

Original citation:

Al-Bawani, Kamal, Englert, Matthias and Westermann, Matthias (2018) Online packet scheduling for CIOQ and buffered crossbar switches. *Algorithmica* . doi:[10.1007/s00453-018-0421-x](https://doi.org/10.1007/s00453-018-0421-x)

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/99884>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"The final publication is available at Springer via " <http://dx.doi.org/10.1007/s00453-018-0421-x>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Online Packet Scheduling for CIOQ and Buffered Crossbar Switches *

Kamal Al-Bawani[†] Matthias Englert[‡] Matthias Westermann[§]

Abstract

We consider the problem of online packet scheduling in Combined Input and Output Queued (CIOQ) and buffered crossbar switches. In the widely used CIOQ switches, packet buffers (queues) are placed at both input and output ports. An $N \times N$ CIOQ switch has N input ports and N output ports, where each input port is equipped with N queues, each of which corresponds to an output port, and each output port is equipped with only one queue. In each time slot, arbitrarily many packets may arrive at each input port, and only one packet can be transmitted from each output port. Packets are transferred from the queues of input ports to the queues of output ports through the internal fabric. Buffered crossbar switches follow a similar design, but are equipped with additional buffers in their internal fabric. In either model, our goal is to maximize the number or, in case the packets have weights, the total weight of transmitted packets.

Our main objective is to devise online algorithms that are both competitive and efficient. We improve the previously known results for both switch models, both for unweighted and weighted packets.

For unweighted packets, Kesselman and Rosén (J. Algorithms '06) give an online algorithm that is 3-competitive for CIOQ switches. We give a faster, more practical algorithm achieving the same competitive ratio. In the buffered crossbar model, we also show 3-competitiveness, improving the previously known ratio of 4.

For weighted packets, we give 5.83- and 14.83-competitive algorithms with an elegant analysis for CIOQ and buffered crossbar switches, respectively. This improves upon the previously known ratios of 6 and 16.24.

1 Introduction

In the widely used *Combined Input and Output Queued (CIOQ) switches*, packet buffers (queues) are placed at both input and output ports. An $N \times N$ CIOQ switch has N input ports and N output ports. Each input port is equipped with N queues, each of which corresponds to an output port, and each output port is equipped with only one queue. The switching fabric connects the input ports with the output ports and is used to transfer packets from the queues of input ports to the queues of output ports. Figure 1 depicts an example of a CIOQ switch.

When a packet arrives at a CIOQ switch, it is first tagged with the following information: the value that represents its class of service, i.e., its priority, the input port through which it enters the

*The second and third author's work was supported by ERC Grant Agreement No. 307696. A preliminary version appeared in Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pages 241–250, 2016.

[†]Department of Computer Science, RWTH Aachen University, Germany kbawani@cs.rwth-aachen.de

[‡]DIMAP and Department of Computer Science, University of Warwick, UK englert@dcs.warwick.ac.uk

[§]Department of Computer Science, TU Dortmund, Germany matthias.westermann@cs.tu-dortmund.de

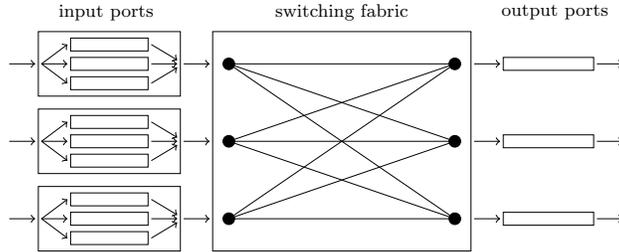


Figure 1: CIOQ switch — An example with $N = 3$

switch, and the output port through which it has to leave the switch. Packets proceed inside the switch in the following way. They are first stored in the queues of the input ports, such that each packet is stored in the queue that corresponds to its output port. After that, they are transferred from input to output ports through the switching fabric, and reside in the queues of the output ports until they are eventually sent out of the switch. However, queues inside the switch are of limited capacities and there may be bursts of packets arriving which exceed the capacities. Thus, queues may overflow. Typically, packets are transferred through the switching fabric with a rate that is \hat{s} times the rate of transmission, i.e., they are transferred through the switching fabric over \hat{s} cycles of speed in each time slot. We call \hat{s} the speedup of the switch. It is worth noting here that we consider non-FIFO queues, i.e., packets can be stored in and released from queues in any arbitrary order.

Closely related to CIOQ switches, another type of switch architecture, the so-called *buffered crossbar switches*, is obtained by adding further queues at the crosspoints of the switching fabric. More specifically, for every queue at the input ports, an additional queue is placed at the switching fabric and dedicated to accommodate packets that are transferred from the input queue before they later on are transferred further to the corresponding output port. The number of those crossbar queues is proportional to the number of crosspoints, i.e., N^2 , but it has been shown that the adoption of crossbar queues significantly decreases the scheduling overhead of CIOQ switches. Figure 2 depicts an example of a buffered crossbar switch.

Packet scheduling in both CIOQ and buffered crossbar switches has been extensively studied in the networking literature (see, e.g., [10, 11]). The design and analysis of scheduling algorithms in that line of research is mostly based on prior assumptions about the traffic distribution, e.g., Poisson-like distributions. However, it has been shown that Internet traffics do not necessarily adhere to such particular distributions (see, e.g., [29, 32]). We do not make any prior assumptions about the arrival behavior of packets, and instead resort to the framework of competitive analysis [31], which is the typical worst-case analysis used to assess the performance of *online* algorithms, i.e., algorithms whose input is revealed piece by piece over time, and the decision they make in each time step is irrevocable.

In competitive analysis, the benefit, in our case the switch's throughput, of an online algorithm is compared to the benefit of an optimal offline algorithm OPT which is assumed to know the entire input sequence in advance. An online algorithm ONL is called *c-competitive* if, for each input sequence σ , the benefit of OPT over σ is at most c times the benefit of ONL over σ . The value c is also called the *competitive ratio* of ONL .

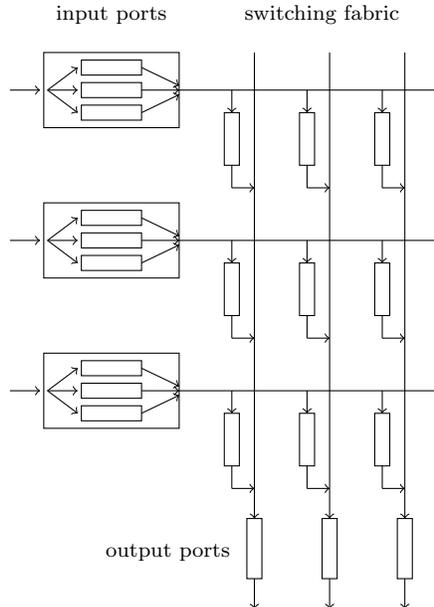


Figure 2: Buffered crossbar switch — An example with $N = 3$

1.1 Our contribution

Our objective in the CIOQ model is twofold: to devise online algorithms that are both competitive and efficient. All online algorithms known for this problem are based on computing a maximum matching in each scheduling cycle, and thus are far from being efficient for real-world switches. We present new algorithms that are significantly more efficient and yet achieve the best competitive ratios known for this problem.

In each scheduling cycle, a bipartite graph is induced from the current configuration of the input and output queues, where the vertices of the left-hand side correspond to the input ports, and the vertices of the right-hand side correspond to the output ports. An edge (i, j) indicates that a packet can be transferred from the i -th input port to the j -th output port. Clearly, a matching in this graph corresponds to an admissible schedule for the current scheduling cycle.

We present two online algorithms in this model: Greedy Matching (GM) for the unit-value case, i.e., where all packets have the same value, and Preemptive Greedy (PG) for the general-value case. Both algorithms are based on greedy maximal matching computations, i.e., we construct a matching incrementally by adding edges, one by one, until no more edges can be added. This is much more efficient than computing the maximum matchings, which have been used in previous works. Moreover, computing maximal matchings complies more with the current practice in distributed systems where packet scheduling has to perform in real time.

With respect to competitiveness, we show in Section 2.1 that GM is 3-competitive for any speedup, and thus it achieves the best competitive ratio known for this problem [23]. In Section 2.2, we show that PG has a competitive ratio of $3 + 2\sqrt{2} \approx 5.83$ for any speedup, which improves upon the previously known competitive ratio of 6 [24].

To obtain these results in an elegant way, we manipulate the queues of an optimal offline algorithm such that certain invariants in relation to our online algorithms are maintained. The techniques we use in the analysis of GM and PG also allow us to achieve improved upper bounds in the related model of buffered crossbar switches. For the unit-value case of this model, Kesselman et al. [21] present a greedy algorithm, which we call Crossbar Greedy Unit (CGU), with a competitive

ratio of 4 for any speedup. We improve on this result and show that CGU is indeed 3-competitive. For the general-value case, they give an algorithm that is 16.24-competitive for any speedup. We present a slightly different algorithm, Crossbar Preemptive Greedy (CPG), and show that it achieves a competitive ratio of ≈ 14.83 for any speedup.

A similar analysis technique has been successfully used by Jež et al. [19] for other packet scheduling related problems. However, in that work the buffer is manipulated in such a way that the optimal algorithm and the online algorithm always have an identical buffer content. In our proofs, we maintain different invariants.

1.2 Related work

For the general-value case of CIOQ switches with FIFO queues, i.e., packets are stored and released in order of their arrival, Kesselman and Rosén [23] give two algorithms with competitive ratios of $4 \cdot S$ and $8 \cdot \min\{k, 2 \log \alpha\}$, where k is the number of distinct packet values and α is the ratio between the largest and the smallest packet value. The latter result was improved by Azar and Richter [7] who give an algorithm with a competitive ratio of 8 for any speedup. Kesselman et al. [22] show that this algorithm is 7.47-competitive. For the buffered crossbar model with FIFO queues, Kesselman et al. [20] give a 19.95-competitive algorithm for any speedup.

A simpler model called input queued switches (IQ) consists of m input queues of the same capacity B and only one output port. It is worth noticing that both the CIOQ and buffered crossbar models generalize this model, e.g., the CIOQ model reduces to the IQ model if the speedup is 1 and only one input port is in use. Therefore, all lower bounds in the IQ model carry over to the CIOQ and buffered crossbar models. In the following, we cite the most known results on the IQ model.

