to traffic characteristics. Certain parameters of this mechanism need to be adjusted to cope with either random or self-similar traffic sources. On the other hand, WFQ was shown to be completely insensitive to this problem, and worked equally well with both random and self-similar traffic without any need for adjustment.

We conclude from these findings that generally, continuous media will be best served by a simple CBQ system using a non-work-conserving rate controller and computer data would be best controlled by WFQ. Fortunately, it is possible to implement a hybrid solution, and thereby meet the needs of both types of traffic in a single system. Figure 2 shows the basic structure of a hybrid CBQ-WFQ system that is very similar to the model we have built using the OPNET simulator.

In this particular example we consider five basic traffic classes: guaranteed video (G_VIDEO); guaranteed audio (G_AUDIO); adaptive video (A_VIDEO); adaptive audio (A_AUDIO); and computer data (DATA). Although the diagram only shows one instance of each continuous media class, the model will allow for multiple instances of any particular class. Furthermore, new classes can be added, or existing classes deleted, as required.
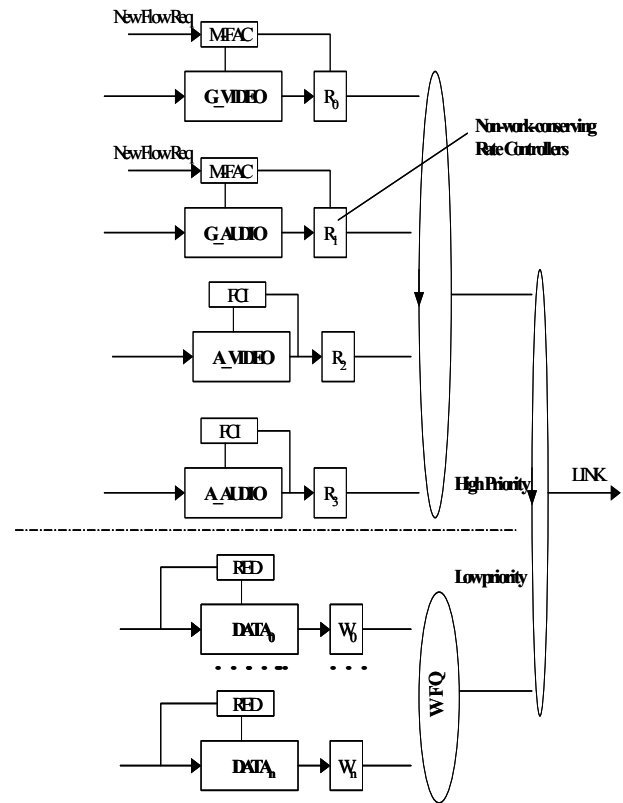
Figure 2: Bandwidth Sharing and Resource Control Model.

In this particular model all the continuous media classes are regulated by a simple non-work-conserving rate control mechanism. The rate control mechanism for each of the continuous media classes is assigned a rate, which may be changed later by certain events. However, this rate will remain fixed between any two successive events that may effect such a change. In the case of the guaranteed continuous media classes M-FAC is used to control the admission of new flows. When a request is made for a new flow M-FAC will predict what rate is required to accommodate the flow in addition the existing flow. If this new rate is below the agreed maximum for the particular class, M-FAC will update the rate allocation for that class and accept the new flow.

In the case of the adaptive continuous media classes there is no admission test for new flows. However, flows belonging to these classes will be required to adapt their output according to network conditions [Ball, 99b]. Signals requiring the sources to reduce their output will be sent via.

a Forward Congestion Indication (FCI) mechanism. FCI messages will be triggered by a new monitoring technique that was developed in our previous work [Tater, 2000]. This work found that whilst certain aspects of feedback control developed for computer data may offer some benefits to adaptive continuous media, generally they are inappropriate to its needs. Techniques such as Random Early Detection (RED) and its variants have been shown to provide a more optimal control with TCP sources [Floyd, 93]. However, our work has shown that they will not offer the same benefits to continuous media sources. Generally, these methods focus on controlling the average delay at each node, whereas controlling a higher percentile of delay is more appropriate for continuous media. We have developed a new congestion monitoring technique that monitors and controls the higher percentiles of delay and uses an explicit congestion feedback signal from the network

