Impact of Traffic Handling on Internet Capacity

Towela P.R. Nyirenda-Jere[1], Victor S. Frost[1], Nail Akar[2]

[1] Information and Telecommunication Technology Center,
University of Kansas

[2] Electrical and Electronics Eng. Dept, Bilkent University

ABSTRACT

This paper describes the impact of traffic handling mechanisms on network capacity for support of Quality of Service (QoS) in the Internet. The emergence of applications with diverse throughput, loss and delay requirements requires a network that is capable of supporting different levels of service as opposed to the single best-effort service that was the foundation of the Internet. As a result the Integrated Services (Intserv) and Differentiated Services (Diffserv) models have been proposed. The Intserv model requires resource reservation on a per-flow basis. The Diffserv model requires no explicit reservation of bandwidth for individual flows and instead relies on a set of pre-defined service types to provide QoS to applications. Flows are grouped into aggregates having the same QoS requirements and the aggregates are handled by the network as a single entity with no flow differentiation. We refer to this type of handling as semi-aggregate. The Best-Effort model does not perform any differentiation and handles all traffic as a single aggregate. Each of these traffic handling models can be used to meet service guarantees of different traffic types, the major difference being in the quantity of network resources that must be provided in each case. In this paper, we consider the issue of finding the cross-over point at which the three approaches of aggregate traffic management, semi-aggregate traffic management and per-flow traffic management become equivalent. Specifically, we determine the network capacity required to achieve equivalent levels of performance under these three traffic management approaches. We use maximum end-to-end delay as the QoS metric and obtain analytic expressions for network capacity based on deterministic network analysis. One key result of this work is that on the basis of capacity requirements, there is no significant difference between semi-aggregate traffic handling and per-flow traffic handling. However Best-Effort handling requires capacity that is several orders of magnitude greater than per-flow handling.

Keywords: Traffic handling, Quality of Service, Aggregation

1. INTRODUCTION AND MOTIVATION

During the last decade the Internet has evolved from a closed community of researchers into a commercial entity and has experienced tremendous growth in the volume of traffic as well as diversity in the type of traffic carried. The emergence of applications with diverse throughput, loss and delay requirements has created a need for a network that is capable of supporting different levels of service as opposed to the single best-effort service that was the foundation of the Internet. Notable results of the effort to incorporate Quality of Service in the Internet are the definition of the Integrated Services (Intserv) and Differentiated Services (Diffserv) models by the IETF [1,2].

The Intserv model parallels the ATM model [3] and is based on the idea that bandwidth must be explicitly managed in order to meet application requirements therefore resource reservation and admission control are a must. Intserv and ATM rely on the reservation of resources based on dynamic signaling using resource reservation protocols. One of the concerns with this model is that it requires each node in the network to maintain state on a per-flow basis and thus poses scalability problems for high-speed links supporting a large number of concurrent flows.

The Diffserv model requires no explicit reservation of resources and relies on mechanisms called Per-Hop Behaviors to provide QoS to a small number of pre-defined service types. Diffserv relies on packet classification according to

This work was supported by a grant from Sprint Corporation

T. Nyirenda-Jere: towela@ittc.ukans.edu, Information and Telecommunication Technology Center, University of Kansas, 2335 Irving Hill Rd, Lawrence, KS 66045; Tel. (785) 864 7761; Fax. (785) 864 0387

V. Frost: frost@ittc.ukans.edu, Tel. (785) 864 4833; Fax. (785) 864 0387

N. Akar: akar@ee.bilkent.edu.tr, Electrical and Electronics Eng. Dept., Bilkent University, 06533 Bilkent, Ankara, Turkey; Tel. (90) 312 290 2337; Fax. (90) 312 266 4192

desired service type at the edges of the network. This aggregation of traffic at the edges of the networks reduces the need for nodes in the network core to maintain per-flow state. The Diffserv effort represents a renewed interest and focus on simple QoS assurances by defining services that map to different levels of sensitivity to loss and delay rather than explicit values of these parameters. The potential for aggregation provided by Diffserv may prove to be beneficial in the backbone of the Internet by reducing the amount of per-flow state that is maintained.

Both approaches to providing service guarantees have translated into a debate on per-flow resource reservation which is the Intserv/ATM model versus aggregate resource reservation which is the Diffserv model versus no resource reservation at all which is the current best-effort model. The biggest argument for per-flow resource reservation is that it allows for more controlled usage of network resources and can be used to provide very strict service guarantees. Advocates of the per-flow model claim that high fidelity interactive audio and video applications need higher quality and more predictable service than that provided by the best-effort Internet and that this can only be achieved through explicit resource reservation [4].

Advocates of Diffserv maintain that with proper network engineering and provisioning, a simple priority structure will be sufficient to meet the needs of real-time traffic. This requires careful definition of classes or priority levels as well as some type of call admission control since an increase in the number of high priority real-time transmissions may degrade the performance of flows in other classes.

