Fair Uni- and Multicasting in Ring Metro WDM Network

Chiara Piglione (i), Martin Reisslein (ii), and Fabio Neri (i)

(i) Dipartimento di Elettronica, Politecnico di Torino, Corso duca degli Abruzzi 24, Torino, Italy

(ii) Corresponding author, Arizona State University, Goldwater Center, MC 5706, Tempe AZ 85287–5706, phone: 480-965–8593, fax: 480-965-8325, http://www.fulton.asu.edu/~mre

{*chiara.piglione, fabio.neri*}*@polito.it, reisslein@asu.edu*

Packet-switched wavelength division multiplexing (WDM) ring networks have been extensively studied as solutions to the increasing amount of traffic in metropolitan area networks, which is widely expected to be a mix of unicast and multicast traffic. In this paper we study the fairness between unicasting and multicasting in slotted packet-switched WDM ring networks that employ a tunable transmitter and fixed tuned receiver at each node and *a posteriori* buffer selection. We find that single-step longest queue buffer selection generally results in unfairness between unicasting and multicasting or a fixed relative priority for multicast vs. unicast traffic. We propose and evaluate dual-step buffer selection policies that achieve fairness and allow for a range of relative priorities of multicast vs. unicast traffic.

1. Introduction

Metropolitan area networks, which interconnect access networks, such as EPONs, with each other and with the high-speed wavelength division multiplexing (WDM) backbone networks, are expected to experience a surge in traffic with the ever increasing speed of the access network technologies. Ring WDM networks have been extensively studied as solutions to the increasing traffic in metropolitan area networks, see for instance [1, 2, 3, 4, 5, 6, 7, 8, 9]. These studies have primarily focused on the efficient transport of point-to-point (unicast) traffic. Point-to-multipoint (multicast) traffic however is widely expected to account for an increasing portion of the traffic in metropolitan area networks as applications that rely on multicasting, such as telepresence/teleconference, telemedicine, multimedia content distribution, software update distribution, become more prevalent. Multicasting in WDM ring networks has received relatively little attention to date, as detailed in Section 1.A. It is widely expected that traffic in future metropolitan area networks will typically be a mix of unicast and multicast traffic and it is therefore important to understand the issues involved in transmitting these two types of traffic together over a packet-switched WDM ring network.

Many different types of ring networks have been studied in the literature, with the single-fiber ring network with a tunable transmitter and fixed tuned receiver (TT-FR) node structure being the most commonly considered type of ring network. We focus on this type of ring network in our study, which to the best of our knowledge is the first to examine the fairness issues involved when transmitting a mix of unicast and multicast traffic over a packet-switched WDM ring network. Unfairness tends to arise when transmitting a mix of unicast and multicast traffic in the common TT-FR single-fiber ring network mainly because (*i*) unicast traffic requires the transmission of a single packet whereas multicast traffic typically requires the transmission of multiple packet copies, and (*ii*) a transmitted unicast packet is received by one destination node, whereas a transmitted multicast packet copy is typically received by several destination nodes.

In this paper we first examine the transmission of a mix of unicast and multicast traffic when different virtual output queue (VOQ) structures are employed in conjunction with the longest queue buffer selection, which determines the queue from which to transmit a packet in a single step and is widely employed in the studied ring networks in the literature. We find that this single-step buffer selection in conjunction with VOQ architectures with separate queues for unicast and multicast traffic generally lead to significant unfairness. Typically the multicast traffic experiences an uneven throughput, or packet loss, or both, compared to the unicast traffic, as detailed in Section 5. Buffering unicast and multicast traffic in the same VOQs, on the other hand, results in fairness in the sense that the ratio of transmitted multicast packet copies to transmitted unicast packets is approximately equal to the ratio of generated multicast packet copies to generated unicast packets. Also, both unicast and multicast traffic experience approximately the same delay and packet loss.

In many application scenarios unicasting is used to transmit time-sensitive messages of relatively high importance, whereas multicasting is used for the non-real-time transfer of bulk data (e.g., content distribution, software updates). To allow for such different delay and loss priorities for unicast and multicast traffic and more generally to allow for a range of multicast to unicast throughput ratios we develop dual-step buffer selection policies. The dual-step policies combine longest queue selection to determine the channel on which to transmit a packet with a probabilistic selection policy to determine whether to transmit a unicast or multicast packet.

This paper is structured as follows. In the following subsection we review related work. In Section 2, we describe the architecture and medium access control (MAC) protocol of the considered single-fiber ring WDM network. We also describe the considered conventional node vs. node fairness mechanism, which ensures that the different ring nodes have approximately equivalent transmission opportunities, irrespective of their location on the ring. In Section 3, we introduce the considered traffic model and the throughput and delay metrics. In Section 4, we introduce the metrics used to assess the fairness between the treatment of unicast and multicast traffic. In Section 5, we present the simulation results for the networks employing separate or joint VOQs for unicast and multicast traffic in conjunction with the single-step longest queue buffer selection policy. In Section 6, we develop and evaluate the dual-step buffer selection policies. We summarize our conclusions in Section 7.

1.A. Related Work

Most closely related to our work are the lines of work on (*i*) multicasting in WDM ring networks, (*ii*) fairness control in unicasting over WDM ring networks, and (*iii*) the joint transport of unicast and multicast traffic in communications networks.

Multicasting in WDM ring networks has received relatively little attention to date [10]. The photonics level issues involved in multicasting over ring WDM networks are explored in [11]. A node architecture suitable for multicasting in WDM ring networks is studied in [12]. The general network architecture and MAC protocol issues arising from multicasting in packet-switched WDM ring networks are addressed in [13, 14]. These studies do not address the fairness issues arising from transmitting a mix of unicast and multicast traffic over a WDM ring network, which is the main focus of this paper. We note for completeness that the wavelength assignment for multicasting in circuit-switched WDM ring networks, which are fundamentally different from the packet-switched networks considered in this paper, has been studied in [15, 16, 17, 18, 19].

The fair transmission of unicast traffic in packet-switched optical bus and ring networks has received considerable attention, see for instance [1, 4, 20, 21, 22, 23, 24, 25]. These studies consider the problem of assuring that each node is provided with fair opportunities for packet transmissions irrespective of the location of the nodes along the ring. This problem is largely orthogonal to the problem of providing fair opportunities for the transmission of unicast and multicast traffic, which we focus on in this paper.

The fairness issues arising when transmitting a mix of unicast and multicast traffic over general packet-switched networks has been examined in a number of studies, see for instance [26, 27, 28, 29, 30], which do not consider the specific properties of optical WDM networks. The problem of scheduling a mix of unicast and multicast traffic in broadcast-and-select WDM networks, which are fundamentally different from the ring WDM network considered in this paper, is studied in [31].

2. Slotted Ring WDM Network

2.A. Network Architecture

In this section, we describe the considered architecture of an all-optical WDM ring network with N nodes and Λ logical wavelength channels. We consider a single-fiber ring network, in which successive nodes are connected with a single unidirectional fiber. The fiber bandwidth is divided into Λ wavelength channels. Each channel is divided into fixed-length time slots whose boundaries are synchronized across all wavelengths. The slot duration equals the transmission time of a fixed-size packet. Each node is equipped with one tunable transmitter and one fixed tuned receiver (i.e., a node can send packets on any wavelength and it is able to receive packets only on a preassigned wavelength). For $N = \Lambda$ each node has its own separate *home channel* for reception. For $N > \Lambda$ each wavelength is shared by several nodes for the reception of packets. In particular, the destination nodes $j = i+n\cdot\Lambda$ with $n \in \{0, 1, \ldots, \lceil N/\Lambda \rceil - 1\}$ share the same drop wavelength i, $i \in \{1, 2, \ldots, \Lambda\}$, i.e., have channel i as their home channel. A given node receiver terminates the wavelength channel on which it is homed. Nodes sharing the same wavelength may have to forward packets toward the destination node, resulting in *multi hopping*. For unicast traffic, the destination node removes the packet from the ring. For multicast traffic, when a node receives a packet, it checks if there are additional destinations downstream; if so, it forwards the packet to the other destinations; otherwise, the node is the last destination and removes the packet from the ring. With this destination release (stripping), wavelengths can be spatially reused by downstream nodes, leading to an increased network capacity.

To avoid head-of-the-line (HOL) blocking each node is typically equipped with at least Λ virtual output queues (VOQs), one for each wavelength on the ring. We refer to the basic buffer architecture with exactly Λ VOQs at a node as Λ*VOQ architecture*. Note that in the ΛVOQ architecture the unicast and the multicast packets to be transmitted on a given wavelength are queued in the same buffer. Often the unicast and multicast packets are buffered separately [13, 28, 32, 33, 34, 35]. The main advantage of separate buffers for unicast and multicast packets is that unicast and multicast packets can be transmitted with different priorities. In this paper we refer to these architectures as $(u+m)$VOQ architectures, where u and m are the number of unicast and multicast buffers, respectively. To avoid HOL blocking both u and m have to be larger than or equal to Λ. We consider the $((N-1)+\Lambda)$VOQ and the $(\Lambda+\Lambda)$VOQ buffer architectures. In the $((N-1)+\Lambda)$VOQ architecture there is a separate unicast buffer for each possible destination node for unicast traffic. In the $(\Lambda+\Lambda)$VOQ architecture, unicast packets to be transmitted on the same wavelength channel are stored in the same unicast buffer. In both architectures, one of the Λ multicast buffers is assigned to each of the Λ wavelength channels. The fanout set of a given multicast packet is partitioned into up to Λ destination groups according to the different home channels of the destination nodes. A copy of the multicast packet is generated for each group of destination nodes and placed in the corresponding multicast buffer. If a multicast packet has destinations on each of the home channels, then Λ packet copies are generated, and one each is placed in the Λ multicast buffers. If all nodes in the fanout set of a multicast packet share the same home channel, then only one packet copy is generated and placed in the corresponding multicast buffer. All the buffers are First In First Out (FIFO).

2.B. MAC Protocol

On each wavelength, each slot consists of a payload field and corresponding control information. The control information gives the slot availability status (empty or busy) and the destination address (fanout set of the packet copy in the case of multicast packet) of the packet transmitted in the slot. The control information may be transmitted on a separate control channel (e.g., as in [9, 25, 36]) or in a subcarrier multiplexed header (e.g., as in [37, 38]). Each node monitors all wavelengths simultaneously and detects the channel availability information in every slot.

We consider an *a posteriori* access strategy, i.e., a node first checks the availability status for a given slot on all wavelength channels and then selects the appropriate buffer for transmission. The node has to wait until an empty slot arrives on at least one wavelength channel. When an arriving slot is empty on one (or more) wavelength channel(s), the node can use this slot to transmit a packet from one of the corresponding queues. In the ΛVOQ architecture, buffer selection is necessary if multiple channels have an empty slot. In the $(u+m)$VOQ architecture, buffer selection is necessary when one or more channels have an empty slot, since at least two queues (the unicast and the multicast queues) store packets for the same home channel and a node can only transmit one packet at any given time with its single transmitter. We consider different buffer selection schemes which we examine in Sections 5 and 6. All considered buffer selection schemes are based on the *Longest Queue* (LQ) buffer selection principle. With the LQ principle, the longest queue is chosen. The motivation for the LQ buffer selection principle is load balancing among the queues in the network, which increases the node and network throughput for an acceptable system complexity [1].

2.C. Node Vs. Node Fairness Control

Due to the ring symmetry and the applied destination release each node has a better-than-average access to channels leading to certain destination nodes and a worse-than-average access to channels leading to other destinations [23]. Spatial reuse may cause starvation, which occurs when a node is constantly being covered by up-stream ring traffic and thus is not able to access the ring for very long periods of time [4]. This fairness problem has received considerable attention in the literature, as detailed in Section 1.A.

In this work we use the fairness control described in [1], which is a modified form of ATMR [39]. This fairness control represents a credit allocation scheme and provides fair channel access by means of a distributed credit mechanism and a cyclic reset scheme based on a monitoring approach. Initially, each node is allocated a predefined credit, called *window size* (W), for each wavelength channel and is set to the active state. The node status (active or inactive) for a channel is included in the control information in each slot. Each node decreases the window size whenever it uses a free slot to send a packet. If the node is still in the active state, i.e., the remaining window size is larger than zero, the node sets the busy address field to the node's address. When the window size reaches zero, the node sets its state as inactive, i.e., the node is not allowed to send any more packets, and leaves the busy address field unchanged. Thus, a node can see whether there are any other nodes in the active state. If a node receives a slot with the busy address field set to the node itself, it knows that all the other nodes are in the inactive state. Then the node immediately sends a reset message to all other nodes by setting the *reset-request field* in the control information and reset its windows size to the predefined value W. The node sends the message only once and waits for the reset message to circulate around the entire ring network. When the reset message is received by the node that sent it, the node strips the message from the ring. When a node receives a reset-request, it sets its status to the active state, sets the window size for the channel to the predefine value W and forwards the reset-request. This algorithm is invoked on all Λ wavelength channels at each node.

3. Unicasting and Multicasting Model

In this section we describe the considered model of the unicast and multicast traffic and define the considered performance metrics. In our traffic model each of the N nodes generates traffic independently of the other nodes. We consider self-similar traffic with a Hurst parameter of 0.75 which we generate from ON/OFF processes with Pareto distributed on-duration and geometrically distributed off-duration [40]. We denote σ, $0 \leq \sigma \leq 1$, for the long run average probability that a given node generates a new packet at the beginning of a given slot. We consider a balanced traffic scenario, i.e., the packet generation probability σ is the same for all nodes. We consider a mix of multicast and unicast traffic with a fraction p_m of multicast traffic. A given newly generated packet is a unicast packet with probability $1-p_m$; a generated packet is a multicast packet with the complementary probability p_m. We consider uniform traffic, i.e., a given unicast packet generated by a node is destined to any one of the other $(N-1)$ nodes with equal probability $1/(N-1)$. For a given generated multicast packet, we draw the fanout (number of destination nodes) independently from a uniform distribution over $[1, N-1]$, and we draw the fanout set (set of destination nodes) randomly and uniformly from the other $N-1$ nodes. We note that the balanced uniform traffic model is generally a good model for the traffic in metro core ring networks. Metro edge ring networks, on the other hand, experience typically strongly hubbed (unbalanced and non-uniform) traffic that is collected from several access networks and forwarded to a core ring. We focus on the balanced uniform traffic model in this paper, but note that our investigation of multicast vs. unicast fairness for balanced uniform traffic is a starting point. In future work it is important to examine the unicast vs. multicast fairness for unbalanced non-uniform traffic models.