Flows of computer data are regulated by a WFQ mechanism, this in turn will be given a share of the link bandwidth by the priority based general scheduler of the CBQ system. Each data class will be assigned a weight, which represents the minimum throughput it can expect to receive. In general the data sources will be able to make use of whatever bandwidth is unused by the continuous media classes. However, the continuous media classes will be limited to a certain percentage of the link bandwidth to ensure that data sources are always given their minimum throughput. A resource reservation protocol such as RSVP could be used for establishing new flows in the guaranteed continuous media class, and to assign and reassign weights to the data classes.

3. MEASUREMENT BASED FLOW ADMISSION CONTROL (M-FAC)

Although CBQ can provide bulk resource allocation to a particular class of traffic by providing isolation from other classes, it cannot directly provide a guarantee to individual flows within that class. In order to achieve this it is also necessary to control admission to the individual class. Admission test are generally based on certain parameters of the traffic's characteristics and usually involves some form of analytical model. A major problem associated with such tests is the difficultly in obtaining accurate a priori estimates of the traffics characteristic. One

solution to this problem is to employed measurement based techniques.

We are developing a measurement based flow admission control mechanism that operates in conjunction with CBQ [Ball, 98], [Ball, 99a]. This mechanism uses a combination of prediction and measurement techniques in order to maintain the Quality of Service (QoS) requirements of continuous media traffic. It can be used to provide admission tests to individual streams, or alternatively, it may be used to dimension the resources needed to support aggregated streams. It is similar in operation to a measurement based ATM Call Admission Control (CAC) scheme proposed by Crosbie et. al. [Crosby, 97] except that it is designed to work with variable length packets, rather than fixed length cells.

The rate assigned to a class is initially predicted from a traffic specification of the aggregated arrival stream and the desired delay/loss characteristics. The prediction is obtained using an analytical model of a $GE/GE/1$[1] queue to estimate the 99% percentile response time of the queue [Ball, 98,99a]. The traffic parameters used by the model are the first and second moments of the packet inter-arrival time, and packet length, distributions. When a request is made for a new stream, its traffic specification, as provide by the source, is combined with the traffic specification of the existing aggregated stream to form an updated traffic specification. This updated specification is then used to predict a new value for the assigned rate. If this new value is below the permitted rate threshold then the new stream is accepted, otherwise it is rejected.

In the event of a new stream being admitted into a class, that class is assigned a new rate as outlined above. However,this rate may be adjusted later if either the actual traffic characteristics are found to differ from those that were determined at the time of the request, or the agreed QoS levels are not being met. Between the acceptances of new streams, a monitoring process will control adjustment of the rate. This monitoring process will not only adjust rate, but will also update the aggregated traffic specification. We have completed development of the prediction mechanism and the remained of this paper focuses on the current state of development of the monitoring process.

---

[1] GE – Generalized Exponential

At this stage we have not determined if it will be necessary to monitor both the arrival process (packet inter-arrival times and packet lengths) and the queue occupancy, or if monitoring only the arrival process will be sufficient. Therefore, in our initial experiments we will monitor both in order to gain a better understanding of the dynamics of the system. In previous work we have developed a method for monitoring a high percentile of queue occupancy, which we have used in a feedback mechanism to control the rates of adaptive continuous media sources. More recently, we have also developed a method to simultaneously monitor the mean and variance of the packet inter-arrival times and packet lengths. These two parameters have been used in the analytical work as input parameters and they give more general characteristics of the arrival process. In the following sections we will outline the percentile monitoring technique and introduce the mean and variance monitoring mechanism.

