Content-based Packet Marking for Application-Aware Processing in Overlay Networks

Panho Lee, Tarun Banka, Anura P. Jayasumana, V. Chandrasekar
{leepanho,tarunb,anura,chandra}@engr.colostate.edu
Department of Electrical and Computer Engineering, Colorado State University,
Fort Collins, CO 80523

Abstract: *In many emerging sensing applications, random network losses may lead to drop of critical information, rendering partially received data useless for the end users. A packet-marking scheme based on the application content is proposed that enables application-aware processing of the data within the overlay network. A token-bucket based rate control algorithm in conjunction with the proposed packet-marking scheme enables on-the-fly selection of data for forwarding/drop to a particular end user at the desired transmission rate. We demonstrate the effectiveness of the packet-marking and token-bucket based rate control scheme in simultaneously meeting heterogeneous QoS requirements of the multiple end users for content quality and bandwidth for a weather monitoring sensing application.*

1. Introduction

Many critical systems, such as those for weather monitoring, industrial environment monitoring, and emergency management are increasingly relying on the Internet. With many of these applications, a variety of data must be transmitted in real time to multiple end users at distributed geographical locations. Unlike audio and video multicast applications in which a source node transmits data to all receivers with a single send, many end users of Distributed Collaborative Adaptive Sensing Systems (DCAS) such as Collaborative Adaptive Sensing of the Atmosphere (CASA) [5,11] use data generated by same set of sensing nodes in different ways to concurrently meet different system specific goals, e.g., tornado detection and rainfall estimation. Thus two different end-users may need different subsets of data from the total data set generated by sensing nodes in the same period of time. A CASA system consists of an array of weather radars generating high-bandwidth streams of data, processing nodes, and multiple end-users distributed over a wide geographic area, all working as a collaborative real-time system [11]. Depending on the end-user QoS requirement, the computation resources available, etc. each user may also be associated with a different rate at which information is delivered. Quality of the data is based on the usefulness of the received data for a particular end user application. In a best-effort environment, there is no guarantee to consistently meet such application specific requirements due to random packet losses within the network because of bandwidth constraints. However, many end users can operate with varying QoS within bounds for data quality and bandwidth. The end users may receive an acceptable performance with the minimum QoS for data quality and bandwidth; yet the performance will improve with reception of higher QoS level, i.e., higher data quality, and bandwidth. There have been considerable efforts since 90's to provide QoS support on the Internet. IntServ and DiffServ [12] architectures, proposed to meet the QoS requirements such as delay guarantees and packet prioritization, have seen only limited deployments in the Internet. While such schemes will enhance the data delivery among nodes of a DCAS system, application aware operations at CASA nodes, operating as an overlay network, could provide significant enhancements in scalability, performance and reliability.

In DCAS applications such as CASA for weather monitoring, end users have critical QoS requirements. However, under extreme weather conditions the underlying network infrastructure may be only partially available. Under such conditions, CASA like systems are required to make best use of the partially available network infrastructure to meet application specific QoS constraints. Moreover it is desired that sensing applications adapt to take advantage of additional network resources that may become available as the network load changes. Different overlay networks based solutions have been proposed to provide a range of useful services such as bandwidth guarantee for enhancing QoS of applications [2,3,14]. Active networks [9,15] introduced the concept of in-network processing, where network nodes perform customized computations on the messages being forwarded. Overlay networks enable deployment of application-specific functionalities at intermediate nodes in the network to enhance QoS received by the end users under given network conditions. In this paper we propose an application-aware content based packet marking and a token-bucket based rate control algorithm to meet content quality and bandwidth requirements of the DCAS applications using overlay networks. The proposed algorithm enables transfer of most suitable subset of the application content to the

---

This work is supported by the Engineering Research Center program of the National Science Foundation under NSF award number 0313747.

1-4244-0419-3/06/$20.00 ©2006 IEEE

end user while providing soft bandwidth guarantees within bounds under available network infrastructure and dynamic network congestion conditions. We also demonstrate the effectiveness of the proposed scheme in overlay multicast where multiple end users receive weather radar data with different QoS requirements.

Most of the prior work on packet marking and application aware data drops has focused on video streaming. OverQoS[14] allows applications to prioritize packets within a stream at the intermediate nodes of overlay networks for enhancing video streaming quality. Kassler et al. in [8] propose and evaluate the use of filtering mechanisms as an application-aware processing for supporting multiple heterogeneous receivers. Keller et al. in [9] propose active router architecture for multicast video distribution using a wavelet-based high scalable video codec and packet tagging. The packet tag enables routers in the network to drop packets belonging to a particular quality set during network congestion.

Section 2 provides the motivation for packet-marking and token-bucket based rate-control algorithm for the DCAS applications. Section 3 describes the proposed application aware packet-marking scheme. Section 4 describes the application of packet-marking in token-bucket based application-aware rate control algorithm. Section 5 discusses experimental results. Conclusions are presented in Section 6.

