An Enhanced Algorithm for Fair Traffic Conditioning in Differentiated Services Networks

Abiola Adegboyega, Rupinder Makkar, Kayvan Mosharaf, Ioannis Lambadaris
Broadband Networks Laboratory
Carleton University Ottawa, Ontario, Canada, K1S 5B6
{agboyega, rup, mosharaf, ioannis}@sce.carleton.ca

Abstract—**Fair bandwidth sharing among traffic flows with different characteristics in Differentiated Services (Diffserv) networks is the focus of the current research. This paper examines and enhances an algorithm developed to enforce fairness among disparate TCP flows in the Assured Forwarding (AF) service in Diffserv. Equation Based Marking (EBM) [1] was introduced by K. Shin et al to enforce fairness in AF by monitoring existent network conditions used in marking decisions. The estimation of packet losses by the algorithm is integral to marking. The loss rates of different connections were demonstrated to converge hence enforcing a fair marking regardless of the metrics of individual flows. In this paper, EBM is analyzed for fairness and enhanced by implementing a more efficient technique for loss rate estimation. Comparison is made between EBM and the enhanced technique with results showing appreciable improvements in the maintenance of fairness. Furthermore, a service definition required by QoS standards is met with the implementation of the additional algorithm to EBM.**

Keywords-Diffserv, QoS, TCP, Fairness, loss rate

I. INTRODUCTION

Differentiated Services (Diffserv) [2] offers a scalable service provisioning architecture in IP networks by the allocation of network resources based on application requirements. Diffserv philosophy, in providing end-to-end Quality of Service (QoS) moves architectural complexity to the 'edge' of the Diffserv domain where ingress traffic is conditioned and marked for differentiation into classes while the 'core' forwards aggregated traffic streams and enforces QoS during network congestion.

The Assured Forwarding (AF) Per-Hop Behavior (PHB) [3] is a Diffserv forwarding mechanism that has been standardized by the Internet Engineering Task Force (IETF). AF maintains traffic profiles for ingress traffic and marks packets at the network edge as IN-profile provided they adhere to the service profile of the network or OUT-of-profile otherwise. The basis of this marking is the comparison of the subscription rate of incoming streams with the Committed Information Rate (CIR) and Peak Information Rate (PIR) assigned for the Diffserv domain. If the ingress traffic rate is lower than the CIR, it is marked as IN and assigned the lowest drop precedence AFx1. When the traffic rate exceeds the CIR but not the PIR, packets are marked as OUT with drop precedence AFx2. Ingress traffic with rates greater than the PIR are marked as OUT with the highest drop precedence AFx3. The AF PHB can be achieved in core routers by a tiered dropping scheme that ensures service differentiation by dropping packets during congestion beginning with the highest drop precedence AFx3, then AFx2 and eventually AFx1 only when all other precedences are dropped. A common implementation of this dropping scheme is done by an active queue management (AQM) scheme like Random Early Detection (RED) [4] with some variations e.g. RED with IN and OUT (RIO), with two thresholds one each for conformant and non-conformant traffic [5].

It was shown in [6] that AF is unable to provide bandwidth assurance when competing TCP traffic streams are different based on the flow metrics of packet size, Round Trip Time (RTT), target rate and the number of microflows aggregated in a flow. Furthermore, while TCP flows are able to slow down their sending rates during congestion, non-responsive flows (e.g. UDP traffic) maintain their network subscription rates regardless of congestion.

One approach of mitigating unfairness is by using 'TCP-aware' traffic conditioners, employing one or more of TCP's characteristics [7][8][9]. This paper presents an enhanced version of the traffic conditioner in [1] that employs the inverse of a modeled steady-state TCP equation in marking functions. Our contributions to the existing algorithm include the implementation of a more efficient loss rate estimation technique, which is integral to packet marking. Loss History Discounting provides enhancements to EBM with more details given in section IV.

The rest of the paper is organized as follows. Section II discusses related work in mitigating unfairness in AF services. Section III presents EBM [1] and its use of loss rate estimation in marking decisions. Section IV presents our version of the modified conditioner. We compare the original and enhanced algorithms with simulations in section V and conclusions are made in section VI.

II. RELATED WORK

Intelligent traffic conditioning for the provision of fairness in Diffserv networks has been researched extensively. In [6], Seddigh et al employed simulations to illustrate the impact of the flow metrics of RTT, packet size, microflows and target on fairness and the distribution of surplus bandwidth. In [10] simulations were employed to demonstrate that TCP/UDP interactions and RTT have a direct impact on achievable throughput of competing flows.