4. MONITORING PERCENTILE QUEUE OCCUPANCY

In separate but related work we have developed a method for monitoring a high percentile of queue occupancy, which we have used in a feedback mechanism to control the rates of adaptive continuous media sources [Tater, 2000]. As yet we have not had an opportunity to test our percentile monitoring method in our M-FAC model, however, we believe that the method would give very similar results when deployed with M-FAC.

The first step in monitoring a high percentile is to choose a target delay that relates to the jitter requirements of the traffic class. Queuing delay is a function of link rate and queue occupancy where queue occupancy changes continuously due to changing load patterns. The target delay can therefore be mapped to a queue occupancy threshold. As the waiting time of a queue can be estimated from the queue occupancy, the threshold can be set to a desired delay value. The desired value itself will depend on the QoS requirements of the traffic class Likewise, the actual percentile can be chosen to give the required level of guarantee. Whilst this method will allow us to monitor the level of QoS it does not provide us with information that can be used directly by the predictive mechanism, therefore we also developed the mean and variance

monitoring mechanisms introduced in the next section.

5. MONITORING MEAN AND VARIANCE

The purpose of this monitoring process is to ensure that the estimated values of the arrival process parameters are close to the true long-term values. To do this we will need repeatedly take sequences of n packets from the incoming packet stream and use these as samples from which to obtain some measure of the true long-term values of the mean and variance for both the inter-arrival times and packet lengths.

Fortunately, there are well know techniques that can be used to obtain confidence intervals for the true values of both the mean and variance of a population given the mean and variance of a sample from that population. Using these techniques we can construct a two-dimensional confidence region for a mean-variance pair, or mean-SD pair. This confidence region (CR) is an area bounded by a trapezium as shown below in figure 3.
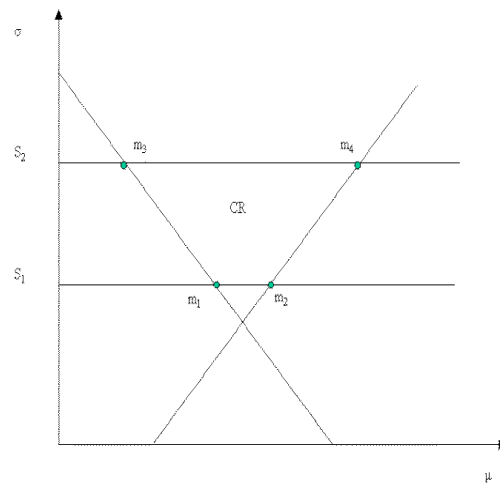

Figure 3 Confidence Region for a mean-SD pair

We are interested in monitoring the mean and SD of both the packet lengths and the packet inter-arrival times of the traffic arriving at a guaranteed service queue (generally, either audio or video). Although both the packet lengths and inter-arrival times will be sampled simultaneously and a separate CR will be constructed for each to simplify discussion we will only refer to packet inter-arrival times in the remainder of this paper.

From information provided by a traffic descriptor, or through previous measurement, we will have estimated values for the both mean and SD of the arrival process at the queue. We will denote the estimated mean and estimated SD as μ_e and σ_e respectively.

For each sample there will be a sample mean and sample SD that will be used to construct the CR for that sample. We will denote the mean and SD of the i^{th} sample as $\mu_{e,i}$ and $\sigma_{e,i}$ respectively. It will be important to carefully consider the sample size and the level of significance that will be use when forming the CR. If the sample size is sufficiently large, normal approximations are valid, however, too large a sample size may reduce responsiveness. Too low a level of significance could lead to a higher number of occasions when the true mean-SD pair is falsely considered to differ from the assumed mean-SD pair by monitoring process, we refer to this condition as a false-rejection of the assumed mean-SD pair or simply false-rejection. Conversely too high a level of significance could reduce sensitivity in detecting a true-reject condition. Therefore the level of confidence will affect the behavior of any control algorithm that is based on this monitoring method. These effects will be considered later in section 6.