2. Motivation

In a CASA application, weather radars periodically generate block of data known as digitized radar signal (DRS), at a constant rate every 'heart-beat' interval. A DRS data block consists of 64 samples each for 256 gates (In radar terminology, a gate refers to a volume in the atmosphere at a particular distance from the radar source node for which data is collected by a radar) [4]. Under ideal conditions, all required samples in a DRS block should be delivered to the end users for processing. However, data may get randomly dropped in the network because of bandwidth constraints. In many DCAS applications such as CASA, different end user algorithms have different bandwidth requirements and show different levels of tolerance to the missing information in a DRS block. Therefore, under bandwidth-constrained scenarios, i.e., during network congestion when all the information cannot be transmitted to the end user in real-time, a subset of the information may be selected by source node for transmission to provide somewhat less accurate, yet acceptable end-product value. The subset of information to be transmitted is determined by knowing the characteristics of the application data, tolerance of end users to the missing information, and bandwidth requirements of the end users. For example, some radar data end user algorithms such as reflectivity computation [4,6] can tolerate uniformly distributed random drops of samples from a DRS block of data generated by the radar node. Alternatively, an algorithm for Doppler velocity computation [4,6] can tolerate bursty drop of samples instead of uniform drop. Thus application-aware data selection at the intermediate overlay nodes can more effectively and efficiently meet the needs of the end user and the DCAS system.

Most of the prior work on application-aware data selection mechanisms relied on end-host applications to adapt to network conditions [1,16]. However, network conditions are dynamic and end-hosts cannot predict a change in the available bandwidth at intermediate nodes during the transmission leading to random losses of the subset of data selected by end-hosts. Thus we cannot assure that most important subset of the data will be delivered to the end-users by using the end-host application-aware data selection scheme alone. Therefore, whenever there is a need to drop packets at intermediate nodes, it is desired that packets be dropped in an application-aware manner. Overlay networks enable deployment of such application-aware data dissemination services over the Internet [14].

Fig. 1 shows an overlay network for application-aware data dissemination to multiple end users. An overlay network consists of different nodes, such as forwarding nodes, multicast nodes, and fusion nodes, each configured to perform distinct application-specific task to best meet QoS requirements of different end users. For example, a source node may perform packet marking based on the properties of the data for a particular application. In Fig. 1, the multicast node is responsible for accepting connection requests from multiple end users with different data quality and bandwidth requirements. Multicast node then forwards the aggregated request to the source node. Source node sends data to the multicast node as per the requested bandwidth and data quality requirements. It is important to note that in DCAS applications, multicast node may simultaneously send different subset of the data to each end user as per their bandwidth and data quality requirements. The multicast nodes may perform functions such as independent flow and congestion control for each end user in an application-aware manner considering their distinct bandwidth requirements and fusion of data from different sources prior to multicast. Existing multicast solutions such as RLM [10], are required to scale to millions of end users, which is significantly higher than the scalability requirements of the DCAS applications. Moreover, unlike RLM, each of the end user in DCAS may have

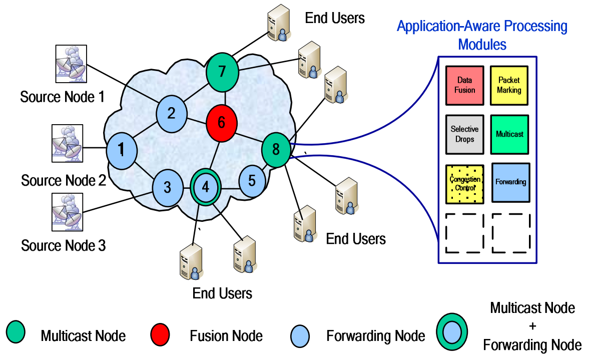

Figure 1. Overlay network for application-aware data dissemination

distinct bandwidth and data quality requirements that need to be satisfied by a single multicast server. A token-bucket based rate control may be implemented at the multicast node to achieve the desired rate for a particular end user under existing network conditions. Alternatively, forwarding nodes can use packet marking to select packets for forwarding or drop while considering available output link bandwidth. Forwarding node can be considered as a special case of multicast node where only one end user is requesting data at a rate equal to the available output link bandwidth.

3. Application-Aware Packet-Marking

Aforementioned, a packet marking may be used at the intermediate overlay nodes to select packets for forwarding such that subset of the data delivered to the end users meets their data quality requirements.

It is considered that source node is aware of the properties of the data generated by sensors such as cameras and weather radars. The source node marks packets at transmission time according to the end users data quality requirements and different rates at which data may be delivered to the end users. Consider an example as shown in Fig. 2, where a sensor node generates 8 application data units (ADU) within the bounded time at rate R1. The ADU is defined as a fundamental application data entity that can be used by an end user algorithm for processing. Each row in Fig. 2 shows the subset of ADUs that are selected for transmission at a lower transmission rate when a higher rate cannot be supported because of bandwidth constraints. The subset of data selected at lower rate depends on the end user data quality requirements. For example, certain end users need uniformly spaced ADUs when only a subset of the data can be selected for transmission. Alternatively, other end users prefer a contiguous group of ADUs when bandwidth is constrained. Consider the case when the source node transmits data at rate R1, and as seen in the figure the data transmitted at lower rates is a subset of the data transmitted at rate R1 and ADUs are selected uniformly at lower rates. The packet containing ADU 1 is marked with different color flags corresponding to different rates, i.e., rates R1-R7. Similarly packet containing ADU 3 is marked with different colors corresponding to different rates, i.e., R1, R2, R4, and R5. As shown in the Fig. 2, every packet contains a flag for each rate for which it is transmitted indicated by different colors. Note that multiple flags can be set to indicate suitability of the packet for multiple transmission rates. It is important to note that packets are encoded at the time of transmission such that there is no dependency between different packets in the stream; end users can decode the packets on-the-fly without waiting for later packets to arrive. Following are the two key advantages of the packet marking when used at intermediate nodes:

(i) On-the-fly selection of packets to one or more end users for forwarding at different rates: Consider the case when marked packets are received at the multicast node from the source node. In this case multicast node may select data on-the-fly for forwarding using packet marking to multiple end users at their respective transmission rates which are determined based on the network congestion for each end user. Note that in DCAS applications, multicast node may send different subset of the received data to each end user concurrently while considering their individual bandwidth requirements. Packets with marking corresponding to end user's transmission rate are selected for forwarding for a particular end user. Fig. 3(a) shows a case when multicast node receives all packets transmitted at rate R1 from the source node. First four packets are shown in the figure. Two end users, i.e., end user 1, and end user 2 are considered that need data at different rates, i.e., rate R5 and rate R1 respectively. As seen in Fig. 3(a), rate R1 packets are marked with green color and rate R5 packets are marked with yellow color. In this case, packets corresponding to rate R5 are subset of the packets transmitted at rate R1. Multicast node selects packets on-the-fly for forwarding to end users 1 and 2 based on the marking flag corresponding to rate R1 and R5 in the packet. Note that multicast node creates a copy of the packet to be forwarded and replaces destination address to the address of an end user for which it is selected for transmission. As seen in Fig. 3(a), out of first 4 packets received at multicast node, packets with ADU 1 and 3 are forwarded to end user 1 and packets with ADUs 1-4 are forwarded to end user 2.

***(ii) On-the-fly compensation for missing marked packets to maintain receiver data quality*:** It is

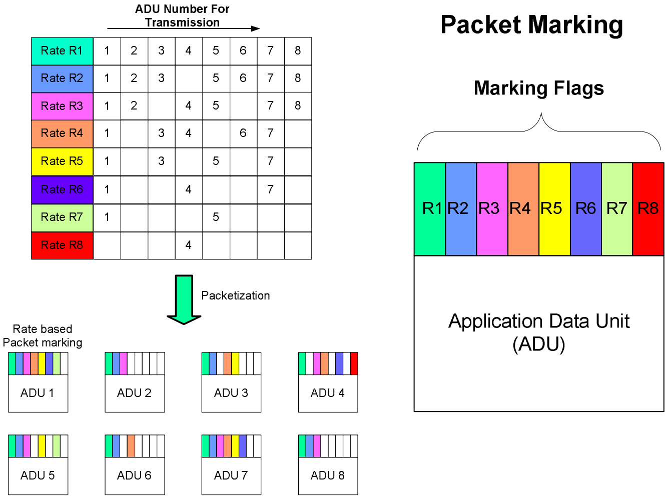

Rate R1	1	2	3	4	5	6	7	8
Rate R2	1	2	3		5	6	7	8
Rate R3	1	2		4	5		7	8
Rate R4	1		3	4		6	7	
Rate R5	1		3		5		7	
Rate R6	1			4			7	
Rate R7	1				5			
Rate R8				4				

Figure 2. Rate based packet marking. Each non-white color represent rate for which packet is marked ,i.e., rate R1-R8.

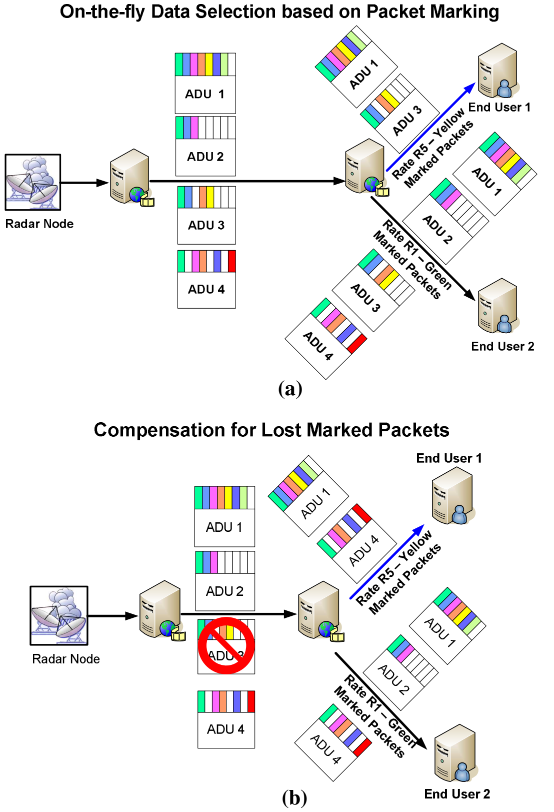

Figure 3. Applications of packet marking. (a) On-the-fly data selection based on packet marking (b) On-the-fly compensation for missing marked packets to meet bandwidth and data quality requirements.

possible that some of the packets with desired marking are dropped or suffers significantly delay in the network when data is sent from a source node to the multicast node for further distribution. If the multicast node further distributes the partially received data then this has the potential to degrade the performance of the end user application. Therefore it is desirable to compensate for the missing packets by selecting packets with the markings corresponding to higher rates than the current transmission rate for a particular end user. When the number of missing packets exceeds some application-specific threshold then multicast node initiate compensation process by selecting packets with marking corresponding to higher transmission rates for forwarding for a particular end user. For example, in Fig. 3(b), a packet with ADU 3 is lost. In this case multicast node may not meet the rate R1 and rate R5 requirements for two end users as both rates need packet with ADU 3. Alternatively, to meet rate R5, multicast node can select packets for forwarding with marking corresponding to rate higher than rate R5, i.e. packets with marking corresponding to any rate between rate R1 and rate R4 may be selected for forwarding to compensate for the missing sample. In Fig. 3(b), multicast node decides to compensate for the missing packet with ADU 3 by selecting a packet with ADU 4 who's marking corresponds to rate R1, rate R3, and rate R4 indicated by Green, Pink and Orange marking flags. Note that compensation for the missing packets with the desired marking is performed only for packets with rates lower than the rate at which data is transmitted by the source node. By performing compensation within the network at intermediate nodes, retransmission of the data may be avoided and thus associated delay is reduced. Note that we have assumed that all nodes that perform marking based packet selection are aware of rate to marking mapping. Each node maintains a static table of flag and the corresponding rate it represents in the packet. Subsequent performance results show that little computation penalty paid for performing application-

