

A Delay Analysis of Maximal Matching Switching with Speedup

Randy Cogill¹ Sanjay Lall²

Abstract

In this paper we analyze the average queue lengths in a combined input-output queued switch using a maximal size matching scheduling algorithm. We compare these average queue lengths to the average queue lengths achieved by an optimal switch. We model the cell arrival process as independent and identically distributed between time slots and uniformly distributed among input and output ports. For switches with many input and output ports, the backlog associated with maximal size matching with speedup 3 is no more than $3\frac{1}{3}$ times the backlog associated with an optimal switch. Moreover, this performance ratio rapidly approaches 2 as speedup increases.

1 Introduction

Although packet switches vary in their internal construction, the most common architecture for high performance switches is the *crossbar* switch. A crossbar switch contains N input lines and N output lines, where each input line meets each output line at a *crosspoint*. This is depicted in Figure 1. When a crosspoint connecting an input line and an output line is closed, cells may be transferred between this input and output. Crossbar switches operate with the constraint that, when routing cells from inputs to outputs, each input may only be connected to a single output, and each output may only be connected to a single input.

Switches are generally analyzed under a model where time is slotted, and only one cell may arrive at each input and depart from each output per time slot. Each arriving cell has a destination output port to which it must eventually be sent. Since multiple cells with the same output destination may arrive simultaneously at the input ports, switches require some form of buffering to store the cells which can not be immediately output. Buffered crossbar switches vary in their architecture, with the simplest being the output queued switch. In an output queued switch, buffers are placed at each output. All arriving cells are placed in their respective output queues in each time slot, and the output queues are served on a first come-first served (FCFS) basis. The average queue backlogs achieved by output queueing are minimum among all buffered crossbar switches. However, for a switch with N inputs and N outputs, output queueing requires that the as many as N

¹Department of Electrical Engineering, Stanford University, Stanford, CA 94305, U.S.A.

Email: rcogill@stanford.edu

²Department of Aeronautics and Astronautics, Stanford University, Stanford CA 94305-4035, U.S.A.

Email lall@stanford.edu

¹The first author was partially supported by a Stanford Graduate Fellowship.

^{1,2}Partially supported by the Stanford URI *Architectures for Secure and Robust Distributed Infrastructures*, AFOSR DoD award number 49620-01-1-0365.




Figure 1: A depiction of a crossbar switch, where input lines are connected to output lines by closing crosspoints.

rounds of scheduling must be performed by the switch in each time slot (consider the case when N cells destined for a single output arrive simultaneously). It is this requirement that makes output queueing infeasible for switches with many input and output ports.

Another alternative is the input queued switch, where cells are placed in buffers at the input ports in which they arrive. Delay performance of an input queued switch is heavily dependent on the service discipline used to serve the queues. It was shown in [1] that if input queues are served FCFS, the switch can only achieve 58% throughput. That is, suppose cells destined for output j arrive at input i at an average rate of λ/N per time slot for all i, j . Then average backlogs are bounded for all $\lambda < 1$ under output queueing, but average backlogs are only bounded for all $\lambda < 0.58$ when input queues are served FCFS. However, if we use a service discipline which schedules cells in each input queue based on their destinations, it is possible to achieve 100% throughput with input queueing. In particular, it was shown in [7] that 100% throughput is achieved by using a service discipline based on constructing maximum weight matchings (MWM) between inputs and outputs in each round of scheduling. This has the advantage over output queueing that only one round of scheduling is required per time slot. However, the algorithms required for computing maximum weight matchings are computationally expensive to implement. Also, the only known bounds on the average backlog under MWM are $O(N^2)$ [6], as opposed to output queueing which has an average backlog which increases as $O(N)$.

Combined input-output queued (CIOQ) switches are an alternative to purely input queued or purely output queued switches. Combined input-output queued switches place buffers at both the input ports and the output ports, and perform some moderate number $s \ll N$ rounds of scheduling per time slot. The number of rounds of scheduling s is commonly referred to the *speedup* of the switch. It was shown in [4] that 100% throughput can be achieved by using speedup $s = 2$ and a simple service discipline based on greedily constructing maximal size matchings between input ports and output ports in each round of scheduling. Unlike pure input queueing with MWM scheduling, maximal size matching schedules can be computed with low computational cost. Also, unlike pure output queueing, the speedup requirements do not increase with the size of the switch.