Proponents of the best-effort model suggest that in the future bandwidth will be practically infinite, therefore there is no need to manage and reserve bandwidth resulting in reduced network complexity. It is believed that by adequate provisioning, a best-effort network can achieve the same performance as a reservation-based network and any requirements for increased service quality can be met by increasing the capacity of network links [4]. The problem with this approach is that the higher the quality of guarantee, the more over-provisioning that must be done for the same level of user satisfaction and hence the lower the efficiency of network utilization. The case for over-provisioning is that declining prices in bandwidth will make the extra capacity required in a best-effort Internet more economical than the complexity of supporting reservations.

Network providers are thus faced with three main options in their quest to provide QoS. The first is the best-effort approach which makes inefficient use of network bandwidth by employing no traffic management. This approach assumes that bandwidth is abundant and cheap and thus the expense associated with traffic management is not needed. The second approach is to employ simple traffic management such as that proposed for Diffserv resulting in moderately efficient use of network bandwidth. Lastly, carriers can use complex traffic management such as that proposed in the Intserv/ATM model which results in the most efficient use of bandwidth. In view of these alternatives, there is a need for a clear understanding of the issues surrounding the provision of QoS in IP-based networks as well as guidelines on how traffic management and network capacity can be used to provide QoS. In this paper, we consider the issue of finding the cross-over point at which the three approaches of no traffic management, simple traffic management and complex traffic management become equivalent. Specifically we determine the network capacity required to achieve equivalent levels of performance under a variety of traffic management schemes. Knowledge of this crossover point will help network engineers and decision-makers determine the suitability of IP QoS traffic management as well as the type of traffic management to use.

In Section 2, we discuss related work and some of the questions that need to be addressed in comparing traffic management strategies. Sections 3 and 4 and describe the analysis and results of analytic study that was undertaken to illustrate how the issues raised in Section 2 could be addressed using a single-link network for illustration. We end our paper with conclusions in Section 5 and describe how we propose to extend this work.

2. RELATED WORK AND MOTIVATION

There are three key issues that have been addressed by the literature to shed more light on the Best Effort, Intserv and Diffserv debate. The first issue has to do with whether the Internet should retain its best-effort architecture or whether it should adopt a reservation-based architecture. In [4] the authors compare a best-effort Internet with one that uses per-flow handling and reservations. They consider the incremental bandwidth that is required to make a best-effort network perform as well as a reservation capable network. Their results indicate that the incremental bandwidth depends on whether the applications are adaptive or non-adaptive with adaptive applications requiring less incremental bandwidth. The general conclusion is that providing a definite answer to the choice between reservation and best-effort will depend on how adaptive applications are and the load patterns in the future Internet.

The second issue addressed in the literature, deals with how aggregation affects network performance. In most cases this is addressed through studies and analyses that compare the performance of per-flow schedulers, class-based schedulers and simple First-in-First-Out (FIFO) queues for applications having diverse QoS characteristics. The literature on scheduling algorithms is extensive and we citenum here a sample of results that are relevant to this work.

In [5] and [6] the authors provide analytical results on end-to-end delay bounds for networks of arbitrary topologies using strict priority schedulers. They conclude that in order to meet delay objectives of high priority traffic, the utilization of traffic in the high priority queue is severely limited by the maximum hop count of the network as well as by the ratio of input to output interfaces at a network node.

In [7] the authors compare the delay performance of FIFO scheduling to Weighted Fair Queueing (WFQ) for sources generating Constant Bit Rate traffic. They find that for high bandwidth flows the delays with FIFO are two orders of magnitude larger than with WFQ and delays for FIFO decrease significantly with a decrease in utilization whereas WFQ is not affected. At low levels of utilization, the difference between FIFO and WFQ decreases and is not very significant.

The work in [8] compares Generalized Processor Sharing (GPS) which is a form of Weighted Fair Queueing, strict priority and FIFO in terms of the admissible region of each policy. Their results suggests that for the performance of FIFO to match that of GPS when the traffic is comprised mostly of traffic with looser delay requirements requires much more bandwidth with FIFO than with GPS. They also find that the strict priority outperforms GPS.

The authors in [10] investigate the use of priority scheduling and FIFO with threshold dropping to provide loss and delay guarantees. They find that FIFO requires 30-70% more bandwidth than priority scheduling to provide the same delay performance.

Lastly, the degree to which traffic should be aggregated in terms of how many service levels or classes should be used has also been addressed in the literature. In [11] the authors address the question of whether to provide a single class of relaxed real-time service using FIFO or multiple levels differentiated by their delay characteristics using priority queueing. From their results, at low load levels, the priority scheme offers no advantages over FIFO. With increasing load, the benefits of priority scheduling increase. In general the conclusion is that multiple service levels increase the load levels at which the network can satisfy the needs of all classes. The work in [9] also addresses the issue of levels of aggregation and the main conclusion is that the division of traffic into two classes, a Real-Time class for audio and video and a non-Real-Time class for data is adequate to meet the stringent delay QoS requirements of the audio and video.