5.1 Adjusting Estimates Of Traffic Parameters

The monitoring process will continually construct the CR for each sample and check that the point represented by μ_e and σ_e lies within the CR. The position of the CR may move for each successive sample as shown below in figure 4.

● Current Estimated Values (μe, σe)
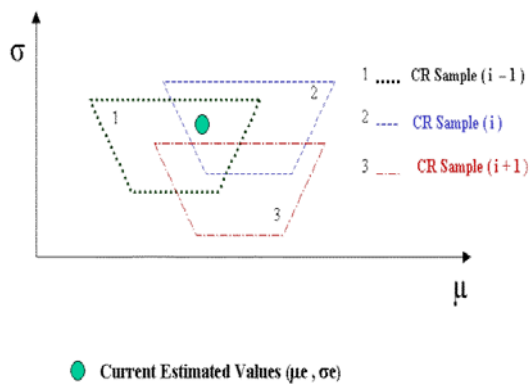
**Figure 4. Change in Position of CR for Different Samples**

In cases where the point represented by μ_e and σ_e is within the CR no action will be taken. However, in cases where the point lies outside the CR the estimated values will need to be updated. As we believe the true value is some where inside the CR we need to move the estimated value to some point within the CR. There are number of actions that could be taken but we limit our self to the three following cases: Firstly the Nearest-Point action in which to move the estimated value by the least distance towards CR, by this action we expect to have less fluctuations in the control process but less accurate tracking of the true values; Secondly, Central-Point action, in which to move the estimated value to the center of CR. In this case we expect a more accurate tracking but greater fluctuations than in the previous case; And finally, a more extreme action is to move the estimate to a furthest-point on the opposite side of the CR. We will call last action as Furthest-Point action. With this action we expect to experience a significantly greater number of fluctuations but to obtain reasonable accuracy in tracking the true values of mean and variance

To evaluate the effectiveness of each of these actions we have carried out a number of experiments to compare their different control behavior.

6. EXPERIMENTAL SIMULATION MODEL

Using the OPNET simulation package we have built a model of the Bandwidth Sharing and Resource Control System shown in figure 2 that is extended to include the multimode case. This model has been designed to be of general use for evaluation in all our resource control work, however, in this instance we utilize only the M-FAC sub-system part of the model, as shown in figure 5.
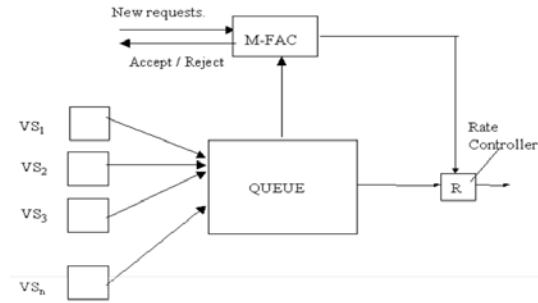
Figure 5. M-FAC Subsystem Model

For these experiments with have chosen to consider the aggregated arrival from a number of video sources, although audio sources would have served equally well. The packet inter-arrival times from these sources are generated randomly, however, the packet lengths are taken from a trace of an MPEG-1 video sequence. Each source can be set independently to have different mean and variance for packet inter-arrival times, and is supplied with its own trace file. Sources can be activated, or deactivated, at any time during the simulation. The main focus for this particular set of experiments is to measure the parameters of the aggregated traffic as it arrives at the queue. The Queue is assumed to be infinite, so that packet loss will not affect the results

We ran an initial experiment to test the effects of sample size. This involved setting the source generators to produce a long sequence packets with known mean and variance, sampling this sequence using a particular sample size, and noting the number of occasions when the know mean-SD pair were outside the CR. Monitoring was considered to be accurate if the proportion of these occasions to the total number of sample

were in proportion to the chosen level of confidence, e.g. for a 95% level of confidence we would expect that for 5% of the samples the known mean-SD pair would be outside the CR. The experiment was repeated over a range of different sample sizes. We found that for sample sizes lower than 200 there were some inaccuracies, a size of 200 provided the required level of accuracy, and increasing the size beyond 200 had little effect. Therefore we decided to use a sample size of 200 in all the following experiment.