We now describe the main implications of the considered traffic model. The probability that a given node generates a unicast packet at the beginning of a slot is

$$\sigma^u = \sigma \cdot (1 - p_m), \tag{1}$$

and the probability that the node generates a multicast packet at the beginning of the slot is:

$$\sigma^m = \sigma \cdot p_m. \tag{2}$$

Note that for a given multicast packet, we generate a copy for each wavelength that leads to destinations of the multicast. In other words, the number of generated packet copies is equal to the number of distinct drop wavelengths among the set of destination nodes. We denote $E[\Delta]$ for the expected number of packet copies generated for a given multicast packet. Noting the equivalence between the nodes attached to an AWG output port in the star network analyzed for multicast traffic in [41] and the nodes sharing a common drop wavelength in the ring network considered here, we obtain

$$\begin{aligned} E[\Delta] &= \frac{\Lambda}{(N-1)-1}\left[\sum_{n=1}^{N-N/\Lambda}\left(1-\frac{(N-n)!(N-N/\Lambda)!}{(N-n-N/\Lambda)!N!}\right)+\sum_{N-N/\Lambda+1}^{N-1}1\right] \\ &= \frac{\Lambda}{N-2}\left[\sum_{n=1}^{N-N/\Lambda}\left(1-\frac{(N-n)!(N-N/\Lambda)!}{(N-n-N/\Lambda)!N!}\right)+\frac{N}{\Lambda}-1\right]. \end{aligned} \tag{3}$$

Thus, a given node generates packet copies with a long run average rate of

$$\sigma^c = \sigma^m \cdot E[\Delta] = \sigma \cdot p_m \cdot E[\Delta] \tag{4}$$

packet copies per slot. Noting that each multicast packet has on average $(N-1+1)/2 = N/2$ destination nodes, we obtain that a given multicast packet copy has an average fanout of $E[f] = N/(2 \cdot E[\Delta])$. That is, a given multicast packet copy reaches on average $E[f]$ destinations by transmission on its wavelength. Note that when $N \gg \Lambda$, then the average fanout approaches the number of wavelength channels, i.e., $E[\Delta] \approx \Lambda$, resulting in an average fanout of approximately $E[f] \approx N/(2 \cdot \Lambda)$ for a given multicast copy on the corresponding wavelength channel.

To determine the stability limit of the network, note that with the considered destination stripping with spatial wavelength reuse, a unicast packet transmission traverses on average half of the circumference of the ring. A multicast packet copy, on the other hand, needs to traverse the ring far enough to reach all of its uniformly randomly distributed destination nodes on the wavelength channel. Noting that the expected number of destination nodes on the wavelength channel $E[f]$ approaches half of the nodes on the wavelength channel, we may reasonably approximate the ring span travelled by a multicast packet copy by the full ring circumference. Thus, the network is stable for $N \cdot \sigma^u/2 + N \cdot \sigma^c \leq \Lambda$, or equivalently $\sigma \leq \Lambda/[(0.5 - 0.5 \cdot p_m + p_m \cdot E[\Delta])N]$.

We proceed to calculate the mean arrival rates of packets to the individual buffers in a given node. In this analysis we initially focus on the $(\Lambda + \Lambda)$ buffer architecture. With the considered uniform traffic scenario, a unicast packet is generated and placed in unicast buffer i, $i = 1, \ldots, \Lambda$, at a given node at the beginning of a slot with probability

$$\sigma_i^u = \frac{\sigma^u}{\Lambda} = \frac{\sigma \cdot (1 - p_m)}{\Lambda}. \tag{5}$$

Similarly, a multicast packet copy is placed in multicast buffer i, $1 \leq i \leq \Lambda$, at a given node at the beginning of a slot with probability

$$\sigma_i^c = \frac{\sigma^c}{\Lambda} = \frac{\sigma \cdot p_m \cdot E[\Delta]}{\Lambda}. \tag{6}$$

Note that for non-uniform traffic patterns the arrival rate at different buffers would be different. Also note that the mean buffer arrival rate for the unicast buffers in the $((N-1)+\Lambda)$ VOQ buffer architecture is obtained by multiplying the arrival rate in the $(\Lambda+\Lambda)$ VOQ architecture by $\Lambda/(N-1)$. The multicast buffer arrival rates are the same in both architectures.

In our performance evaluation we consider the transmitter throughput, the receiver throughput, the packet delay, and the packet loss probability as defined in detail as follows. We also consider a fairness index, which we define in Section 4. We define the mean transmitter throughput of unicast buffer i at node n as the probability that a packet is transmitted from this buffer in a given slot and denote this throughput by $T_{n,i}^u$. We define the mean transmitter throughput of multicast buffer i at node n in analogous fashion and denote it by $T_{n,i}^m$. Note that the mean aggregate unicast (multicast) transmitter throughput in the network is $T^u = \sum_{n=1}^{N} \sum_{i=1}^{\Lambda} T_{n,i}^u$ ($T^m = \sum_{n=1}^{N} \sum_{i=1}^{\Lambda} T_{n,i}^m$). The total transmitter throughput in the network is $T = \sum_{n=1}^{N} \sum_{i=1}^{\Lambda} (T^u_{n,i} + T^m_{n,i})$ and gives the average number of source nodes transmitting a packet (copy) in steady state. Note that it is not sufficient to consider only these absolute throughput levels when assessing the fair treatment of unicast and multicast packets. This is because the mean arrival rates differ for unicast and multicast buffers. In order to fairly assess the achieved throughput for unicast and multicast packets, we normalize the absolute throughput levels by the mean arrival rates. More formally, we define the effective unicast (multicast) transmitter throughput of unicast (multicast) buffer i as $\tau^u_{n,i} = T^u_{n,i}/\sigma^u_i$ ($\tau^m_{n,i} = T^m_{n,i}/\sigma^c_i$).

We define the mean receiver throughput of unicast buffer i at node n as the average number of destination nodes that receive a packet (copy) from buffer i in node n in a given slot in steady state. Note that the number of destination nodes receiving a packet from a unicast buffer in a slot is upper bounded by N/Λ. We denote the receiver throughput of unicast buffer i in node n by $R^u_{n,i}$. We define the mean receiver throughput of multicast buffer i at node n in analogous fashion and denote it by $R^m_{n,i}$. Note that the instantaneous receiver throughput of multicast buffer in a slot is also upper bounded by N/Λ. Also, note that a single multicast packet copy can count up to N/Λ times toward the receiver throughput, which is the case when the packet copy is destined to all N/Λ nodes sharing the drop wavelength channel associated with the buffer. The mean aggregate receiver throughput in the network is $R = \sum_{n=1}^{N} \sum_{i=1}^{\Lambda} (R^u_{n,i} + R^m_{n,i})$ and gives the average number of destination nodes receiving a packet (copy) in steady state. Similar to the effective unicast (multicast) transmitter throughput we define the effective unicast (multicast) receiver throughput of unicast (multicast) buffer i as $\rho^u_{n,i} = R^u_{n,i}/\sigma^u_i$ ($\rho^m_{n,i} = R^m_{n,i}/\sigma^c_i$).

We define the mean packet delay in unicast buffer i at node n as the time period elapsed from the generation of a packet to the complete reception of the packet in slots in steady state and denote it by $D^u_{n,i}$. For a multicast packet, we consider the individual delays until the complete reception of the individual packet copies by the individual receivers. (For instance, for a multicast packet with three packet copies, each with a fanout of two, there are six delay samples.) This counting of the delay accounts for the delays experienced by the individual receivers until they receive their copies of the multicast packet. We denote the mean packet delay by $D^m_{n,i}$.

We define the relative packet loss L as the ratio of the total number of dropped packets to the total number of generated packets in the network. A generated multicast packet with m destination nodes counts as m generated packets. A multicast packet copy directed to n destination nodes that finds the buffer full and is consequently dropped, counts as n dropped packets. We denote L^m and L^u for the relative packet loss of multicast packet copies and unicast packets, respectively.

We estimate these performance metrics from discrete event simulation. Each simulation was run for 10^6 time slots (including a transient phase of 10^5 time slots for warming up each simulation).

4. Unicast vs. Multicast Fairness Metrics

In this section we define the considered metrics for assessing the fair treatment of unicast and multicast packets. In general, fairness can be defined in many ways. We approach the problem of finding a sensible fairness metric by considering a linear combination of the performance metrics defined in the preceding section. We consider initially the $(\Lambda + \Lambda)$ VOQ architecture; adapting the developed metrics to the other architectures is straightforward. We initially define the *fairness index* for unicast buffer i at node n as

$$\alpha^u_{i,n} = \frac{w_1}{D^u_{i,n}} + w_2 T^u_{i,n} + w_3 R^u_{i,n}, \tag{7}$$

with $w_1 + w_2 + w_3 = 1$. The fairness index $\alpha^m_{n,i}$ for the corresponding multicast buffer is defined in analogous fashion. By adjusting the weights w_1, w_2, and w_3 we can adjust the emphasis placed on delay, transmitter throughput, and receiver throughput. Note that the thus defined fairness index captures the absolute throughput and delay levels.

In fairness studies it is quite common to consider the throughput provided to a traffic flow in proportion to the amount of traffic of the flow. In our context, we can express this in a *proportional* fairness index, which employs the effective throughput measures,

$$\pi^u_{n,i} = \frac{w_1}{D^u_{i,n}} + w_2 \tau^u_{i,n} + w_3 \rho^u_{i,n}. \tag{8}$$

We can assess the fairness of a considered packet transmission scheme by comparing the fairness indices achieved by the individual unicast and multicast buffers. With a perfectly fair packet transmission scheme the fairness indices associated with all buffers should be (approximately) equal. At the same time, it is desirable to maximize the throughput-delay performance of the network, which corresponds to maximizing the sum of all the fairness indices. More formally, it is desirable to

$$\text{Maximize } \sum_{n=1}^{N} \sum_{i=1}^{\Lambda} (\pi^u_{n,i} + \pi^m_{n,i}) \text{ and Minimize } \delta \tag{9}$$

subject to

$$|\pi^\theta_{l,i} - \pi^\phi_{n,j}| \le \delta \text{ for } 1 \le l, n \le N, \ 1 \le i, j \le \Lambda, \ \theta, \phi = u, m. \tag{10}$$

For the homogeneous network with balanced uniform traffic considered in this study, the delay and transmitter throughput for each buffer are related by the same inversely proportional relationship. We can thus limit our fairness index to the transmitter and receiver throughput, i.e., we consider

$$\pi^\theta_{n,i} = w \tau^\theta_{n,i} + (1 - w) \rho^\theta_{n,i}, \quad \theta = u, m \tag{11}$$

with $0 \le w \le 1$.

In our performance evaluations we consider the difference (unfairness) metrics defined as follows.

$$|\pi^u_{l,i} - \pi^u_{n,j}| < \delta^u \qquad 1 \le l, n \le N \qquad 1 \le i, j \le \Lambda \tag{12}$$

$$|\pi^m_{l,i} - \pi^m_{n,j}| < \delta^m \qquad 1 \le l, n \le N \qquad 1 \le i, j \le \Lambda \tag{13}$$

$$|\pi^u_{l,i} - \pi^m_{n,j}| < \delta^{u/m} \qquad 1 \le l, n \le N \qquad 1 \le i, j \le \Lambda \tag{14}$$

$$\max\{\delta^u, \ \delta^m, \ \delta^{u/m}\} < \delta. \tag{15}$$

Our main focus is on the difference metric $\delta^{u/m}$, which indicates how fairly unicast traffic is treated with respect to multicast traffic and vice versa. The smaller $\delta^{u/m}$, the fairer the relative treatment of unicast and multicast traffic.

Note that conventional fairness mechanisms designed for unicast traffic ideally provide each node n fair access to packet transmission slots on each wavelength i, i.e., the mechanisms ensure that all the $(\tau^u_{n,i} + \tau^m_{n,i})$ are (approximately) equal.

We briefly outline some implications of the definition of our difference metrics. Note that our fairness definition enforces

$$\pi^m_{l,i} \approx \pi^u_{n,j}, \tag{16}$$

or equivalently,

$$w \frac{T^m_{l,i}}{\sigma^c_i} + (1-w)\frac{R^m_{l,i}}{\sigma^c_i} \approx w \frac{T^u_{n,j}}{\sigma^u_j} + (1-w)\frac{R^u_{n,j}}{\sigma^u_j} \tag{17}$$

Setting the weight to $w = 1$ forces the effective transmitter throughput levels to be (approximately) equal, which implies that the absolute transmitter throughput levels satisfy

$$\frac{T^m_{l,i}}{T^u_{n,j}} \approx \frac{p_m \cdot E[\Delta]}{1 - p_m}. \tag{18}$$

Noting that each multicast packet copy reaches on average $E[f]$ destinations and that $E[\Delta] \cdot E[f] = N/2$, the transmitter throughput ratio in (18) implies for the receiver throughput

$$\frac{R^m_{l,i}}{R^u_{n,j}} \approx \frac{p_m \cdot N/2}{1 - p_m}. \tag{19}$$

Intuitively, setting $w = 1$ enforces that each unicast packet and multicast packet *copy* experiences the same opportunities for transmission in a slot on a wavelength and hence the same delay.

On the other hand, setting the weight to $w = 0$ forces the effective receiver throughput levels to be (approximately) equal, which in turn implies for the ratio of the absolute receiver throughput levels

$$\frac{R^m_{l,i}}{R^u_{n,j}} \approx \frac{p_m \cdot E[\Delta]}{1 - p_m}. \tag{20}$$

This in turn implies for the transmitter throughput

$$\frac{T^m_{l,i}}{T^u_{n,j}} \approx \frac{p_m \cdot E[\Delta]/E[f]}{1 - p_m}. \tag{21}$$

Intuitively, setting $w = 0$ favors the unicast traffic, which is desirable in scenarios where relatively important and delay sensitive data is sent via unicast and bulk data distribution (e.g., software distribution, off-line pushing of web and multimedia content into proxy servers) is carried out via multicast. Note that when the network is lightly loaded and there is negligible loss, both unicast traffic and multicast traffic are served at the rates at which they are generated. When the load and losses increase the fairness policy should ideally serve all unicast packets and drop only multicast packet copies until the throughput ratios (20) and (21) are met.

In our performance studies we consider both the difference metrics and the more intuitive ratio of the absolute throughput levels.