In this paper we enhance prior research by considering the issue of how network capacity is affected by the particular traffic handling strategy employed. As previously stated, there are three options to providing service quality which can be related to the level of aggregation of flows used by traffic handling mechanisms within the network. In the best-effort model, all flows are enqueued in the same buffer and share the buffer and link resources. We call this a (total) aggregation environment. This is the simplest and most prevalent form of traffic handling. The link must be configured with enough capacity to meet the most stringent QoS and the typical approach to maintaining QoS in this situation is to add more capacity to the link - "throwing more bandwidth".

In the Diffserv model which we call a partial aggregation (or semi-aggregate) environment, flows are divided into classes based on some criteria, the most obvious one being to group flows with similar QoS requirements. In this way, the QoS needs of a class of flows can be ensured in isolation from other classes.

The per-flow model represents an environment with zero aggregation in which each flow is assigned its own set of resources and thus attains its QoS independent of other flows. This is the best means of ensuring QoS but it is also the most complex to administer. This environment corresponds to the dedicated resources solution and has been referred to as "throwing complexity at the network". The common term for zero aggregation is per-flow queueing.

Isolation between traffic aggregates can be achieved through a combination of per-aggregate buffer admission mechanisms and per-aggregate schedulers. For the sake of simplicity, we assume in this paper that all flows which belong to the same aggregate join a FIFO queue that is sized to accommodate all bursts corresponding to that aggregate. Total aggregation can be achieved with a single queue with FIFO scheduling while for partial aggregation Priority Queueing (PQ) and Class Based Queueing (CBQ) are typical approaches. Priority Queueing imposes a strict service order by assigning each queue to a fixed priority level and serving the queues accordingly. With Class-Based Queueing, flows are mapped to classes based on some predefined attribute and service weights are assigned to each class. Per-flow queueing can be implemented using (Weighted) Fair Queueing, (Weighted) Round Robin and their many variants.

Given the levels of aggregation and the associated scheduling mechanisms which we couple under the umbrella term of traffic handling [12], the question facing the network engineer is that of determining the equivalence of the different traffic handling mechanisms in terms of their ability to support traffic with varying QoS requirements. Of particular interest is the trade-off between the complexity of traffic handling mechanisms and the network capacity required to support QoS.

It is widely accepted that the use of aggregate schemes may necessitate the provisioning of more network capacity than per flow schemes but it is not clear just how much more capacity is needed nor is it clear how the complexity of per-flow management measures up against the cost of additional capacity with aggregate traffic handling. In particular, very little is known about how semi-aggregate schemes compare to per-flow schemes. In the next two sections we describe a methodology and some results that have been obtained as part of on-going research to address these issues.

3. ANALYTIC STUDY OF TRAFFIC AGGREGATION IN A SINGLE NETWORK LINK

In order to obtain results that are easily understood and verified we focused on the simplest model of a network with a single link. For characterization of the traffic sources we used the burstiness constraint model of Cruz [13] in which traffic is characterized by two parameters, a burstiness parameter σ and an average rate parameter ρ. We assume that the network uses regulator elements or shapers to ensure that the traffic entering it conforms to these parameters. We chose to use this bounded model for the traffic processes so that the results obtained are general and applicable to a variety of situations and do not depend on specific traffic assumptions. The model is very appealing because both the IETF and ATM Forum have defined network elements which can convert an arbitrary traffic process into a process that is bounded in this way. [3], [14]

We chose four applications that are representative of current Internet usage and which provide diversity in their attributes as shown in Table 1.

Table 1. Traffic Classes and Parameters

Traffic Class (Index-Type)	**Rate** ρ_k **(Mbps)**	**Burstiness** σ_k **(bytes)**	**Packet Size** L_k **(Bytes)**	**Max. Delay** D_k **(sec)**
1-Voice	0.064	64	64	0.002
2-Video	1.5	8000	512	0.005
3-E-mail	0.128	3072	512	0.5
4-WWW	1	40960	1500	0.5

The service metric that we use is the maximum queuing delay. Typically, e-mail and WWW traffic are considered to be elastic or adaptive in that they do not have stringent delay requirements and can adjust their rates according to network conditions. We thus choose to assign delays to them that are an order of magnitude higher than the inelastic voice and video to emphasize the fact that although their delay requirements may not be stringent, bandwidth may still need to be provisioned for them to prevent total starvation.

We considered four different traffic handling strategies: Weighted Fair Queueing(WFQ), Weighted Class-Based Queueing(CBQ), Strict Priority Queueing(PQ) and FIFO Queueing. For CBQ and PQ, two classes/priority levels were used: Real Time(RT) and non-Real-Time(NRT) with voice and video classified as RT and e-mail and WWW as NRT traffic. With WFQ, since each flow is assigned its resources independently of other flows based on its own maximum delay requirement. For CBQ and PQ, in order to meet the most stringent constraint in a class, the minimum delay over all flows in a class is used to determine resources allocated for that class. For FIFO, the minimum delay over all flows determines the resources allocated.

In order to have a unified basis for comparison, we chose WFQ as the reference mechanism and for different sets of load values we calculated the number of sources that could be supported for each traffic type using WFQ. We then used this same number of sources to find the capacity required by CBQ, PQ and FIFO. The following paragraphs describe the methodology in more detail for an OC-3 link with a total load of 80% divided as 40% voice, 10% video, 15% e-mail and 15% WWW.