The purpose of this paper is to show that average backlog performance of a CIOQ switch using maximal matching scheduling with low speedup is comparable to that of an output queued switch. Several previous papers have addressed the problem of analyzing backlogs in combined input-output queued switches with speedup. In [2], it was shown that under

any traffic, an output queued switch can be *exactly* emulated by a CIOQ switch operation with speedup 2. However, the queueing discipline used in each round of scheduling has quite high computational cost. In [6], an upper bound on average backlog was proven for maximal matching scheduling with speedup 2 assuming IID Bernoulli traffic with uniform loading on input and output ports. Unlike the best known bound for MWM scheduling, the ratio between this upper bound and a lower bound on the backlog for an output queued switch is constant as N increases. However, this ratio becomes arbitrarily large as the arrival rate λ approaches 1. The same problem was considered and another upper bound on backlog was computed in [9]. There it was shown that the average backlog associated with maximal matching with speedup 2 is no more than 5 times the backlog associated with an output queued switch. In this paper we also consider switches under uniformly loaded IID traffic. We show that average backlog associated with maximal matching with speedup s gets arbitrarily close to 2 times the backlog associated with an output queued switch as s increases. Specifically, for a switch with many input and output ports, we show that for speedup $s = 3$, the backlog associated with maximal matching with speedup 3 is no more than $3\frac{1}{3}$ times the backlog associated with an output queued switch. This performance ratio rapidly approaches 2 as s increases.

2 Preliminaries

2.1 Maximal Size Matchings

Performing a round of scheduling in a crossbar switch can be thought of as constructing a matching in a bipartite graph G . This is shown in Figure 2. The vertices in G represent input and output ports, and there is an edge between vertices i and j if the queue at input port i contains a cell to be sent to output port j . Scheduling corresponds to choosing a collection of edges in G . That is, edge (i, j) is chosen if a cell is to be sent from input port i to output port j . The connectivity constraint imposed by the crossbar requires that the scheduled transfers correspond to a matching in the graph. A matching is a subgraph of G with the defining property that no two edges are incident on the same vertex.




Figure 2: A bipartite graph, where the heavy lines show a maximal size matching.

Scheduling algorithms for input and combined input-output queued switches essentially amount to various criteria for selecting matchings. In this paper we consider *maximal size matchings*. The main advantage to scheduling using maximal size matchings is that these

matchings can be computed very efficiently using a simple greedy algorithm. A maximal size matching is a subgraph $H \subset G$ with the property that if we add any edge in $G - H$ to H , then H is no longer a matching. The key property of maximal size matchings which is used in our later proofs is that if edge (i, j) is in G , then there is an edge in H incident to either vertex i or vertex j .

2.2 The Markov Chain Switch Model

Here we assume a traffic model in which at most one cell may arrive at each input in a single time slot, and that cell arrivals at all time slots are independent and identically distributed. We let $A_{ij}(t) \in \{0, 1\}$ be the random variable giving the number of cells arriving at input i destined for output j in time slot t . For simplicity, here we consider the case where arrivals are uniformly and independently distributed across inputs and outputs. This implies that the first and second moments of $A_{ij}(t)$ are

$$\begin{aligned} E[A_{ij}(t)] &= \frac{\lambda}{N} && \text{for all } i, j \\ E[A_{ij}(t)^2] &= \frac{\lambda}{N} && \text{for all } i, j \\ E[A_{kj}(t)A_{lj}(t)] &= \frac{\lambda^2}{N^2} && \text{for all } j \text{ and all } k \neq l \\ E[A_{ik}(t)A_{il}(t)] &= 0 && \text{for all } i \text{ and all } k \neq l, \end{aligned}$$