Azar and Richter [6] show that any work-conserving policy for the IQ model is 2-competitive. In the unit-value case, they provide a lower bound of $2 - 1/m$ on the competitive ratio of any deterministic algorithm. Albers and Schmidt [3] improve this result and give a policy called SEMI-GREEDY that is $17/9 \approx 1.89$ -competitive, for any B with $m \gg B$, and $13/7 \approx 1.86$ -competitive, for $B = 2$. They also give a lower bound of $2 - 1/B$ on the competitive ratio of any greedy algorithm. Bienkowski [8] presents a lower bound of $e/(e - 1) \approx 1.58$ on the competitive ratio of any (even randomized) algorithm. Azar and Richter [6] give an optimal randomized policy and Azar and Litichevsky [4] give an optimal deterministic policy matching this lower bound for large B . For $m = 2$, Schmidt [30] presents a lower bound of $16/13 \approx 1.23$ on the competitive ratio of any (even randomized) algorithm, and Bienkowski and Madry [9] give an optimal randomized policy and Kobayashi et al. [25] give an optimal deterministic policy matching this lower bound.

For the general-value case of the IQ model with m FIFO queues, Azar and Richter [6] give a generic technique that transforms any single-queue algorithm with a competitive ratio c into an m -queue algorithm with a competitive ratio $2c$. Given the results of Englert and Westermann [12] on FIFO single-queue switches, the technique of [6] leads to a competitive ratio of $\sqrt{13} - 1 \approx 2.61$ for the special case in which each packet can take only the values 1 and $\alpha > 1$ and a competitive ratio of $2\sqrt{3} \approx 3.47$ for arbitrary packet values. For two packet values, Kobayashi, Miyazaki, and Okabe [26] give improved upper bounds for large enough values of B . For arbitrary packet values, Azar and Richter [5] give an improved upper bound by presenting the 3-competitive Transmit Largest Head algorithm (TLH). Itoh and Takahashi [18] refine this result and show that TLH is in fact $(3 - 1/\alpha)$ -competitive when the packet values are from the interval $[1, \alpha]$.

The problem of packet scheduling (also known as buffer management) has also been studied under several other models. For example, the multi-queue model with shared memory [1, 15, 16], the multi-queue model with class segregation [2], the single-queue (FIFO) model [12], and the

bounded delay model, where packets have deadlines besides their values [13, 27]. Comprehensive and up-to-date surveys on this problem and its variants can be found in [14, 17, 28].

1.3 Models and notations

We consider a CIOQ switch with N input ports and N output ports. Each input port has N queues and each output port has one queue. We call the queues at the input ports the *input queues* and those at the output ports the *output queues*. An input queue that is placed at input port i ($i = 1, \dots, N$) and corresponds to output port j ($j = 1, \dots, N$) is denoted by Q_{ij} . An output queue that is placed at output port j ($j = 1, \dots, N$) is denoted by Q_j . For any input or output queue Q , the capacity of Q , i.e., the number of packets that can be stored in Q , is denoted by $B(Q)$, and $Q(t)$ denotes the set of packets that reside in Q at time t . All queues in the switch are non-FIFO, i.e., packets may be stored in and released from queues in any arbitrary order.

An input instance of this problem is a sequence of packets arriving at the switch in an online manner, i.e., packets that arrive at time t are not known before t . All packets have the same size. For each packet p in the input sequence, $v(p)$, $\text{arr}(p)$, $\text{in}(p)$, and $\text{out}(p)$ denote p 's value, arrival time, input port, and output port, respectively, where $\text{in}(p)$ and $\text{out}(p)$ take on values between 1 and N .

We denote the arrival of a new packet as an *arrival event*, the transfer of a packet from an input queue to an output queue as a *scheduling event* and the sending of a packet from an output queue as a *transmission event*. Therefore, an input sequence σ can be seen as a sequence of arrival, scheduling and transmission events. The time that precedes the first arrival event of the sequence is denoted as time 0. We assume that the queues of any algorithm are all empty at time 0.

Continuous time is divided into slots of unit length, and each of these time slots is further divided into three phases; namely, arrival, scheduling, and transmission phases. For simplicity, we assume that for any given queue, all events in arrival, scheduling, and transmission phases occur at different (fractional) times. For example, the arrival events in the first time slot will occur at distinct time points in the interval $(0, 1/3)$, the scheduling events will occur at distinct time points in the interval $(1/3, 2/3)$, and the transmission events will occur at distinct time points in the interval $(2/3, 1)$.

In the *arrival phase*, arbitrarily many packets arrive at the switch. An arriving packet p is either *accepted* and thus inserted in queue Q_{ij} , where $i = \text{in}(p)$ and $j = \text{out}(p)$, or it is *rejected*, i.e., discarded.

In the *scheduling phase*, a set of packets that are stored in input queues are transferred to their corresponding output queues through the switching fabric. These transfers take place in internal time cycles which we call the *scheduling cycles*. We say that a switch has a speedup \hat{s} when it is capable of performing \hat{s} scheduling cycles within a single time slot. We denote the s -th cycle of time slot T by $T[s]$, for $s = 1, \dots, \hat{s}$. In any scheduling cycle, a matching between input and output ports is computed, such that at most one packet is released from each input port and at most one packet is admitted to each output port. More specifically, when a packet p is transferred from queue Q_{ij} in scheduling cycle $T[s]$, it is forwarded through the switching fabric to queue Q_j , and no packet except p is released from input port i or forwarded to output port j in $T[s]$.

Finally, in the *transmission phase*, at most one packet is sent out from each output queue, i.e., transmitted to its next destination on the network.

Preemption is allowed, i.e., a packet that was previously inserted into a queue can be *preempted*, i.e., discarded, before it is sent. Therefore, a packet may be lost in one of two occasions: rejection upon its arrival, or preemption after getting stored in a queue.

The benefit made by an online algorithm ONL on an input sequence σ is denoted by $\text{ONL}(\sigma)$,

and is defined as the total value of packets that ONL sends from the output queues. We aim at maximizing this benefit. An algorithm that knows the entire input beforehand and makes the maximum benefit on any sequence is denoted as OPT. An online algorithm ONL is c -competitive if $\text{OPT}(\sigma) \leq c \cdot \text{ONL}(\sigma)$ for any input sequence σ .

Buffered crossbar switches are obtained by adding further queues at the crosspoints of the switching fabric. A crossbar queue that is placed at the crosspoint of input port i ($i = 1, \dots, N$) and output port j ($j = 1, \dots, N$) is denoted by C_{ij} . Again, all queues in the switch are non-FIFO, i.e., packets may be stored in and released from queues in any arbitrary order.

All other notations and conventions of the CIOQ model hold also for the buffered crossbar model. However, each cycle of the scheduling phase in the buffered crossbar model is divided into two subphases: the *input subphase* and the *output subphase*. In the input subphase, packets can be transferred from any input queue Q_{ij} to its corresponding crossbar queue C_{ij} , such that at most one packet is transferred from each input port i . In the output subphase, packets can be transferred from any crossbar queue C_{ij} to its corresponding output queue Q_j , such that at most one packet is transferred to each output port j .

2 CIOQ switches

2.1 Unit-value case

In this case, all packets have unit value. Thus, our goal is to maximize the number of transmitted packets. In the following, we present the Greedy Matching algorithm (GM).

- **Arrival phase:** For every arriving packet p with $\text{in}(p) = i$ and $\text{out}(p) = j$, accept p if Q_{ij} is not full; otherwise, reject p .
- **Scheduling phase:** In every scheduling cycle $T[s]$, a bipartite graph $G_{T[s]} = (U, V, E)$ is induced from the current configuration of the switch, where $U = \{u_1, \dots, u_N\}$, $V = \{v_1, \dots, v_N\}$, and an edge $(u_i, v_j) \in E$ if and only if the input queue Q_{ij} is not empty and the output queue Q_j is not full at $T[s]$.

A greedy matching $M_{T[s]}$ is then computed on $G_{T[s]}$ in the following way: Start with an empty matching and iterate over all edges of E . Add an edge e to the current matching if e does not violate the matching property.

After $M_{T[s]}$ is computed, for each edge $(u_i, v_j) \in M_{T[s]}$, the head packet of Q_{ij} is transferred to Q_j .

- **Transmission phase:** For every non-empty output queue Q_j , send the packet at the head of Q_j .

The next theorem shows that GM is 3-competitive for any speedup.

Theorem 1. *The competitive ratio of GM is at most 3 for any speedup.*

From now on, we fix an input sequence σ , and, for any input or output queue Q , we reserve the notation Q for the online algorithm and use Q^* to denote the corresponding queue of the offline algorithm OPT.

First, without loss of generality, we assume that OPT is greedy in transmission events, i.e, it sends a packet from an output queue as long as its queue is not empty. Obviously, as OPT knows in advance which packets it is going to send, holding packets back in output queues, rather than sending them as early as possible, cannot improve its benefit.

Now, we modify OPT in a way that does not decrease its benefit of σ . Specifically, at the end of each scheduling cycle $T[s]$, i.e., immediately after OPT has performed its scheduling policy, we apply the following two modifications on the configuration of OPT in the given order:

Modification 2.1.1. *Suppose that GM transfers a packet from Q_{ij} and OPT does not transfer any packet from Q_{ij}^* in $T[s]$. If Q_{ij}^* is not empty in $T[s]$, OPT sends a packet p from Q_{ij}^* directly out of the switch, i.e., through an imaginary channel. In this case, p is called a privileged packet of Type 1 and contributes to the benefit of OPT.*

Modification 2.1.2. *Suppose that OPT transfers a packet p to Q_j^* and GM does not transfer any packet to Q_j in $T[s]$. If Q_j is not full in $T[s]$, OPT sends p directly out of the switch. In this case, p is called a privileged packet of Type 2 and contributes to the benefit of OPT.*

Clearly, these modifications do not decrease the benefit of OPT. They can only make it stronger by allowing it to send packets directly from input ports to outside the switch without being enqueued in output ports. The input and output queues will respectively become shorter in this case and thus the optimal algorithm may accept more new packets.

Before we continue, we introduce further notations. We call packets that OPT schedules through the normal channels, i.e., they are not privileged, *normal* packets. We use S^* and P^* to denote the sets of OPT's normal and privileged packets, respectively. Clearly, the benefit of OPT is given by $|P^*| + |S^*|$. We also use S to denote the set of packets sent by GM. Thus, we want to show that $|P^*| + |S^*| \leq 3|S|$.

We now show how to derive the competitive ratio of 3. First, we show in Lemma 1 how Modifications 2.1.1 and 2.1.2 are used to preserve the following invariant: At any time, each queue in GM is not shorter than its counterpart in OPT. Therefore, for any transmission event at time t and output port j , if OPT sends a packet from Q_j^* at t , GM must also send a packet from Q_j at t . Hence, $|S^*| \leq |S|$. After that, we show by Lemma 3 that $|P^*| \leq 2|S|$. Thus, the proof of Theorem 1 follows directly from these two lemmas.

