
The Odd-Even ATM Switch [†]

Christos KOLIAS^{††} and Leonard KLEINROCK^{††}, *Nonmembers*

SUMMARY This paper introduces and studies the performance of an $N \times N$ space-division, single-stage ATM switch with dual input-queueing. Each input port has two separate FIFO queues, an “odd” and an “even” queue. An incoming cell is stored at the input at either of two FIFOs according to its output port destination (output ports are also labeled as “odd” or “even”). Hence we call this scheme the *Odd-Even* switch. We compare the Odd-Even switch to the ordinary input-buffered switch and we find that it can achieve a remarkably higher performance, in terms of throughput, mean delay and cell loss. This is due to the fact that the Head-of-Line effect becomes less problematic under the Odd-Even switch. Our results are based on various traffic models. Finally, we compare the Odd-Even scheme to the Look-ahead (input “window”) policy.

key words: *Odd-Even, ATM Switching, Input-Queueing, Head-of-Line Blocking.*

1. Introduction

Asynchronous Transfer Mode (ATM) wide area networks are currently being deployed by telecommunications companies in order to support the forthcoming B-ISDN services. Various blocking and non-blocking, time-division and space-division, single stage and multi-stage ATM switches have been proposed and studied extensively, in order to provide high performance packet switching for integrated ATM transport. Some of these have been developed and tested under real-time traffic requirements and have eventually become available as commercial products.

Switch architectures are classified [7] according to their buffering methods (e.g., input buffers, output buffers), according to their switch fabric (i.e. crossbar, Batcher-Banyan, Starlite) or according to their internal switching (non-blocking, blocking).

In this paper we study an input-buffered, non-blocking, crossbar, ATM switch with N input ports and N output ports. Incoming cells are stored in (one of the two) dedicated buffers at the input ports, according to

their output port destination.

The model we present for this switch ignores any hardware design and specific implementation details; we are mainly concerned with the performance evaluation of the switching system. Achieved throughput and delay are the two main performance metrics of interest. That allows us to compare it in terms of performance with other switches that have different characteristics (i.e. arbitration policy) but with the same philosophy (i.e. input buffering).

The paper is organized as follows. Section 2 describes the Odd-Even switch along with our modeling assumptions. In section 3, we present some performance evaluation results for the Odd-Even switch while we compare it to the ordinary (simple) input-queueing switch - to which we refer as “regular” switch - under various traffic conditions. Section 4 extends the comparison of the Odd-Even model by considering also the Look-ahead scheme. Concluding remarks are provided in section 5.

2. The Odd-Even Switch

2.1 Switch Architecture

The switch under consideration uses input-queueing to buffer a cell before it accesses the switch fabric in the case where there is no available connection to the output port and there are buffers only at the inputs. Cell arbitration among the input buffers is needed in order to control the switch and initiates the contention resolution cycle. The switch has N inputs and N outputs, therefore the switch is of size N . What is novel about this switch (Figure 1) is the fact that each input is split (expanded) into two FIFO queues which we call *Odd* and *Even*; hence we call this switching system the *Odd-Even* ATM switch[‡]. By convention, we refer to outputs numbered as 1,3,5,... as odd, and to 2,4,6,... as even ports. An incoming cell destined to an odd (even) numbered output port joins the odd (even) queue of the input port to which it was fed and awaits its turn to be switched to its output. When more than one cell are present at the heads of the queues contending for the same output address only one can access

Manuscript received April 30, 1997.

Manuscript revised July 24, 1997.

[†]This research was supported by ARPA/CSTO under grant MDA-972-91-J-1011, Advanced Networking and Distributed Systems. The first author was also supported by the Lilian Voudouri Fellowship (Greece). Part of this paper was presented at the International Conference on Communications (ICC) '96, Dallas, Texas, June 23-27, 1996. Web URL: <http://millennium.cs.ucla.edu/~ck>

^{††}The authors are with the Computer Science Department, UCLA, Los Angeles, CA 90095, U.S.A.

[‡]We would like to thank Larry Roberts, CTO of *Connectware Inc.*, for introducing the Odd-Even model to us.