We begin by finding for each traffic type k, the guaranteed rate g_k^{WFQ} required under WFQ given by [15]:

$$g_k^{WFQ} = \max\left\{\frac{\sigma_k + L_k}{D_k - \frac{L_{max}}{C}}, \rho_k\right\} \quad (1)$$

$$g_k^{WFQ} \approx \max\left\{\frac{\sigma_k + L_k}{D_k}, \rho_k\right\} \quad (2)$$

where $L_{max} = \max\{L_k\}$ and C is the link capacity. The number of connections for type k is then given by:

$$N_k = \left\lfloor \frac{w_k * C}{g_k^{WFQ}} \right\rfloor \quad (3)$$

where $\lfloor x \rfloor$ is x rounded down to the nearest integer. The results are shown in Table 2.

Table 2. Guaranteed Rates and Connections

Traffic Class (Index-Type)	**Guaranteed Rate g_k(Mbps)**	**Number of Connections N_k**
1-Voice	0.512	131
2-Video	13.6	1
3-E-mail	0.128	181
4-WWW	1	23

We then determine how much capacity would be required to support the same traffic using the other three schemes. This is done based on the deterministic analysis of maximum queueing delay presented in [13,16]. For CBQ with P classes, the required bandwidth C_{CBQ} is found as :

$$C^{CBQ} = \sum_{p=1}^{P}\sum_{k\in p}\frac{N_k\sigma_k + L_{max}}{D_{class\ p}} \quad p = 1, 2, ...P \quad (4)$$

where $D_{class\ p} = \min_{k\in p}\{D_k\}$

where we have assumed that the factor $\frac{L_{max}}{C}$ does not contribute significantly to the delay. For Priority Queueing with P priority levels such that $1 > 2 > .....P$, the required capacity C^{PQ} is found as:

$$C^{PQ} = \max_p\left\{\sum_{j=1}^{p}\sum_{k\,\in\,class\ j}\frac{N_k\sigma_k + L_{max}}{D_{class\ p}} + \sum_{j=1}^{p-1}\sum_{k\,\in\,class\ j}N_k\rho_k\right\} p = 1, 2, ...P \quad (5)$$

For FIFO, the capacity C_{FIFO} is given by:

$$C^{FIFO} = \sum_{k=1}^{K}\frac{N_k\sigma_k}{D_{min}} \quad (6)$$

where $D_{min} = \min_k\{D_k\}$. Applying the formulas yields the results in Table 3 which shows the actual capacity in Mbps and the capacity quantized to the minimum number of OC-3 links.

The capacities calculated using this analysis should be understood to be the minimum capacities that will ensure that the delay objectives for each traffic type are met. An additional constraint that should be factored in to ensure stability is that the capacity should always be greater than $\sum_k N_k\rho_k$.

Table 3. Example of Capacity Requirements

Scheme	Capacity (Mbps)	Capacity (OC-3 links)
WFQ	124	1
CBQ	89	1
PQ	68	1
FIFO	757	5