where $0 \leq \lambda < 1$ is a parameter describing the traffic intensity. Let $D_{ij}(t) \in \{0, \dots, s\}$ denote the number of cells sent from input queue i to output queue j in time slot t , and let $E_j(t) \in \{0, 1\}$ be the number of cells served from output queue j in time slot t . Also, we let $X_{ij}(t)$ denote the number of cells in input queue i destined for output j in time slot t , and let $Y_j(t)$ denote the number of cells in output queue j in time slot t . These random variables satisfy

$$\begin{aligned} X_{ij}(t+1) &= X_{ij}(t) + A_{ij}(t) - D_{ij}(t) \\ Y_j(t+1) &= Y_j(t) + \sum_{i=1}^N D_{ij}(t) - E_j(t). \end{aligned}$$

Throughout this paper, we will occasionally write these quantities in lowercase when simply referring to feasible values that they may take.

We will consider the problem of controlling the system to regulate the steady-state average per-period backlog in the input and output queues,

$$\lim_{t \rightarrow \infty} \frac{1}{t+1} \sum_{k=0}^t E \left[\sum_{i=1}^N \sum_{j=1}^N X_{ij}(k) + \sum_{j=1}^N Y_j(k) \middle| X(0), Y(0) \right].$$

Under a maximal matching scheduling policy, $D(t)$ and $E(t)$ depend only on $X(t)$ and $Y(t)$. When this is the case, this system evolves as a Markov chain and we can use the following lemma to bound the average per-period backlog. This lemma is a special case of a more general result shown in [3]. Results similar to the lemma below also appear, for example, in [8].

Lemma 1. Consider a Markov chain X with state space \mathcal{X} . For the cost function $r : \mathcal{X} \rightarrow \mathbb{R}$, let

$$J(z) = \lim_{t \rightarrow \infty} \frac{1}{t+1} \sum_{k=0}^t E[r(X(k)) | X(0) = z].$$

For any $h_U : \mathcal{X} \rightarrow \mathbb{R}$ such that $\inf_{x \in \mathcal{X}} \{h_U(x)\} > -\infty$,

$$J(z) \leq \sup_{x \in \mathcal{X}} \{r(x) + E[h_U(X(t+1)) | X(t) = x] - h_U(x)\}$$

for all $z \in \mathcal{X}$.

Proof. Let $y = \inf_{x \in \mathcal{X}} \{h_U(x)\}$,

$$\Delta_U(x) = E[h_U(X(t+1)) | X(t) = x] - h_U(x),$$

and

$$\beta_U = \sup_{x \in \mathcal{X}} \{r(x) + \Delta_U(x)\}.$$

For all $t \geq 0$,

$$\begin{aligned} \frac{1}{t+1} \sum_{k=0}^t E[r(X(k)) | X(0)] &= \frac{1}{t+1} \sum_{k=0}^t E[r(X(k)) + \Delta_U(X(k)) | X(0)] \\ &\quad + \frac{h_U(X(0)) - y}{t+1} - \frac{E[h_U(X(t+1)) | X(0)] - y}{t+1} \\ &\leq \beta_U + \frac{h_U(X(0)) - y}{t+1}. \end{aligned}$$

Clearly

$$\lim_{t \rightarrow \infty} \frac{h_U(X(0)) - y}{t+1} = 0$$

for all $z \in \mathcal{X}$, hence

$$\lim_{t \rightarrow \infty} \frac{1}{t+1} \sum_{k=0}^t E[r(X(k)) | X(0)] \leq \beta_U. \quad \blacksquare$$

3 Main Result

Our overall goal is to show that maximal matching scheduling with a speedup of s keeps the average backlog relatively close to the backlog achieved by an output queued switch. Specifically, we will: (i) compute an upper bound on the backlog associated with maximal matching with speedup s , (ii) compute the backlog associated with an output queued switch, and (iii) compute a bound on the ratio of these quantities.