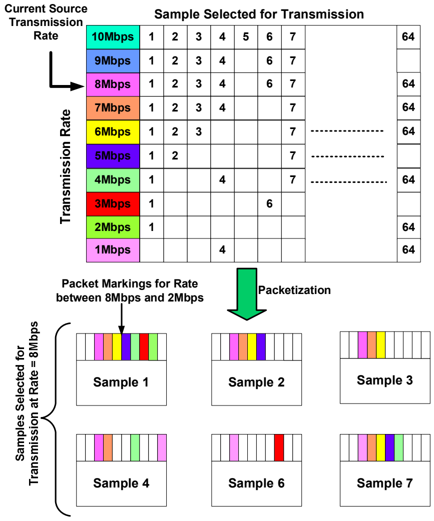

10Mbps	1	2	3	4	5	6	7		64
9Mbps	1	2	3	4		6	7		
8Mbps	1	2	3	4		6	7		64
7Mbps	1	2	3	4			7		64
6Mbps	1	2	3				7	----------------	64
5Mbps	1	2					7	----------------	
4Mbps	1			4			7	----------------	64
3Mbps	1					6			
2Mbps	1								64
1Mbps				4					64

Figure 4. Packet marking for radar data when current transmission rate is 8Mbps. Sample is application data unit (ADU) for the radar data

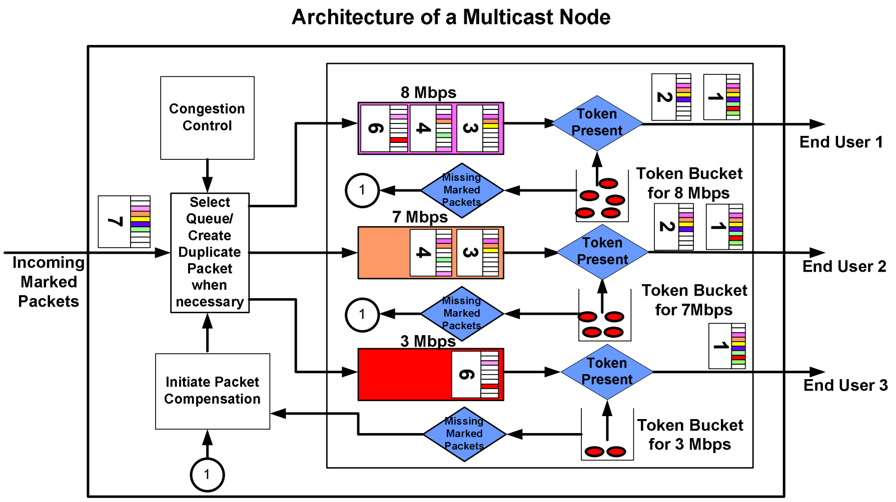

Figure 5. Architecture of a multicast node. Three end user's current transmission rates are 8Mbps, 7Mbps, and 3Mbps and arrival rate of the data at multicast node is 8Mbps. Initial number of tokens in the bucket depends on the end user current transmission rate.

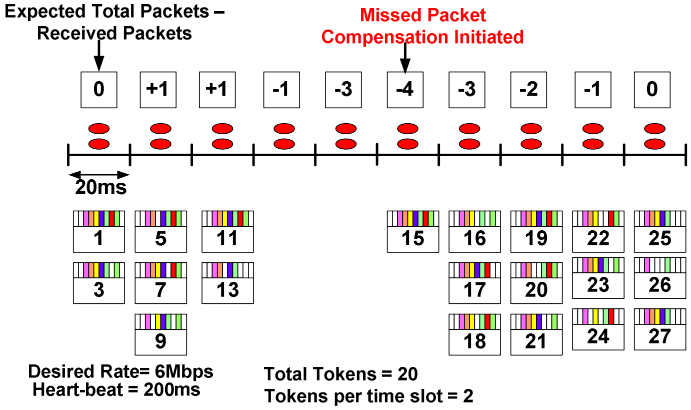

Figure 6. Token bucket based packet compensation to meet bandwidth and data quality requirement

aware processing at each node helps in significantly enhancing the quality of the content delivered to the end users during network congestion. We now explain the application of this packet marking strategy for streaming weather radar data in CASA environment.

3.1. Packet-Marking for Radar Data – An Example

Fig. 4 shows the packet marking of the DRS block of data generated by a radar node. In this example it is assumed that radar data is generated at 10Mbps and current source transmission rate is 8Mbps. A subset of samples is selected at rate 8Mbps for transmission. It is assumed that end users can tolerate uniform loss of samples. Therefore, as seen in Fig. 4, sample drops are uniformly distributed among 64 samples generated by radar for different transmission rates. In Fig. 4, data may be transmitted at ten different rates between 1Mbps and 10Mbps indicated by ten marking flags in each packet. Note that flag is marked when packet is suitable for the current transmission rate.

4. Applications of Packet-Marking

In Fig. 1, different overlay nodes use packet marking to perform application-aware processing within the network to meet the QoS requirements of the DCAS application. Following section describes the application of packet marking in multicast node and the forwarding node.

4.1. Token bucket based rate control: At intermediate nodes such as forwarding nodes, and multicast nodes, packets are selected for transmission based on the current transmission rate and the corresponding packing marking. A token bucket based rate control algorithm is used to achieve the desired transmission rate at these intermediate nodes.