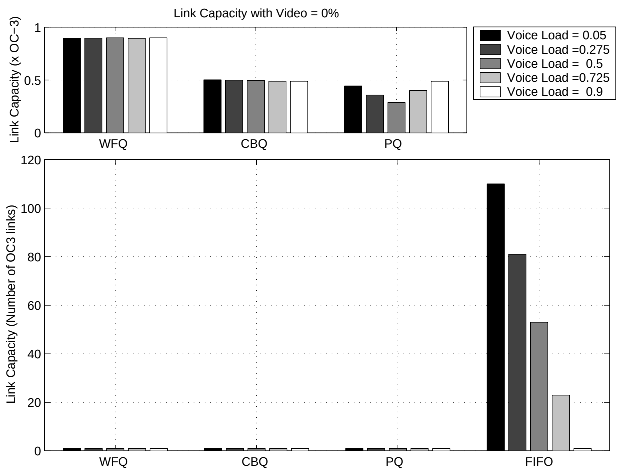

Figure 1. Capacity Requirement with No Video and Varying Voice Load

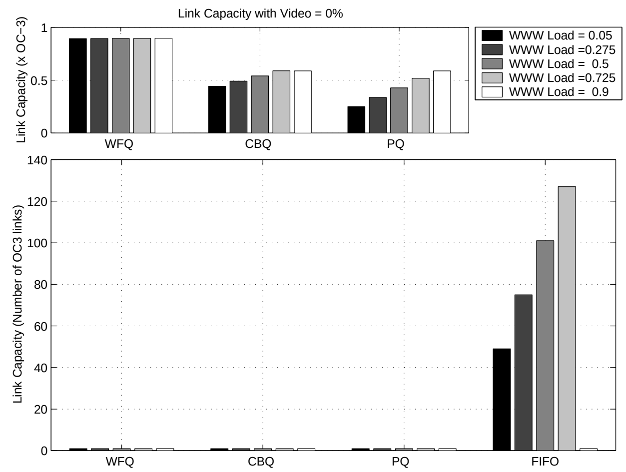

Figure 2. Capacity Requirement with No Video and Varying WWW Load

4. ANALYTIC RESULTS

4.1. Comparison of Bandwidth Requirements

In this section we present results on the difference in bandwidth requirements of the four schemes under varying load conditions. Using the notation w_T for the total load on the link and w_k for the fraction of link capacity allocated to traffic type k, we used three different values for video load: $w_2 = 0, 0.1, 0.2$. In one case, for each value of video load, we varied the voice load from 0.05 to $w_T - w_2$, setting the e-mail and WWW load to half of the remaining bandwidth. In the second case we reversed the roles of the voice and WWW traffic. We plot the capacity requirements in terms of the number of the minimum number of OC-3 links required by each scheme. Figure 1 shows the capacity requirements when there was no video traffic and voice load was varied.

We observe that CBQ, PQ and WFQ are not affected by the volume of voice traffic and are able to meet the delay guarantees for all cases with one OC-3 link. For FIFO, the amount of voice traffic significantly affects the bandwidth requirements: when the proportion of voice traffic is small, the bandwidth requirements are higher and vice versa. This is because when the voice load is small, the e-mail and WWW traffic proportions increase and more capacity is required to equalize the performance of the e-mail and WWW to that of voice in order to guarantee the delay objectives of voice traffic. The capacity for FIFO is more than 100 times that of the other schemes when voice is 5% and equal when voice is 90%. When the WWW load is varied, Figure 2 shows that WFQ, CBQ and PQ are still able to support all the traffic types with one OC-3 link. With FIFO, increasing the WWW load increases the required capacity when voice traffic is present. When there is no voice traffic and no video traffic, the capacity requirements of FIFO decrease significantly and one OC-3 link is sufficient.

In Figures 3 and 4 we plot the capacity requirements when the video load is 20%. WFQ, CBQ and PQ are still able to support the traffic using one OC-3 link when either the voice or WWW traffic is varied. For FIFO, the effect of the video traffic is to reduce the bandwidth requirements compared to the case with no video load, since now the proportion of e-mail and WWW traffic is reduced thus decreasing the capacity needed to equalize the performance of the e-mail and WWW traffic to that of voice. Increasing the voice load reduces the FIFO capacity while increasing the WWW load increases the capacity.

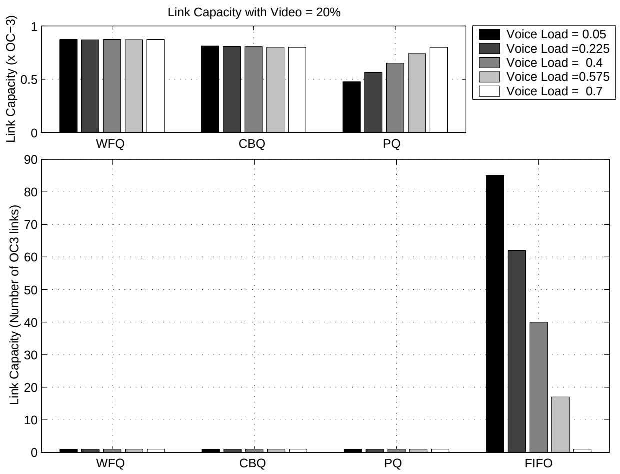

Figure 3. Capacity Requirement with 20% Video and Varying Voice Load

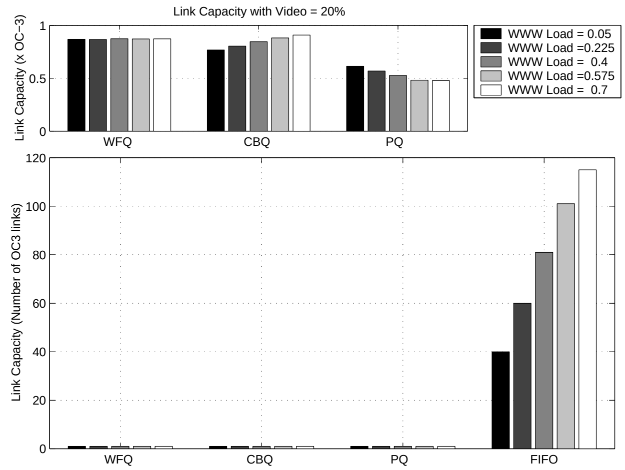

Figure 4. Capacity Requirement with 20% Video and Varying WWW Load

4.2. Sensitivity to Design Point

The goal of this analysis was to explore the ability of the three schemes to provide acceptable delay QoS guarantees when the traffic submitted exceeded the traffic for which the network was designed. We considered two scenarios: one in which voice is the dominant traffic type and another in which WWW traffic is dominant. For each scenario, the capacity required by each of the four schemes was calculated using the procedures in Section 4. The number of sources, the link capacities and the delay performance are collectively referred to as the design point. For each scenario, the volume of either voice or WWW traffic was varied and the delay for each traffic type was calculated using the design point capacities. We note in practice that call admission procedures would be used to restrict the number of flows admitted but since we are testing the sensitivity of the traffic handling schemes we assume no call admission control. Instead we consider two approaches for bandwidth allocation under WFQ. In the first method which we call WFQ1, an increase in the traffic of a particular class is handled by re-distributing the bandwidth share of that class (as determined by the load at the design point) equally among the sources (old and new) of that class. In the second approach called WFQ2, an increase in voice traffic is accommodated by "stealing" bandwidth from the e-mail and WWW classes to guarantee the voice traffic its delay QoS. We present our results in the form of plots of the ratio of actual delay to desired delay as a function of the % change in voice or WWW load. We focus on the QoS of voice since it is the most stringent.