Lemma 1. *For any $i, j \in \{1, \dots, N\}$ and any time t , the following inequalities hold:*

I1. $|Q_{ij}^*(t)| \leq |Q_{ij}(t)|$

I2. $|Q_j^*(t)| \leq |Q_j(t)|$

Proof. Inequalities I1 and I2 can be shown by a simple induction over the event sequence. Let the induction base be at time 0, i.e., before the sequence starts. All queues are empty at this time and thus I1 and I2 hold. Assume now that they hold for any time up to some arrival, scheduling, or transmission event τ . Then we have to show that they still hold right after the event τ .

Assume τ is an arrival event. Clearly, output queues do not change in arrival events and thus I2 holds for this case. For I1, the only critical case is when the arriving packet is rejected by GM and accepted by OPT. However, the input queue of GM must be full in this case and thus I1 still holds.

Now, let τ be a scheduling event. Here, the only critical case for I1 is when GM transfers a packet from Q_{ij} while OPT does not transfer anything from Q_{ij}^* . However, either Q_{ij}^* is empty in this case or it cannot happen due to Modification 2.1.1. For I2, the only critical case is when OPT inserts a packet into Q_j^* while GM does not insert anything into Q_j . However, either Q_j is full in this case or it cannot happen due to Modification 2.1.2.

Finally, assume τ is a transmission event. Clearly, the input queues do not change in transmission events and thus I1 holds for this case. For I2, the only critical case is when GM sends a packet from Q_j while OPT does not send anything from Q_j^* . However, since we assume that OPT is greedy at sending, its output queue must be empty in this case and thus I2 still holds. \square

The following lemma shows that if Modification 2.1.2 takes place, GM must transfer a packet from the same input port.

Lemma 2. *Suppose that, in $T[s]$, OPT transfers a packet p from Q_{ij}^* to Q_j^* and GM does not transfer any packet to Q_j . If Q_j is not full in $T[s]$, then GM transfers a packet p' from $Q_{ij'}$ in $T[s]$, where $j' \neq j$.*

Proof. Recall the bipartite graph $G_{T[s]}$ and the corresponding matching $M_{T[s]}$ which are induced from the configuration of GM right before performing the scheduling cycle $T[s]$.

Assume that Q_j is not full in $T[s]$. By Inequality I1 of Lemma 1, since OPT transfers p from Q_{ij}^* in $T[s]$, GM must have at least one packet in Q_{ij} . Therefore, an edge (u_i, v_j) must be in E . Nevertheless, since GM does not transfer any packet to Q_j , (u_i, v_j) is not in $M_{T[s]}$. Since $M_{T[s]}$ is a maximal matching, there must exist an edge $(u_i, v_{j'})$, for $j' \neq j$, such that $(u_i, v_{j'}) \in M_{T[s]}$. Hence, a packet p' is transferred from $Q_{ij'}$ in $T[s]$. \square

Lemma 3. *The following inequality holds:*

$$|P^*| \leq 2|S| .$$

Proof. We carry out the following mapping scheme from P^* to S in each scheduling cycle $T[s]$.

1. Let p be a privileged packet of Type 1 that is sent by OPT from Q_{ij}^* in $T[s]$. By Modification 2.1.1, GM transfers a packet p' from Q_{ij} in $T[s]$. Map p to p' .
2. Let p be a privileged packet of Type 2 that is sent by OPT from Q_{ij}^* . By Lemma 2, GM transfers a packet p' from $Q_{ij'}$ in $T[s]$, where $j' \neq j$. Map p to p' .

Clearly, this mapping scheme is feasible, i.e., each packet $p \in P^*$ is mapped to a packet $q \in S$. Furthermore, at most two privileged packets can be mapped to each packet $q \in S$. To see that, let q be a packet transferred by GM from Q_{ij} in a scheduling cycle $T[s]$. Clearly, q can get mapped only in $T[s]$, provided that OPT sends privileged packets at this time. By Modifications 2.1.1 and 2.1.2, OPT can send at most 2 privileged packets from input port i in $T[s]$: one of Type 1 if OPT's queue of Q_{ij}^* is not empty, and one of Type 2 if it transfers a packet from another queue $Q_{ij'}^*$. Thus, these two privileged packets are mapped to q . \square

2.2 General-value case

For the case of arbitrary packet values, we present the Preemptive Greedy algorithm (PG) that is a variant of a 6-competitive algorithm given by Kesselman and Rosén [24]. We show next that PG has a competitive ratio of $3 + 2\sqrt{2} \approx 5.83$ for any speedup.

Before we describe PG formally, we introduce further notations. Let $g_{ij}(t)$ denote the packet with the greatest value in Q_{ij} at time t , and $l_{ij}(t)$ (resp. $l_j(t)$) denote the packet with the least value in Q_{ij} (resp. Q_j) at time t . Additionally, let $\beta \geq 1$ be a parameter of the algorithm that will be determined later.

- **Arrival phase:** If a packet p arrives at time t with $\text{in}(p) = i$ and $\text{out}(p) = j$, accept p if

$$|Q_{ij}(t)| < B(Q_{ij}) \bigvee v(l_{ij}(t)) < v(p) ;$$

otherwise, reject p . If p is accepted while $|Q_{ij}(t)| = B(Q_{ij})$, then $l_{ij}(t)$ is preempted. In this case we also say that p causes the preemption of $l_{ij}(t)$ or, to be more concise, that p preempts $l_{ij}(t)$.

- **Scheduling phase:** In every scheduling cycle $T[s]$, a weighted bipartite graph $G_{T[s]} = (U, V, E, w)$ is induced from the current configuration of the switch, where $U = \{u_1, \dots, u_N\}$, $V = \{v_1, \dots, v_N\}$, an edge $(u_i, v_j) \in E$ if and only if

$$|Q_{ij}(T[s])| > 0 \wedge \left(|Q_j(T[s])| < B(Q_j) \vee v(g_{ij}(T[s])) > \beta v(l_j(T[s])) \right),$$

and the weight of (u_i, v_j) is given by $w(u_i, v_j) = v(g_{ij}(T[s]))$.

A greedy matching $M_{T[s]}$ is then computed on $G_{T[s]}$ in the following way: Start with an empty matching and iterate over all edges of E in a descending order of their weights. Add an edge e to the current matching if e does not violate the matching property.

After $M_{T[s]}$ is computed, for each edge $(u_i, v_j) \in M_{T[s]}$, the packet $g_{ij}(T[s])$ is transferred to Q_j . If $g_{ij}(T[s])$ is transferred while $|Q_j(T[s])| = B(Q_j)$, then $l_j(T[s])$ is preempted. Again, in this case we also say that $g_{ij}(T[s])$ preempts $l_j(T[s])$.

- **Transmission phase:** For every non-empty output queue Q_j , send the packet with the greatest value in Q_j .

As described above, unlike the algorithm given in [24], PG computes a maximal weighted matching in each scheduling cycle rather than a maximum weighted matching.

Theorem 2. *For $\beta = \sqrt{2} + 1$, the competitive ratio of PG is at most $3 + 2\sqrt{2} \approx 5.83$ for any speedup.*

First, we fix an input sequence σ . Without loss of generality, we make the following assumptions about OPT:

- A1. OPT is greedy in scheduling and transmission events, i.e., when it transfers or sends a packet p from an input or output queue, it chooses p as the one with the greatest value in the queue.
- A2. OPT is work-conserving at output ports, i.e., it sends a packet from every non-empty output queue in each transmission event.

Obviously, as OPT knows in advance which packets it is going to send, it does not matter for OPT in which order these packets are released from queues or when they are transmitted from output queues. Hence, based on the greediness of PG, we make another harmless assumption:

- A3. In all input and output queues, PG and OPT store packets in the order of their values, where the packet with the greatest value is at the queue's head and the one with the least value is at the queue's tail. Ties are broken arbitrarily but consistently.

Assumption A3 is made for ease of exposition, essentially in the statement and proof of Lemma 4 where packets of any queue in PG and OPT are shown to be consistently aligned.

Similarly to the unit-value case, we modify OPT without decreasing its benefit. Specifically, at the end of each scheduling cycle $T[s]$, i.e., immediately after OPT has performed its scheduling policy, we apply the following modifications on the configurations of OPT:

Modification 2.2.1. *Suppose that PG transfers a packet from Q_{ij} and OPT does not transfer any packet from Q_{ij}^* in $T[s]$. If Q_{ij}^* is not empty in $T[s]$, OPT sends the head packet p of Q_{ij}^* , i.e., the packet with the greatest value in Q_{ij}^* , directly out of the switch. In this case, p is called a privileged packet of Type 1 and contributes to the benefit of OPT.*

Modification 2.2.2. If OPT transfers a packet p to Q_j^* and PG transfers a packet q to Q_j in $T[s]$ with $v(q) < v(p)$, OPT sends p directly out of the switch. In this case, p is called a privileged packet of Type 2 and contributes to the benefit of OPT.

Modification 2.2.3. Suppose that OPT transfers a packet p to Q_j^* and PG does not transfer any packet to Q_j in $T[s]$. If Q_j is not full in $T[s]$ or $v(p) > \beta v(l_j(T[s]))$, OPT sends p directly out of the switch. In this case, p is called a privileged packet of Type 3 and contributes to the benefit of OPT.

Note that Modifications 2.2.2 and 2.2.3 are closely related and dealing with them separately is only for ease of exposition.

Let $\delta_{ij}(k, t)$ (resp. $\delta_j(k, t)$) denote the packet at position k in Q_{ij} (resp. Q_j) at time t , where position 1 corresponds to the head of the queue. Let $\delta_{ij}^*(k, t)$ and $\delta_j^*(k, t)$ be the corresponding notations for OPT. The following lemma shows that each packet in an OPT's input queue is aligned to a packet of the same or greater value in the corresponding input queue of PG, and each packet p in an OPT's output queue is aligned to a packet q in the corresponding output queue of PG, where $v(p) \leq \beta v(q)$.

Lemma 4. For any $i, j \in \{1, \dots, N\}$ and any time t , the following holds:

- I1. $|Q_{ij}^*(t)| \leq |Q_{ij}(t)|$ and $v(\delta_{ij}^*(k, t)) \leq v(\delta_{ij}(k, t))$, for $k = 1, \dots, |Q_{ij}^*(t)|$
- I2. $|Q_j^*(t)| \leq |Q_j(t)|$ and $v(\delta_j^*(k, t)) \leq \beta v(\delta_j(k, t))$, for $k = 1, \dots, |Q_j^*(t)|$

Proof. We show I1 and I2 by induction over the event sequence. Let the induction base be at time 0, i.e., before the sequence starts. All queues are empty at this time and thus I1 and I2 hold. Assume now that they hold for any time up to some arrival, scheduling, or transmission event τ . Then we have to show that they still hold right after the event τ . Let t' be a time just before event τ (but after the preceding event) and let t be a time just after τ (but before the following event). In other words, we assume that I1 and I2 hold up to time t' and want to argue that they also hold at time t . In the following, we will argue only for I2. The argument for I1 is analogous, and we will put the main differences between [] at the respective positions.