5. Simulation Results for Single-Step Buffer Selection

In this section we examine the throughput-delay as well as fairness performance when a single-step longest queue buffer selection is employed. Throughout we consider a network with $N = 64$ nodes and $\Lambda = 8$ wavelength channels, which results in a mean number of $E[\Delta] = 7.22$ copies for each multicast packet. We set the probability of a packet being a multicast packet to $p_m = 0.3$, which implies that the network is stable for $\sigma < 0.0497$. Following [35] we set the window size for fairness control to $W = 500$. We first consider the $((N-1)+\Lambda)$ VOQ architecture as a representative of an architecture with separate buffers for unicast and multicast packets. Subsequently we examine the Λ VOQ architecture in which unicast and multicast packets are buffered together. These two architectures are illustrated in Fig 1. Initially we employ the following counting policy to determine the longest queue. We count each unicast packet as one and each multicast packet as its fanout, i.e., the number of intended destination nodes. This is motivated by the fact that the transmission of a multicast packet copy contributes to the receiver throughput according to its fanout.

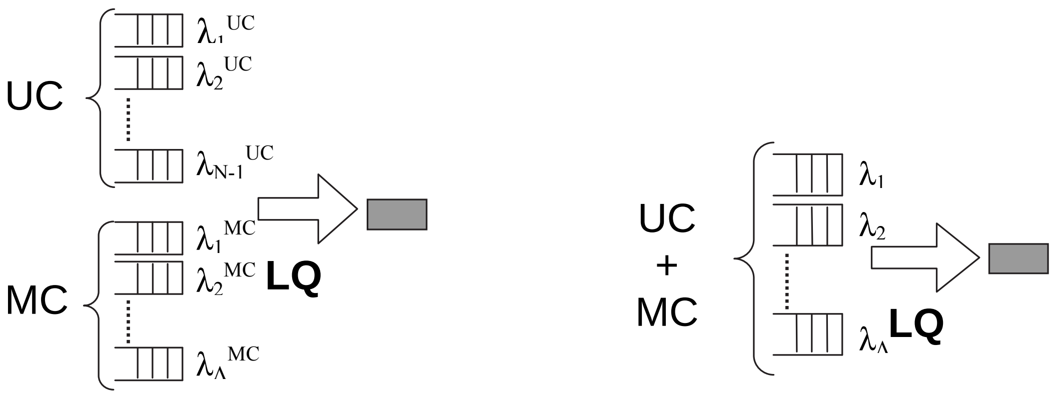

(a) $((N-1)+\Lambda)$VOQ architecture.

(b) ΛVOQ architecture.

Fig. 1. Illustration of single-step longest queue buffer selection.

5.A. *Separate Buffers for Unicast and Multicast Traffic:* $((N-1)+\Lambda)$*VOQ Architecture*

As described in Section 2.A, with the $((N-1)+\Lambda)$VOQ architecture, each node is equipped with Λ buffers for multicast traffic and $(N-1)$ buffers for unicast traffic. Initially we set the capacity of each multicast buffer B^m and the capacity of each unicast buffer B^u to the same value $B = 64$. In Fig. 2 we plot the performance metrics as a function of the packet generation probability (mean packet arrival rate) σ.

We observe from the plots a number fairness issues. First, we observe from Fig. 2(b) that unicast packets experience significantly larger delays than multicast packets. This is primarily caused by the LQ buffer selection giving priority to the multicast buffer, which have typically longer queue counts due to the counting of the fanout of each multicast packet copy. On the other hand, we observe from Fig. 2(c) that multicast packet experience significantly larger loss probabilities for arrival rates up to about 0.4. This is primarily due to the multiple packet copies generated for each multicast packet and the fact that there are only Λ multicast buffers whereas there are $N-1$ unicast buffers in each node. These effects can be mitigated by setting the unicast and multicast buffer capacities such that the available buffer capacities are proportional to the packet (copy) generation probabilities as detailed shortly.

We observe from Figs. 2(a), (d), and (e) that for large arrival rates the network tends to transport only multicast packets. This is primarily due to the LQ buffer selection giving priority to transmissions from the multicast buffers. For a closer inspection of the throughput ratio we plot in Fig. 2(f) the throughput ratio for a range of smaller arrival rates. We observe that the throughput ratio initially decreases and then increases. This is primarily due to the significantly lower packet loss probability for unicast packets up to an arrival rate of approximately 0.1.

Next, we examine the performance when the buffer capacities are set proportional to the packet generation probabilities. Specifically, we use a base buffer capacity of $B = 92$ packets and set $B^u = \lfloor (1-p_m) \cdot B \rfloor = 64$ and $B^m = \lfloor p_m \cdot E[\Delta] \cdot B \cdot (N-1)/\Lambda \rfloor = 1569$. We observe from Fig. 2(g) that this buffer setting exacerbates the unfairness in terms of the throughput ratio. This is due to the LQ policy now giving even more priority to the multicast buffers since they can now grow significantly longer than the unicast buffers. We also observe from Fig. 2(h) that the packet loss behavior is now reversed compared to Fig. 2(c), which is again due to the relatively larger buffer length counts achieved with the large multicast buffers.

In Fig. 3 we have also examined the performance when counting each unicast packet and multicast packet copy as one in the separate unicast and multicast buffers. We have found that when the unicast and multicast buffers have the same capacity, then there is unfairness in terms of the delay, packet loss, and throughput ratio similar to the situation depicted in Figs. 2(b), (c), and (f). On the other hand, when the buffer capacities are set in proportion to the traffic generation rates of unicast packets and multicast packet copies, then there is unfairness in terms of the throughput ratio similar to the situation depicted in Fig. 2(g).

5.B. Joint Buffers for Unicast and Multicast Traffic: ΛVOQ architecture

In this section we examine the behavior of the ring for the ΛVOQ buffer architecture. As described in Sect. 2.A, each node is equipped with Λ buffers, one for each wavelength channel. A given buffer is used both for multicast and unicast traffic; i.e., unicast packets and packet copies (generated from a multicast packet) are stored in the same queue, according to the home channel of the destination node(s). To make the overall buffer capacity at each node the same as in the Section 5.A, we set the capacity of each buffer to $B = 64(1 + \frac{N-1}{\Lambda})$.

In Figs. 4 we plot the performance metrics as a function of the mean arrival rate σ. We observe from Figs. 4 (b) and (c) that the unicast and multicast packets are treated fairly in terms of delay and packet loss. This is intuitive as unicast and multicast packets are not distinguished in the ΛVOQ architecture. We observe from Fig. 4(a) that the unicast and multicast transmitter throughputs satisfy $T^m/T^u = p_m \cdot E[\Delta]/(1 - p_m) = 3.1$ with reasonable accuracy and thus meet our $w = 1$ fairness criterion. We furthermore observe from Figs. 4(d) and (e) that the receiver throughputs satisfy $R^m/R^u = p_m \cdot N/[2 \cdot (1 - p_m)] = 13.7$ with relatively good accuracy, which corresponds to $w = 1$ fairness.

5.C. Conclusions from Single-Step Buffer Selection Simulations

In this section we have presented the performance results obtained for unicast and multicast traffic when employing the LQ buffer selection scheme in different buffer architectures. We found that the buffering of unicast and multicast packets in different queues in conjunction with the single-step longest queue buffer selection generally results in significant unfairness and fluctuating relative performance for unicast and multicast traffic. We obtained fair treatment of unicast vs. multicast traffic only for the ΛVOQ architecture, where each of the Λ buffers stores both the unicast and multicast traffic for a given wavelength. More specifically, the ΛVOQ architecture in conjunction with LQ buffer selection gives the $w = 1$ fairness, where unicast packets and multicast packet copies are provided with the same transmission opportunities and experience the same delay.

As noted in Section 3, for the common traffic scenario with delay sensitive unicast traffic and delay insensitive multicast traffic it is desirable to provide more transmission opportunities to unicast traffic. Also, generally it is desirable to control the relative priority of multicast vs. unicast traffic over a range of priorities. Toward this goal we introduce and examine dual buffer selection policies for buffer architectures with separate unicast and multicast buffers next.

6. Dual-Step Buffer Selection

In this section we develop and evaluate dual step buffer selection policies that are designed to allow for the consistently fair treatment of unicast and multicast traffic with different preferences for achieving a range of ratios of multicast to unicast transmitter or receiver throughput.

Our dual-step buffer selection policies employ the *longest queue (LQ)* buffer selection to decide on which wavelength *channel*, among the channels with an empty slot, to transmit a packet. We refer to this selection step as *LQ-C selection*. In addition, we employ a buffer selection scheme that in each time slot determines whether to transmit a unicast packet or a multicast packet. We refer to this step as *UC/MC selection*.

We now proceed to explain these selection steps in detail. Throughout this discussion we focus on the $\Lambda + \Lambda$ VOQ architecture; the selection steps can be adapted to other VOQ architectures with separate queues for unicast and multicast traffic in a straightforward fashion. In Fig. 5 we illustrate the dual step buffer selection.

As illustrated in the figure there are two possible permutations for executing the buffer selection steps. We refer to the sequence LQ-C selection followed by UC/MC selection as *channel priority* policy, as it gives priority to selecting the wavelength channel for transmission. On the other hand, we refer to the sequence UC/MC selection followed by LQ-C selection as *traffic priority* policy, as it gives priority to selecting the type of traffic (unicast or multicast) to be transmitted.

For the formal definition of the LQ-C selection we initially focus on the channel priority policy. We let β_i^u, $i = 1, \ldots, \Lambda$, denote the occupancy of unicast buffer i. Note that this occupancy is equal to the number of packets stored in this buffer. Similarly, we let β_i^m denote the occupancy of multicast buffer i storing multicast packet copies destined to nodes that receive on channel i. This occupancy is weighted by the fanout of the multicast packet copies stored in the buffer, i.e., it is equal to the sum of the fanouts of the stored packets. We let $\beta_i = \beta_i^u + \beta_i^m$, $i = 1 \ldots \Lambda$, denote the total buffer occupancy (backlog) for wavelength channel i. Among the channels with a free slot, the LQ-C selection chooses the channel with the largest backlog for transmission. Unless either the unicast buffer or the multicast buffer for the chosen channel is empty, the UC/MC selection (as detailed shortly) is then invoked to determine whether to transmit a unicast packet or a multicast packet copy.

With the traffic priority policy, the LQ-C selection is done analogously considering either only the unicast buffer occupancies β_i^u or the multicast buffer occupancies β_i^m depending on the type of traffic selected by the UC/MC selection.

In each time slot, the UC/MC selection decides whether to transmit a unicast packet or a multicast packet. We transmit a unicast packet with probability P^u and a multicast packet copy with probability P^m. We calculate these selection probabilities as a function of the input traffic characteristics (the probability σ_i^u of generating a unicast packet for a buffer and the probability σ_i^c of generating a multicast packet copy for a buffer) and of the weight w in the fairness metrics as follows. We note that the transmission of a unicast packet and a multicast packet copy contribute the same toward the transmitter throughput. On the other hand, the transmission of a multicast packet copy contributes to the receiver throughput on average $E[f]$ times more than the transmission of a unicast packet. Thus, $\tau_{n,i}^u$ and $\rho_{n,i}^u$ are proportional to P^u, $\tau_{n,i}^m$ is proportional to P^m, and $\rho_{n,i}^m$ is proportional to $E[f] \cdot P^m$. Hence we can rewrite (17) in terms of the transmission probabilities to obtain

$$w \frac{P^m}{\sigma_i^c} + (1-w) \frac{E[f] \cdot P^m}{\sigma_i^c} \approx w \frac{P^u}{\sigma_i^u} + (1-w) \frac{P^u}{\sigma_i^u}. \tag{22}$$

From (22) in conjunction with $P^u + P^m = 1$ we obtain

$$P^u = \frac{1}{1 + \frac{p_m \cdot E[\Delta]}{1-p_m} \frac{1}{w+(1-w) \cdot E[f]}} \tag{23}$$

$$P^m = \frac{1}{1 + \frac{1-p_m}{p_m \cdot E[\Delta]} (w + (1-w) \cdot E[f])}. \tag{24}$$

Table 1. Transmission probabilities for a network with $N = 64$ nodes, $\Lambda = 8$ wavelength channels, and a fraction of $p_m = 0.3$ of multicast traffic.

w	P^u	P^m
0	0.59	0.41
0.5	0.42	0.58
1	0.25	0.75

The unicast and multicast transmission probabilities for different values of the weight w for the considered network are given in Table 1.

In summary, the channel priority policy first applies LQ-C selection to determine among the wavelength channels with an empty slot the channel i with the largest cumulative buffer occupancy $\beta_i^u + \beta_i^m$. If only one of the buffers for the selected channel i (either the unicast buffer or the multicast buffer) is non-empty, the node transmits one packet from the non-empty buffer. If both buffers are non-empty, the node employs the UC/MC selection and transmits a unicast packet with probability P^u, with the complementary probability P^m the node transmits a multicast packet.

The traffic priority policy first applies the UC/MC selection, i.e., with probability P^u the unicast buffers are considered in the subsequent LQ-C selection; with the complementary probability P^m the multicast buffers are considered in the LQ-C selection. If all unicast buffers (multicast buffers) are empty, then the UC/MC selection is not invoked and only the multicast buffers (unicast buffers) are considered in the LQ-C selection step. The LQ-C selection chooses the buffer with the largest occupancy β_i^u or β_i^m (according to the choice made in the UC/MC selection step) for transmission.

The motivation for the described dual-step buffer selection is as follows. The LQ-C selection achieves load balancing among the queues in the system, i.e., it strives to keep δ^u (12) and δ^m (13) small and at the same time strives to maximize the transmitter/receiver throughputs. The UC/MC buffer selection, on the other hand, strives to keep $\delta^{u/m}$ (14) small, i.e., to enforce fairness between unicast and multicast traffic.

6.A. *Simulation Results*

In this section we present simulation results for our dual-step buffer selection policies. We initially focus on the channel priority policy employed in the $(\Lambda + \Lambda)$VOQ buffer architecture. To avoid unfairness among unicast and multicast packets in term of packet loss we set the unicast and multicast buffer capacities in proportion to the unicast packet and multicast packet copy generation probabilities. In particular, we set $B^u = (1 - p_m)B$ and $B^m = p_m \cdot B \cdot E[\Delta]$ with $B = 460$ which implies $B^u = 322$ and $B^m = 996$.

In Fig. 6 we plot the performance metrics of the $(\Lambda + \Lambda)$VOQ architecture with channel priority policy for $w = 1$. We observe that the performance metrics in these plots are roughly equivalent to the plots in Fig. 4, which confirms that the channel priority policy employed in the $(\Lambda + \Lambda)$VOQ architecture is able to achieve the same $w = 1$ fairness as the ΛVOQ architecture in conjunction with LQ buffer selection. In particular, the ratio of multicast to unicast transmitter throughput T^m/T^u is close to $p_m \cdot E[\Delta]/(1-p_m) = 3.1$ as required for $w = 1$ fairness, see (18). Similarly, the ratio of multicast to unicast receiver throughput R^m/R^u is close to $(p_m \cdot N/2)/(1-p_m) = 13.7$, as is required for $w = 1$ fairness, see (19).