4.1.1. Multicast Node: Fig. 5 shows architecture of the multicast node. In a CASA application, each end user may specify its rate requirement and content-quality requirement in terms of tolerance towards bursty losses or uniform losses within the DRS block at the time of request for the data. The rate at which the data is transmitted to the end users can vary from user to user based on the available bandwidth. A congestion control protocol determines the transmission rate for each end user independently. We consider TRABOL (TCP Friendly Rate Adaptation Based on Losses) congestion control protocol to determine transmission rate of each end user. During network congestion, TRABOL performs rate adaptation while considering end user specific minimum rate (MR) and target rate (TR) requirements [4]. When a packet with application-specific markings arrives at the multicast node, multicast node determines the users to which this packet must be transmitted. Decision to transmit a packet to a particular end user depends on the packet markings and the rate requirements of the end users. The multicast node maintains separate output queues for sending data to each end user. Fig. 5 shows a case for on-the-fly selection of the data for transfer to multiple end users when radar data shown in Fig. 4 is received at the multicast node. In Fig. 5, three end users, i.e., end user 1-3 are considered with current transmission rate of 8Mbps, 7Mbps, and 3Mbps respectively. Multicast node maintains the mapping between marking flag and the corresponding transmission rate. Note that in Fig. 5 the packet with sample 1 is forwarded to all three end users because packet is marked for transmission rate 8Mbps, 7Mbps, and 3Mbps, i.e., the current transmission rates of the end user 1, end user 2, and end user 3 respectively. Alternatively, packet with sample 2 is not forwarded to end user 3 because the packet is not marked for transmission rate 3Mbps.

A token-bucket based scheme is used to maintain the required average transmission rate for each end user. A token bucket size for each end user is determined periodically every 'heart-beat' interval based on its transmission rate. As explained in Section 2, 'heart-beat' interval is the periodic interval after which radar node generates DRS block of data at a constant rate. We consider a case where 1 packet is removed from the end user output queue for transmission for every 1 token present in the bucket. Therefore, the token bucket size gives the upper bound on the number of packets that the multicast node can transmit during a 'heart-beat' interval to a particular end user. Fig. 6 shows the token-bucket scheme and the packet compensation process to meet the current transmission rate and data quality requirement of the end users. Consider an example in Fig. 6 where 20 tokens present in the bucket for a given transmission rate within a 'heart-beat' interval of 200ms are evenly distributed among 10 time slots of 20ms each. A counter is maintained to track if number of received packets with desired marking is equal to the expected number of packets within each time slot. If the number of packets received with the desired marking is less than the number of expected packets then it indicates that some of the desired marked packets are lost or significantly delayed in the network. When this difference exceeds some threshold then it indicates the start of a compensation process. For example, in Fig. 6, compensation for missing packets start when the difference between arriving and expected packets falls below threshold (=-3). As seen in the figure, after compensation is initiated, packets with marking corresponding to a rate higher than the current transmission rate are used to meet the transmission rate requirements. At the end of 'heat-beat' interval, if the difference between total arriving packets and total expected packets becomes 0 then it indicate that compensation process succeeded in meeting the transmission rate requirements.

In Fig. 7, packet-marking algorithm primarily consists of two key phases: (i) Processing of Arriving Packets, and (ii). Processing Packets for Transmission. As seen in the figure, at the beginning of every *heart_beat* interval, algorithm determines the token bucket size *N[END_USER]* for the *END_USER* based on its current transmission rate. The *heart_beat* interval is divided into *S* time slots of equal duration. Let *B[END_USER,i]* denote the number of tokens allocated to the *END_USER* in the time slot *i*. Therefore, *N[END_USER]/S* tokens are evenly allocated to each time slot, i.e., *B[END_USER,i]=N[END_USER]/S*. When a new packet arrives, it is pushed to the appropriate end user queue. At the time of transmission, for each time slot *i* algorithm checks for the number of tokens present in the bucket for each end user. If the tokens are present in the bucket *N[END_USER]* and corresponding queue is not empty, then the packet is popped from the queue and transmitted to the *END_USER*, and *B[END_USER,i]* and *N[END_USER]* are decremented by 1. At the end of time slot *i*, if *B[END_USER,i]* is greater than 0, it is the indication of missing packets with desired marking for a particular *END_USER*. Since data may also arrive in bursts, therefore *B[END_USER,i]* can be less than 0. For e.g., if *B[END_USER,i]=2* and then in the i^{th} time slot 3 packets with required marking arrives then *B[END_USER,i]* is -1 at the end of i^{th} slot. At the end of the i^{th} slot, *B[END_USER,i]* is subtracted from the *L[END_USER]* to track the difference between total expected packets and the actual received packets until the end of the i^{th} time slot. *L[END_USER] >0* indicates that more number of packets with the desired marking arrives than the expected number of packets at the overlay node until that instant. Alternatively, *L[END_USER]<0* indicates that missing or delayed packets with the desired marking. When

L[END_USER]<0 and is below an acceptable threshold value *THRESHOLD* then the protocol initiates marked packet compensation process in the $(i+1)^{th}$ time slot. During the packet compensation process, multicast node also start selecting packets with marking corresponding to rates higher than the current transmission rate *R[END_USER]* of the *END_USER*. Therefore queue for *END_USER* may contain packets with marking corresponding to rates greater than *R[END_USER]*. This process continues as long as *L[END_USER]<0* or until the expiration of the *heart_beat* interval. At the end of the *heart_beat* interval, there may be some marked packets left in the end user queue that couldn't be transmitted within *heart_beat* interval. In that case all queues are cleared leading to drop of marked packets at the overlay node. The motivation for dropping remaining packets in the queues after the *heart_beat* interval is that these packets are likely to arrive late at the end user destination; therefore they may be useless for the end user application. It is important to note that performance of the token bucket based rate control is sensitive to the duration of the *heart_beat* interval. Token bucket based algorithm will be able to select the most appropriate marked packets for forwarding if delay jitter is small compared to the *heat_beat* interval.