Before we start, we say that a packet $p \in Q_j^*(\bar{t})$ is in a *legal alignment*, if p is aligned at time \bar{t} to a packet $q \in Q_j(\bar{t})$ with $v(p) \leq \beta v(q)$. Clearly, it suffices to show that any packet $p \in Q_j^*(t)$ is in a legal alignment. We distinguish between two cases:

Case I2.1 $p \in Q_j^*(t')$. Thus, by induction, p is aligned at t' to a packet $q \in Q_j(t')$ with $v(p) \leq \beta v(q)$ [resp. $v(p) \leq v(q)$]. We need to show in this case that p either remains in the same alignment at t or it changes to another legal alignment. Assumption A3 implies that any packet p from t' either remains in its position at time t , moves one step ahead (if a packet that is in front of p is sent from the queue) or moves one step back (if a new packet is inserted in front of p).

Assume now that p remains in its position at t but q moves. Note that neither q nor any packet in front of it can be released from the queue at time t ; otherwise, by Assumption A2 [resp. Modification 2.2.1], some packet would be also released from Q_j^* , which makes p move one step ahead. Thus, q can only move back at t . In this case, however, the packet q' that is directly in front of q is aligned with p . Since $v(q) \leq v(q')$, p is again in a legal alignment.

Next, assume that p moves one step ahead at t . In this case, p either remains in a legal alignment with q (in case q moves ahead as well) or it aligns with a packet that is in front of q at t' and thus makes again a legal alignment.

Finally, assume that p moves one step back at t . Thus, a packet p' must be inserted in front of p , implying that $v(p) \leq v(p')$. Note that the insertion of p' happens only in one of two cases: (i) if

a packet r with $v(r) \geq v(p')$ is inserted into Q_j (by Modification 2.2.2), or (ii) if Q_j is full at t and $v(p') \leq \beta v(l_j(t))$ (by Modification 2.2.3). Let k denote the position of the alignment (p, q) at time t' . In case (i), either (1) r is inserted in a position $k' \leq k$, and thus p will be aligned again with q at t , or (2) r is inserted in a position $k' > k$, and thus p will be aligned with some packet q' at t . By Assumption A3, the second case implies that $v(r) \leq v(q')$. Since $v(p) \leq v(p') \leq v(r)$, then $v(p) \leq v(q')$. Hence, p is in a legal alignment in either case.

In case (ii), since Q_j is full at t , p must be aligned with some packet q' at t . By Assumption A3, $v(l_j(t)) \leq v(q')$. Moreover, since $v(p') \leq \beta v(l_j(t))$, $v(p) \leq v(p') \leq \beta v(q')$. Thus, p makes a legal alignment with q' . [The respective cases for I1 are: case (i) p' is also inserted into Q_{ij} , thus $r = p'$ in the above argument, and case (ii) Q_{ij} is full at t and $v(l_j(t)) \geq v(p')$.]

Case I2.2 $p \notin Q_j^*(t')$. Thus, p is a new packet that is inserted in the queue at time t . Again, note that the insertion of p into Q_j^* happens only in one of two cases: (i) if a packet r with $v(r) \geq v(p)$ is inserted into Q_j (by Modification 2.2.2), or (ii) if Q_j is full at t and $v(p) \leq \beta v(l_j(t))$ (by Modification 2.2.3). In case (ii), since Q_j is full at t , p must be aligned with a packet q at t . Since $v(p) \leq \beta v(l_j(t))$, $v(p) \leq \beta v(q)$. Thus, p makes a legal alignment with q .

Now, consider case (i). Let k denote the position at which p is inserted. If $k = 1$, p is aligned with the most valuable packet in Q_j at t . Since r is in Q_j at time t , p must be aligned with a packet of value at least $v(r) \geq v(p)$. Now suppose $k > 1$. Let p' be the packet that is directly in front of p at t . Clearly, $p' \in Q_j^*(t')$ and $v(p) \leq v(p')$. Furthermore, let q' be the packet aligned with p' at time t' . Thus, $v(p) \leq v(p') \leq \beta v(q')$. Additionally, let q be the packet at position k in Q_j at time t' (assume $q = \emptyset$, if this is an empty position in Q_j).

Note that (1) r is inserted in position k , and thus p will be aligned with r at t , (2) r is inserted in a position $k' < k$, and thus p will be aligned with q' at t , or (3) r is inserted in a position $k' > k$, and thus p will be aligned with q at t . Clearly, the last case implies that $q \neq \emptyset$ and that $v(q) \geq v(r) \geq v(p)$. Therefore, we have $v(p) \leq v(r)$ in the first case, $v(p) \leq \beta v(q')$ in the second, and $v(p) \leq v(q)$ in the third. Hence, p is in a legal alignment in any case.

[The respective cases for I1 are: case (i) p is also inserted into Q_{ij} , thus $r = p$ in the above argument, and case (ii) Q_{ij} is full at t and $v(l_j(t)) \geq v(p)$.] \square

Similarly to the analysis of the unit-value case, granting OPT with privileged packets must be done carefully, so that the total value of privileged packets remains within a certain factor of the total value of packets that PG sends. Obviously, each privileged packet of Type 1 can be paired with a packet that PG transfers from the same input queue. In the following two lemmas, we show that such a pairing is feasible for privileged packets of Types 2 and 3 as well. Of course, as packets of PG may be preempted after being transferred to output queues, some pairs can be destructed. However, we will show in Lemma 7 how to fix this problem.

Lemma 5. *If OPT transfers a packet p from Q_{ij}^* to Q_j^* and PG transfers a packet q to Q_j in $T[s]$ with $v(q) < v(p)$, then PG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p') \geq v(p)$.*

Proof. Recall the bipartite graph $G_{T[s]}$ and the corresponding matching $M_{T[s]}$ which are induced from the configuration of PG right before performing the scheduling cycle $T[s]$.

By I1 of Lemma 4, since OPT transfers p from Q_{ij}^* in $T[s]$, PG must have at the head of Q_{ij} a packet r with $v(r) \geq v(p)$. Obviously, $v(r) > v(q)$ and thus $q \neq r$. As a result, q must be transferred from an input queue $Q_{i'j}$ with $i' \neq i$. Moreover, since q is inserted in Q_j , the edge $(u_{i'}, v_j) \in E$, and either $|Q_j(T[s])| < B(Q_j)$ or $v(q) > \beta v(l_j(T[s]))$. Thus, it holds also for r that either $|Q_j(T[s])| < B(Q_j)$ or $v(r) > \beta v(l_j(T[s]))$. Hence, the edge $(u_i, v_j) \in E$ as well, and clearly $w(u_i, v_j) > w(u_{i'}, v_j)$. This implies that (u_i, v_j) is considered before $(u_{i'}, v_j)$ during the

computation of $M_{T[s]}$. However, since (u_i, v_j) is not in the matching, the node u_i must have been matched before considering (u_i, v_j) , and thus there exists an edge $(u_i, v_{j'})$, for $j' \neq j$, that is inserted in the matching before considering (u_i, v_j) . As a result, a packet p' is transferred from $Q_{ij'}$, and it must hold that $w(u_i, v_{j'}) \geq w(u_i, v_j)$. Hence, $v(p') \geq v(r) \geq v(p)$. \square

The proof of the following lemma is analogous to that of Lemma 5.

Lemma 6. *Suppose that, in $T[s]$, OPT transfers a packet p from Q_{ij}^* to Q_j^* and PG does not transfer any packet to Q_j . If Q_j is not full in $T[s]$ or $v(p) > \beta v(l_j(T[s]))$, then PG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p') \geq v(p)$.*

Now, recall I2 of Lemma 4. It implies that if OPT sends a packet of value v from some output queue at some time, PG must send a packet of value at least v/β from the same output queue at the same time. Let S (resp. S^*) denote the set of all packets that PG (resp. OPT) sends from output queues. Thus,

$$\sum_{p \in S^*} v(p) \leq \beta \sum_{p \in S} v(p) .$$

Moreover, let P^* denote the set of all privileged packets, of all types, that OPT sends directly out of the switch. The next lemma shows that

$$\sum_{p \in P^*} v(p) \leq \frac{2\beta}{\beta - 1} \sum_{p \in S} v(p) .$$

Thus, we can conclude the competitive ratio of PG as follows

$$\begin{aligned} \text{OPT}(\sigma) &= \sum_{p \in S^*} v(p) + \sum_{p \in P^*} v(p) \\ &\leq \beta \sum_{p \in S} v(p) + \frac{2\beta}{\beta - 1} \sum_{p \in S} v(p) \\ &= \left(\beta + \frac{2\beta}{\beta - 1} \right) \text{PG}(\sigma) . \end{aligned}$$

Finally, it is easy to verify that the optimal value for β is $\sqrt{2} + 1$, resulting in a competitive ratio of $3 + 2\sqrt{2} \approx 5.83$.

Lemma 7. *The following inequality holds:*

$$\sum_{p \in P^*} v(p) \leq \frac{2\beta}{\beta - 1} \sum_{p \in S} v(p) .$$

Proof. We consider the following mapping scheme:

1. Let p be a privileged packet of Type 1 that is sent by OPT from Q_{ij}^* in scheduling cycle $T[s]$. By Modification 2.2.1, PG transfers a packet p' from Q_{ij} in $T[s]$, and by I1 of Lemma 4, $v(p) \leq v(p')$. Map p to p' .
2. Let p be a privileged packet of Type 2 that is sent by OPT from Q_{ij}^* in scheduling cycle $T[s]$. By Lemma 5, PG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p) \leq v(p')$. Map p to p' .

3. Let p be a privileged packet of Type 3 that is sent by OPT from Q_{ij}^* in scheduling cycle $T[s]$. By Lemma 6, PG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p) \leq v(p')$. Map p to p' .
4. Let q be a packet that is preempted by PG from an output queue Q_j due to accepting another packet p' . For each privileged packet p that is mapped to q , re-map p to p' .

As shown above, this mapping scheme is feasible, i.e., each packet $p \in P^*$ is mapped to a packet $p' \in S$. Now, it remains to show that the total value of privileged packets that are mapped to each packet $p' \in S$ is at most $\frac{2\beta}{\beta-1}v(p')$.

For any packet $p' \in S$, p' can get mapped in two events: when it is scheduled and when it preempts a packet from an output queue.