We also observe that the values of the difference metric δ, which gives the largest difference between the relative transmitter throughput for any two different unicast or multicast buffers for the considered case of $w = 1$, are relatively large for small mean arrival rates. These relatively large δ are primarily due to the high burstiness of the considered self-similar traffic. We have found that for Bernoulli traffic, the δ is always below 0.11 (see Fig. 7(f)).

Next, we examine the performance of the channel priority policy for $w = 0$, for which we plot the performance metrics in Fig. 8.

A number of important observations are in order. First, we observe from Fig. 8(a), (d), and (e) that for mean arrival rates larger than 0.2, which correspond to heavy overload, the throughput ratios are very close to the ratios required for $w = 0$ fairness. In particular, the ratio of the multicast to unicast transmitter throughput T^m/T^u is close to $(p_m \cdot E[\Delta]/E[f])/(1-p_m) = 0.72$, see (21). At the same time, the ratio of the multicast to unicast receiver throughput R^m/R^u is close to $p_m \cdot E[\Delta]/(1-p_m) = 3.1$, see (20). This indicates that the channel priority policy is effective in enforcing $w = 0$ fairness with consistent throughput ratios for multicast vs. unicast traffic under heavy overload conditions. We also observe from Fig. 8(b) that the unicast traffic experiences consistently significantly smaller delays than the multicast traffic. This indicates that with the $w = 0$ setting, the channel priority policy is quite effective in ensuring small delays for time-sensitive unicast traffic while bulk multicast traffic experiences larger delays.

Next we observe from Fig. 8(e) that for very small traffic loads, i.e., a mean arrival rate close to zero, the channel priority policy does not achieve the desired multicast to unicast receiver throughput ratio R^m/R^u of 3.1. Instead, the throughput ratio R^m/R^u is close to the ratio of 13.7, which was desired for the $w = 1$ setting. This is because for very small arrival rates there is not enough unicast traffic to take advantage of the transmission opportunities offered to the unicast traffic with the $w = 0$ setting, i.e., the UC/MC selection is frequently not invoked as either only the unicast buffer or the multicast buffer for the channel selected by the LQ-C selection holds a packet. Also, observe from Fig. 8(c) that essentially no packet is lost, i.e., essentially all the generated traffic is served. Hence, the ratio of multicast to unicast transmitter throughput is essentially equal to the ratio of the generation rates of multicast packet copies and unicast packets, which is $p_m \cdot E[\Delta]/(1-p_m)$. Noting that each multicast packet copy reaches on average $E[f]$ destination nodes, the corresponding ratio of the receiver throughputs is $p_m \cdot E[\Delta] \cdot E[f]/(1-p_m) = (p_m \cdot N/2)/(1-p_m)$, which corresponds to the receiver throughput ratio required for $w = 1$ fairness and is 13.7 with our parameter settings.

As the mean arrival rate increases from close to zero up to a level of about 0.2, the increasing amount of unicast traffic can take advantage of the transmission opportunities offered by the channel priority policy with $w = 0$, and the throughput ratios approach the values desired with the $w =$

0 setting. In particular, we observe from Fig. 8(c) that the loss probability for unicast traffic is significantly smaller than the loss probability for multicast traffic as the unicast traffic is given priority to achieve the desired ratio of $(1-p_m)/(p_m \cdot E[\Delta]/E[f]) \approx 1.4$ unicast packet transmissions for each transmission of a multicast packet copy. We observe that the channel priority policy does allow for some minuscule loss of unicast packets in the range from $\sigma = 0.025 - 0.1$, even though the desired ratio of unicast packet to multicast packet copy transmissions is not yet reached. Ideally, the fairness mechanism should not drop any unicast packet as long as the desired ratio of unicast packet to multicast packet copy transmissions is not yet reached. The unicast packet losses that do occur are primarily due to the highly bursty nature of the considered self-similar traffic. In more extensive simulations (see Fig. 9) we have found that with Bernoulli traffic no unicast packet is dropped up to an arrival rate of $\sigma = 0.1$.

The results for the difference metric δ in Fig. 8(f) also reflect that it is impossible to achieve the throughput ratios desired for $w = 0$ when the arrival rates are small. Indeed, we observe that $\delta \approx E[f] \approx 4.4$, which reflects the fact that for low traffic loads all generated packets are served and consequently the relative multicast receiver throughput is $E[f]$ times the relative unicast receiver throughput.

Setting w to intermediate values achieves throughput ratios as well as delay and packet loss behaviors that lie between the extremes observed above for $w = 0$ and $w = 1$. (see Fig. 10).

We have also examined the traffic priority policy and found that it gives essentially the same results as the channel priority policy (see Figs. 11, 12, 13, and 14). Furthermore, we have developed and evaluated a traffic priority policy with memory, which keeps a memory of missed transmission opportunities for unicast or multicast traffic due to instances where all unicast or multicast buffers were empty. We have found that the traffic priority policy with memory gives essentially the same performance as the policies without memory (see Figs. 15, 16, 17, and 18).

7. Conclusion

We have examined the fairness issues arising when transmitting a mix of unicast and multicast traffic in a packet-switched WDM ring network. We have considered a single-fiber ring network with TT–FR node structure, *a posteriori* buffer selection, and destination stripping for uniform balanced traffic. We have found that VOQ architectures with separate queues for unicast and multicast traffic give generally rise to unfairness when the conventional single-step longest queue buffer selection is employed. We have also found that a VOQ architecture with one queue buffering both unicast and multicast traffic for a given wavelength channel achieves fairness in the sense that the relative transmitter throughput levels achieved by unicast and multicast traffic are approximately equal, or equivalently the ratio of the transmitter throughput for unicast traffic to the transmitter throughput for multicast traffic is approximately equal to the ratio of generated unicast packets to generated multicast packet copies.

To allow for a range of throughput ratios, we have developed and evaluated dual-step buffer selection policies which combine longest queue selection for determining the wavelength channel to transmit on with UC/MC buffer selection for determining whether to transmit a unicast or multicast packet. We have demonstrated that the dual-step buffer selection policies achieve a range of throughput ratios. The policies cover the range from being fair in the sense that unicast and multicast traffic experience approximately the same effective transmitter throughput to being fair in the sense that unicast and multicast traffic experience approximately the same effective receiver throughput.

There are several exciting and important avenues for future work. One important direction is to consider the transmission of a mix of unicast and multicast traffic in conjunction with non-uniform unbalanced traffic patterns. Another interesting direction for future work is to examine the fairness issues arising when transmitting a mix of unicast and multicast traffic over different ring network architectures, such as networks with different node structures, dual-fiber ring networks, meshed ring networks [42, 43], or DWADM virtual circle networks [44].

Acknowledgement

We are grateful to Prof. Andrea Bianco of the Politecnico di Torino for his careful review of an earlier version of this paper. We are grateful to Hyo-Sik Yang of Arizona State University for assisting in building the simulation models used in this work. This work was conducted while Chiara Piglione was visiting Arizona State University in 2003 and 2004, the support of ASU for her visit is gratefully acknowledged.

References and Links

[1] K. Bengi and H. R. van As, “Efficient QoS Support in a Slotted Multihop WDM Metro Ring,” IEEE Journal on Selected Areas in Communications 20, no. 1, 216–227 (2002).

[2] J. Cai, A. Fumagalli, and I. Chlamtac, “The Multitoken Interarrival Time (MTIT) Access Protocol for Supporting Variable Size Packets Over WDM Ring Network,” IEEE Journal on Selected Areas in Communications 18, no. 10, 2094–2104 (2000).

[3] Y. Cai, R. M. Fortenberry, and R. S. Tucker, “Demonstration of Photonic Packet–Switched Ring Network with Optically Transparent Nodes,” IEEE Photonics Technology Letters 6, no. 9, 1139–1141 (1994).

[4] I. Cidon and Y. Ofek, “MetaRing – A Full–Duplex Ring with Fairness and Spatial Reuse,” IEEE Transactions on Communications 41, no. 1, 110–120 (1993).

[5] J. Fransson, M. Johansson, M. Roughan, L. Andrew, and M. A. Summerfield, “Design of a Medium Access Control Protocol for a WDMA/TDMA Photonic Ring Network,” in *Proceedings of IEEE Globecom* (Institute of Electrical and Electronics Engineers, 1998), pp. 307–312.

[6] D. Guo and A. S. Acampora, “Scalable Multihop WDM Passive Ring with Optimal Wavelength Assignment and Adaptive Wavelength Routing,” IEEE/OSA Journal of Lightwave Technology 14, no. 6, 1264–1277 (1996).

[7] C. S. Jelger and J. M. H. Elmirghani, “Photonic Packet WDM Ring Networks Architecture and Performance,” IEEE Communications Magazine 40, 110–115 (2002).

[8] A. Smiljanić, M. Boroditsky, and N. J. Frigo, “High–Capacity Packet–Switched Optical Ring Network,” IEEE Communications Letters 6, no. 3, 111–113 (2002).

[9] I. M. White, M. S. Rogge, K. V. Shrikhande, and L. G. Kazovsky, “A summary of the HORNET project: A next-generation metropolitan area network,” IEEE Journal on Selected Areas in Communications 21, no. 9, 1478–1494, (2003).

[10] J. He, S.-H. G. Chan, and D. H. Tsang, “Multicasting in WDM networks,” IEEE Communications Surveys and Tutorials 4, no. 1 (2002).

[11] M. Boroditsky, C. F. Lam, S. L. Woodward, N. N. J. Frigo, and M. M. D. Feuer, “Power management for enhanced system performance of passive optical rings,” IEE Proceedings Optoelectronics 150, no. 3, 229–234 (2003).

[12] S. Aleksic and K. Bengi, “Multicast-capable access nodes for slotted photonic ring networks,” in *Proceedings of European Conference on Optical Communications (ECOC)* (2000), pp. 83–84.

[13] A. Carena, V. D. Feo, J. Finochietto, R. Gaudino, F. Neri, C. Piglione, and P. Poggiolini, “RingO: a Demonstrator of WDM Optical Packet Network on a Ring Topology,” IEEE Journal on Selected Areas in Communications (to be published).

[14] E. Shimada, S. Fujiwara, K. Okazaki, and I. Sasase, “A transmission system for multicast traffic with preallocation scheme in WDM ring networks,” in *Proceedings of IEEE Pacific Rim Conference on Communications, Computers and Signal Processing* (Institute of Electrical and Electronics Engineers, 2003), pp. 490–493.

[15] D.-R. Din, “Genetic algorithms for multiple multicast on WDM ring network,” Computer Communications 27, no. 9, 840–856 (2004).

[16] X. Jia, X.-D. Hu, L. Ruan, and J. Sun, “Multicast routing, load balancing, and wavelength assignment on tree of rings,” IEEE Communications Letters 6, no. 2, 79–81 (2002).

[17] R. Libeskind-Hadas, J. R. Hartline, K. Dresner, E. Drucker, and K. Ray, “Multicast virtual topologies in WDM paths and rings with splitting loss,” in *Proceedings of Eleventh International Conference on Computer Communications and Networks* (2002), pp. 318–321.

[18] J. Wang, B. Chen, and R. Uma, “Dynamic wavelength assignment for multicast in all-optical WDM networks to maximize the network capacity,” IEEE Journal on Selected Areas in Communications 21, no. 8, 1274–1284 (2003).

[19] C. Zhou and Y. Yang, “Wide–sense Nonblocking Multicast in a Class of Regular Optical WDM Networks,” IEEE Transaction on Communications 50, no. 1, 126–134 (2002).

[20] J. Chen, I. Cidon, and Y. Ofek, “A Local Fairness Algorithm for the MetaRing and its performance study,” IEEE Journal on Selected Areas in Communications 11, 1183–1192 (1993).

[21] I. Cidon, L. Georgiadis, R. Guerin, and Y. Shavitt, “Improved fairness algorithms for rings with spatial reuse,” IEEE/ACM Transactions on Networking 5, no. 2, 190–204 (1997).

[22] S. Kasemlonnapa, J. Meditch, and B. Mukherjee, “A delay-throughput performance analysis of the pi-persistent protocol for unidirectional broadcast bus networks,” in *Proceedings of IEEE Infocom* (Institute of Electrical and Electronics Engineers, Ottawa, Canada, 1989) , pp. 834–840.

[23] M. Ajmone Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri, “MAC Protocols and Fairness Control in WDM Multirings with Tunable Transmitters and Fixed Receivers,” IEEE Journal of Lightwave Technology 14, no. 6, 1230–1244 (1996).

[24] N. Sarkar and K. Pawlikowski, “A delay-throughput performance improvement to the pi-persistent protocol,” in *Proceedings of IEEE Int. Symposium on Computers and Communications* (Institute of Electrical and Electronics Engineers, 2001), pp. 615–620.

[25] I. M. White, M. S. Rogge, K. Shrikhande, and L. G. Kazovsky, "Design of a Control-Based Media-Access-Control Protocol for HORNET," Journal of Optical Networking 1, no. 12, 460–473 (2002).

[26] D. M. Chiu, "Some observations on fairness of bandwidth sharing," in *Proceedings of IEEE Int. Symposium on Computers and Communications* (Institute of Electrical and Electronics Engineers, 2000), pp. 125–131.

[27] S. Deb and R. Srikant, "Congestion control for fair resource allocation in networks with multicast flows," in *Proceedings of IEEE Conference on Decision and Control* (Institute of Electrical and Electronics Engineers, 2001), pp. 1911–1916.

[28] F. Filali and W. Dabbous, "Fair Bandwidth Sharing Between Unicast and Multicast Flows in Best-Effort Networks," Computer Communications, Special Issue on Quality of Future Internet 27, no. 4, 330–344 (2004).

[29] A. Legout, J. Nonnenmacher, and E. W. Biersack, "Bandwidth-allocation policies for unicast and multicast flows," IEEE/ACM Transactions on Networking 9, no. 4, 464–478 (2001).

[30] B. Prabhakar, N. McKeown, and R. Ahuja, "Multicast Scheduling for Input-Queued Switches," IEEE Journal on Selected Areas in Communications 15, no. 15, 855–866 (1997).