Figure 5 shows the results obtained for a design point with voice as the dominant traffic type corresponding to $w_1 = 40\%$, $w_2 = 10\%$, $w_3 = 15\%$ and $w_4 = 15\%$ with changing voice traffic. We find in general that WFQ1, CBQ and PQ exhibit the greatest sensitivity to increasing voice load, FIFO exhibits marginal sensitivity while WFQ2 is not sensitive to increases in the voice load. WFQ2 allows us to meet the delay requirements for the voice traffic while degrading the e-mail and WWW performance. Figure 6 shows that when we increase the WWW traffic, FIFO is now the most sensitive and we cannot meet the delay objectives for voice. WFQ, CBQ and PQ do not affect the voice performance since the voice traffic is isolated from the impact of the WWW traffic.

For a network initially dominated by WWW traffic and design load $w_1 = 15\%$, $w_2 = 10\%$, $w_3 = 15\%$ and $w_4 = 40\%$ the results show a similar trend to the case of a network dominated by voice. The picture emerging from these results is that the traffic handling schemes are both sensitive to the type of traffic that dominates the network at the design point as well as to the type of traffic that increases the load on the network. For a network designed with voice as the dominant class, FIFO is the least sensitive to increases in voice traffic and the most sensitive to increases in WWW traffic when considering the delay objectives of voice. WFQ, CBQ and PQ are both sensitive to increases in the voice load and if the goal is to maintain the delay objectives of voice at all costs, the use of a scheme like WFQ2 can achieve this with a corresponding exponential increase in the delay of e-mail and WWW traffic. The value of these results is best demonstrated by taking into account the permissible variances in the delay objectives which means using statistical objectives as opposed to deterministic ones and this will be explored in extensions to this research.

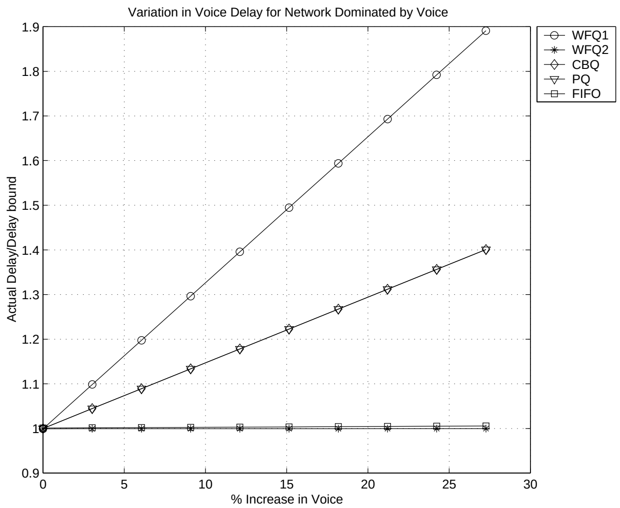

Figure 5. Variation in Voice Delay with increase in Voice load

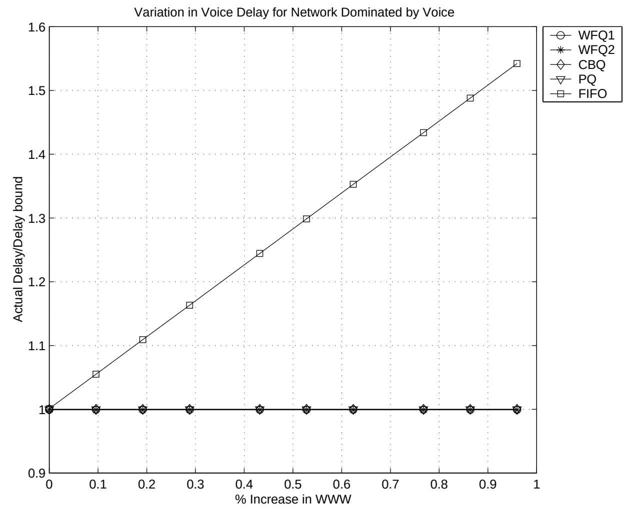

Figure 6. Variation in Voice Delay with increase in WWW load

4.3. Required Capacity with Projections on Traffic Growth

In this part of the analysis we calculate the capacity required to support yearly projections on growth in voice and WWW traffic. Current industry estimates are that voice traffic on the Internet will grow at a rate of 5-15% each year. The trend in WWW traffic has been almost a 100% increase in traffic per year. We assume the two scenarios in Section 4.2 of either voice or WWW being the dominant traffic type. Using the same procedures as before, we calculated the capacity required over a 5 year period assuming a 15% growth in voice traffic per year and a 100% growth in WWW traffic per year. The results obtained are shown in Figures 7 and 8 with the capacity expressed in terms of the minimum number of OC-3 links.