Assume that p' is scheduled from $Q_{ij'}$ to $Q_{j'}$ during scheduling cycle $T[s]$. Now, assume that OPT transfers a packet from Q_{ij}^* to Q_j^* during $T[s]$. Clearly, we can only send one privileged packet p_1 of Type 1 from $Q_{ij'}$ in $T[s]$ (in case $j \neq j'$). Furthermore, we can only send from Q_{ij}^* either a privileged packet p_2 of Type 2 (in case PG transfers a packet q to Q_j with $v(q) < v(p_2)$), or a privileged packet p_3 of Type 3 (in case PG does not transfer any packet to Q_j). Hence, at most two privileged packets may be sent during $T[s]$ from each input port. Since privileged packets are mapped only to packets that are transferred by PG from the same input port during the same scheduling cycle, at most two packets from $\{p_1, p_2, p_3\}$ can be mapped to p' . Furthermore, as shown in the mapping scheme above, the value of any of these privileged packets is at most the value of p' . Thus, the total value of privileged packets that are mapped to p' when it is scheduled is at most $2v(p')$.

Assume now that p' is the m -th packet in a chain of packets q_0, \dots, q_m in which packet q_n preempts packet q_{n-1} , for $1 \leq n \leq m$. Let $x(q_n)$ denote the total value of privileged packets that are mapped to a packet q_n after it preempts q_{n-1} . Thus, the total value of privileged packets that are mapped to p' is given by $x(q_m)$. Note that q_0 does not preempt any packet and thus the total value of privileged packets that are mapped to q_0 is at most $2v(q_0)$. Thus, $x(q_m)$ can be given by the following recursion:

$$\begin{aligned} x(q_0) &\leq 2v(q_0) \quad \text{and} \\ x(q_n) &\leq 2v(q_n) + x(q_{n-1}) \quad , \quad \text{for } 0 < n \leq m \quad . \end{aligned}$$

Solving this recursion, we obtain that

$$x(q_m) \leq 2 \sum_{n=0}^m v(q_n) \quad .$$

Note also that $v(q_{n-1}) \leq v(q_n)/\beta$, for $1 \leq n \leq m$. Hence, we can rewrite $x(q_m)$ as follows:

$$\begin{aligned} x(q_m) &\leq 2v(q_m) \sum_{n=0}^m \frac{1}{\beta^n} \\ &< \frac{2\beta}{\beta-1}v(q_m) \quad . \end{aligned}$$

□

3 Buffered crossbar switches

3.1 Unit-value case

For the case where all packets have unit value, Kesselman et al. [21] considered the following algorithm, which we call Crossbar Greedy Unit (CGU). Arrival and transmission phases of CGU are the same as ones of GM (Section 2.1). In a scheduling phase, CGU works as follows.

- **Scheduling phase:** We divide every scheduling cycle $T[s]$ into two subphases.

- **Input subphase:** For each input port i , choose an arbitrary input queue Q_{ij} that satisfies

$$|Q_{ij}(T[s])| > 0 \bigwedge |C_{ij}(T[s])| < B(C_{ij}) ,$$

and transfer its head packet.

- **Output subphase:** For each output queue Q_j , choose an arbitrary crossbar queue C_{ij} that satisfies

$$|Q_j(T[s])| < B(Q_j) \bigwedge |C_{ij}(T[s])| > 0 ,$$

and transfer its head packet.

The next theorem shows that CGU is 3-competitive for any speedup.

Theorem 3. *The competitive ratio of CGU is at most 3 for any speedup.*

First, we fix an input sequence σ . Again, we modify OPT in a way that does not decrease its benefit over σ . Specifically, at the end of each scheduling cycle $T[s]$, i.e., immediately after OPT has performed its scheduling policy, we apply the following modifications on the configuration of OPT in the given order:

Modification 3.1.1. *Suppose that CGU transfers a packet from Q_{ij} and OPT does not transfer any packet from Q_{ij}^* in $T[s]$. If Q_{ij}^* is not empty in $T[s]$, OPT transfers a packet p from Q_{ij}^* in $T[s]$. If C_{ij}^* is not full in $T[s]$, p is transferred to C_{ij}^* . Otherwise, p is sent directly out of the switch. In either case, p is called a privileged packet and contributes to the benefit of OPT.*

Modification 3.1.2. *Suppose that CGU transfers a packet from Q_{ij} and OPT does not transfer any packet from Q_{ij}^* in $T[s]$. If Q_{ij}^* is empty in $T[s]$ and C_{ij}^* is not full in $T[s]$, a new packet is created and inserted into C_{ij}^* . Such a new packet is called an extra packet of Type 1 and contributes to the benefit of OPT.*

Modification 3.1.3. *Suppose that OPT transfers a packet from C_{ij}^* and CGU does not transfer any packet from C_{ij} in $T[s]$. If C_{ij} is not empty in $T[s]$, a new packet is created and inserted into C_{ij}^* . Such a new packet is called an extra packet of Type 2 and contributes to the benefit of OPT.*

Note that extra packets are not used in the analysis of the algorithms presented in Section 2 for the CIOQ model. Furthermore, note that Modification 3.1.1 takes place if Q_{ij}^* is not empty in $T[s]$ and Modification 3.1.2 takes place if Q_{ij}^* is empty in $T[s]$. Therefore, Modifications 3.1.1 and 3.1.2 cannot take place together in the same scheduling cycle.

Next, we show how the above modifications are used to show invariants that are different from the invariants shown in Section 2.1.

Lemma 8. *For any time t and any $i, j \in \{1, \dots, N\}$, the following inequalities hold:*

$$I1. |Q_{ij}(t)| \geq |Q_{ij}^*(t)|$$

$$I2. |C_{ij}^*(t)| \geq |C_{ij}(t)|$$

Proof. We show Inequalities I1 and I2 by a simple induction over the event sequence. Let the induction base be at time 0, i.e., before the sequence starts. All queues are empty at this time and all the inequalities hold. Assume now that they hold for any time up to some arrival, scheduling, or transmission event τ . Then we have to show that they still hold right after the event τ . Input and crossbar queues can change only in arrival and scheduling events. So, we assume that τ is either an arrival or a scheduling event.

Assume τ is an arrival event. Clearly, crossbar queues do not change in arrival events and thus I2 holds for this case. For I1, the only critical case is when the arriving packet is rejected by CGU and accepted by OPT. However, the input queue of CGU must be full in this case and thus I1 still holds.

Now, let τ be a scheduling event. Here, the only critical case for I1 is when CGU transfers a packet from Q_{ij} while OPT does not transfer any packet from Q_{ij}^* . However, either Q_{ij}^* is empty in this case or it cannot happen due to Modification 3.1.1. For I2, the first critical case is when CGU inserts a packet into C_{ij} while OPT does not insert any packet into C_{ij}^* . However, either C_{ij}^* is full in this case or it cannot happen due to Modification 3.1.2. The second critical case for I2 is when OPT transfers a packet from C_{ij}^* while CGU does not transfer any packet from C_{ij} . However, either C_{ij} is empty in this case or the size of C_{ij}^* does not decrease due to Modification 3.1.3. \square

In the following, we use $S_{T[s]}^*$ to denote the set of OPT's *normal* packets in the input subphase of $T[s]$. These are packets that OPT schedules through the normal channels, i.e., they are not privileged, and are part of the original input sequence, i.e., they are not extra. On the other hand, we use $S_{T[s]}$ to denote the set of packets that CGU schedules in the input subphase of cycle $T[s]$, i.e., from input queues to crossbar queues.

Lemma 9. *For any scheduling cycle $T[s]$, $|S_{T[s]}^*| \leq |S_{T[s]}|$.*

Proof. We want to show that in the input subphase of any scheduling cycle $T[s]$, if OPT transfers a normal packet from an input port i , CGU also transfers a packet from i .

Assume that OPT transfers a normal packet p from Q_{ij}^* (to C_{ij}^*) in $T[s]$. Thus, by I1 and I2 of Lemma 8, Q_{ij} is not empty and C_{ij} is not full in $T[s]$. (Note that OPT would not schedule a packet to a full crossbar queue, as all packets are of the same value.) Hence, CGU transfers a packet from either Q_{ij} or another $Q_{ij'}$ in $T[s]$. \square

Let $P_{T[s]}^*$ denote the set of OPT's privileged and extra packets (of either type) that occur in a scheduling cycle $T[s]$.

We consider the following mapping scheme from $P_{T[s]}^*$ to $S_{T[s]}$. For packets that are inserted in CGU's output queues, we use the notion of a *marked* packet. Initially, all packets are unmarked.

1. Let p be a privileged packet that is transferred by OPT from Q_{ij}^* in scheduling cycle $T[s]$. By Modification 3.1.1, CGU transfers a packet q from Q_{ij} in $T[s]$. Map p to q .
2. Let p be an extra packet of Type 1 that is inserted into C_{ij}^* in the input subphase of scheduling cycle $T[s]$. By Modification 3.1.2, CGU transfers a packet q into C_{ij} in $T[s]$. Map p to q .
3. Let p be an extra packet of Type 2 that is inserted into C_{ij}^* in the output subphase of scheduling cycle $T[s]$. By Modification 3.1.3, C_{ij} is not empty in $T[s]$, and OPT transfers a

packet p' to Q_j^* . Thus, Q_j^* is not full right before $T[s]$. Now, let q be the first unmarked packet in Q_j , i.e., the nearest to the queue's head. Map p to q and then mark q . Note that q can be the packet that CGU may insert into Q_j in $T[s]$.

Next, we show that the mapping scheme is feasible, i.e., each packet $p \in P_{T[s]}^*$ is mapped to a packet $q \in S_{T[s]}$. Clearly, Steps 1 and 2 are feasible. We show now that Step 3 is feasible as well. Let $M_j(t)$ denote the set of marked packets in Q_j at time t , for any $1 \leq j \leq N$. We first show the following lemma.

Lemma 10. *At any time t , $|M_j(t)| \leq |Q_j^*(t)|$.*

Proof. We show the lemma by induction over scheduling and transmission events. Clearly, $M_j(t)$ and $Q_j^*(t)$ can change only in these events.

Assume first that a transmission event occurs at t . The only critical scenario in this event is that OPT sends a packet from Q_j^* while CGU does not send a marked packet. If that happens, then either CGU sends an unmarked packet or it does not send any packet at all. The first case cannot happen while $|M_j(t)| > 0$ since marked packets are always at the front of the queue. The second case is safe because it implies that Q_j is empty and thus $|M_j(t)| = 0$.

Now, assume that a scheduling event occurs at t . The only critical scenario in this event is that a packet q is marked in Q_j while OPT does not insert any packet into Q_j^* . However, according to Step 3 of the mapping scheme, marking q implies that OPT transfers a packet p' from C_{ij}^* to Q_j^* . Thus, this scenario cannot happen in scheduling events. \square

Now, to show that Step 3 of the mapping scheme is feasible, we need to show that at least one packet is unmarked in Q_j in the scheduling cycle $T[s]$. For the sake of contradiction, assume that all packets in Q_j are marked or Q_j is empty in $T[s]$. The first thing that follows from this assumption is that CGU does not insert any packet into Q_j in $T[s]$ (because otherwise the inserted packet would be initially unmarked).