[31] Z. Ortiz, G. N. Rouskas, and H. G. Perros, "Scheduling combined unicast and multicast traffic in broadcast WDM networks," Photonic Network Communication 2, no. 2, 135–153 (2000).

[32] M. Andrews, S. Khanna, and K. Kumaran, "Integrated Scheduling of Unicast and Multicast Traffic in an Input-Queued Switch," in *Proceedings of IEEE Infocom* (Institute of Electrical and Electronics Engineers, New York, NY, 1999), pp. 1144–1151.

[33] S.-H. Lee and D.-H. Cho, "Packet-Scheduling Algorithm Based on Priority of Separate Buffers for Unicast and Multicast Services," Electronics Letters 39, no. 2, 259–260 (2003).

[34] N. McKeown, M. Izzard, A. Mekkittikul, B. Ellersick, and M. Horowitz, "Tiny Tera: a packet switch core," IEEE Micro 17, no. 1, 26–33 (1997).

[35] H.-S. Yang, M. Herzog, M. Maier, and M. Reisslein, "Metro WDM Networks: Performance Comparison of Ring and Star Topologies," IEEE Journal on Selected Areas in Communications (to be published).

[36] I. M. White, E. S.-T. Hu, Y.-L. Hsueh, K. V. Shrikhande, M. S. Rogge, and L. G. Kazovsky, "Demonstration and system analysis of the HORNET," IEEE/OSA Journal of Lightwave Technology 21, no. 11, 2489–2498 (2003).

[37] K. Shrikhande, A. Srivatsa, I. M. White, M. S. Rogge, D. Wonglumsom, S. M. Gemelos, and L. G. Kazovsky, "CSMA/CA MAC Protocols for IP–HORNET: An IP over WDM Metropolitan Area Ring Network," in *Proceedings of IEEE Globecom* (Institute of Electrical and Electronics Engineers, San Francisco, CA, 2000), pp. 1303–1307.

[38] K. V. Shrikhande, I. M. White, D.-R. Wonglumsom, S. M. Gemelos, M. S. Rogge, Y. Fukashiro, M. Avenarius, and L. G. Kazovsky, "HORNET: A Packet–Over–WDM Multiple Access Metropolitan Area Ring Network," IEEE Journal on Selected Areas in Communications 18, no. 10, 2004–2016 (2000).

[39] K. Imai, T. Ito, H. Kasahara, and N. Morita, "ATMR: Asynchronous Transfer Mode Ring Protocol," Computer Networks and ISDN Systems 26, 785–798 (1994).

[40] K. Park and W. Willinger, *Self–Similar Network Traffic and Performance Evaluation* (John Wiley & Sons Inc., 2000).

[41] C. Fan, M. Reisslein, and S. Adams, "The $FT^{\Lambda} - FR^{\Lambda}$ AWG Network: a Practical Single-Hop Metro WDM Network for Efficient Uni- and Multicasting," in *Proceedings of IEEE Infocom*, (Institute of Electrical and Electronics Engineers, Hong Kong, 2004).

[42] I. Rubin and H.-K. Hua, "SMARTNet: An All–Optical Wavelength–Division Meshed–Ring Packet–Switching Network," in *Proceedings of IEEE Globecom* (Institute of Electrical and Electronics Engineers, Singapore,1995), pp. 1756–1760.

[43] I. Rubin and H.-K. H. Hua, "Synthesis and Throughput Behavior of WDM Meshed–Ring Networks Under Nonuniform Traffic Loading," IEEE/OSA Journal of Lightwave Technology 15, no. 8, 1513–1521 (1997).

[44] W. Cho and B. Mukherjee, "Design of MAC Protocols for DWADM–Based Metropolitan–Area Optical Ring Networks," in *Proceedings of IEEE Globecom* (Institute of Electrical and Electronics Engineers, San Antonio, TX, 2001), pp. 1575–1579.

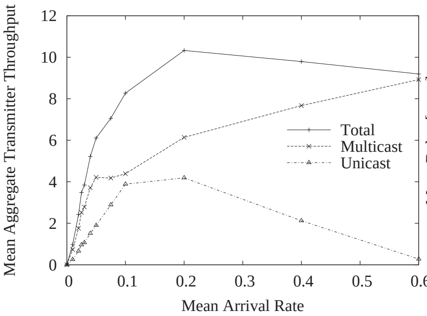

(a) Aggregate transmitter throughput T; $B_U = B_M = 64$.

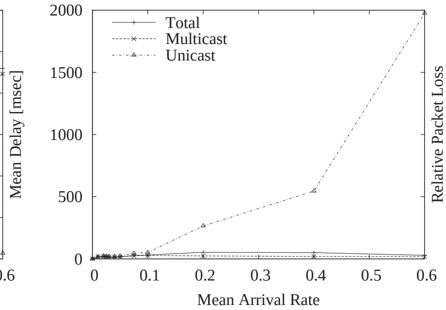

(b) Mean aggregate delay D; $B_U = B_M = 64$.

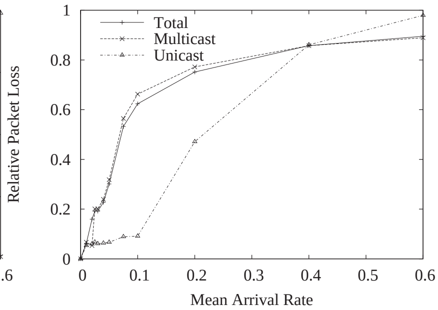

(c) Relative packet loss L; $B_U = B_M = 64$.

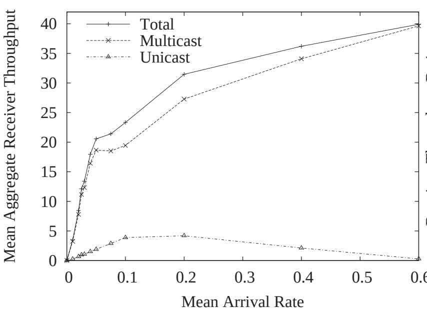

(d) Aggregate receiver throughput T; $B_U = B_M = 64$.

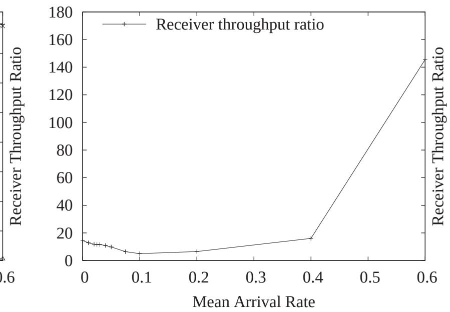

(e) Ratio of multicast to unicast receiver throughput R^m/R^u; $B_U = B_M = 64$; $\sigma \in \{0.0, 0.6\}$.

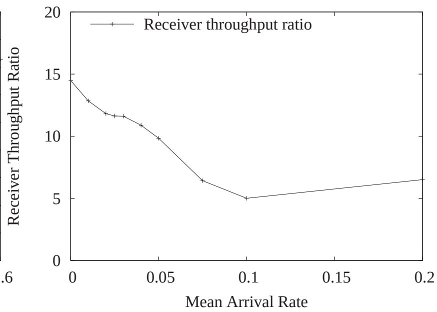

(f) Ratio of multicast to unicast receiver throughput R^m/R^u; $B_U = B_M = 64$; $\sigma \in \{0.0, 0.2\}$.

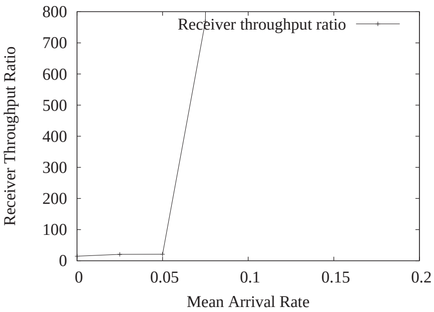

(g) Ratio of multicast to unicast receiver throughput R^m/R^u; $B_U = 64, B_M = 1569$; $\sigma \in \{0.0, 0.2\}$.

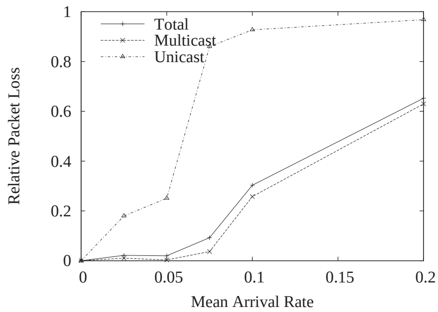

(h) Relative packet loss L; $B_U = 64, B_M = 1569$.

Fig. 2. Performance for $((N-1)+\Lambda)$VOQ architecture with single-step longest queue buffer selection.

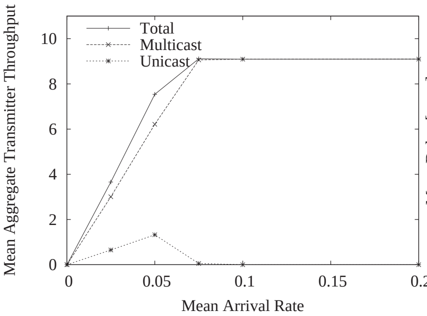

(a) Aggregate transmitter throughput T; $B_U = 64, B_M = 1569$.

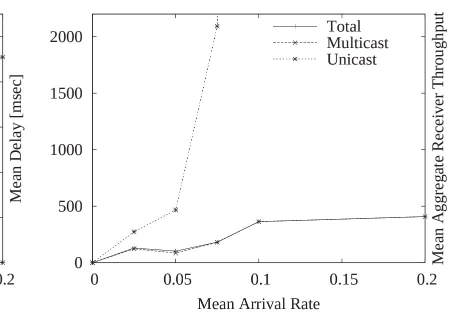

(b) Mean aggregate delay D; $B_U = 64, B_M = 1569$.

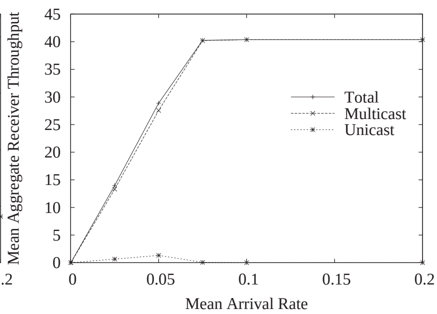

(c) Aggregate receiver throughput T; $B_U = 64, B_M = 1569$.

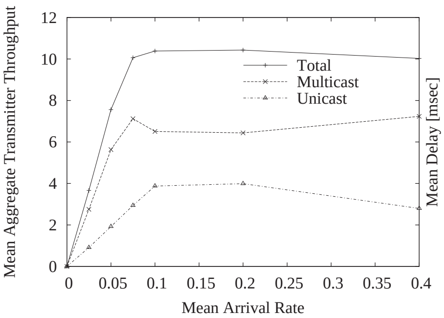

(d) Aggregate transmitter throughput T.

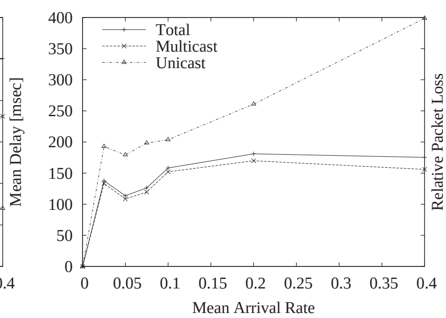

(e) Mean aggregate delay D.

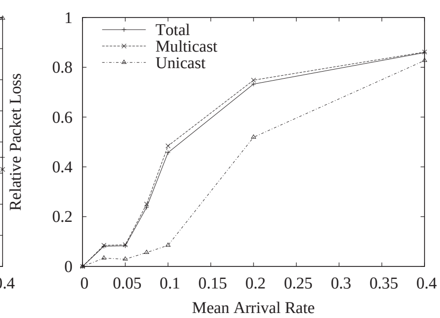

(f) Relative packet loss L.

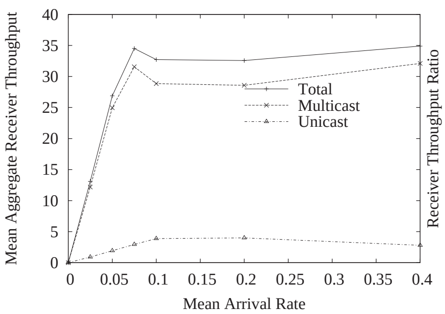

(g) Aggregate receiver throughput T.

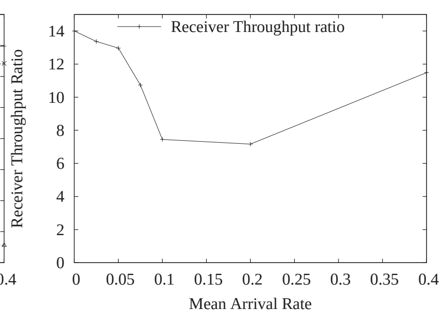

(h) Ratio of multicast to unicast receiver throughput R^m/R^u.

Fig. 3. Performance for $(\Lambda + \Lambda)$VOQ architecture with single-step longest queue buffer selection, when counting each unicast packet and multicast packet copy as one.

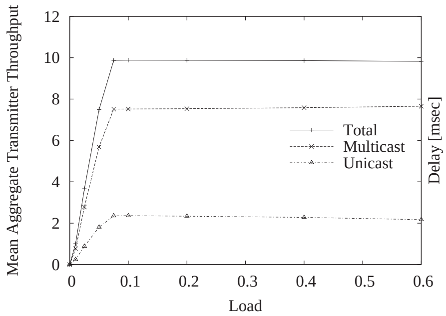

(a) Aggregate transmitter throughput T.

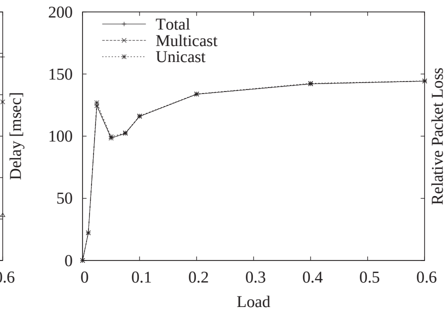

(b) Mean aggregate delay D.

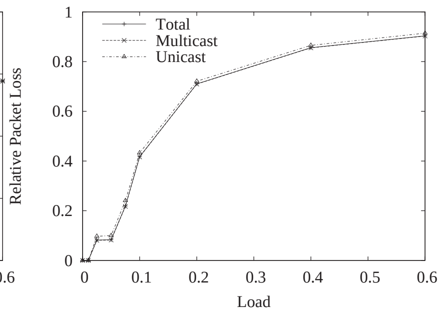

(c) Relative packet loss L.