0-7803-8794-5/04/$20.00

Intelligent traffic conditioners can employ local knowledge (at the edge of the DiffServ domain) in enforcing fairness [11]. The algorithms implemented in [7-9] fall under this category. In [8] the traffic conditioner protects high RTT traffic by increasing their amount of IN-profile traffic over short RTT connections since connections with shorter RTTs send faster updates from end-to-end and thus obtain more than their fair share of the available bandwidth. An enhanced Time Sliding Window (TSW) marker [7] remedies the drawback of the traditional TSW, which estimates ingress traffic throughput inefficiently with a profile meter that either keeps a long or a short history of packet arrivals. The enhanced TSW increases the sensitivity to TCP throughput dynamics by employment of watermarks that track the oscillation of TCP transmission rates.

Another approach to intelligent traffic conditioning is the computation of arrival rates based on measurement of all flows aggregated at ingress nodes. The drawback of tracking individual flows is insensitivity to the number of microflows aggregated in a flow [11]. In [12], a Packet Marking Engine (PME) is designed which can either be integrated with end hosts or external at the network edge. The PME maintains the local state and passively observes the throughput of a connection or an aggregate to make marking decisions. While it eliminates the disadvantage of monitoring individual flows, its feedback mechanism is inaccurate and causes fluctuations in performance [1]. Performance fluctuations are mitigated in [13], which uses a Time Sliding Window-Token Bucket (TSW-TC) marker. It enhances arrival throughput estimation with a mechanism that tracks the variations in TCP congestion window and RTT, which is kept in memory as a variable that is compared with the current average rate estimation. This enables tracking of flow rate change and implicitly, the variations in RTT and congestion window.

The approach in EBM [1] periodically tracks the TCP dynamics through a widely accepted model of the TCP equation [14]. The estimation of packet loss rates in EBM used in marking decisions differentiates it from all the mechanisms discussed. Our enhancement of EBM is in the implementation of a more efficient scheme for loss rate estimation. The guiding principles of EBM are discussed in the next section.

III. EQUATION BASED MARKING

The Equation Based Marker [1] subsequently referred to as EBM, was developed to enforce fairness among disparate AF traffic flows by employing the parameters of the steady-state TCP rate equation in [14] (Appendix A) to determine packet marking. The conditioning function of EBM is modeled on the inverse of this equation. While other implementations have employed approximate steady-state models, the motivation for EBM is its estimation of packet losses at the domain edge used in marking decisions while able to accurately predict the required bandwidth share of heterogeneous connections from the rate equation. EBM utilizes TCP's ability to manage congestion by reducing transmission rates through feedback signaling with packet loss indication. This feedback is integrated into the EBM traffic conditioner as a loss rate estimation mechanism at the edge of the network. The computed loss rate is compared with two thresholds derived from the equation's inverse to determine packet marking. The TCP rate equation is given in (1) as a function of the parameters (Appendix A) that govern packet transmission.

$$B = F(p, RTT, W_{\max}, T_0) \qquad (1)$$

Where B is the sending rate of the TCP source, p is the packet loss probability, RTT is the two-way latency through the network, W_{max}, the maximum congestion window is the sender limit on the number of packets that can be sent at once and T_0 is the retransmission timeout.

The reasoning behind the employment of the loss probability p from (1) in EBM is due to the duality between throughput and loss indications as shown in Fig. 1 [1]. The Loss probability[1] is the feedback signal between end hosts for congestion management. EBM uses this feedback signal to calculate two thresholds used in marking. Equation (1) is inverted with respect to p giving Equation (2).

$$p = F^{-1}(B, RTT, W_{\max}, T_0) \qquad (2)$$

With the sending rate B replaced with the CIR and the PIR of an incoming flow, we obtain:

$$p_{cir} = F_p^{-1}(CIR, RTT, W_{\max}, T_0) \qquad (3)$$

$$p_{pir} = F_p^{-1}(PIR, RTT, W_{\max}, T_0) \qquad (4)$$

These are target loss probabilities with respect to its CIR and PIR respectively[2]. These two thresholds are computed for each flow and then compared with the flow's estimated loss rate to determine marking. From (4), CIR and PIR are easily obtainable from profile meters. The EBM modules for the estimation of RTT and marking are discussed. We discuss loss rate estimation under section IV.