Recall from Step 3 that C_{ij} is not empty in $T[s]$. Thus, since no packet is inserted into Q_j , Q_j must be full in $T[s]$. Hence, since all packets are marked by assumption, $|M_j(t)| = B(Q_j)$, where t is the time right before $T[s]$. Thus, by Lemma 10, $|Q_j^*(t)| = B(Q_j^*)$ as well. However, this contradicts with the fact that OPT inserts p' into Q_j^* in $T[s]$. Hence, at least one packet is unmarked in Q_j in the scheduling cycle $T[s]$, and thus Step 3 is feasible.

Lemma 11. *For any scheduling cycle $T[s]$, $|P_{T[s]}^*| \leq 2|S_{T[s]}|$.*

Proof. As shown above, the mapping scheme is feasible for each scheduling cycle $T[s]$. So, it remains to show that at most two packets from $P_{T[s]}^*$ are mapped to any packet $q \in S_{T[s]}$.

Consider a packet $q \in S_{T[s]}$. Let Q_{ij} be the input queue from which q is transferred in the scheduling cycle $T[s]$. By the above mapping scheme, q may get mapped in at most 3 occasions during its entire lifespan: (i) by a privileged packet p that is transferred from Q_{ij}^* in $T[s]$, (ii) by an extra packet p' of Type 1 that is inserted into C_{ij}^* in the input subphase of $T[s]$, and (iii) by an extra packet of Type 2 that is inserted into $C_{i'j}^*$ in the output subphase of $T[s]$, with $i \neq i'$. However, as noted above, Modifications 3.1.1 and 3.1.2 cannot take place together in the same scheduling cycle and thus p and p' cannot exist together. Hence, at most two packets are mapped to q . \square

Now, as CGU does not preempt packets, each packet which CGU schedules in an input subphase must be eventually sent, and thus it contributes to the benefit of CGU. Hence, $\text{CGU}(\sigma) = \sum_{T[s]} |S_{T[s]}|$. Furthermore, note that $\text{OPT}(\sigma) = \sum_{T[s]} |S_{T[s]}^*| + |P_{T[s]}^*|$. Therefore, the proof of Theorem 3 follows immediately from Lemmas 9 and 11.

3.2 General-value case

For the case of arbitrary packet values, we present the Crossbar Preemptive Greedy algorithm (CPG) that is a variant of a 16.24-competitive algorithm given by Kesselman et al. [21].

Recall the notations $g_{ij}(t)$, $l_{ij}(t)$, and $l_j(t)$ that we used with algorithm PG (Section 2.2). Let $gc_{ij}(t)$ and $lc_{ij}(t)$ be the corresponding notations for crossbar queue C_{ij} , i.e., the packet with the greatest value and the packet with the least value, respectively, in C_{ij} at time t . Additionally, let $\beta \geq 1$ and $\alpha \geq 1$ be two parameters of the algorithm that will be determined later. If $\beta = \alpha$, our algorithm will be the same as the algorithm given in [21]. However, we show that to minimize the competitive ratio for this algorithm, these two parameters must take on different values.

Arrival and transmission phases of CPG are the same as ones of PG. In a scheduling phase, CPG works as follows.

- **Scheduling phase:** We divide every scheduling cycle $T[s]$ into two subphases.

- **Input subphase:** For each input port i , let J be defined as follows:

$$J = \left\{ j : |Q_{ij}(T[s])| > 0 \wedge \left(|C_{ij}(T[s])| < B(C_{ij}) \vee v(g_{ij}(T[s])) > \beta v(lc_{ij}(T[s])) \right) \right\} .$$

If $J \neq \emptyset$, choose Q_{ij} such that for all $j' \in J$,

$$j \in J \wedge v(g_{ij}(T[s])) \geq v(g_{ij'}(T[s])) .$$

Transfer $g_{ij}(T[s])$ to C_{ij} . If $|C_{ij}(T[s])| = B(C_{ij})$, preempt $lc_{ij}(T[s])$ first. In this case we also say that $g_{ij}(T[s])$ causes the preemption of $lc_{ij}(T[s])$ or, to be more concise, that $g_{ij}(T[s])$ preempts $lc_{ij}(T[s])$.

- **Output subphase:** For each output queue Q_j , choose a crossbar queue C_{ij} such that for all $i' \neq i$,

$$|C_{ij}(T[s])| > 0 \wedge v(gc_{ij}(T[s])) \geq v(gc_{i'j}(T[s])) .$$

If the following condition is satisfied

$$|Q_j(T[s])| < B(Q_j) \vee v(gc_{ij}(T[s])) > \alpha v(l_j(T[s])),$$

transfer $gc_{ij}(T[s])$ to Q_j . If $|Q_j(T[s])| = B(Q_j)$, preempt $l_j(T[s])$ first. Again, in this case we also say that $gc_{ij}(T[s])$ preempts $l_j(T[s])$.

Note that all ties in CPG are broken arbitrarily.

Theorem 4. *There is a choice for β and α such that the competitive ratio of CPG is at most 14.83 for any speedup.*

The analysis of CPG is carried out in a similar way as PG. We extend Assumptions A1 - A3 to include crossbar queues as well, and modify OPT in a slightly different way. Specifically, at the end of each scheduling cycle $T[s]$, i.e., immediately after OPT has performed its scheduling policy, we apply the following modifications on the configurations of OPT:

Modification 3.2.1. *Suppose that CPG transfers a packet from Q_{ij} and OPT does not transfer any packet from Q_{ij}^* in $T[s]$. If Q_{ij}^* is not empty in $T[s]$, OPT sends the head packet p of Q_{ij}^* , i.e., the packet with the greatest value in Q_{ij}^* , directly out of the switch. In this case, p is called a privileged packet of Type 1 and contributes to the benefit of OPT.*

Modification 3.2.2. Suppose that OPT transfers a packet p to C_{ij}^* and CPG does not transfer any packet to C_{ij} in $T[s]$. If C_{ij} is not full in $T[s]$ or $v(p) > \beta v(lc_{ij}(T[s]))$, OPT sends p directly out of the switch. In this case, p is called a privileged packet of Type 2 and contributes to the benefit of OPT.

Modification 3.2.3. Suppose that CPG transfers a packet from C_{ij} and OPT does not transfer any packet from C_{ij}^* in $T[s]$. If C_{ij}^* is not empty in $T[s]$, OPT sends the head packet p of C_{ij}^* , i.e., the packet with the greatest value in C_{ij}^* , directly out of the switch. In this case, p is called a privileged packet of Type 3 and contributes to the benefit of OPT.

Notice that Modifications 3.2.1 and 3.2.2 occur in the input subphase of $T[s]$, while Modification 3.2.3 occurs in the output subphase.

The following lemma extends Lemma 4 to include crossbar queues. We similarly use $\gamma_{ij}(k, t)$ (resp. $\gamma_{ij}^*(k, t)$) to denote the packet at position k in C_{ij} (resp. C_{ij}^*) at time t .

Lemma 12. For any $i, j \in \{1, \dots, N\}$ and any time t , the following holds:

- I1. $|Q_{ij}^*(t)| \leq |Q_{ij}(t)|$ and $v(\delta_{ij}^*(k, t)) \leq v(\delta_{ij}(k, t))$, for any position $k = 1, \dots, |Q_{ij}^*(t)|$
- I2. $|C_{ij}^*(t)| \leq |C_{ij}(t)|$ and $v(\gamma_{ij}^*(k, t)) \leq \beta v(\gamma_{ij}(k, t))$, for any position $k = 1, \dots, |C_{ij}^*(t)|$
- I3. $|Q_j^*(t)| \leq |Q_j(t)|$ and $v(\delta_j^*(k, t)) \leq \alpha\beta v(\delta_j(k, t))$, for any position $k = 1, \dots, |Q_j^*(t)|$

Proof. We show I1–I3 by induction over the event sequence. Let the induction base be at time 0, i.e., before the sequence starts. All queues are empty at this time and thus I1–I3 hold. Assume now that they hold for any time up to some arrival, scheduling, or transmission event τ . Then we have to show that they still hold right after the event τ . Let t' be a time just before event τ (but after the preceding event) and let t be a time just after τ (but before the following event). In other words, we assume that I1–I3 hold up to time t' and want to argue that they also hold at time t . In the following, we will argue only for I2 and I3. The argument for I1 is the same as in the proof of Lemma 4.

Before we start with I2, we say that a packet $p \in C_{ij}^*(\bar{t})$ is in a *legal alignment* if p is aligned at time \bar{t} to a packet $q \in C_{ij}(\bar{t})$ with $v(p) \leq \beta v(q)$. Clearly, it suffices to show that any packet $p \in C_{ij}^*(t)$ is in a legal alignment. We distinguish between two cases:

Case I2.1 $p \in C_{ij}^*(t')$. Thus, by induction, p is aligned at t' to a packet $q \in C_{ij}(t')$ with $v(p) \leq \beta v(q)$. We need to show in this case that p either remains in the same alignment at t or it changes to another legal alignment.

The only critical case is when p moves one step back at t (other cases are the same as in the proof of Lemma 4). In this case, a packet p' must be inserted in front of p , implying that $v(p) \leq v(p')$. Here, we distinguish between two cases: (i) CPG transfers a packet r to C_{ij} and (ii) CPG does not transfer any packet to C_{ij} . Let k denote the position of the alignment (p, q) at time t' . In case (i), due to I1 of this lemma, it must hold that $v(p') \leq v(r)$. Now, notice that either (1) r is inserted in a position $k' \leq k$, and thus p will be aligned again with q at t , or (2) r is inserted in a position $k' > k$, and thus p will be aligned with some packet q' at t . Clearly, the second case implies that $v(r) \leq v(q')$. Since $v(p) \leq v(p') \leq v(r)$, $v(p) \leq v(q') \leq \beta v(q')$. Hence, p is in a legal alignment in either case.

In case (ii), C_{ij} must be full at t and $v(p') \leq \beta v(lc_{ij}(t))$; otherwise, due to Modification 3.2.2, OPT would not insert any packet in C_{ij}^* . Thus, p must be aligned with some packet q' at t . Clearly, $v(lc_{ij}(t)) \leq v(q')$. Thus, $v(p) \leq v(p') \leq \beta v(q')$. Hence, p makes a legal alignment with q' .

Case I2.2 $p \notin C_{ij}^*(t')$. Thus, p is a new packet that is inserted in C_{ij}^* at time t . Again, we distinguish between two cases: (i) CPG transfers a packet r to C_{ij} , and (ii) CPG does not transfer any packet to C_{ij} . In case (ii), C_{ij} must be full at t and $v(p) \leq \beta v(l_{C_{ij}}(t))$; otherwise, due to Modification 3.2.2, OPT would not insert any packet into C_{ij}^* . Thus, p must be aligned with a packet q at t . Clearly, $v(l_{C_{ij}}(t)) \leq v(q)$. Thus, $v(p) \leq \beta v(q)$. Hence, p makes a legal alignment with q .