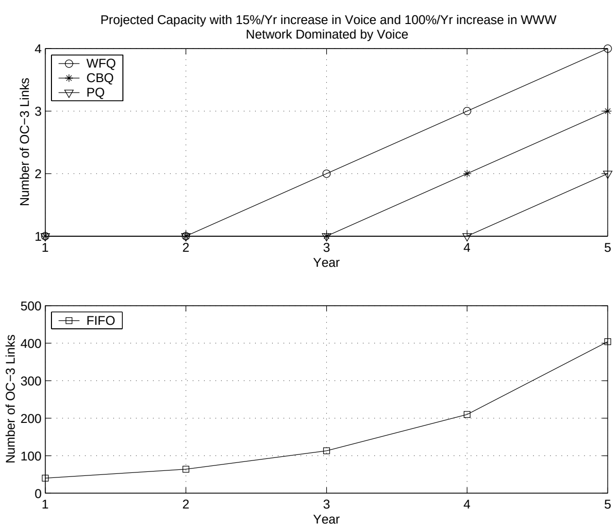

Figure 7. Network Capacity with Projections on Voice and WWW for Voice-dominated Network

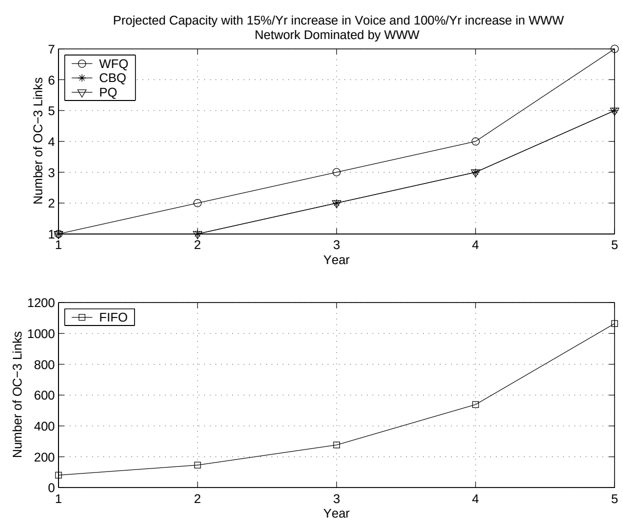

Figure 8. Network Capacity with Projections on Voice and WWW for WWW-dominated Network

For the network dominated by voice we find that the capacity required for WFQ increases to 4 times the initial capacity, CBQ by a factor of 3 and PQ by a factor of 2 after the 5 year period. FIFO capacity increases to 8 times the initial capacity, reaching 400 OC-3 links after 5 years. When we start with a network dominated by WWW traffic, the capacity of WFQ increases by a factor of 7 while CBQ and PQ capacity increases by 5 after the 5 year period. FIFO capacity increases by a factor of 13. In both cases FIFO is affected the most by the increase in traffic especially since we are increasing the volume of WWW traffic by a substantial amount. We note that the increase in capacity for WFQ is faster than that of CBQ and PQ and that WFQ capacity after the first year is larger than CBQ or PQ capacity. To complete the picture we consider a hypothetical future situation in which the growth of WWW traffic is 15% and that of voice is 100%. This corresponds to the hypothesis that eventually growth in voice traffic will outpace growth in WWW traffic. The projections on capacity in this case are shown in Figures 9 and 10.

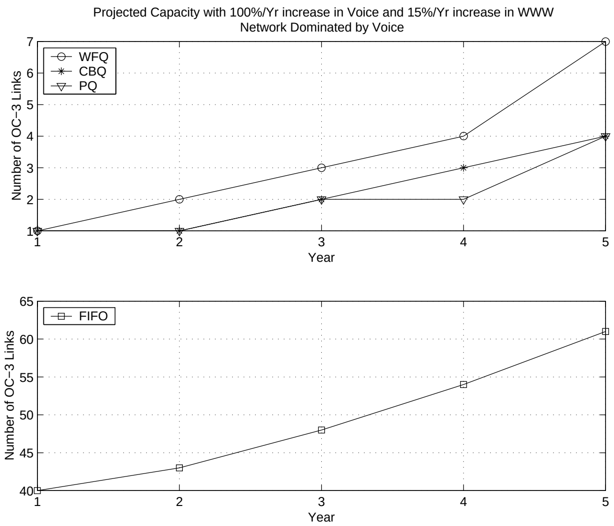

Figure 9. Network Capacity with Projections on Voice and WWW for Voice-dominated Network