The main experiment involved starting the simulation with a number of sources set to be active and the remainder left inactive. Throughout the simulation run certain sources would be switched from the active to the inactive state, or vice versa. The monitoring method was applied to traffic as it arrived at the queue and changes in the estimated values for both the mean and SD were recorded. Simultaneously, the inter-arrival time and packet length of each packet were logged. This allowed the true mean and SD for each sub-sequence of packets to be calculated. A sub-sequence being a sequence of packets occurring between any two changes in traffic activity. Updates to the estimated values were then plotted against the true values to show the behavior of the monitoring method.

The experiment was repeated using the Nearest-Point, Central-Point and Furthest-Point options, and in each case was also repeated for an 80% and a 99% level of confidence. Results were plotted for both mean and SD of inter-arrival times and it was noted that there was little difference the set of results relating to the mean and those relating to the SD. Therefore, we chose to report only those results relating to the SD
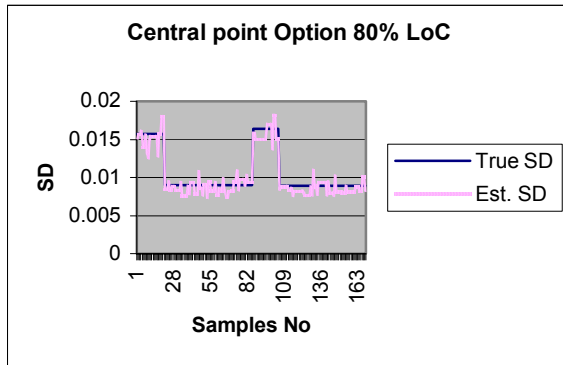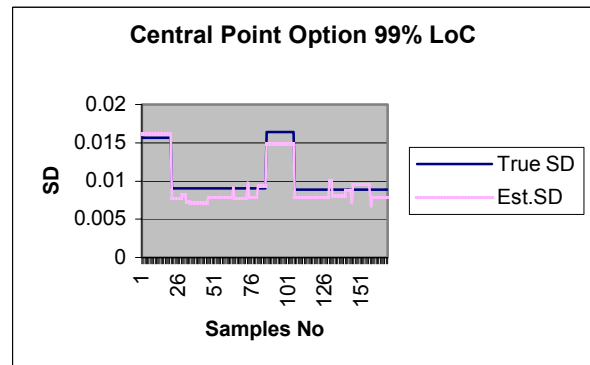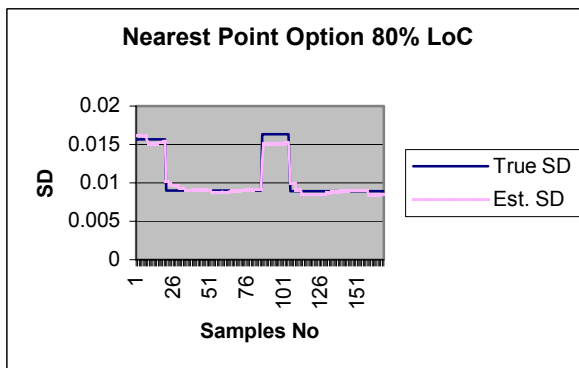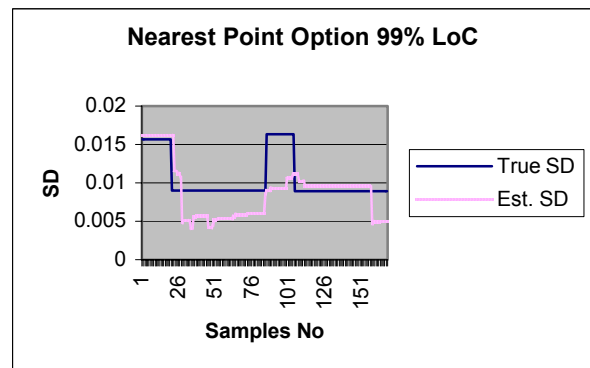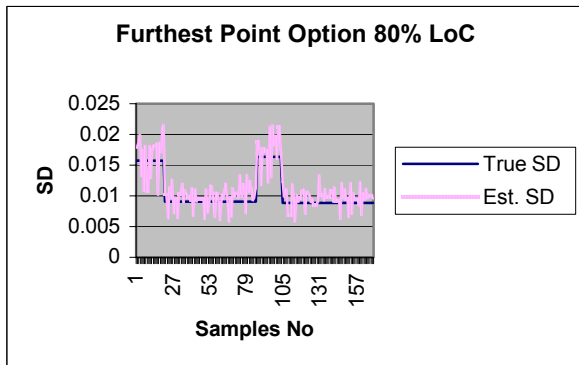
Fig 6a