Fig. 1 An 8×8 Odd-Even ATM switch (two FIFO queues per input port).

that output while the others remain blocked until they are finally released in future time slots. This blocking phenomenon is known as *head-of-line (HOL)* blocking, a problem from which all input-buffered packet switches suffer. A queued cell that moves up to the head of its queue, after the previous HOL is switched, is called a *fresh* cell. Time is slotted and since cells are packets of fixed size, service time is deterministic and equal to a time slot (which is determined by the actual speed of the switch fabric) and which we assume here to be one time unit.

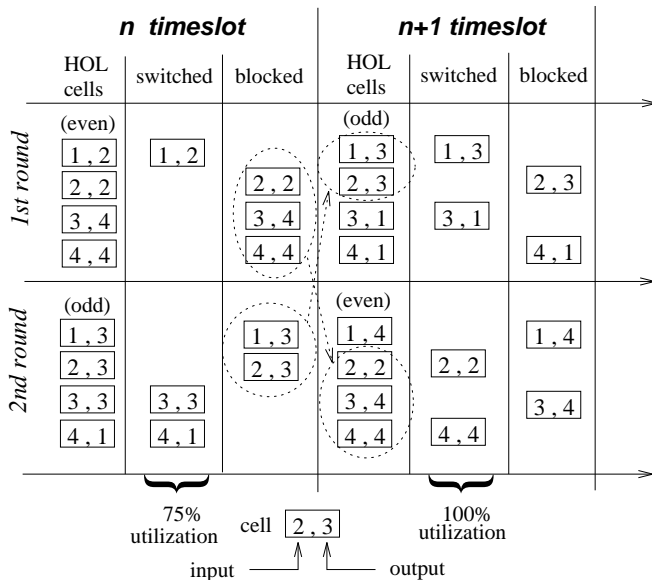


Fig. 2 Arbitration and contention resolution (random selection) example in a 4×4 Odd-Even switch.

As mentioned above, cell arbitration is employed in order to resolve any potential contention among the input buffers. The full contention resolution process occurs at the beginning of each time slot, immediately following the cell arrivals, and consists of two (very short) rounds. Arbitration during the first round involves the HOL cells at the even input queues. In the second round cells at the HOL of the odd input queues

contend for the odd output addresses. However, those input ports whose even queues could not access an output port in the first round - because they either lost the contention (and therefore were blocked) or simply did not have any HOL cells present - are allowed to participate in a contention among their odd queues in the subsequent second round (see example in Figure 2). Consequently, at each time slot, an input port always gets the opportunity to send a cell either from an even or an odd queue, but not more than one cell is allowed to be cleared, from any given input port, in both arbitration phases during the course of a time slot.

In order to maintain fairness we allow arbitration cycles to interchange order in every time slot (if it is even-odd in one time slot it is odd-even in the next one), otherwise it would be as if we introduced some kind of priorities. We use a random HOL selection policy*. Under a balanced traffic assumption, this arbitration strategy can guarantee stability and fairness. Other possible policies are the “longest queue selection”** and the “longest waiting HOL cell selection”*** - if there is a tie, the lowest numbered input queue is the winner. Scheduling among the HOL cells can further improve the throughput. We do not consider any such mechanisms or algorithms in our paper.

Clearly, the proposed architecture has no speed-up features. However, as we demonstrate in the next sections, comparing this switch architecture to the “regular” input-buffered switch it can increase rather significantly the output trunk utilization and decrease delays and switch cell loss. This is achieved as the Odd-Even scheme introduces a certain degree of parallelism. We now proceed with the assumptions supporting our model.

*When k cells contend for the same output at the beginning of a time slot one of the k is randomly chosen to be switched.

**The buffer with the longest queue (most stored cells) is selected.

***The HOL cell that has been waiting the longest at the HOL position is selected. This obviously requires timestamping.

2.2 Modeling Approach

The proposed switch is modeled as a multi-queue (input links), multi-server (output trunks), discrete-time queueing system. Cells arriving at the input links are simultaneously and independently routed to the appropriate output ports. Cells arrivals occur at the beginning of the time slot and incoming cells are queued accordingly in the Odd-Even FIFO queues. Cell departures from the switch occur in a concurrent, independent fashion.