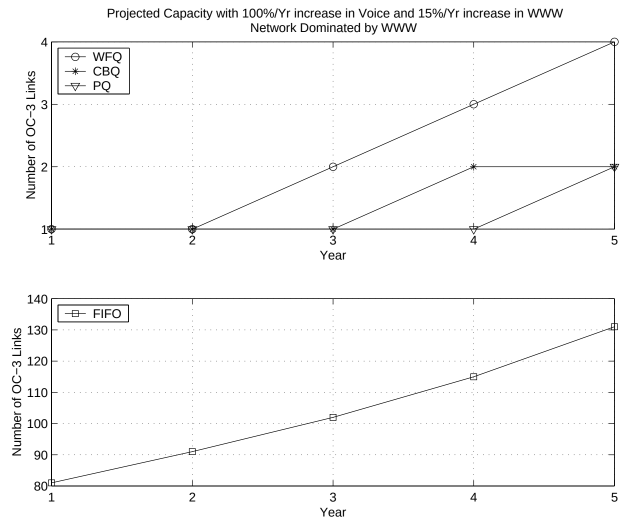

Figure 10. Network Capacity with Projections on Voice and WWW for WWW-dominated Network

We find in this case that for a Voice-dominated network, CBQ and PQ capacity increase the least by a factor of 4 while WFQ increases by a factor of 7. FIFO capacity increases by only a factor of 1.5. For a WWW-dominated network, WFQ capacity increases by a factor of 4 and CBQ and PQ by 2 while the FIFO capacity increases by a factor of 2. We conclude that WFQ is affected more by the volume of voice traffic than the aggregate schemes while FIFO is affected most by the volume of WWW traffic when voice traffic is present in the network. We also observe that when the volume of voice traffic is high as in Figures 9 and 10, CBQ and PQ require slightly less capacity than WFQ, illustrating the multiplexing gains possible with aggregate schemes.

5. CONCLUSION AND FUTURE WORK

While the analysis and methodology presented in this paper is straightforward, it clearly demonstrates that it is possible to quantify the trade-off between network capacity and traffic management. Although the results presented apply to a single link we anticipate that the general trend of the results will be preserved when we analyze networks of arbitrary size and topology. One of the most encouraging results from this work is that on the basis of network capacity there is no significant difference between using semi-aggregate traffic handling and using per-flow traffic handling. It is still an open issue how to capture the complexity associated with the three traffic handling methods and determine how that would influence the trade-off between total aggregation on the one hand and partial aggregation and per-flow handling on the other. There are several ways in which we intend to apply and extend our analysis in order to fully address the trade-off between complexity of traffic management and network capacity. We are currently extending the analysis to networks of arbitrary size and topology and investigating the use of statistical descriptions for the delay objectives.

REFERENCES

1. IETF Differentiated Services Working Group, *www.ietf.org/html.charters/diffserv-charter.html*
2. IETF Integrated Services Working Group, *www.ietf.org/html.charters/intserv-charter.html*
3. ATM Forum, *Traffic Management Specification version 4.1, AF-TM-0121.000*, Mar 1999.
4. Breslau L., S. Shenker, “Best Effort versus Reservations: A Simple Comparative Analysis” in *Proc. ACM SIGCOMM* 1998, p.3-16.

5. Charny A., "Delay Bounds in a Network with Aggregate Scheduling", ftp:ftpeng.cisco.com/ftp/acharny/aggregate_delay_v4.ps, Feb. 2000
6. Charny A., J. Leboudec, "Delay Bounds in a Network with Aggregate Scheduling", *1st International Workshop on Quality of future Internet Services, QofIS 2000*, Sept. 2000.
7. Sahni J., P. Goyal, H. Vin, "Scheduling CBR Flows: FIFO or Fair Queueing?", in *Proc. 9th Int'l Workshop on Network and Operating Systems Support for Digital Audio and Video, NOSSDAV '99*, p. 13-27, June 1999.
8. Presti F.L., Z. Zhang, D. Towsley, "Bounds, Approximations and Application for a Two-Queue GPS System" in *Proc. IEEE Infocom 1996*, p.1310-1317.
9. Dolzer K.,W. Payer, M. Eberspacher, "A Simulation Study of Traffic Aggregation in Multi-Service Networks" in *Proc. IEEE Conference on High Performance Switching and Routing*, Heidelberg, June 2000
10. Sahu S., D. Towsley, J. Kurose, *A Quantitative Study of Differentiated Services for the Internet*, Umass CMPSCI Tech Report 99-09, University of Massachusetts, 1999.
11. Bajaj S., L. Breslau, S. Shenker, "Is Service Priority Useful in Networks?" in *Proc. ACM SIGCOMM 1998*, p66-77.
12. Bernet Y., "The Complementary Roles of RSVP and Differenitiated Services in the Full QoS Network" in *IEEE Communications Magazine*, Feb. 2000, p.154-162.
13. Cruz R.L., "A Calculus for Network Delay, Part I: Network Elements in Isolation" in *IEEE Transactions on Information Theory*, Vol 37-1, Jan, 1991.
14. IETF RFC 2215, General Characterization Parameters for Integrated Services Network Elements,*www.ietf.org/rfc/rfc2215.txt*
15. Parekh A., R. Gallager, "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case" in *IEEE/ACM Transactions on Networking*, Vol.1-3, June 1993.
16. Cruz R.L., "Quality of Service Guarantees in Virtual-Circuit Switched Networks" in *IEEE JSAC 13(6)* pp.1048-1056, 1995