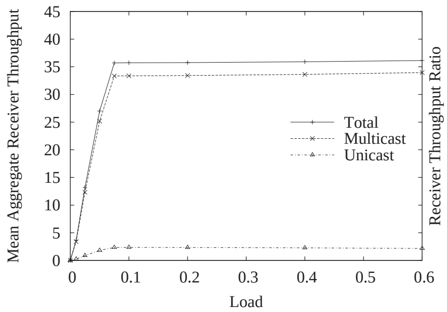

(d) Aggregate receiver throughput T.

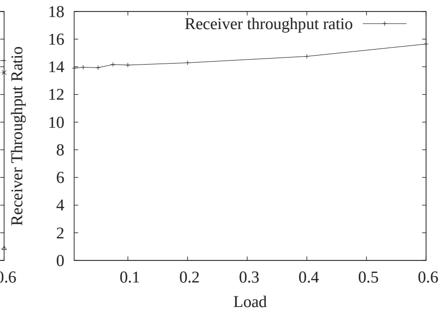

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

Fig. 4. Performance for ΛVOQ architecture with single-step longest queue buffer selection.

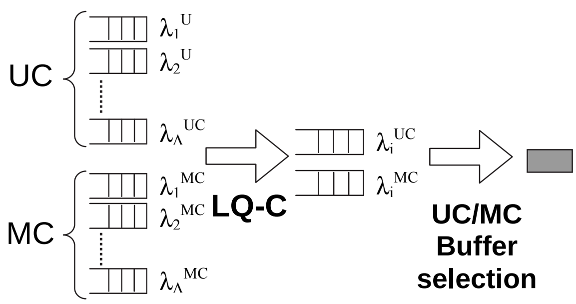

(a) Channel priority policy: LQ-C selection followed by UC/MC selection.

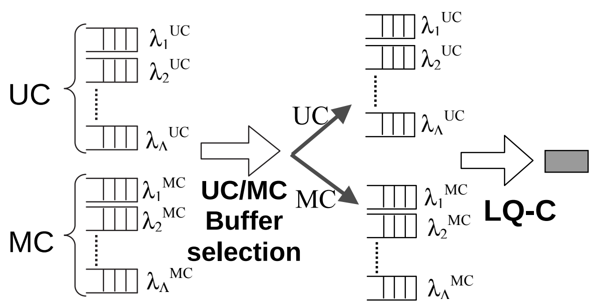

(b) Traffic priority policy: UC/MC selection followed by LQ-C selection.

Fig. 5. Dual-step buffer selection schemes for $(\Lambda + \Lambda)$VOQ architectures.

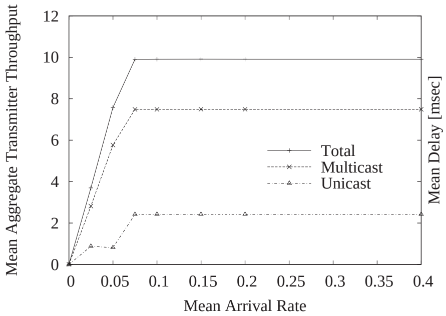

(a) Aggregate transmitter throughput T.

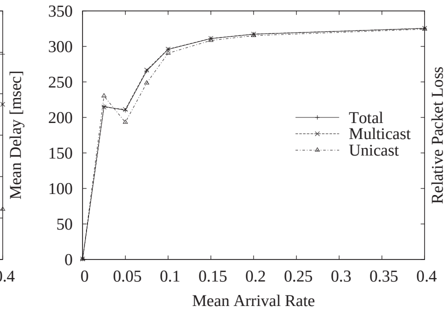

(b) Mean aggregate delay D.

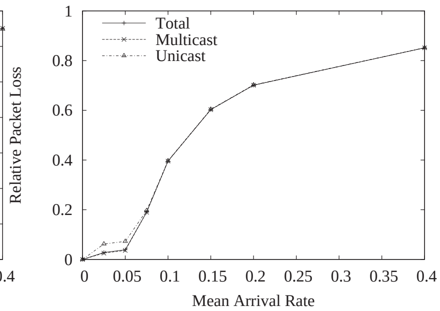

(c) Relative packet loss L.

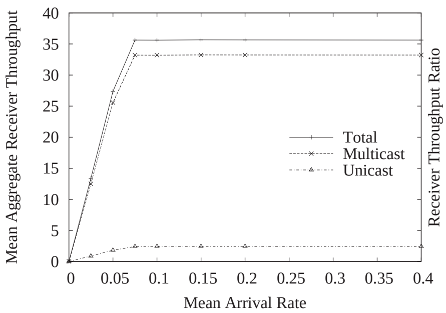

(d) Aggregate receiver throughput T.

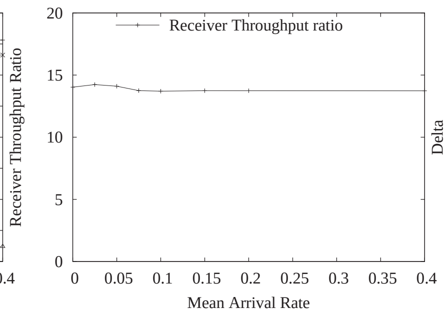

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

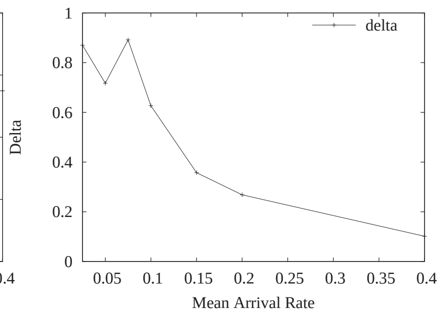

(f) Fairness performance: δ.

Fig. 6. Performance for $(\Lambda + \Lambda)$VOQ architecture with channel priority policy with $w = 1$.

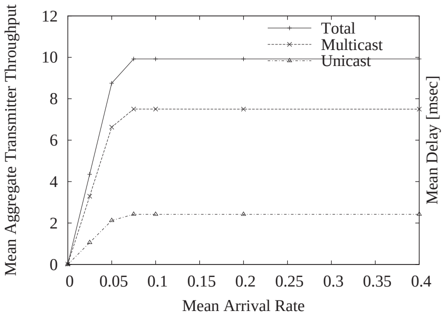

(a) Aggregate transmitter throughput T.

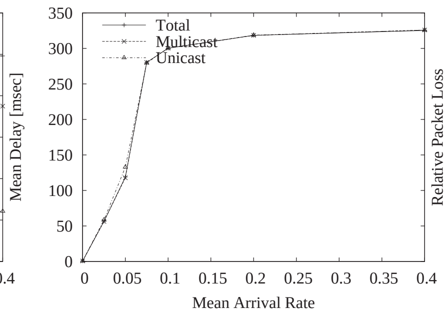

(b) Mean aggregate delay D.

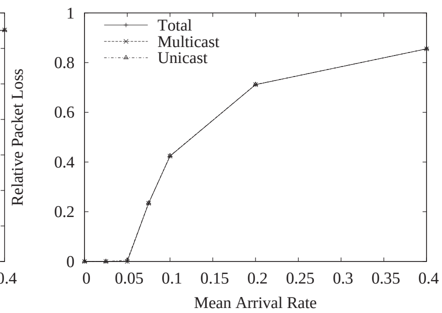

(c) Relative packet loss L.

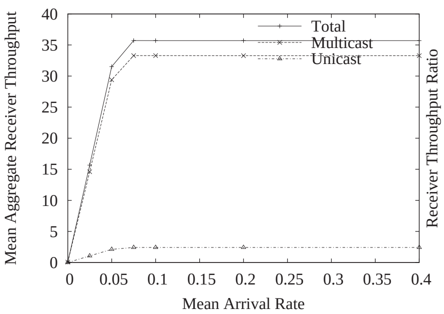

(d) Aggregate receiver throughput T.

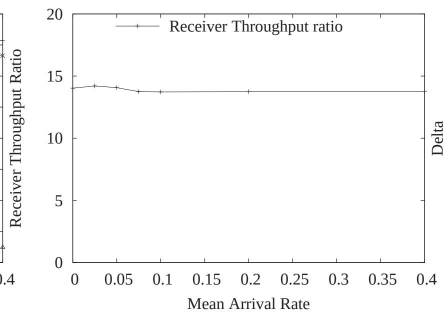

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

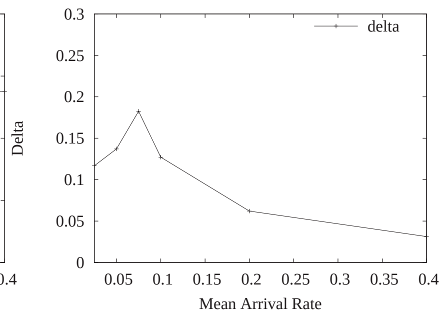

(f) Fairness performance: δ.

Fig. 7. Performance for $(\Lambda + \Lambda)$VOQ architecture with channel priority policy with $w = 1$ with Bernoulli traffic.

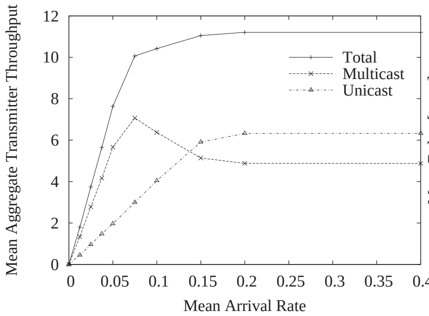

(a) Aggregate transmitter throughput T.

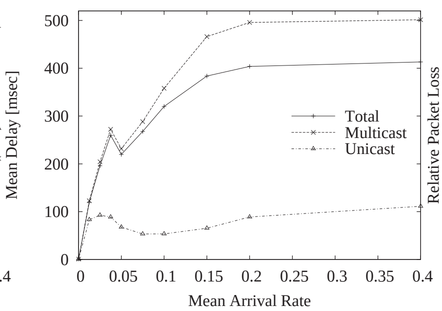

(b) Mean aggregate delay D.

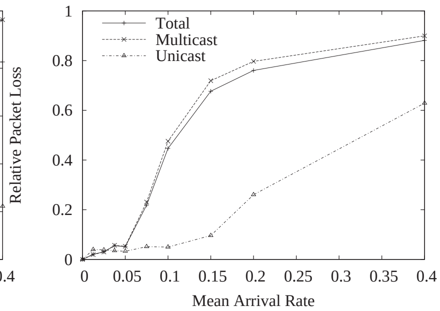

(c) Relative packet loss L.

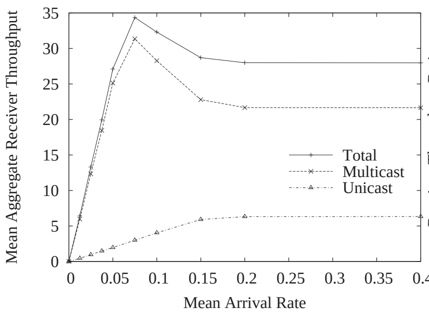

(d) Aggregate receiver throughput T.

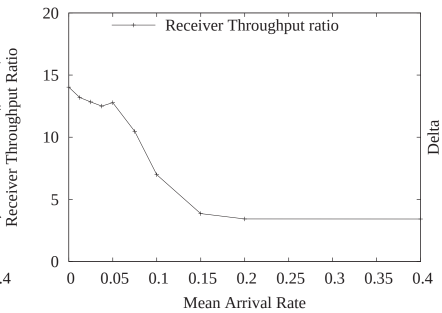

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

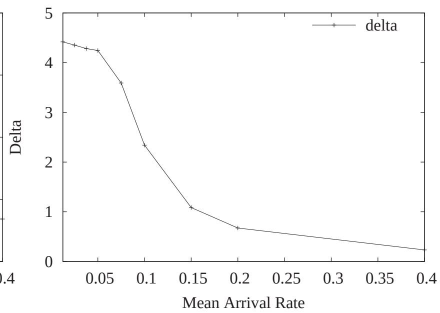

(f) Fairness performance: δ.

Fig. 8. Performance for $(\Lambda + \Lambda)$VOQ architecture with channel priority policy with $w = 0$.

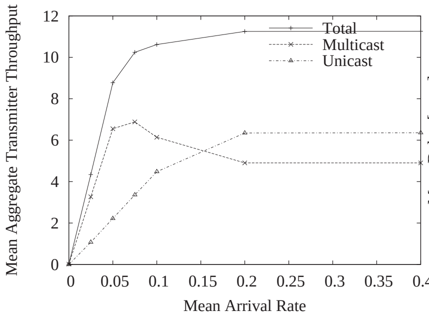

(a) Aggregate transmitter throughput T.

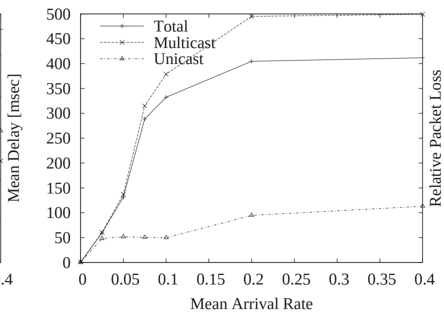

(b) Mean aggregate delay D.

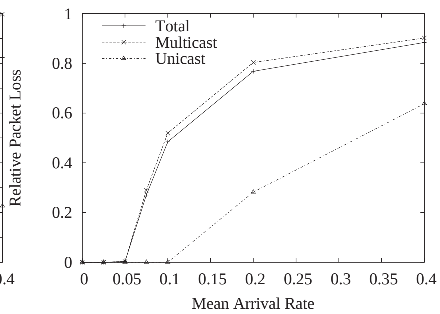

(c) Relative packet loss L.

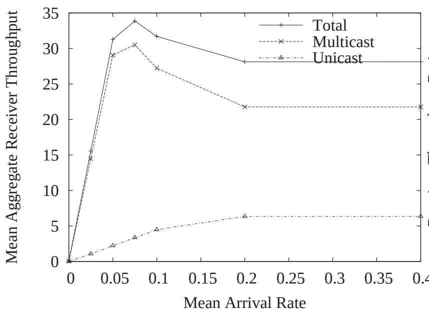

(d) Aggregate receiver throughput T.