4.1.2. Forwarding Node: Fig. 1 shows forwarding nodes in the overlay path. Main task of the forwarding node is to select the appropriately marked packets for forwarding according to the available output link bandwidth. Consider a case, where a high-bandwidth upstream flow relays packets through a forwarding node to a low-bandwidth output link. In this case forwarding node may either buffer or selectively discard packets received from the upstream node. The implementation of a forwarding node is similar to the implementation of a multicast node as explained in Fig. 5-7. However, forwarding node may determine output link bandwidth using bandwidth measuring tools instead of congestion control algorithm used by the multicast node.

```
                         Packet-Marking Algorithm
#heart_beat:  Periodic time interval to generate block of sensor data
#S:           Total number of time slots within heart-beat interval
#i:           Index of the time slot within heat-beat interval
#PACKET:      Packet received at the node
#PKT_SIZE:    Size of the PACKET
THRESHOLD:            Application defined threshold to initiate packet compensation
#R[END_USER]:     Current transmission rate of END_USER
#N[END_USER]:      Number of tokens in the bucket for END_USER
#B[END_USER,i]:    Number of tokens assigned to  time slot i
#L[END_USER]:     Number of missing marked packets for END_USER
#get_transmission_rate(END_USER): Determine current transmission rate for END_USER
#push(PACKET,END_USER):           Push PACKET on a queue for END_USER
#pop_and_transmit(END_USER):      Pop packet from queue for END_USER and transmit it
#clear(END_USER):                 Clear queues for END_USER

#Initialization phase every heart-beat interval
FOR each END_USER {
   R[END_USER] = get_transmission_rate(END_USER) #Determine transmission rate for each end
user
   N[END_USER] = (R[END_USER]*heart_beat)/PKT_SIZE  #Determine token bucket size for each end
user
   FOR each time slot i=0; i <S;i++   #Assign token to each slot of heat-beat interval for END_USER
     B[END_USER,i] = N[END_USER]/S
      L[END_USER]=0   #Initialize missing packets to 0 for each END_USER
  }
#Process New Arriving Packets
FOR each arriving PACKET
   FOR each END_USER {
      #Check if the desired packet arrives based on packet marking
      IF (PACKET MARKING==END_USER MARKING) {
         push(PACKET, END_USER) #Push packet to queue for the end user
         }
     #Check if packet with higher marking than the desired arrives
     ELSE IF (PACKET MARKING > END_USER MARKING) {
        IF (L[END_USER]<THRESHOLD) { #Check if compensation phase
           push(PACKET, END_USER)  #Push packet to the queue for the end user
           }
        }
     }
 #Process Packet Transmission
  FOR each time slot i=0;i < S;i++ {
     FOR each END_USER { # Repeat following loop multiple time until current time slot expires
        IF  (N[END_USER] >0) {  # Check for presence of tokens in the bucket for end user
           IF  (!queue_empty(END_USER)) {
              pop_and_transmit(END_USER)  #Pop queue and transmit
              B[END_USER, i)  =  B[END_USER,i] -1  # Track number of packets received in time slot i

              N[END_USER]     = N[END_USER]-1  # Number of tokens left in the bucket for
END_USER
            }
        }
     }
   # After time slot i expires, update number of missing or excess packets until time slot i
   L[END_USER]     =  L[END_USER] - B[END_USER,i]
   }
  #Clear all queues at the end of heat-beat interval
  FOR each  END_USER
       clear(END_USER)
```

Figure 7. Implementation of the packet-marking algorithm for multicast node and forwarding node.

5. Performance Results

Performance of the packet-marking scheme and token-bucket based rate control algorithm is evaluated in a network emulation environment as shown in Fig. 8. The NIST Net [7] based network emulator along with TCP cross-traffic is used to control the bandwidth between source node and the multicast node, and to control bandwidth between multicast node and end users. In all experiments, we consider the case where source node generates data at a constant rate of 10Mbps with different bandwidth requirements for different end users. The multicast node receives a single copy of the packet from the source node and is transmitted to multiple end users with different round trip delays, and with different bandwidth and data quality requirements. Experiments were conducted to investigate: (i) Performance in meeting bandwidth requirements of end users, (ii) Impact of packet-marking scheme on data quality, and (iii) Impact of packet compensation algorithm on data quality.