Our first step will be to compute the upper bound on the average backlog. The following lemma will be used in the proof of the bound. Here we will define the quantities

$$\begin{aligned} \alpha_1(s) &= \frac{1}{(s-1)^2 - 1}, \\ \alpha_2(s) &= \frac{s-1}{(s-1)^2 - 1}, \end{aligned}$$

and

$$Q_{ij}(\lambda, d, e, s) = (1 + \alpha_1(s) + \alpha_2(s))\lambda - \left(\alpha_1(s) \sum_{l=1}^N d_{lj} + e_j + \alpha_2(s) \sum_{l=1}^N d_{il} \right)$$

which will be used throughout the rest of this paper.

Lemma 2. *For a CIOQ switch operating at speedup $s \geq 3$ using maximal matching scheduling,*

$$Q_{ij}(\lambda, d, e, s)x_{ij} \leq (\lambda - 1)x_{ij}$$

for all $\lambda \leq 1$ and all feasible values of x_{ij} , d , and e .

Proof. When operating at speedup s , s rounds of scheduling occur in each time slot. When using maximal matching scheduling, if there is a cell in input queue i destined for output j at the start of a round of scheduling, then either a cell is removed from input i or a cell is sent to output j in that round. Also, if a cell is sent to output queue j , then output queue j is served at the end of the time slot.

It is clear that the lemma holds if $x_{ij} = 0$. To prove the lemma for $x_{ij} > 0$, we will consider three cases:

1. When $\sum_{l=1}^N d_{lj} = s$ and $\sum_{l=1}^N d_{il} = 0$,

$$\begin{aligned} \alpha_1(s) \sum_{l=1}^N d_{lj} + e_j + \alpha_2(s) \sum_{l=1}^N d_{il} &= \alpha_1(s)s + 1 \\ &= \frac{s-1+1}{(s-1)^2-1} + 1 \\ &= 1 + \alpha_1(s) + \alpha_2(s). \end{aligned}$$

2. When $\sum_{l=1}^N d_{lj} = 0$ and $\sum_{l=1}^N d_{il} = s$,

$$\begin{aligned} \alpha_1(s) \sum_{l=1}^N d_{lj} + e_j + \alpha_2(s) \sum_{l=1}^N d_{il} &\geq \alpha_2(s)s \\ &= \frac{(s-1)(s-1+1) - 1 + 1}{(s-1)^2-1} \\ &= 1 + \alpha_1(s) + \alpha_2(s). \end{aligned}$$

3. When $\sum_{l=1}^N d_{lj} > 0$ and $\sum_{l=1}^N d_{il} > 0$, it is sufficient to consider the case where $\sum_{l=1}^N d_{lj} = \sum_{l=1}^N d_{il} = 1$ since $\alpha_1(s) \geq 0$ and $\alpha_2(s) \geq 0$. In this case,

$$\alpha_1(s) \sum_{l=1}^N d_{lj} + e_j + \alpha_2(s) \sum_{l=1}^N d_{il} = 1 + \alpha_1(s) + \alpha_2(s).$$

Note that if $x_{ij} > 0$ and

$$\sum_{l=1}^N x_{lj} + \sum_{l=1}^N x_{il} - x_{ij} < s,$$

then the total number of cells either in input queue i or destined for output queue j is less than s . However, in this case at least one cell must be sent from input queue i to output queue j , implying that $\sum_{l=1}^N d_{lj} > 0$ and $\sum_{l=1}^N d_{il} > 0$. ■

Now we are ready to prove the upper bound on the backlog associated with maximal matching scheduling. We will let J_{MMs} denote the average per-period backlog associated with maximal matching scheduling with speedup s .

Theorem 3. *A CIOQ switch operating with speedup s using a maximal matching scheduling policy has average per-period backlog satisfying*

$$J_{MMs} \leq \frac{(k_1(s) \left(1 - \frac{1}{N}\right) \lambda^2 + k_2(s)\lambda - k_3(s)\lambda^2) N}{2(1 - \lambda)},$$

where

$$\begin{aligned} k_1(s) &= 1 + \alpha_1(s) \\ k_2(s) &= 2 + (\alpha_1(s) + \alpha_2(s))(s + 1) \\ k_3(s) &= 2 + 2\alpha_1(s) + 2\alpha_2(s). \end{aligned}$$