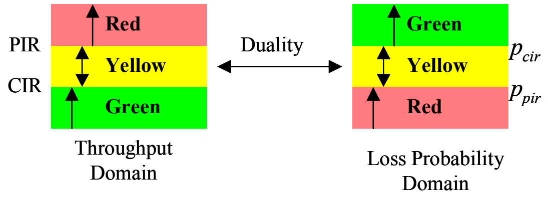

Fig. 1. Duality between target rate and loss probability

---

[1] The steady-state TCP model is derived with p as the probability of a packet loss while it is generally understood as the actual packet loss which is computed as the loss rate in EBM

[2] F_p^{-1} is the inverse of the TCP equation with respect to p

A. Estimation of RTT

EBM uses the RTT estimation module written into available implementations of TCP. TCP senders and receivers both use sequence numbers to measure RTT. Every time the receiver sends an ACK to the sender, it includes the sequence number from the most recent data packet. The difference between the present clock and the timestamp based on the received sequence number is an estimate of the RTT. This is smoothened with an exponentially weighted moving average according to the formula:

$$srtt = w * rtt + (1 - w) * srtt \tag{5}$$

Where *srtt* is the smoothened RTT from the current estimate. The weight *w* is used to determine the responsiveness of the transmission rate to changes in the round trip time [15]. The complete module is given in [1].

B. Marking

Marking is based on proportionality to the target loss probabilities p_{cir} and p_{pir} derived in (3) and (4). The drop precedences AFx1, AFx2 and AFx3 are described as Green, Yellow and Red respectively. With the computation of p_{cir} and p_{pir}, a linear function is used to calculate P_{yellow} (probability of marking a packet as yellow) and P_{red} (probability of marking a packet as red) when a three-color marking scheme is used. The algorithm is given in Fig. 2. The *YScale* and *RScale* are design parameters based on current network conditions. A complete description of the operational analysis of EBM is found in [1]. We proceed with our contributions to EBM.

```
If (loss_rate) > $p_{cir}$
        mark packet as GREEN
Else if ( $p_{cir}$ > loss_rate > $p_{pir}$ )
        $P_{yellow}$ = ($p_{cir}$ -  loss_rate)/$p_{cir}$ x ($p_{cir}$ x YScale)
        with probability $P_{yellow}$ mark packet as YELLOW and
        with probability (1- $P_{yellow}$) mark packet as GREEN
Else if (loss_rate < $p_{pir}$  )
        $P_{red}$ = ($p_{pir}$ – loss_rate)/$p_{pir}$ x ($p_{pir}$ x  RScale)
        $P_{yellow}$ = ($p_{cir}$ – $p_{pir}$)/$p_{pir}$ x ($p_{pir}$ x  YScale)
        with probability $P_{red}$ mark packet as RED and
        with probability $P_{yellow}$  mark packet as YELLOW and
        with probability (1- ($P_{yellow}$ + $P_{red}$)) mark packet as GREEN
```

Fig. 2. EBM marking algorithm

IV. LOSS HISTORY DISCOUNTING

Three methods of loss rate estimation were discussed in [15]. The Dynamic History Window method uses a history window of packets with window length determined from the current transmission rate. This method is not resilient to noise. Even with a periodic loss rate, loss events entering and exiting the window cause changes to the estimated loss rate adding unnecessary noise to the loss signal.

The Exponentially Weighted Moving Average (EWMA) method uses an exponentially weighted moving average of the number of packets spanning two loss events. Depending on weighting, this can either add too much weight to the most recent interval or take too much history into account in computing the loss rate.

The average loss interval method is able to track smoothly in the presence of noise while its weight distribution is stable to current loss rates without accounting for a long sequence in the loss history. However, it has its limitations given certain network conditions that could exacerbate marking functions for EBM. A *loss interval* is defined as the number of packets between loss events. The average loss interval $\hat{s}_{(1,n)}$ is calculated as the weighted average of the last *n* intervals. From [15], this is given as:

$$\hat{s}_{(1,n)} = \frac{\sum_{i=1}^{n} w_i s_i}{\sum_{i=1}^{n} w_i} \tag{6}$$

For weights w_i:

$$w_i = 1,\ 1 \le i \le n/2,$$

$$w_i = 1 - \frac{i - n/2}{n/2 + 1}, \quad n/2 < i \le n$$

In practice [15], $n = 8$ is an accepted value which is the lower bound that achieves stability with a reasonable balance between resilience and noise. The given loss event rate is then given as $1/\hat{s}$. Equation (6) is implemented in EBM.

The interval since the most recent loss, depending on its length may not reflect the current loss rate of the network. Secondly, there can be sudden changes in estimated rate because unrepresentative loss intervals leave the number *n* of intervals being observed [16]. In view of these factors, *loss history discounting* is implemented to adapt the weights for loss rate estimation.