Fig 6b

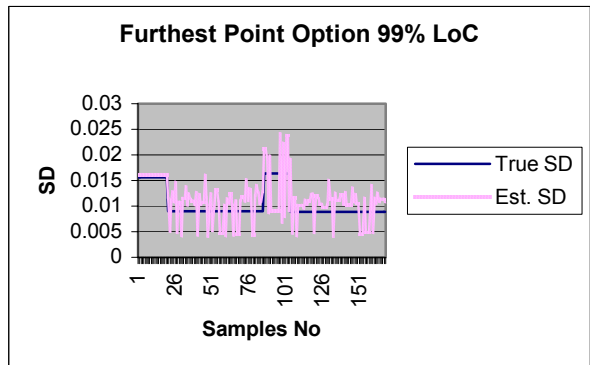Fig 6c

Fig 6d

Fig 6e

Fig 6f

Figure 6 Comparative Behavior of Nearest-Point, Central-Point, and Furthest-Point

From figure 6 it can be seen that the general behavior of the different options is as expected. However, the results clearly show that two of the six options (Fig 6a and 6c) offer significantly better behavior over the remaining four. Of these two best options Central-Point with an 80% level of confidence (Fig 6a) offers more accurate tracking but a greater degree of fluctuation than Nearest-Point with an 80% level of confidence (Fig 6c). In order to determine which of these two options will be most beneficial to the complete M-FAC system, we need to carry out further experiments that include the simultaneous operation of the prediction and monitoring processes. We are currently building the prediction mechanism into our simulation model and will soon be in a position to evaluate the complete system.

7.CONCLUSION AND FURTHER WORK

This paper has present work that is focused on the development of a Measurement Based Flow Admission Control (M-FAC) mechanism that will support guaranteed services in mutliservice packet switched networks. We have described the relationship between M-FAC and CBQ and discussed the role of CBQ in isolating traffics in classes. We have introduced Admission Control as a unit responsible for admitting flows to a guaranteed class, and discussed the importance of providing real traffic parameters rather than a simply using the characteristics stated by the source. We have outlined the monitoring and predicting algorithms that form the basic component of the M-FAC mechanism, and discussed in detail the development and evaluation of the monitoring mechanism. We have presented results from experiments to test the behaviour of the monitoring process in the presence of dynamically changing traffic streams, and have shown how M-FAC can use this to adjust its estimate of the current load vector. The next stage in our work is to incorporate both the predictive and monitoring mechanisms into the simulation model and to experiment with different control strategies for adjusting rate allocation and maintaining QoS.

REFERENCES

[Ball, 98] F. Ball, P. Callinan, D.D. Kouvatsos and C Skianis, "Dynamic Dimensioning for Guaranteed Services in Packet Switched Networks." **IFIP 6th Workshop on Performance Modeling and Evaluation of ATM Networks**, (July 1998) Ilkley, UK.