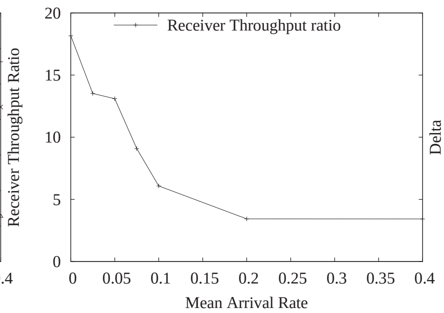

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

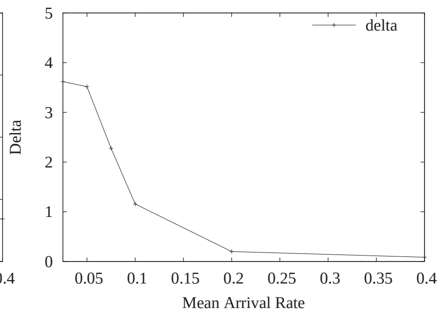

(f) Fairness performance: δ.

Fig. 9. Performance for $(\Lambda + \Lambda)$VOQ architecture with channel priority policy with $w = 0$ with Bernoulli traffic.

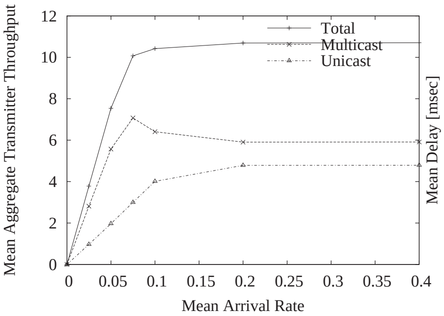

(a) Aggregate transmitter throughput T.

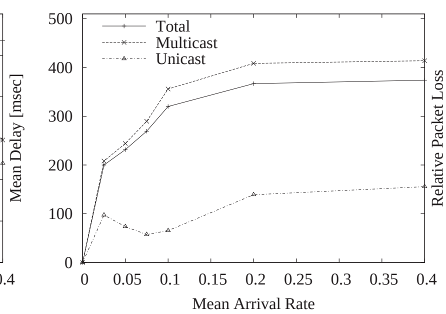

(b) Mean aggregate delay D.

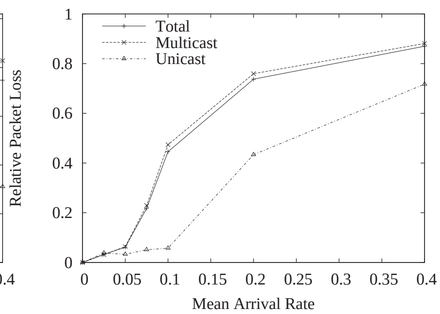

(c) Relative packet loss L.

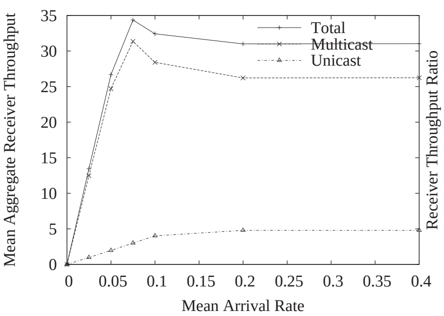

(d) Aggregate receiver throughput T.

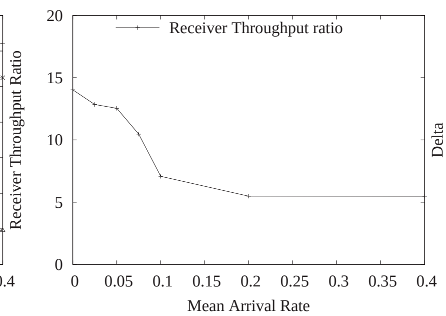

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

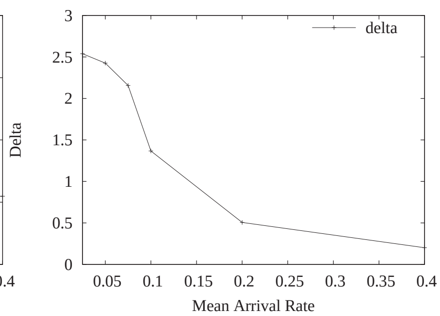

(f) Fairness performance: δ.

Fig. 10. Performance for $(\Lambda + \Lambda)$VOQ architecture with channel priority policy with $w = 0.5$.

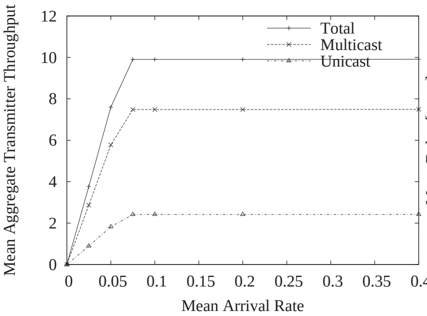

(a) Aggregate transmitter throughput T.

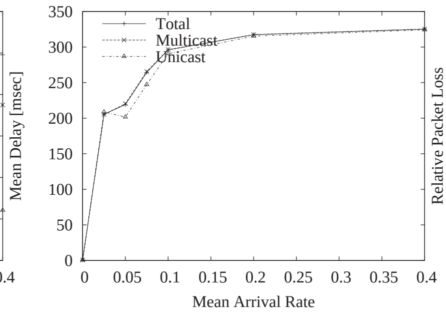

(b) Mean aggregate delay D.

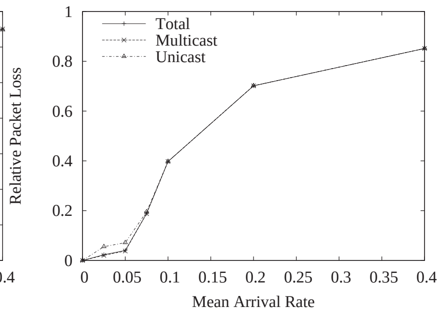

(c) Relative packet loss L.

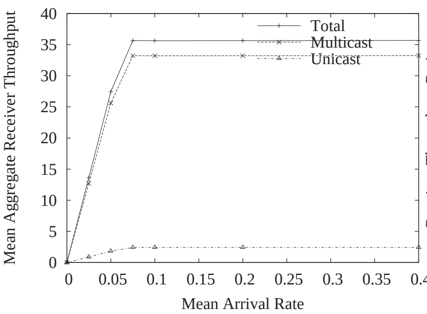

(d) Aggregate receiver throughput T.

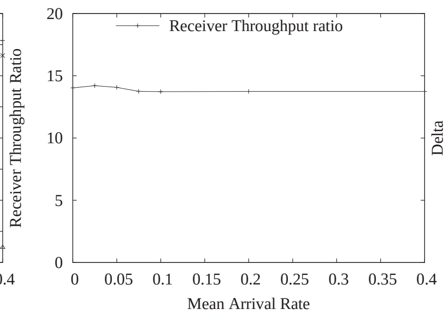

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

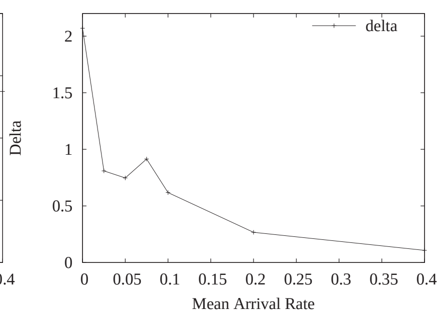

(f) Fairness performance: δ.

Fig. 11. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with $w = 1$.

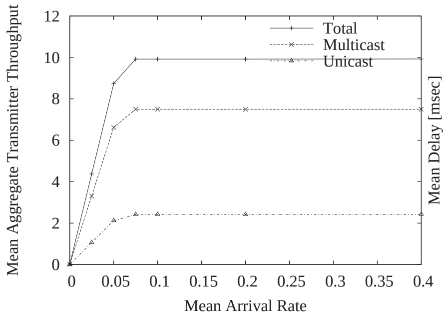

(a) Aggregate transmitter throughput T.

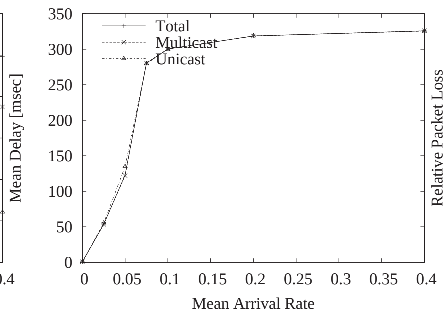

(b) Mean aggregate delay D.

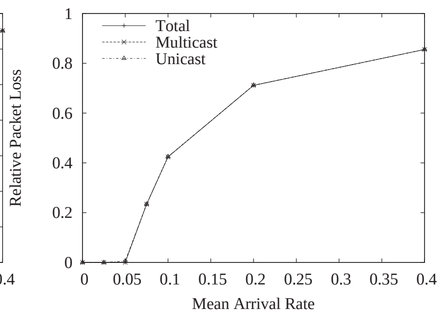

(c) Relative packet loss L.

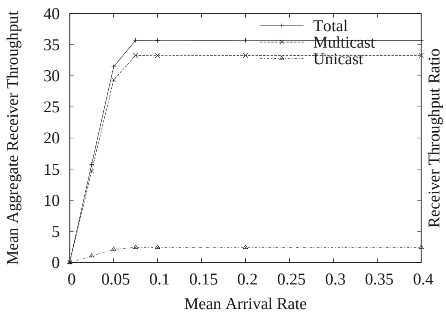

(d) Aggregate receiver throughput T.

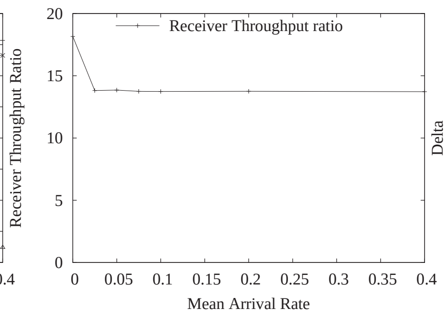

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

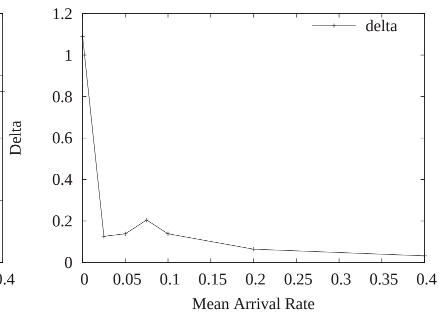

(f) Fairness performance: δ.

Fig. 12. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with $w = 1$ with Bernoulli traffic.

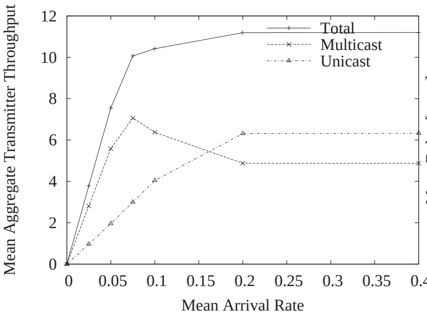

(a) Aggregate transmitter throughput T.

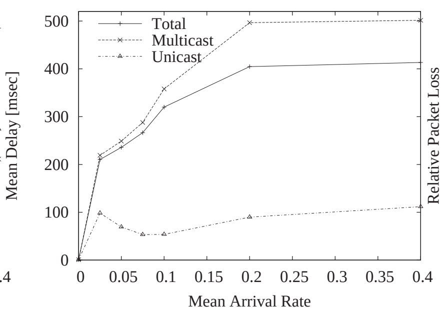

(b) Mean aggregate delay D.

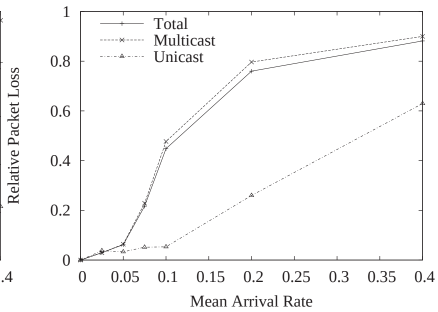

(c) Relative packet loss L.

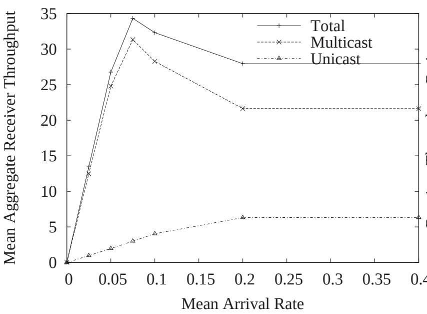

(d) Aggregate receiver throughput T.

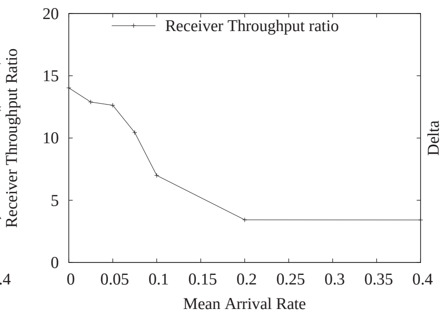

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

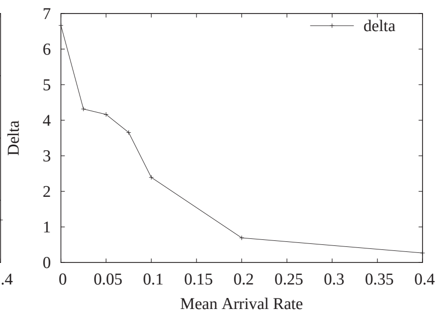

(f) Fairness performance: δ.

Fig. 13. Performance for $(\Lambda{+}\Lambda)$VOQ architecture with traffic priority policy with $w = 0$.

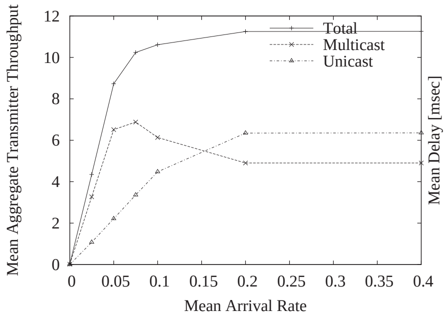

(a) Aggregate transmitter throughput T.

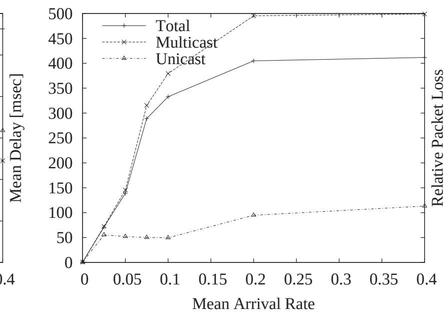

(b) Mean aggregate delay D.