The model ignores any flow or congestion control mechanism or implications by assuming that such control is accounted for in other parts of the network. We assume that the traffic management is efficient enough so that there is no cell loss inside the switch (though it may occur elsewhere as part of a policing mechanism) and that input buffers are infinitely large (we will later relinquish this assumption). Buffers are considered to be FIFO queues, thus preventing any resequencing of cells. It is clear from the arbitration policy, that the throughput of the Odd-Even model is at least as high as that of the regular switch. A packet switch where the input traffic intensity is the same for every input link and the output address of each incoming cell is assigned with equal probability to any output port is said to be characterized as a homogeneous system.

The *throughput* of the switch, denoted by γ and defined as the average number of cells successfully (i.e. not blocked) switched per unit of time (timeslot), will be used as the main performance evaluation measure. Switch throughput is obtained as the average utilization of the N servers, where each server's utilization is obtained as the percentage of time the server is busy. What actually presents interest is finding the maximum achievable switch throughput, γ_{max} . The maximum throughput (also known as saturation throughput) is the maximum applied input load than can be handled by the switch without causing any instability (the switch then operates at its maximum attainable limit). Input traffic exceeding that critical load threshold drives the system to unstable and overloaded situations (e.g., infinitely large queues and delays), while for input load less than γ_{max} throughput is equal to the input load. A throughput of, e.g., $\gamma_{max} = 0.5$ means that on average only half of the input ports route cells in a typical time slot and that throughput gives the probability that a cell at an HOL position gets served during a time slot, provided of course that all HOL positions are occupied, i.e. a heavy traffic situation. Another interpretation is that an HOL cell has to wait, again on average, for two ($1/\gamma_{max}$) timeslots, to be switched, assuming again a heavy traffic case. *Delay* is obtained as the total time experienced by a cell when it is fed into an input port until it is finally switched out from its desired output port. Thus, delay comprises of any

waiting time (queueing delay) and the (fixed) switching time (equal to a time slot). The time spent at the HOL position is considered part of the queueing delay.

Our performance study is carried under a variety of traffic patterns: non-bursty, bursty, and non-uniform. Since exact analysis is rather intractable, we present simulation results in order to evaluate the Odd-Even model^{††}. In our simulation runs we considered various switch sizes, namely, $N = 1, 2, 4, 8, 16, 32, 64, 128, 256$ and 512. Throughput results, presented as a function of N (on a logarithmic scale), refer to saturation situations. As we have indicated, we are concerned with those the cases where input queues become saturated (saturation may be exhibited at different input load levels for switches of different sizes). For the balanced traffic case we expect all input queues, including odd and even queues in our model, to become saturated at the same input load.

3. Performance Results

3.1 Non-bursty traffic

For a quantitative evaluation and comparison of the Odd-Even switch we consider a homogeneous system. Cell arrivals to each of the N input ports are independent (thus we consider non-bursty traffic) and identically distributed according to a Bernoulli process with parameter λ , that is a cell arrives at an input port with probability λ .

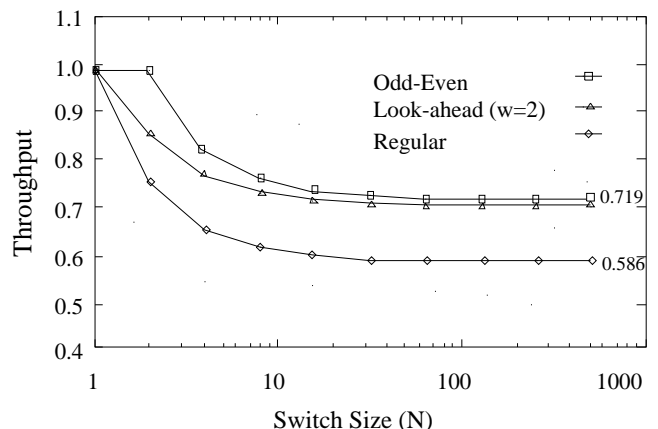


Fig. 3 Throughput of the Odd-Even switch vs. Regular switch (Bernoulli arrivals).