History discounting (HD) [17] identifies a particularly long interval since the last drop and smoothly discounts the weights given to the older loss intervals. The current loss interval ppl_i, the most recent of the last *i* intervals is considered to be unrepresentative of the loss event rate when it exceeds a multiple *m* of the average size ppl^{avg} over the *i*-1 intervals. In HD a factor d_i is determined for the unrepresentative interval to mitigate its impact on the estimation of the loss rate. This is given as:

$$d_i = \begin{cases} 1 & \text{for } ppl_i \le m.ppl_{i-1}^{avg} \\ \dfrac{m.ppl_{i-1}^{avg}}{ppl_i} & \text{for } ppl_i > m.ppl_{i-1}^{avg} \end{cases} \tag{7}$$

With a greater discrepancy between the size of the average loss interval and the current loss interval, the smaller is the deweighting factor and thus the weight assigned to the older loss intervals. The new calculation of the loss estimate is given as:

$$\hat{s} = \frac{\sum_{i=0}^{n-1} d_i w_{i+1} s_i}{\sum_{i=1}^{n} d_{i-1} w_i} \qquad (8)$$

While history discounting adapts to long intervals unrepresentative of the loss rate in the network, it also needs to adapt to a sudden increase in loss rate *after deweighting*. To this end, a lower bound can be imposed on the deweighting factor. This prevents a shift in all the weight to the current loss interval.

The discounting mechanism that was added to EBM is given as follows: If the current loss interval $s_0 > 2\hat{s}_{(i\geq1)}$, then the current loss interval s_0 is considerably longer than the recent calculated average. The weights for the older intervals are discounted by using the discount factor:

$$d_i = \max(0.5, \frac{2\hat{s}_{(i\geq1)}}{s_0}) \quad for\ i > 0;\ d_0 = 1 \qquad (9)$$

The lower bound of 0.5 ensures that past losses are not completely discarded in the calculation of the loss rate and improves adaptability with a sudden increase in losses after deweighting is applied. When a loss occurs, the old interval s_0 is shifted to s_1. The corresponding discount factors are shifted so that once an interval is discounted it is not un-discounted. Without discounting, $d_i = 1$ for all values of i.

V. EVALUATION

We make a performance comparison of three conditioners: EBM, EBM with History Discounting (EBM-HD) and the Time Sliding Window Three Color Marker (TSWTCM). Simulations were done using the simulation tool ns-2 [18]. The simulation network setup [1] is given in Fig. 3. There are ten sources (S1-S10) that send TCP traffic to the corresponding destinations (D1-D10). Markers are inserted between each source and the network edge router E1. The bottleneck link connects three core devices, with a capacity of 10 Mbps and a latency of 10 ms between each device as shown. The access links (between each source and network edge E1 and between E2 and each destinations) have a capacity of 10 Mbps with a latency of 10ms. The core routers achieve the AF PHB using RIO with three-drop precedences. The threshold parameter values for RIO in non-overlapped mode are given in table 1. A TCP agent at each source generates unlimited FTP bulk data. Based on traffic conditions, a CIR and PIR common to all sources is specified. This topology presents two departures from the simulations in [1]. There are no UDP sources (background traffic). Investigation is done solely with TCP sources. The second departure is the definition of subscription level. In [1], total subscription is the sum of two rates: Half is the sum of the CIRs of all assured sources (TCP) and the other half is the bit-rate of the background UDP sources. The latter is varied by either incrementing beyond one-half defined as *oversubscription* or decremented as a definition of *undersubscription*. In our experiments, 100% bandwidth subscription is the sum of all CIRs. With the sum greater than the bottleneck we have oversubscription while we have undersubscription when it is less.

The goodput (the total packets successfully received at the destinations) was measured over a simulation run (800 s) with the experiment repeated 20 times for a confidence interval of 95 %. The metric of fairness *fi* was computed with Jain's fairness index [19], which indexes fairness among *n* competing flows, each with a bandwidth share of x_i, as:

$$fi = \frac{(\sum_{i=1}^{n} x_i)^2}{n\sum_{i=1}^{n} x_i^2} \qquad (10)$$

The investigations are based on three basic traffic metrics that affect fairness: Packet Size, RTT and Target rate.