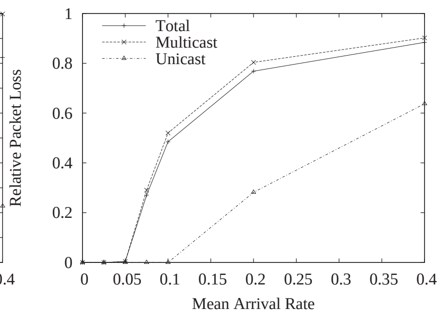

(c) Relative packet loss L.

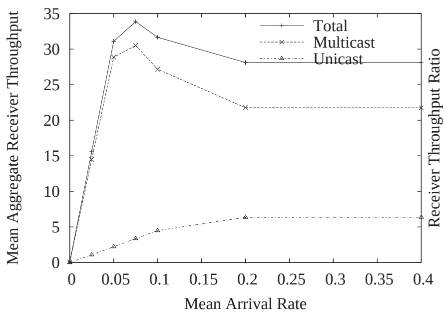

(d) Aggregate receiver throughput T.

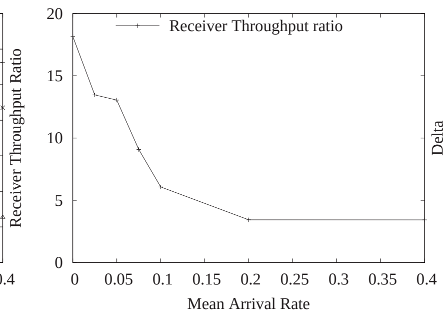

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

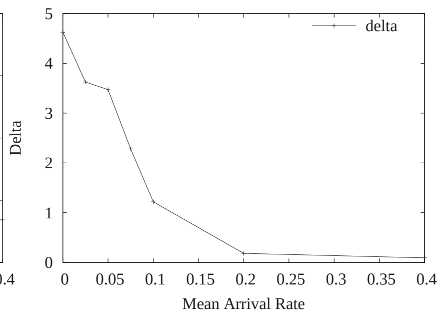

(f) Fairness performance: δ.

Fig. 14. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with $w = 0$ with Bernoulli traffic.

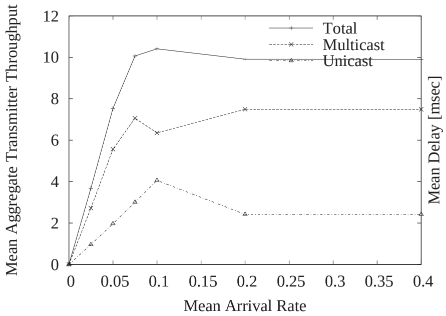

(a) Aggregate transmitter throughput T.

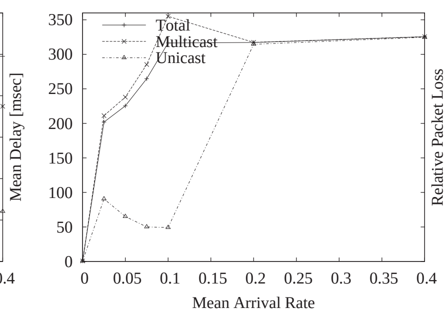

(b) Mean aggregate delay D.

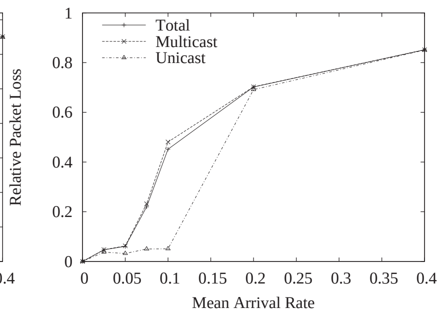

(c) Relative packet loss L.

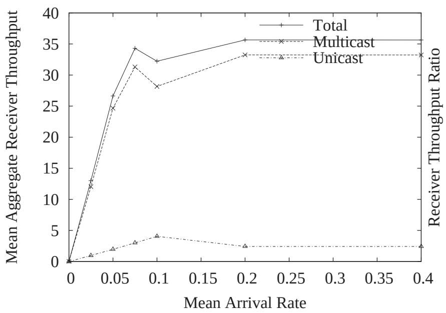

(d) Aggregate receiver throughput T.

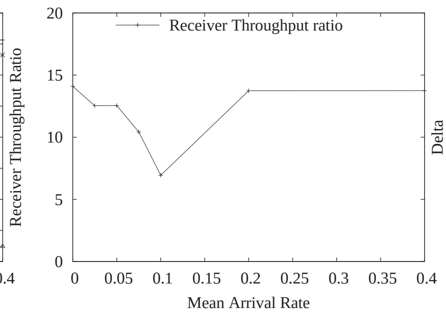

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

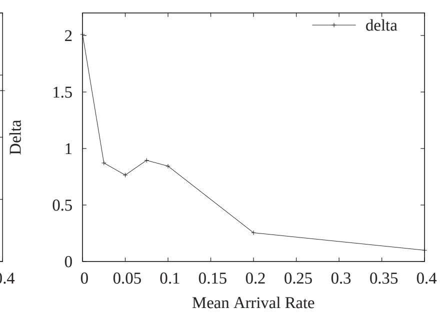

(f) Fairness performance: δ.

Fig. 15. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with memory with $w = 1$.

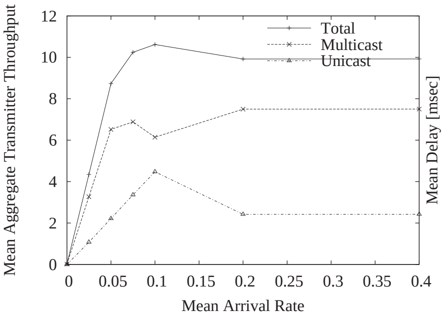

(a) Aggregate transmitter throughput T.

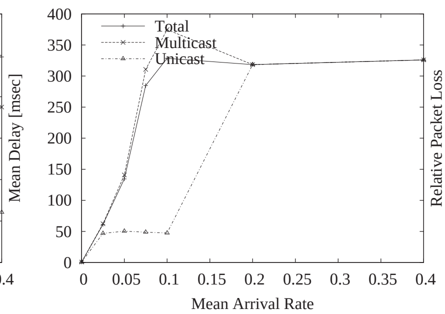

(b) Mean aggregate delay D.

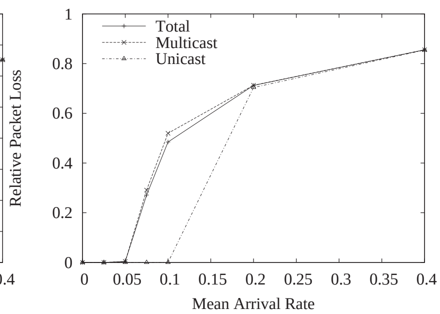

(c) Relative packet loss L.

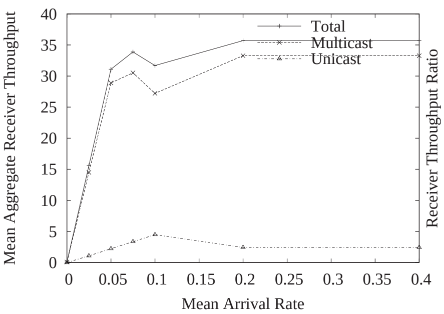

(d) Aggregate receiver throughput T.

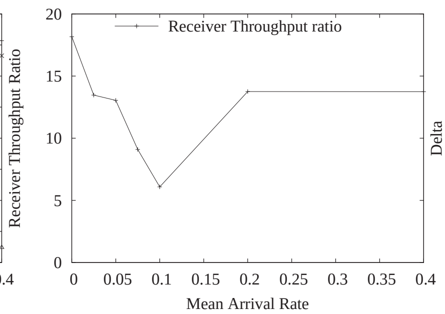

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

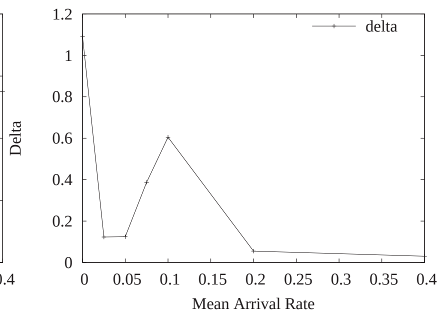

(f) Fairness performance: δ.

Fig. 16. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with memory with $w = 1$ with Bernoulli traffic.

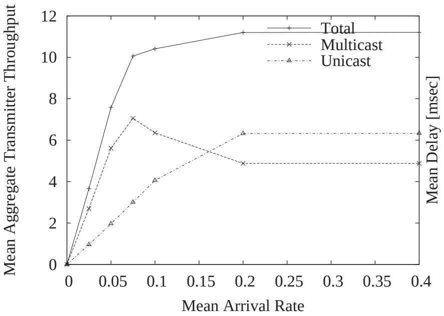

(a) Aggregate transmitter throughput T.

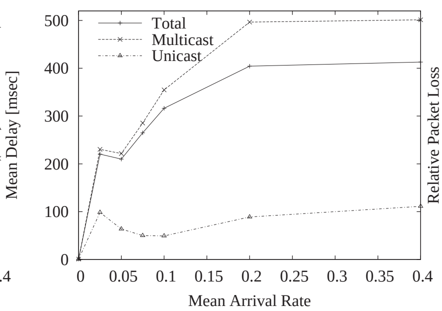

(b) Mean aggregate delay D.

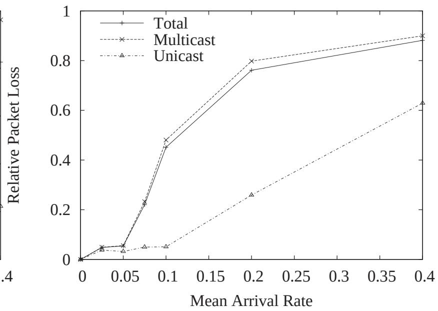

(c) Relative packet loss L.

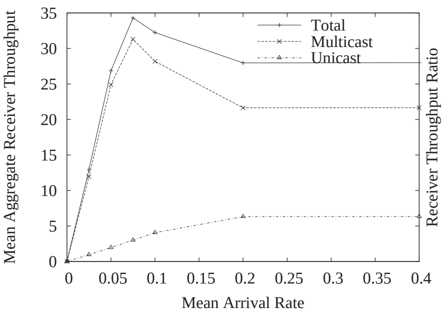

(d) Aggregate receiver throughput T.

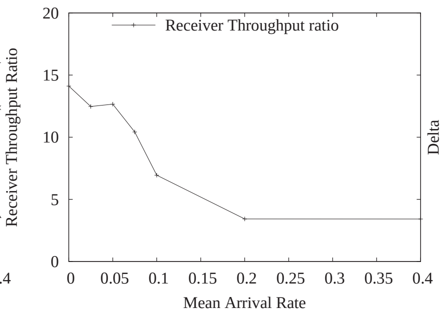

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

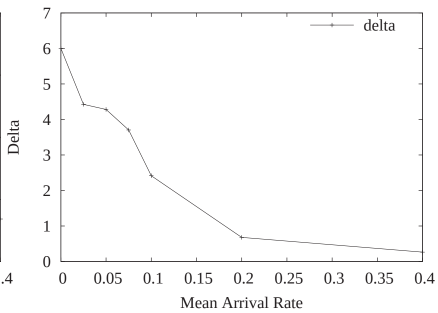

(f) Fairness performance: δ.

Fig. 17. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority with memory policy with $w = 0$.

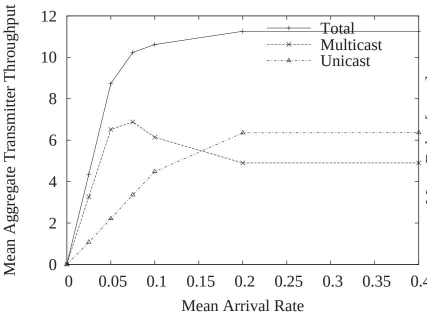

(a) Aggregate transmitter throughput T.

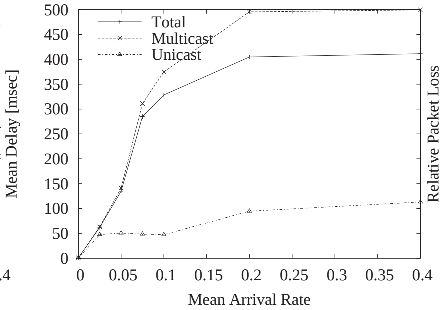

(b) Mean aggregate delay D.

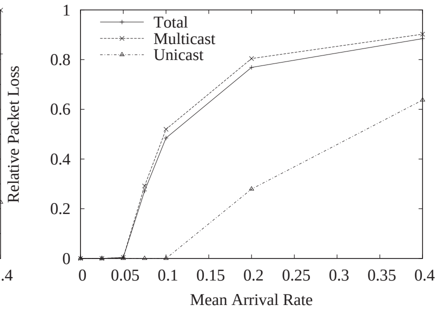

(c) Relative packet loss L.

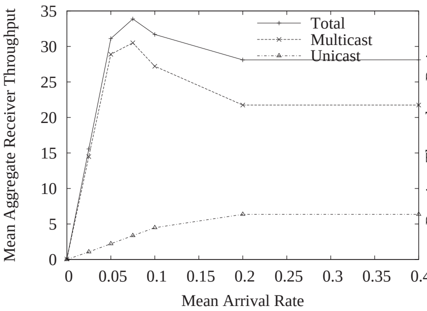

(d) Aggregate receiver throughput T.

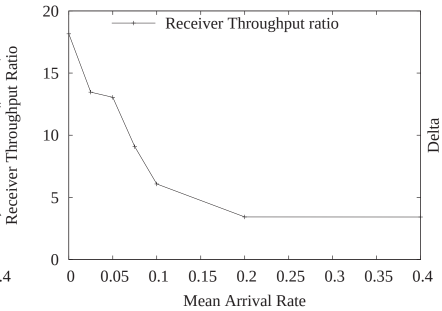

(e) Ratio of multicast to unicast receiver throughput R^m/R^u.

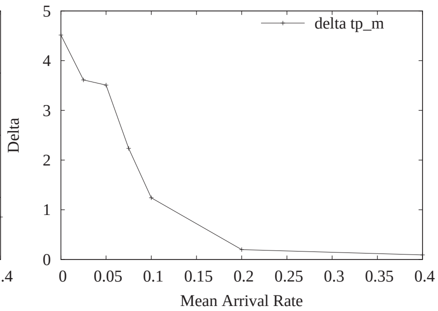

(f) Fairness performance: δ.

Fig. 18. Performance for $(\Lambda+\Lambda)$VOQ architecture with traffic priority policy with memory with $w = 0$ with Bernoulli traffic.