Proof. We prove this bound using Lemma 1 with

$$h_U(x, y) = h_1(x) + h_2(x) + h_3(x, y),$$

where

$$\begin{aligned} h_1(x) &= \frac{\alpha_1(s)}{2(1 - \lambda)} \sum_{j=1}^N \left(\left(\sum_{i=1}^N x_{ij} \right)^2 + (s - 2\lambda) \sum_{i=1}^N x_{ij} \right) \\ h_2(x) &= \frac{\alpha_2(s)}{2(1 - \lambda)} \sum_{i=1}^N \left(\left(\sum_{j=1}^N x_{ij} \right)^2 + (s - 2\lambda) \sum_{j=1}^N x_{ij} \right) \\ h_3(x, y) &= \frac{1}{2(1 - \lambda)} \sum_{j=1}^N \left(\sum_{i=1}^N x_{ij} + y_j \right)^2 + \frac{1 - 2\lambda}{2(1 - \lambda)} \sum_{j=1}^N \left(\sum_{i=1}^N x_{ij} + y_j \right) \end{aligned}$$

Since h_U is quadratic with positive second order coefficients, it is clear that

$$\inf_{x, y} \{h_U(x, y)\} > -\infty,$$

satisfying the required condition of Lemma 1. Let

$$\Delta_i(x, y, d) = E[h_i(X(t + 1), Y(t + 1)|x, y, d)] - h_i(x, y)$$

denote the expected drift in h_i when in state (x, y) and action d is taken.

$$\begin{aligned}
\Delta_1(x, y, d) &= \frac{\alpha_1(s)}{1-\lambda} \sum_{j=1}^N \left(E \left[\sum_{l=1}^N (A_{lj} - d_{lj}) \right] \right) \sum_{i=1}^N x_{ij} \\
&\quad + \frac{\alpha_1(s)}{2(1-\lambda)} \sum_{j=1}^N E \left[\left(\sum_{l=1}^N (A_{lj} - d_{lj}) \right)^2 \right] \\
&\quad + \frac{\alpha_1(s)(s-2\lambda)}{2(1-\lambda)} \sum_{j=1}^N E \left[\sum_{l=1}^N (A_{lj} - d_{lj}) \right] \\
&\leq \frac{\alpha_1(s)}{1-\lambda} \sum_{i=1}^N \sum_{j=1}^N \left(\lambda - \sum_{l=1}^N d_{lj} \right) x_{ij} + \alpha_1(s) \frac{\left(\left(1 - \frac{1}{N}\right) \lambda^2 + (s+1)\lambda - 2\lambda^2 \right) N}{2(1-\lambda)},
\end{aligned}$$

where we used the fact that

$$\begin{aligned}
E \left[\left(\sum_{l=1}^N A_{lj} \right)^2 \right] &= \sum_{k=1}^N \sum_{l=1}^N E[A_{kj} A_{lj}] \\
&= \left(1 - \frac{1}{N} \right) \lambda^2 + \lambda.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\Delta_2(x, y, d) &= \frac{\alpha_2(s)}{1-\lambda} \sum_{i=1}^N \left(E \left[\sum_{l=1}^N (A_{il} - d_{il}) \right] \right) \sum_{j=1}^N x_{ij} \\
&\quad + \frac{\alpha_2(s)}{2(1-\lambda)} \sum_{i=1}^N E \left[\left(\sum_{l=1}^N (A_{il} - d_{il}) \right)^2 \right] \\
&\quad + \frac{\alpha_2(s)(s-2\lambda)}{2(1-\lambda)} \sum_{i=1}^N E \left[\sum_{l=1}^N (A_{il} - d_{il}) \right] \\
&\leq \frac{\alpha_2(s)}{2(1-\lambda)} \sum_{i=1}^N \sum_{j=1}^N \left(\lambda - \sum_{l=1}^N d_{il} \right) x_{ij} + \alpha_2(s) \frac{\left((s+1)\lambda - 2\lambda^2 \right) N}{2(1-\lambda)},
\end{aligned}$$

where we used the fact that

$$\begin{aligned}
E \left[\left(\sum_{l=1}^N A_{il} \right)^2 \right] &= \sum_{k=1}^N \sum_{l=1}^N E[A_{ik} A_{il}] \\
&= \lambda.
\end{aligned}$$