Figure 3 shows the throughput of the Odd-Even switch (upper curve), whose value is (for large N) $\gamma_{max} = 0.719$, a sizable increase compared to that of the ordinary input buffered switch, which is known to

^{††}See [4] for an analytical treatment of throughput for the non-bursty case.

be $2 - \sqrt{2} \approx 0.586$ (cf. [6], [3], [2]), as also seen in the same figure. In [4] we show that an approximate mathematical analysis results in a switch throughput of 0.705 for the Odd-Even model, which is in fact a good approximation to 0.719. We see that the Odd-Even scheme considerably outperforms the regular switch (by more than 20%).

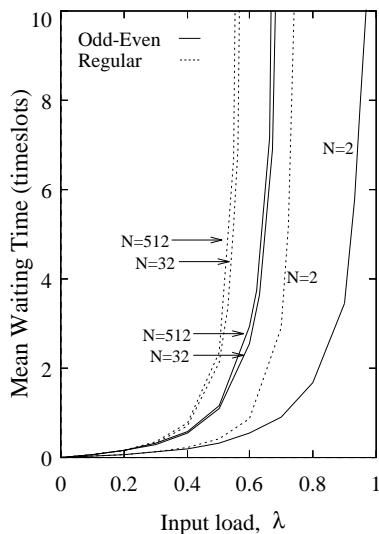


Fig. 4 Mean waiting times (Bernoulli arrivals).

The mean waiting time for the Odd-Even model, as compared to the regular input-queueing switch, is shown in figure 4, as a function of λ and for the values of $N = 2, 32$ and 512 , under a random HOL selection policy. We see again a significant improvement using the Odd-Even switch over the regular one.

Scheme	γ_{max}
Regular - Random	0.586
Odd-Even, Longest Waiting HOL	0.713
Odd-Even, Longest Queue	0.718
Odd-Even, Random	0.719

Table 1 Maximum throughput, γ_{max} , for different selection policies.

Table 1 shows how the Odd-Even model behaves under different policies. While in the regular switch the selection policy has no effect on throughput, this is not true for the Odd-Even. In fact, we notice that the random selection strategy yields a higher throughput than the other two. Similar results (which we do not include in this paper) were obtained with respect to the mean waiting time which showed the random selection policy yields the lowest.

Last but not least we extend our performance evaluation, by considering finite-size buffers as to study the cell drop ratio for both switches, namely the Odd-Even and the regular switch. The *cell drop ratio* is measured as the fraction of incoming cells that are discarded due to insufficient buffer space (buffer overflow) inside the switch. For fairness, we assume the same total buffer capacity, per input port, for both switching systems, i.e., if we assume that a FIFO queue in the regular switch is 32 cells long then the two FIFO queues (per input) in the Odd-Even switch are 16 cells each. Table 2 shows various cell drop ratios for different switch and buffer sizes. The tabulated ratios refer to those saturation cases where the applied input load slightly exceeds the switch's throughput. These input load levels are different for switches of different sizes. We observe from that table that the Odd-Even switch has lower (with a few exceptions) cell drop ratios when the same input load is applied to both switching systems.

3.2 Interrupted Bernoulli Process (IBP) traffic

As Variable Bit Rate (VBR) applications characterize a significant portion of the traffic (e.g., packetized voice and video and large data files) carried by ATM networks, we evaluate our model also under a bursty traffic pattern. In this model, input traffic alternates between active (busy) and silent (idle) periods, while output port addresses are still uniformly distributed. In particular, each of the inputs is described by the same *On-Off* model where both busy and idle periods have a geometrically distributed duration. Cells within the same burst are destined to the same output port (i.e. cells belong to the same fragmented packet). Basically, each time slot is governed, again, by a Bernoulli process, where given an input is in an On (busy) state, it will remain in that state (i.e. during the next time slot) with probability $1 - p$, while it will switch to the Off (idle) state with probability p (see figure 5). An On state corresponds to a burst being transmitted on an input link while the Off state corresponds to a silence period (figure 5). The probability of a burst consisting of k time slots is given by:

$$b(k) = p(1 - p)^{k-1}, \quad k \geq 1 \quad (1)$$

If the input is in the Off (idle) state, where no cells arrive, it will stay in the same state with probability $1 - q$ and then the probability that an idle period lasts k time slots is simply:

$$i(k) = q(1 - q)^k, \quad k \geq 0 \quad (2)$$

which also takes into account the possibility that a burst can be followed immediately by another burst (e.g., two different streams are multiplexed on the same input link) in which case there is no idle period ($k = 0$). This traffic model is referred as the Interrupted Bernoulli Process (IBP). The offered input load