A. Packet Size Experiment

This section explores the fairness among different TCP flows with disparate packet sizes. For the network in Fig. 3, Packet sizes vary from 100 bytes (S1) to 1000 bytes (S10). Each source has a CIR of 1 Mbps and a PIR of 1.2 Mbps. All core links have a capacity *l*. Capacity allocation (100% subscription) is at 10 Mbps. The Queue length parameter in Table 1 is made 1000 packets. Oversubscription (125%) is at 8 Mbps with an expected fair share of 800 kbps per source. Undersubscription (75%) is 13 Mbps and an expected fair share of 1300 kbps per source. Fig. 4 to Fig. 6 give the results.

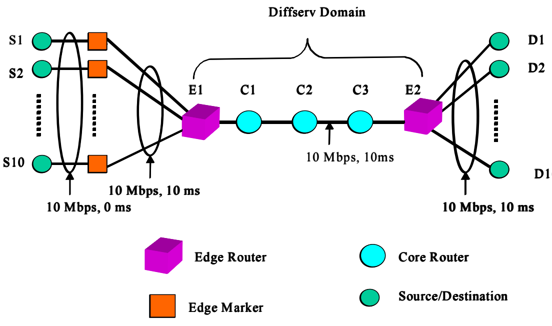

Fig. 3: Simulation Network

Table 1: Parameters for RIO with three-drop precedences

Parameters for:	Green	Yellow	Red
Queue Length(Packets)	L	L	L
max_{th}(Packets)	0.875L	0.625L	0.3125L
min_{th}(Packets)	0.625L	0.3125L	0.025L
max_p	0.02	0.05	0.1
w_q	0.002	0.002	0.002

From Fig. 4, the disparity in achieved goodput with the TSWTCM increases with packet size. Since it does not detect packet sizes at the edge router given that TCP bandwidth is directly proportional to the packet size, the connections with the larger packet sizes hog the available bandwidth. EBM achieves a skewed fairness. Flows with packet sizes varying from 100 to 300 bytes are denied fair bandwidth share (for 100% and 125% subscription). EBM-HD, while not achieving the required 1000 kbps per source sustains fairness with an average goodput of 900 kbps. The effectiveness of history discounting can be explained by the convergence of the loss rate estimate over time. Fig. 7 and Fig. 8 respectively compare the loss rate estimate for EBM and EBM-HD. In Fig. 8, EBM-HD enhances sensitivity of loss rate estimation for convergence to a given range that enforces fairness. It is observed that EBM is not sensitive enough to calculate the loss rate at lower packet sizes (100-300 bytes). EBM-HD is able to achieve fairness at the three given levels of subscription within 90% of the CIR, which is an accepted definition of a Service Level Agreement (SLA) [20] in Diffserv.

B. RTT Experiments

RTT investigations were conducted with the same network topology shown in Fig. 3. Based on the range chosen for RTT research [7][21] we conducted experiments spanning 4 intervals: [40, 160] ms, [40, 240] ms, [40, 320] ms and [40, 400] ms. This follows from the RTT in real networks which allows for a definition of various levels of assurance. For the network in Fig. 3, the RTTs for the sources S1 to S10 were varied according to each interval. Each source has a CIR of 500 kbps and a PIR of 700 kbps. Subscription levels [75,100,125]% was achieved with [6.67,5,4] Mbps respectively. The queue length in table 1 is set to 800 packets. We show the results for the first and last intervals from Fig. 10 to Fig. 15 while Fig. 16 and Fig. 17 show the fairness index and bandwidth assurance from the smallest to the largest interval. The performance of TSWTCM follows from the TCP rate equation, which is inversely proportional to RTT. It shows some degree of fairness (at 100% and 125% subscription) over the short-range [40, 160] ms with performance degrading as the range increases to [40, 400] ms.

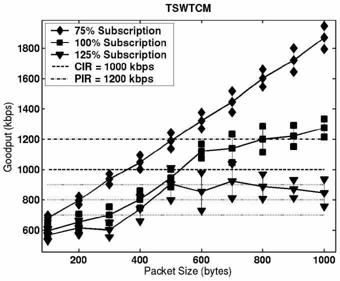

Fig. 4. TSWTCM (Goodput vs. Packet Size)

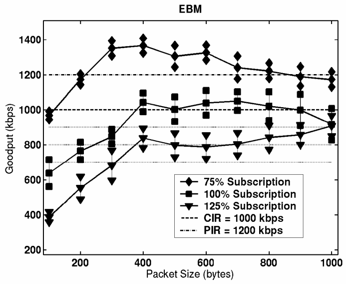

Fig.5. EBM (Goodput vs. Packet Size)