The first set of experiments evaluates the effectiveness of the proposed packet-marking scheme along with token-bucket rate control scheme in meeting bandwidth requirements of the end users. As seen in Fig. 9, when bottleneck is 30Mbps, which lies between sum of target rate requirements and sum of minimum rate requirements of all the end users, then all end users shares the bandwidth while meeting their bandwidth constraints. Next experiment results investigate the effect of application-aware selective drop and packet-

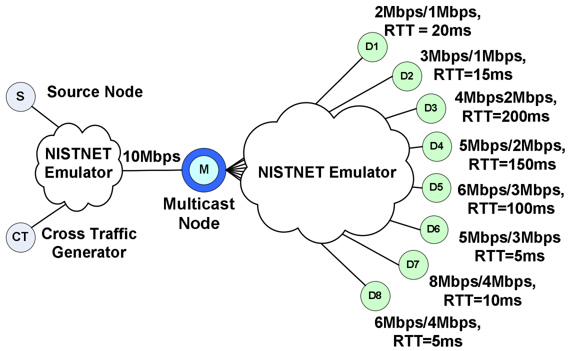

Figure 8. Emulation network for application-aware multicasting of weather radar data

marking on data quality of the end users. Eight different end users with different RTT requests multicast node for the data with their data quality and bandwidth requirements. In this experiment, the multicast node combines the bandwidth requests of all end user and makes an aggregate request to the source node with target rate (TR) = 8Mbps, minimum rate (MR) = 4Mbps requirement. Source node initially selects subset of the data for transmission at 8Mbps from the DRS block. We consider random sample drop and selective sample drop as the sample selection scheme at the source node for determining subset of data for transmission at rates lower than the data generation rates[5]. In the selective drop, the source node considers end user's sensitivity to different components of the DRS block while selecting a subset of samples according to the current transmission rate for a multicast node. Alternatively, in the random drop case; the source node randomly selects subset of the samples from the DRS block for a given transmission rate.

Experiments were conducted to compare quality of the data delivered to the end users when packet-marking is used, i.e., packet-marking based forwarding to the case when no packet marking is used, i.e., random forwarding. In the packet-marking based forwarding, the source node mark packets as explained in Section 3 and the packet marking enables the multicast node to determine packets to be forwarded to the each end user independently based on the available bandwidth for each end user. In the random forwarding case packets are not marked; packets that arrive at the multicast node are transmitted to an end user in FIFO basis as long as there are tokens present in the bucket within 'heart-beat' interval. We compare the performance for three cases based on how data is dropped at the source node and whether packet-marking is supported or not: (1) Random drop, Random forwarding, (2) Selective drop, Random forwarding, and (3) Selective drop, Packet-marking based forwarding.

In case of weather radar data, quality of the received raw data is measured by computing standard deviation in the reflectivity parameter [4] for each end user. Lower standard deviation is a measure of better quality data [4,5,6]. Fig.10 shows the standard deviation of reflectivity parameter of 10 gates (141 ~ 150) of one end user with TR=4Mbps, MR=2Mbps. We compare the results with the baseline case in which all samples generated by a radar node is delivered to the end user application. As seen in Fig. 10, standard deviation is minimum for the baseline case and is highest for a case when data is randomly dropped at source node and no packet-marking is supported. Alternatively, quality of the data improves, i.e., lower standard deviation, with selective drop and packet-marking support. Improvement in data quality at end user is due to end user receiving the required subset of the data at any given transmission rate.

Experiments are conducted to study the effect of the packet compensation scheme on the receiver throughput and data quality of the end users. Bandwidth between source node and multicast node is configured as 10Mbps using NIST Net. TCP cross-traffic is used to introduce random losses in the network between source node and the multicast node. Initially source node transmits marked packets with radar data to the multicast node at 8Mbps but due to competing TCP cross-traffic, random radar data packet may get lost between source node and the multicast node. Alternatively, bandwidth between multicast node and end users is sufficient to support the target rate all end users thus no losses are encountered in network between multicast node and the end user. Performance is compared when marked data is transmitted using packet compensation and without any compensation. In Fig. 11, dotted lines show receiver throughput of the end users without any packet compensation. In Fig. 11, solid lines show end users receiver throughput using packet compensation scheme. As seen in the figure, when network losses increase due to increase in TCP cross-traffic, the throughput of end users decreases. However, when the packet compensation is used end user receives higher throughput compared to no compensation case. Since TR=8 Mbps, MR=4 Mbps end user receives data at the highest rate, the multicast node cannot find any higher stream to compensate for the missing packets, so no compensation is performed for that particular end user. Fig. 12 shows the impact of the packet compensation scheme on the quality of the data received by the end users for different gates. Fig. 12 shows standard deviations of the reflectivity parameter for two end users with bandwidth requirements TR=6 Mbps, MR=3 Mbps, and TR=3 Mbps, MR=2 Mbps. As seen in the figure data quality improves when packet compensation is used at the multicasting node indicated by decrease in standard

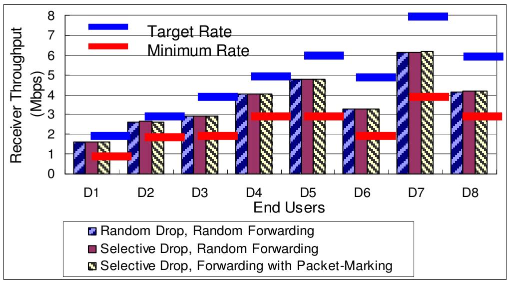

Figure 9. Receiver throughputs of the end users with different rate requirements

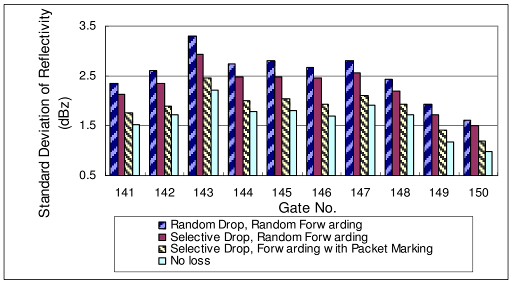

Figure 10. Effect of the application-aware selective drop and packet-marking on data quality for the end user with TR=4Mpbs, MR=2Mbps