Also,

$$\begin{aligned}
\Delta_3(x, y) &= \frac{1}{1-\lambda} \sum_{j=1}^N \left(E \left[\sum_{l=1}^N A_{lj} - e_j \right] \right) \left(\sum_{i=1}^N x_{ij} + y_j \right) \\
&\quad + \frac{1}{2(1-\lambda)} \sum_{j=1}^N E \left[\left(\sum_{l=1}^N A_{lj} - e_j \right)^2 \right] \\
&\quad + \frac{1-2\lambda}{2(1-\lambda)} \sum_{j=1}^N E \left[\sum_{l=1}^N A_{lj} - e_j \right] \\
&= \frac{1}{1-\lambda} \sum_{j=1}^N (\lambda - e_j) \left(\sum_{i=1}^N x_{ij} + y_j \right) + \frac{\left((1 - \frac{1}{N}) \lambda^2 + 2\lambda - 2\lambda^2 \right) N}{2(1-\lambda)}.
\end{aligned}$$

Therefore,

$$\begin{aligned}
r(x, y, d) + \sum_{i=1}^3 \Delta_i(x, y, d) &\leq \sum_{i=1}^N \sum_{j=1}^N \left(1 + \frac{Q_{ij}(\lambda, d, e, s)}{1-\lambda} \right) x_{ij} \\
&\quad + \sum_{j=1}^N \left(1 + \frac{\lambda - e_j}{1-\lambda} \right) y_j \\
&\quad + \frac{(k_1(s) (1 - \frac{1}{N}) \lambda^2 + k_2(s) \lambda - k_3(s) \lambda^2) N}{2(1-\lambda)}.
\end{aligned}$$

From Lemma 2, we have

$$\sum_{i=1}^N \sum_{j=1}^N \left(1 + \frac{Q_{ij}(\lambda, d, e, s)}{1-\lambda} \right) x_{ij} \leq 0$$

for all values of x . Also, since $e_j = 1$ if $y_j > 0$,

$$\sum_{j=1}^N \left(1 + \frac{\lambda - e_j}{1-\lambda} \right) y_j = 0$$

for all values of y . Therefore,

$$\begin{aligned}
J_{MMs} &= \sup_{x, y} \{ r(x, y) + \Delta(x, y, d) \} \\
&\leq \frac{(k_1(s) (1 - \frac{1}{N}) \lambda^2 + k_2(s) \lambda - k_3(s) \lambda^2) N}{2(1-\lambda)}.
\end{aligned}$$

■

The previous theorem established an upper bound on J_{MMs} , the average per-period backlog under maximal matching with some fixed speedup s . We would now like to determine the expected per-period backlog associated with an output queued switch, which we will denote by J_{OQ} . This is a standard result, but is presented here to keep our treatment self-contained.

Lemma 4. *An output queued switch has*

$$J_{OQ} = \frac{\left(\left(1 - \frac{1}{N}\right)\lambda^2 + 2\lambda - 2\lambda^2\right)N}{2(1 - \lambda)}.$$

Proof. Output queue j is a discrete-time queue with queue with arrival process $A_{1j} + \dots + A_{Nj}$. By the Pollaczek-Khintchine formula, (see, for example, [5]) the average steady-state per-period backlog of output queue j is

$$\frac{E\left[\left(\sum_{i=1}^N A_{ij}(t)\right)^2\right] + \lambda - 2\lambda^2}{2(1 - \lambda)}.$$

Using the fact that

$$\begin{aligned} E\left[\left(\sum_{l=1}^N A_{lj}\right)^2\right] &= \sum_{k=1}^N \sum_{l=1}^N E[A_{kj}A_{lj}] \\ &= \left(1 - \frac{1}{N}\right)\lambda^2 + \lambda, \end{aligned}$$

we sum over all output queues to obtain

$$J_{OQ} = \frac{\left(\left(1 - \frac{1}{N}\right)\lambda^2 + 2\lambda - \lambda^2\right)N}{2(1 - \lambda)}.$$

■

The upper bound and the result of the previous lemma are now used to determine a bound on the performance ratio between maximal matching and output queueing.