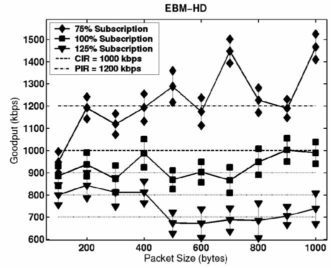

Fig. 6. EBM-HD (Goodput vs. Packet Size)

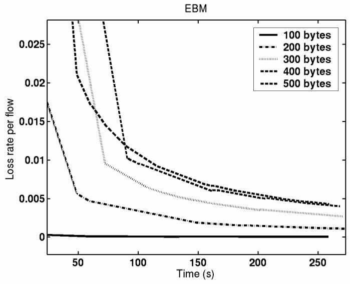

Fig. 7. EBM (loss rate per flow vs. time)

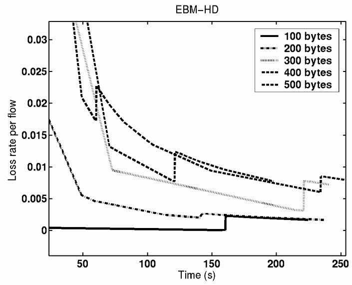

Fig. 8. EBM-HD (Loss rate per flow vs. time)

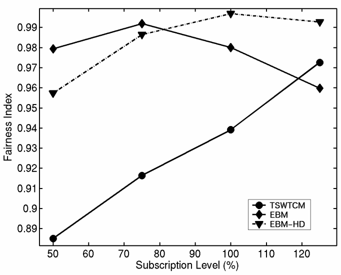

Fig. 9. Fairness Index vs. Subscription level

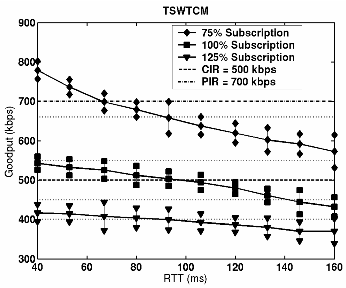

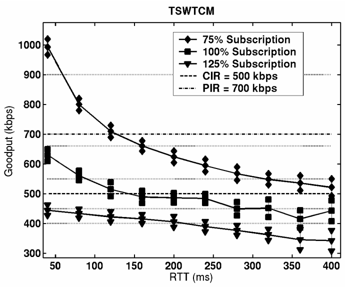

Fig. 10. TSWTCM: Goodput vs. RTT [40,160] ms Fig. 11. TSWTCM: Goodput vs. RTT [40,400] ms