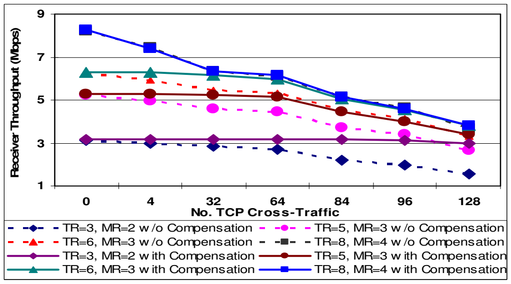

Figure 11. Effect of packet compensation on the receiver throughput at the end users

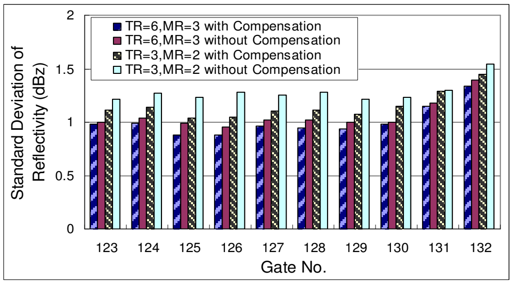

Figure 12. Effect of packet compensation on the data quality at the end users

deviation in presence of packet compensation.

6. Conclusions

This paper presented content based packet marking and a token-bucket rate control scheme to support application-aware transport requirements using overlay networks. Although simpler schemes can be used to meet the bandwidth requirement, the results show that application aware schemes at intermediate nodes can result in better quality of the end result. We demonstrated the effectiveness of the proposed approach by evaluating its performance in a network emulation environment. During network congestion, packet marking is very effective in delivering high quality data to the end user. Moreover, when packet compensation technique is used it further improves the received bandwidth and data quality of the end users.

References

[1] Andersen, D., Bansal, D., Curtis, D., Seshan, S., Balakrishnan, H., "System Support for Bandwidth Management and Content Adaptation in Internet Applications," 4th USENIX OSDI Conf., San Diego, California, Oct. 2000.

[2] Andersen, D., Balakrishnan, H., Kaashoek, M.F., Morris, R., "Resilient Overlay Networks," in Proc. 18th ACM SOSP, Banff, Canada, Oct. 2001

[3] Banerjee, S., Kommareddy, C., Kar, K. Bhattacharjee, B., and Khuller, S. "Construction of an efficient overlay multicast infrastructure for realtime applications," in Proc. IEEE INFOCOM, June 2003.

[4] Banka, T., Maroo, A., Jayasumana, A.P., Chandrasekar, V., Bharadawaj, N., and Chittababu, S.K., "Radar Networking: Considerations for Data transfer Protocols and Network Characteristics," Proc. of AMS IIPS for Meteorology, Oceanography, and Hydrology, American Meteorological Society (AMS), 19.11. Jan. 2005

[5] Banka, T., Lee, P., Jayasumana, A.P., Chandrasekar, V., "Application Aware Overlay One-to-Many Data Dissemination Protocol for High-Bandwidth Sensor Actuator Network," IEEE COMSWARE 2006, New Delhi, India, Jan. 2006

[6] Bringi, V. N., and Chandrasekar, V., "Polarimetric Doppler Weather Radar: Principles and Operations," Camb. Univ. Press, Aug. 2001

[7] Carson, M., and Santay, D., "NIST Net: A Linux-based Network Emulation Tool," ACM SIGCOMM Computer Comm. Review, Vol. 33, Issue 3, pp. 111–126, Aug. 2003

[8] Kassler, A., Neubeck, A., and Schulthess, P., "Classification and Evaluation of Filters for Wavelet Coded Videostreams," Signal Processing: Image Communication, 16(8):795-807, Elsevier May, 2001

[9] Keller, R., Choi, S., Dasen, M., Decasper, D., Fankhauser, G., and Plattner, B., "An Active Router Architecture for Multicast Video Distribution," Proc. IEEE INFOCOM, Mar. 2000

[10] McCanne, S., Jacobson, V., "Receiver-driven Layered Multicast," Proc ACM SIGCOMM'96 Aug. 1996

[11] McLaughlin, D.J., Chandrasekar, V., Droegemeier, K., Frasier, S., Kurose, J., Junyent, F., Philips, B., Cruz-Pol, S., and Colom, J. "Distributed Collaborative Adaptive Sensing (DCAS) for Improved Detection, Understanding, and Prediction of Atmospheric Hazards," Proc. of AMS IIPS for Meteorology, Oceanography, and Hydrology, American Meteorological Society (AMS), 11.3, Jan 2005.

[12] Nicolas, K., Blake, S., F. and D. Black, "Definition of the Differentiated Service Field (DS Field) in the IPv4 and IPv6 Headers", RFC2474, Dec. 1998

[13] Peterson, L., Anderson, T., Culler, D., and Roscoe, T., "A Blueprint for Introducing Disruptive Technology into the Internet," 1st ACM Workshop on Hot Topics in Networks, HotNets-I, Oct. 2002

[14] Subramanian, L., Stoica, I., Balakrishnan, H., Katz, R.,"OverQoS: An Overlay Based Architecture for Enhancing Internet QoS," Proc. 1st Symposium on Networked Systems Design and Implementation (NSDI), San Francisco, CA, Mar. 2004.

[15] Tennenhouse, D.L., Smith, J. M., Sincoskie, W. D., Wetherall, D. J., and Minden, G. J., "A Survey of Active Network Research," IEEE Communications Magazine, 35(1):80–86, 1997.

[16] Zhang, Z.-L., Nelakuditi, S., Aggarwal, R., and Tsang, R. "Efficient Selective Frame Discard Algorithms for Stored Video Delivery across Resource Constrained Networks," Proc. of IEEE INFOCOM, Mar. 1999