Theorem 5. *The ratio of the average backlog under maximal matching scheduling with speedup $s \geq 3$ to the average backlog of an output queued switch satisfies*

$$\frac{J_{MMs}}{J_{OQ}} \leq \frac{N}{N-1} \left(\frac{2(s-1)^2 + (s-1)}{(s-1)^2 - 1} \right). \quad (1)$$

Proof. From Theorem 3 and Theorem 4 we have

$$\frac{J_{MMs}}{J_{OQ}} \leq \frac{k_1(s)\left(1 - \frac{1}{N}\right)\lambda + k_2(s) - k_3(s)\lambda}{2 - \left(1 + \frac{1}{N}\right)\lambda}.$$

By differentiating, it is straightforward to show that for $s \geq 3$ and $N \geq 2$, the previous expression is increasing in λ for $0 \leq \lambda \leq 1$. Therefore,

$$\begin{aligned} \frac{J_{MMs}}{J_{OQ}} &\leq \frac{k_1(s)\left(1 - \frac{1}{N}\right) + k_2(s) - k_3(s)}{1 - \frac{1}{N}} \\ &= \frac{1 + \alpha_1(s)s + \alpha_2(s)(s-1) - (1 + \alpha_1(s))\frac{1}{N}}{1 - \frac{1}{N}} \\ &\leq \frac{N}{N-1}(1 + \alpha_1(s)s + \alpha_2(s)(s-1)) \\ &= \frac{N}{N-1} \left(\frac{2(s-1)^2 + (s-1)}{(s-1)^2 - 1} \right) \end{aligned}$$

■

For large N , the performance ratio approaches 2 as s increases. Table 3 shows the value of this ratio for several low values of speedup.

s	3	4	5	8	15
$J_{\text{MMs}}/J_{\text{OQ}}$	3.36	2.65	2.42	2.20	2.10

Table 1: Values of (1) at several low values of speedup for a 128×128 switch.

4 Conclusions

In this paper we have analyzed the average backlogs in network switches using a maximal size matching scheduling policy with speedup. It is shown that switches using maximal matching with speedup achieve backlogs comparable to an optimal switch. For the sake of simplicity, we have focused on the case of IID arrivals with uniform loading on input and output ports. We believe that the performance bounds proven in this paper can be tightened when arrivals are time correlated, and this is a subject of future research.

References

- [1] T. Anderson, S. Owicki, J. Saxe, and C. Thacker. High speed switch scheduling for local area networks. *ACM Trans. Comp. Sys.*, 11(4):319–351, 1993.
- [2] S.T. Chuang, A. Goel, N. McKeown, and B. Prabhakar. Matching output queueing with a combined input-output queued switch. *IEEE INFOCOM 1999*, 3:1169–1178, 1999.
- [3] R. Cogill and S. Lall. Suboptimality bounds in stochastic control: A queueing example. *To appear in the Proceedings of the 2006 American Control Conf.*, 2006.
- [4] J. Dai and B. Prabhakar. The throughput of data switches with and without speedup. *IEEE INFOCOM 2000*, 2:556–564, 2000.
- [5] D. Gross and C. Harris. *Fundamentals of Queueing Theory*. John Wiley, New York, 1985.
- [6] E. Leonardi, M. Mellia, F. Neri, and A. Marsan. Bounds on delays and queue lengths in input-queued cell switches. *Journal of the ACM*, 50(4):520–550, 2003.
- [7] N. McKeown, V. Anantharan, and J. Walrand. Achieving 100% throughput in an input-queued switch. *IEEE INFOCOM 1996*, 1:296–302, 1996.
- [8] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
- [9] D. Shah. Maximal matching scheduling is good enough. *IEEE Globecom*, 22(1):3009–3013, 2003.