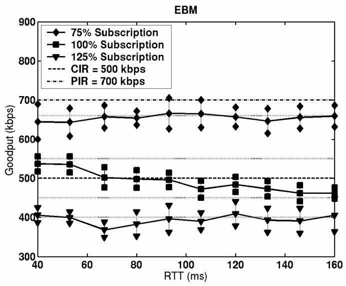

Fig. 12. EBM: Goodput vs. RTT [40,160] ms

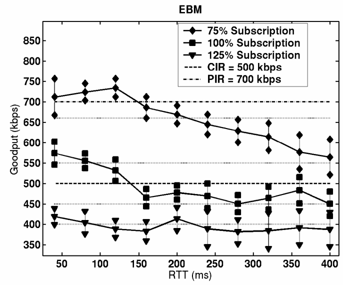

Fig. 13. EBM: Goodput vs. RTT [40,400] ms

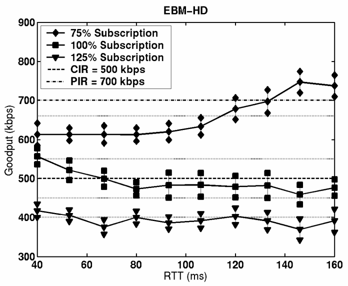

Fig. 14. EBM-HD: Goodput vs. RTT [40,160] ms

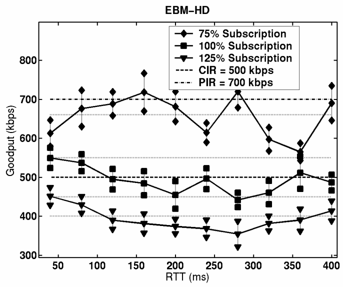

Fig. 15. EBM-HD: Goodput vs. RTT [40,400] ms

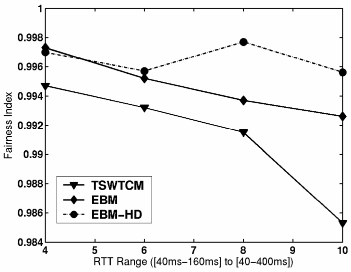

Fig. 15. Fairness Index vs. RTT range

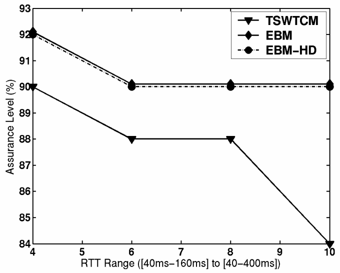

Fig. 16. Assurance Level vs. RTT Range

From Fig. 12 and Fig. 13 (for EBM) and Fig. 14 and Fig. 15 (EBM-HD), at 100% subscription (with an expected goodput of 500 kbps for all sources), the first three connections with the lowest RTTs exceed the CIR of 500 kbps while the other seven connections maintain a bandwidth that is approximately 460 kbps. This allows for an SLA definition of a lower bound of 92% of customer CIR. History Discounting provides no enhancement beyond the regular method of loss rate estimation. EBM-HD however improves fairness among competing flows when the RTT range increases to [40 400] ms as shown in Fig. 15.

C. Target Rate Experiments

To explore fairness according to target rates, the network in Fig. 3 was simulated. Target rates vary from 500 kbps (S1) to 2300 kbps (S10). Each source has a PIR that is twice its CIR. All core links have a capacity of 14 Mbps, which is the sum of the CIRs of all connections. The queue length in table 1 was set to 1000 packets. Testing was done only at capacity allocation (100% subscription). From Fig. 17, EBM offers the required protection based on the target loss rates determined according to both the PIR and CIR. History Discounting works to overprotect the low CIR connections and is therefore inhibited in ensuring fairness according to target rates.

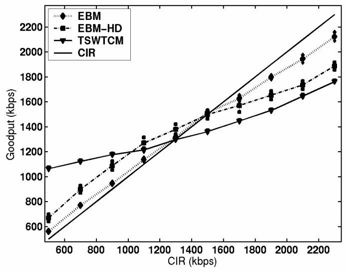

Fig. 17.Goodput vs. CIR for all 3 markers

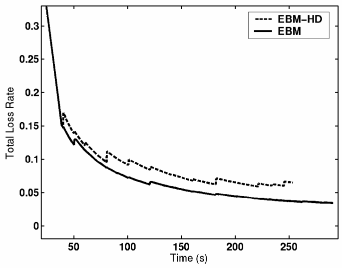

Fig. 18: Loss rate comparison for EBM & EBM-HD

VI. CONCLUSIONS

In this paper we demonstrate the enhancement of EBM by history discounting in providing an acceptable level of fairness and bandwidth assurance when competing connections with different packet sizes share a single Diffserv AF network. EBM-HD provides the same level of bandwidth assurance with disparate RTTs while improving on fairness over a wide range of RTTs. History Discounting should be inhibited when connections differ in target rate due to its overprotection of connections with low CIRs.

EBM-HD shows better performance in the maintenance of fairness from undersubscription through capacity allocation to oversubscription. Overall EBM-HD leads to a higher loss rate as shown in Fig. 18 for all experiments. The tuning parameter that affects the level of discounting can however help to mitigate this.

From other experiments not shown, both EBM and EBM-HD do not to provide bandwidth assurance when connections differ in the number of microflows within an aggregate. The algorithm observes the characteristics of a single flow and is oblivious of the number of flows in an aggregate.

The addition of a module to EBM that monitors aggregated microflows and the adaptation of EBM's techniques to monitor arrival rates at the edge of the Diffserv domain is the subject of future research.

REFERENCES

[1] M. El-Gendy, K. Shin, "Equation Based Marking for Assured Forwarding Services," IEEE INFOCOM, Vol. 2, June 2002.

[2] K. Nichols, V. Jacobson, L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," IETF, RFC 2638, July 1999.

[3] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, "Assured Forwarding Per-Hop-Behavior Group," RFC 2597, June 1999.

[4] S. Floyd, V. Jacobson, "Random early detection gateways for congestion avoidance," IEEE/ACM Transactions on Networking, Vol. 2 No. 4, August 1993.

[5] D.D. Clark, W. Fang, "Explicit allocation of best-effort packet delivery service," IEEE/ACM Transactions on Networking, Vol. 6 No. 4, August 1998.

[6] N. Seddigh, B. Nandy, P. Pieda, "Bandwidth assurance issues for TCP flows in a differentiated services network," IEEE Globecom vol. 3, 1999.