[Ball, 99a] F.Ball, P.Callinan, D.D.Kouvatsos and C Skianis, "A Measurement Based Admission Control Mechanisms for use with CBQ in Packet Switched Networks", **15th Annual UK Performance UKPEW'99**, (22nd-23rd July 1999) Bristol, UK. ISBN 0-9524027-8-5.

[Ball, 99b] F. Ball and S. Tater, "Supporting Adaptive Video Applications in Future IP Networks", **European Workshop on DISTRIBUTED IMAGING**,(18th -19th Nov 1999) London, UK.

[Basu, 2000] K. Basu, F. Ball and D.D. Kouvatsos," IPV6 over ATM: Flow mapping for guaranteed services", To be presented at the **16th UK Teletraffic Symposium UKTS** (22nd -24th May 2000) Harlow, UK

[Callinan, 2000] P. Callinan, M.Witwit, F. Ball, "A Comparitive Evaluation of Sorted Priority Algorithms and Class Based Queueing using Simulation", **Advanced Simulation Technologies Conference ASTC 2000**, (16th-20th April 2000) Washington DC.

[Floyd, 93] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", **IEEE/ACM Transactions on Networking,** August 1993.

[Floyd,93] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", **IEEE/ACM Transactions on Networking,** August 1993.

[Floyd, 95] S. Floyd and V. Jacobson, " Link-Sharing and Resource Management Models for Packet Networks", **IEEE/ACM Transactions on Networking**, 3, 4 August 1995.

[Parekh, 93] Abhay K. Parekh, "A Generalised Processor Sharing Approach to Flow Control in Integrated Services Networks", LIDS-TH-2089, MIT Laboratory, a doctoral dissertation.

[Tater, 2000] S Tater and F.Ball " Supporting Adaptive Video Applications in Future IP Networks", **18th IASTED International Conference Applied Informatics AI'2000** (14th-17th Feb 2000) Innsbruck, Austria.

[Crosby, 97] Simon Crosby, Ian Leslie, John T. Lewis, Raymond Russel, Fergal Toomy, and B McGurk. "Practical Connection Admission Control for ATM networks Based on On-Line measurements", **IEEE ATM'97, June 1997, Lisbon**

[Li, 2000] Xianglin Li and F Ball " A Generic Architecture for Multiservice Routing ", Proceedings of the 2nd annual Postgraduate Symposium on The Convergence of Telecommunication, Networking and Broadcasting, **Liverpool John Moores University (18th-19th June 2001) UK, ISBN 1902560078**.

Frank Ball received a BSc in Computer Science (1988), and was awarded a PhD (1996) for his work on supporting Quality of Service (QoS) in heterogeneous networks; both degrees are from the University of Lancaster UK. He is also a member of the IEE and a Chartered Engineer. For over 12 years he has been involved in research relating to multiservice networking, network performance evaluation, teletraffic modelling and the development of congestion control mechanisms. He is a Senior Lecturer in the School of Computing and Mathematical Science at Oxford Brookes University, where he is involved in the research activities of the School's Distributed Systems Research Group. A major part of his current research is focused on providing multiservice in packet switched networks, with particular emphasis on supporting the QoS requirements of continuous media traffic.

Ali Maqousi received his MSc. Degree in computer engineering from Electrical Engineering Institute, St. Petersburg, Russia 1989. He was a network administrator and part-time teacher assistant at the University of Petra from 1993 to 1997. He joined the faculty of science in 1999 as a full-time teacher assistant and has received a scholarship from University of Petra to continue his higher education towards getting a PhD in computing. He is currently a PhD student at Oxford Brookes University, UK. As a member of Distributed Systems Research Group (DSRG) he is working in the area of Traffic Engineering, queuing theory, and Providing QoS in packet switched networks.