Now, consider case (i). Due to I1 of this lemma, it must hold that $v(p) \leq v(r)$. Let k denote the position at which p is inserted. If $k = 1$, p is aligned with the most valuable packet in C_{ij} at t . Since r is in C_{ij} at time t , p is aligned with a packet of value at least $v(r) \geq v(p)$. Now suppose $k > 1$. Let p' be the packet that is directly in front of p at t . Clearly, $p' \in C_{ij}^*(t')$ and $v(p) \leq v(p')$. Furthermore, let q' be the packet aligned with p' at time t' . Thus, $v(p) \leq v(p') \leq \beta v(q')$. Additionally, let q be the packet at position k in C_{ij} at time t' (assume $q = \emptyset$ if this is an empty position in C_{ij}).

Notice that (1) r is inserted in position k , and thus p will be aligned with r at t , (2) r is inserted in a position $k' < k$, and thus p will be aligned with q' at t , or (3) r is inserted in a position $k' > k$, and thus p will be aligned with q at t . Clearly, the last case implies that $q \neq \emptyset$ and that $v(q) \geq v(r) \geq v(p)$. Therefore, we have $v(p) \leq v(r)$ in the first case, $v(p) \leq \beta v(q')$ in the second, and $v(p) \leq v(q)$ in the third. Hence, p is in a legal alignment in any case.

Before we continue with I3, we say that a packet $p \in Q_j^*(\bar{t})$ is in a *legal alignment* if p is aligned at time \bar{t} to a packet $q \in Q_j(\bar{t})$ with $v(p) \leq \alpha \beta v(q)$. Clearly, it suffices to show that any packet $p \in Q_j^*(t)$ is in a legal alignment. We distinguish between two cases:

Case I3.1 $p \in Q_j^*(t')$. Thus, by induction, p is aligned at t' to a packet $q \in Q_j(t')$ with $v(p) \leq \alpha \beta v(q)$. We need to show in this case that p either remains in the same alignment at t or it changes to another legal alignment.

The only critical case is when p moves one step back at t (other cases are the same as in the proof of Lemma 4). In this case, a packet p' must be inserted in front of p , implying that $v(p) \leq v(p')$. Here, we distinguish between two cases: (i) CPG transfers a packet r to Q_j as well, and (ii) CPG does not transfer any packet to Q_j . Let k denote the position of the alignment (p, q) at time t' . In case (i), due to I2 of this lemma, it must hold that $v(p') \leq \beta v(r)$. Now, notice that either (1) r is inserted in a position $k' \leq k$, and thus p will be aligned again with q at t , or (2) r is inserted in a position $k' > k$, and thus p will be aligned with some packet q' at t . Clearly, the second case implies that $v(r) \leq v(q')$. Since $v(p) \leq v(p') \leq \beta v(r)$, $v(p) \leq \beta v(q')$. Hence, p is in a legal alignment in either case.

In case (ii), recall that p' is transferred by OPT from C_{ij}^* . Thus, due to I2 of this lemma, C_{ij} is not empty. Therefore, since CPG does not transfer any packet to Q_j in this case, Q_j must be full at t and $v(gc_{ij}(t)) \leq \alpha v(l_j(t))$. Since $v(p') \leq \beta v(gc_{ij}(t))$ (again due to I2), $v(p') \leq \alpha \beta v(l_j(t))$. Now, since Q_j is full at t , p must be aligned with some packet q' at t . Clearly, $v(l_j(t)) \leq v(q')$. Thus, $v(p) \leq v(p') \leq \alpha \beta v(q')$. Hence, p makes a legal alignment with q' .

Case I3.2 $p \notin Q_j^*(t')$. Thus, p is a new packet that is inserted in the queue at time t . Again, we distinguish between two cases: (i) CPG transfers a packet r to Q_j , or (ii) CPG does not transfer any packet to Q_j . In case (ii), recall that p is transferred by OPT from C_{ij}^* . Thus, due to I2 of this lemma, C_{ij} is not empty. Therefore, since CPG does not transfer any packet to Q_j in this case, Q_j must be full at t and $v(gc_{ij}(t)) \leq \alpha v(l_j(t))$. Since $v(p) \leq \beta v(gc_{ij}(t))$ (again due to I2), $v(p) \leq \alpha \beta v(l_j(t))$. Now, since Q_j is full at t , p must be aligned with a packet q at t . Clearly, $v(l_j(t)) \leq v(q)$. Thus, $v(p) \leq \alpha \beta v(q)$. Hence, p makes a legal alignment with q .

Now, consider case (i). Due to I2 of this lemma, it must hold that $v(p) \leq \beta v(r)$. Let k denote the position at which p is inserted. If $k = 1$, p is aligned with the most valuable packet in Q_j at t . Since

r is in Q_j at time t , p is aligned with a packet of value at least $v(r) \geq v(p)/\beta$. Now suppose that $k > 1$. Let p' be the packet that is directly in front of p at t . Clearly, $p' \in Q_j^*(t')$ and $v(p) \leq v(p')$. Furthermore, let q' be the packet aligned with p' at time t' . Thus, $v(p) \leq v(p') \leq \alpha\beta v(q')$. Additionally, let q be the packet at position k in Q_j at time t' (assume $q = \emptyset$ if this is an empty position in Q_j).

Notice that (1) r is inserted in position k , and thus p will be aligned with r at t , (2) r is inserted in a position $k' < k$, and thus p will be aligned with q' at t , or (3) r is inserted in a position $k' > k$, and thus p will be aligned with q at t . Clearly, the last case implies that $q \neq \emptyset$ and that $v(r) \leq v(q)$, and thus $v(p) \leq \beta v(q)$. Therefore, we have $v(p) \leq \beta v(r)$ in the first case, $v(p) \leq \alpha\beta v(q')$ in the second, and $v(p) \leq \beta v(q)$ in the third. Hence, p is in a legal alignment in any case. \square

The following lemma extends the claim of Lemma 6 to crossbar queues concerning the feasibility of mapping privileged packets of type 2.

Lemma 13. *Assume that OPT transfers a packet p from Q_{ij}^* to C_{ij}^* in $T[s]$ and CPG does not transfer any packet to C_{ij} . If C_{ij} is not full in $T[s]$ or $v(p) > \beta v(lc_{ij}(T[s]))$, then CPG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p') \geq v(p)$.*

Proof. Assume that C_{ij} is not full in $T[s]$ or $v(p) > \beta v(lc_{ij}(T[s]))$. By I1 of Lemma 12, since OPT transfers p from Q_{ij}^* in $T[s]$, CPG must have at the head of Q_{ij} a packet r with $v(r) \geq v(p)$. Thus, if $v(p) > \beta v(lc_{ij}(T[s]))$, then it must also hold that $v(r) > \beta v(lc_{ij}(T[s]))$. Hence, C_{ij} is not full in $T[s]$ or $v(r) > \beta v(lc_{ij}(T[s]))$, and therefore r is eligible to be transferred to C_{ij} . Nevertheless, as CPG does not transfer any packet to C_{ij} , another eligible packet p' must be transferred from another input queue $Q_{ij'}$, where $j' \neq j$. Obviously, as CPG preferred p' over r , it must hold that $v(p') \geq v(r)$, and hence $v(p') \geq v(p)$. \square

Now, recall I3 of Lemma 12. It implies that if OPT sends a packet of value v from some output queue at some time, CPG must send a packet of value at least $v/(\alpha\beta)$ from the same output queue at the same time. Let S (resp. S^*) denote the set of all packets that CPG (resp. OPT) sends from output queues. Thus,

$$\sum_{p \in S^*} v(p) \leq \alpha\beta \sum_{p \in S} v(p) .$$

Moreover, let P^* denote the set of all privileged packets of all types, which OPT sends directly out of the switch. The next lemma shows that

$$\sum_{p \in P^*} v(p) \leq \frac{2\alpha\beta + \alpha\beta(\beta - 1)}{(\alpha - 1)(\beta - 1)} \sum_{p \in S} v(p) .$$

Thus, we can conclude the competitive ratio of CPG as

$$\text{OPT}(\sigma) \leq \left(\alpha\beta + \frac{2\alpha\beta + \alpha\beta(\beta - 1)}{(\alpha - 1)(\beta - 1)} \right) \text{CPG}(\sigma) .$$

It can be verified that this competitive ratio is minimized when $\beta = \frac{\rho^2 + \rho + 4}{3\rho}$, where $\rho = (19 + 3\sqrt{33})^{1/3}$ and $\alpha = \frac{2}{(\beta - 1)^2}$. The resulting competitive ratio is $\frac{(\chi + 4) \cdot \rho^2 + (\chi + 16) \cdot \rho + 56}{12} \approx 14.83$, where $\chi = (19 - 3\sqrt{33})$.

Lemma 14. *The following inequality holds:*

$$\sum_{p \in P^*} v(p) \leq \frac{2\alpha\beta + \alpha\beta(\beta - 1)}{(\alpha - 1)(\beta - 1)} \sum_{p \in S} v(p) .$$

Proof. We consider the following mapping scheme:

1. Let p be a privileged packet of Type 1 that is sent from Q_{ij}^* in $T[s]$. By Modification 3.2.1, CPG transfers a packet p' from Q_{ij} in $T[s]$, and by I1 of Lemma 12, $v(p) \leq v(p')$. Map p to p' .
2. Let p be a privileged packet of Type 2 that is sent from Q_{ij}^* in $T[s]$. By Lemma 13, CPG transfers a packet p' from $Q_{ij'}$ in $T[s]$ with $j' \neq j$ and $v(p) \leq v(p')$. Map p to p' .
3. Let p be a privileged packet of Type 3 that is sent from C_{ij}^* in $T[s]$. By Modification 3.2.3, CPG transfers a packet p' from C_{ij} in $T[s]$, and by I2 of Lemma 12, $v(p) \leq \beta v(p')$. Map p to p' .
4. Let q be a packet that is preempted by CPG from a crossbar or an output queue due to accepting another packet p' . For each privileged packet p that is mapped to q , re-map p to p' .

As shown above, this mapping scheme is feasible, i.e., each packet $p \in P^*$ is mapped to a packet $p' \in S$. Now, it remains to show that the total value of privileged packets that are mapped to each packet $p' \in S$ is at most $\frac{2\alpha\beta + \alpha\beta(\beta-1)}{(\alpha-1)(\beta-1)}v(p')$.