[7] W. Lin, R. Zheng, and J.C. Hou, "How to make assured service more assured," In Proc. IEEE International Conference on Network Protocols (ICNP 1999), pp. 182-191, November 1999

[8] B Nandy, N Seddigh, P Pieda and J Ethridge, "Intelligent traffic conditioners for assured forwarding based on differentiated services networks," Proc of High Performance Networking 2000 Conference, Paris, France, May 2000

[9] Y. Chait, C. Hollot, V. Misra, D. Towsley, H. Zhang and J. Lui, "Providing throughput differentiation for TCP flows using adaptive two-color marking and two level AQM," IEEE Conference on

Computer Communications (INFOCOM 2002), vol. 2, pp. 837-844, June 2002.

[10] J. Ibanez and K. Nichols, “Preliminary simulation evaluation of an Assured Service,” Internet Draft, draft-ibanez-diffserv-assured-eval-00.txt, August 1998.

[11] X. Chang, J. Muppala, “Adaptive marking threshold for improving bandwidth assurance in a differentiated services network,” IEEE Global Telecommunications Conference, 2003, (GLOBECOM), Vol. 6, Dec. 2003

[12] W. Feng, D. Kandlur, D. Saha and K Shin, “Adaptive packet marking for maintaining end-to-end throughput in a differentiated services Internet,” IEEE/ACM Transactions on Networking, Vol. 7, No. 5, October 1999.

[13] K. Kumar, A. Ananda, and L. Jacob, “A Memory-Based Approach for a TCP-friendly Traffic Conditioner in Diffserv Networks,” IEEE International Conference on Network Protocols (ICNP 2001), pp. 138-145, November 2001.

[14] J. Padhye, V. Firoiu, D. Towsley, J. Kurose, “Modeling TCP Reno Performance: A Simple Model and Its Empirical Validation”, IEEE/ACM transactions on Networking, Vol. 6 Issue 2, April 2000

[15] S. Floyd, J. Padhye and J. Widmer, “Equation-based control for unicast applications”, Proc. Of ACM SIGCOMM ’00, 2000.

[16] S. Floyd, J. Padhye and J. Widmer, “Equation-based control for unicast applications: the Extended Version”, International Computer Science Institute technical report, tr-00-03, March 2000.

[17] Widmer, J. “Equation Based Congestion Control” Diploma Thesis, Department of Mathematics and Computer Science, University of Mannheim, February 2000.

[18] Network Simulator, ns-2, University of California at Berkeley, version 2.1b6, www.isi.edu/nsnam/ns with the EBM implementation suite by El-Gendy available at http://kabru.eecs.umich.edu/qos_network/diffserv/EBM/ebm.html

[19] Raj Jain, “The Art of Computer Systems Performance Analysis”, John Wiley and Sons Inc., 1991

[20] S. Blake et al., “An Architecture for Differentiated Services”, RFC 2475, December 1998

[21] A. Habib, B. Bhargava, S. Fahmy, “A round trip time and timeout aware traffic conditioner for differentiated services networks,” IEEE international conference on communications ICC 2002, Vol. 2, May 2002

Appendix A Steady state TCP Model

The achievable throughput B of a TCP sender is given as:

$$B(p) = M \frac{\frac{1-p}{p} + W(p) + \frac{Q(p,w)}{1-p}}{RTT(\frac{b}{2}W(p)+1) + \frac{Q(p,w)F(p)T_o}{1-p}} \qquad \text{If } W(p) < W_{max}$$

$$B(p) = M \frac{\frac{1-p}{p} + W_{\max} + \frac{Q(p,w_{\max})}{1-p}}{RTT(\frac{b}{8}W_{\max} + \frac{1-p}{pW_{\max}} + 2) + \frac{Q(p,w_{\max})F(p)T_o}{1-p}} \qquad \text{Otherwise}$$

$$W(p) = \frac{2+b}{3b} + \sqrt{\left(\frac{8(1-p)}{3bp} + \left(\frac{2+b}{3b}\right)^2\right)}$$

$$Q(p,w) = \min\left(1, \frac{(1-(1-p)^3)(1+(1-p)^3(1-(1-p)^{(w-3)}))}{1-(1-p)^w}\right)$$

$$F(p) = 1 + p + 2p^2 + 4p^3 + 8p^4 + 16p^5 + 32p^6$$

M is the maximum segments size in bits;
W_{max} is the maximum congestion window;
T_0 is a Timeout;
b is the number of packets acknowledged by an ACK;
p is the loss probability;
RTT is the Round-Trip-Time of the connection;
$W(p)$ is the average window size