For any packet $p' \in S$, p' can get mapped in four cases: (1) when it is scheduled in an input subphase, (2) when it preempts a packet from a crossbar queue, (3) when it is scheduled in an output subphase, and (4) when it preempts a packet from an output queue. We first consider cases (1) and (2). Assume that p' is scheduled from $Q_{ij'}$ to $C_{ij'}$ in the input subphase t . Now, assume that OPT transfers a packet from Q_{ij}^* to C_{ij}^* in the same time. Clearly, if $j \neq j'$, a privileged packet, say p_1 , of Type 1 can be sent from $Q_{ij'}$ in t , and the packet which OPT transfers from Q_{ij}^* can become a privileged packet, say p_2 , of Type 2. Hence, at most two privileged packets may be sent in t from each input port i . Since privileged packets of Types 1 and 2 are mapped only to packets that are transferred by CPG from the same input port during the same input subphase, only p_1 and p_2 can be mapped to p' in t . Furthermore, as shown in the mapping scheme above, the value of any of these privileged packets is at most the value of p' . Thus, the total value of privileged packets that are mapped to p' when it is scheduled in an input subphase is at most $2v(p')$.

Assume now that p' preempts a packet from $C_{ij'}$. Using the same argument of preemption chains in the proof of Lemma 7, we can show that the total value of privileged packets that are mapped to p' when it preempts a packet from $C_{ij'}$ is at most $\frac{2\beta}{\beta-1}v(p')$.

Now, we consider cases (3) and (4). In case (3), p' is scheduled in an output subphase t to the output queue Q'_j . Additionally, assume that OPT does not transfer any packet from $C_{ij'}$ in t . Clearly, a privileged packet, say p_3 , of Type 3 will be sent in this case from $C_{ij'}$. Since privileged packets of Type 3 are mapped only to packets that are transferred by CPG from the same crossbar queue in the same output subphase, only p_3 is mapped to p' in t . Furthermore, as shown in the mapping scheme above, the value of p_3 is at most β times the value of p' . Thus, the total value of privileged packets that are mapped to p' when it is scheduled in an output subphase is at most $(\beta + 2\beta/(\beta-1))v(p')$.

Finally, assume that p' preempts a packet from Q'_j . Again, using the same argument of preemption chains in the proof of Lemma 7, we can show that the total value of privileged packets that are mapped to p' when it preempts a packet from Q'_j is at most

$$\frac{\alpha}{\alpha-1} \cdot \left(\beta + \frac{2\beta}{\beta-1} \right) v(p') = \frac{2\alpha\beta + \alpha\beta(\beta-1)}{(\alpha-1)(\beta-1)} v(p') .$$

□

4 Conclusion

In this paper, we analyze online algorithms that are both competitive and efficient for the problem of packet scheduling in two closely related models of network switches, the CIOQ switches and the buffered crossbar switches. For unweighted packets in the CIOQ model, we give a faster algorithm that achieves the best known competitive ratio of 3. In the buffered crossbar model, we also show 3-competitiveness, improving the previously known ratio of 4. For weighted packets, we show 5.83- and 14.83-competitive algorithms for the CIOQ and buffered crossbar switches, respectively, which improves upon the previously known ratios of 6 and 16.24.

Despite the considerable interest that the switching problem in the CIOQ and buffered crossbar models has received, no result is known on any randomized algorithm in these models. Furthermore, we are not aware of any deterministic lower bounds that are especially constructed for these models.

As we show in Section 1.2, several lower bounds that are all strictly below 2 are known for the IQ model and they apply also to the CIOQ and buffered crossbar models. However, the gap between these lower bounds and the upper bounds we show in this work is still quite significant, especially in the general-value case where we show an upper bound of 5.83 in the CIOQ model and 14.83 in the buffered crossbar model. Due to the complex interaction between input and output queues in these models, we consider this problem as one of the most challenging open problems in the area of buffer management—it is already intriguing for us whether a lower bound that is only $2 + \epsilon$ is attainable, even in the unit-value case of the CIOQ model.

Moreover, it is still an open problem whether the upper bounds shown in this paper are tight for the given algorithms. When applied on the IQ model (i.e., $N \times 1$ switches with speedup 1), our algorithms GM and PG become the same algorithms given by [6] and [5], respectively, and for those algorithms asymptotic lower bounds are known: 2 for the unit value case and 3 for the general value case.

Furthermore, we notice that all the results presented in this paper for the CIOQ and buffered crossbar models can be generalized to an $N \times M$ switch, where N and M are not necessarily equal. The focus in the literature on an $N \times N$ architecture seems to be for practical reasons only.

Finally, from a practical point of view, choosing the right values for the parameters β and α of algorithms PG and CPG can be better based on a prior knowledge of the packet sequence. For example, and roughly speaking, the competitive ratio of PG consists of two terms, β and $2\beta/(\beta - 1)$: The first term corresponds to the scenario where PG admits to output queues packets which are considered negligible by OPT, and thus whenever PG sends a packet of this kind, OPT sends a packet of a larger value (up to β times the value of the packet neglected by OPT). Therefore, if the arrival rate of such large packets is significantly greater than the arrival rate of small packets, then β should be chosen sufficiently small. On the other hand, the second term corresponds to the scenario where PG excessively preempts packets from output queues while buffers can accommodate most of the packets. Thus, if it is more likely to have this scenario in practice, β should be chosen sufficiently large.

References

- [1] William Aiello, Alex Kesselman, and Yishay Mansour. Competitive buffer management for shared-memory switches. *ACM Transactions on Algorithms*, 5(1):Article 3, 2008.
- [2] Kamal Al-Bawani and Alexander Souza. Buffer overflow management with class segregation. *Information Processing Letters*, 113(4):145–150, 2013.

- [3] Susanne Albers and Markus Schmidt. On the performance of greedy algorithms in packet buffering. *SIAM Journal on Computing*, 35(2):278–304, 2006.
- [4] Yossi Azar and Arik Litichevsky. Maximizing throughput in multi-queue switches. *Algorithmica*, 45(1):69–90, 2006.
- [5] Yossi Azar and Yossi Richter. The zero-one principle for switching networks. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pages 64–71, 2004.
- [6] Yossi Azar and Yossi Richter. Management of multi-queue switches in QoS networks. *Algorithmica*, 43:81–96, 2005.
- [7] Yossi Azar and Yossi Richter. An improved algorithm for CIOQ switches. *ACM Transactions on Algorithms*, 2(2):282–295, 2006.
- [8] Marcin Bienkowski. An optimal lower bound for buffer management in multi-queue switches. *Algorithmica*, 68(2):426–447, 2014.
- [9] Marcin Bienkowski and Aleksander Madry. Geometric aspects of online packet buffering: An optimal randomized algorithm for two buffers. In *Proceedings of the 8th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 252–263, 2008.
- [10] Shang-Tse Chuang, Ashish Goel, Nick McKeown, and Balaji Prabhakar. Matching output queueing with a combined input output queued switch. *IEEE Journal on Selected Areas in Communications*, 17:1030–1039, 1999.
- [11] Shang-Tse Chuang, Sundar Iyer, and Nick McKeown. Practical algorithms for performance guarantees in buffered crossbars. In *Proceedings of the 24th IEEE Conference on Computer Communications (INFOCOM)*, pages 981–991, 2005.
- [12] Matthias Englert and Matthias Westermann. Lower and upper bounds on FIFO buffer management in QoS switches. *Algorithmica*, 53(4):523–548, 2009.
- [13] Matthias Englert and Matthias Westermann. Considering suppressed packets improves buffer management in QoS switches. *SIAM Journal on Computing*, 41(5):1166–1192, 2012.
- [14] Leah Epstein and Rob van Stee. Buffer management problems. *SIGACT News*, 35(3):58–66, 2004.
- [15] Patrick Th. Eugster, Alexander Kesselman, Kirill Kogan, Sergey I. Nikolenko, and Alexander Sirotkin. Essential traffic parameters for shared memory switch performance. In *Proceedings of the 22nd International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 61–75, 2015.
- [16] Patrick Th. Eugster, Kirill Kogan, Sergey I. Nikolenko, and Alexander Sirotkin. Shared memory buffer management for heterogeneous packet processing. In *Proceedings of the 34th IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 471–480, 2014.
- [17] Michael H. Goldwasser. A survey of buffer management policies for packet switches. *SIGACT News*, 41:100–128, 2010.
- [18] Toshiya Itoh and Noriyuki Takahashi. Competitive analysis of multi-queue preemptive QoS algorithms for general priorities. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E89-A(5):1186–1197, 2006.

- [19] Łukasz Jeż, Fei Li, Jay Sethuraman, and Clifford Stein. Online scheduling of packets with agreeable deadlines. *ACM Transactions on Algorithms*, 9(1):Article 5, 2012.
- [20] Alex Kesselman, Kirill Kogan, and Michael Segal. Packet mode and QoS algorithms for buffered crossbar switches with FIFO queuing. *Distributed Computing*, 23(3):163–175, 2010.
- [21] Alex Kesselman, Kirill Kogan, and Michael Segal. Best effort and priority queuing policies for buffered crossbar switches. *Chicago Journal of Theoretical Computer Science*, 2012(5):1–14, 2012.
- [22] Alex Kesselman, Kirill Kogan, and Michael Segal. Improved competitive performance bounds for CIOQ switches. *Algorithmica*, 63(1–2):411–424, 2012.
- [23] Alex Kesselman and Adi Rosén. Scheduling policies for CIOQ switches. *Journal of Algorithms*, 60(1):60–83, 2006.
- [24] Alex Kesselman and Adi Rosén. Controlling CIOQ switches with priority queuing and in multistage interconnection networks. *Journal of Interconnection Networks*, 9(1–2):53–72, 2008.
- [25] Koji M. Kobayashi, Shuichi Miyazaki, and Yasuo Okabe. A tight upper bound on online buffer management for multi-queue switches with bicodal buffers. *IEICE Transactions*, 91-D(12):2757–2769, 2008.
- [26] Koji M. Kobayashi, Shuichi Miyazaki, and Yasuo Okabe. Competitive buffer management for multi-queue switches in qos networks using packet buffering algorithms. *Theoretical Computer Science*, 675:27–42, 2017.
- [27] Fei Li, Jay Sethuraman, and Clifford Stein. Better online buffer management. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 199–208, 2007.
- [28] Sergey I. Nikolenko and Kirill Kogan. Single and multiple buffer processing. *Encyclopedia of Algorithms*, pages 1–9, 2014.
- [29] Vern Paxson and Sally Floyd. Wide-area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, 1995.
- [30] Markus Schmidt. Packet buffering: Randomization beats deterministic algorithms. In *Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 293–304, 2005.
- [31] Daniel Sleator and Robert Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
- [32] Andras Veres and Miklós Boda. The chaotic nature of TCP congestion control. In *Proceedings of the 19th IEEE Conference on Computer Communications (INFOCOM)*, pages 1715–1723, 2000.