Switch Size $N \times N$	Cell drop ratio									
	$b = 2$		$b = 16$		$b = 32$		$b = 128$		$b = 512$	
	<i>OE</i>	regular	<i>OE</i>	regular	<i>OE</i>	regular	<i>OE</i>	regular	<i>OE</i>	regular
2×2	1.94E-1	2.48E-1	2.92E-2	2.49E-1	1.40E-2	2.49E-1	2.62E-3	2.49E-1	1.90E-5	2.49E-1
4×4	2.17E-1	2.43E-1	5.45E-2	2.29E-1	3.83E-2	2.29E-1	2.79E-2	2.29E-1	2.73E-2	2.29E-1
8×8	2.28E-1	2.48E-1	6.71E-2	2.27E-1	5.04E-2	2.27E-1	4.08E-2	2.27E-1	4.04E-2	2.27E-1
16×16	2.20E-1	2.31E-1	4.98E-2	1.98E-1	2.99E-2	1.98E-1	1.39E-2	1.98E-1	1.07E-2	1.97E-1
32×32	2.23E-1	2.33E-1	5.28E-2	1.98E-1	3.22E-2	1.98E-1	1.64E-2	1.98E-1	1.32E-2	1.97E-1
64×64	2.27E-1	2.37E-1	5.84E-2	2.03E-1	3.81E-2	2.03E-1	2.32E-2	2.03E-1	2.09E-2	2.03E-1
128×128	2.23E-1	2.32E-1	5.28E-2	1.95E-1	3.19E-2	1.94E-1	1.50E-2	1.94E-1	1.15E-2	1.94E-1

Table 2 Cell drop ratios for the Odd-Even (*OE*) and the regular switches under various buffer capacities and for different switch sizes.

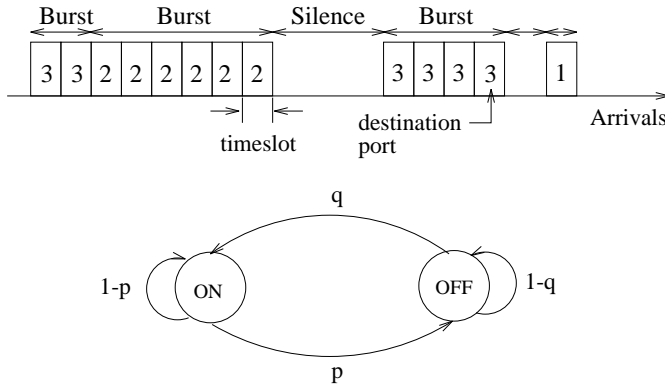


Fig. 5 On(Burst)-Off(Silence) traffic model.

is calculated as:

$$\rho = \frac{E[\text{busy period}]}{E[\text{busy period}] + E[\text{idle period}]} = \frac{q}{q + p - pq} \quad (3)$$

A $p = 0.5$ means that the average burst length is $1/p = 2$ cells. Figure 6 illustrates the saturation throughput for both the Odd-Even and the regular switch, for IBP traffic with $p = 0.05$ (i.e., very bursty), while Figure 7 gives the maximum throughput for different values of p (the probability of being at the On state) for a 512×512 switch. For $\rho = 1$ then from Equation (3) we get $q = 1$ which means that there are no silence periods, regardless of the value of p ($p \neq 0$). This leads the system to overloaded situations. If now $p = 1$ (bursts consist of only one cell) then $q = \rho$ which then has the same effect, in terms of throughput, as the Bernoulli($\lambda = \rho$) case in the previous paragraph. From Figure 7, which gives the asymptotic behavior, we notice that the absolute difference in throughput between the two switches remains almost constant for any value of p . This means that the percentage gain in throughput using the Odd-Even switch becomes larger




Fig. 6 Throughput of the Odd-Even switch vs. Regular switch (IBP arrivals, $p=0.05$).

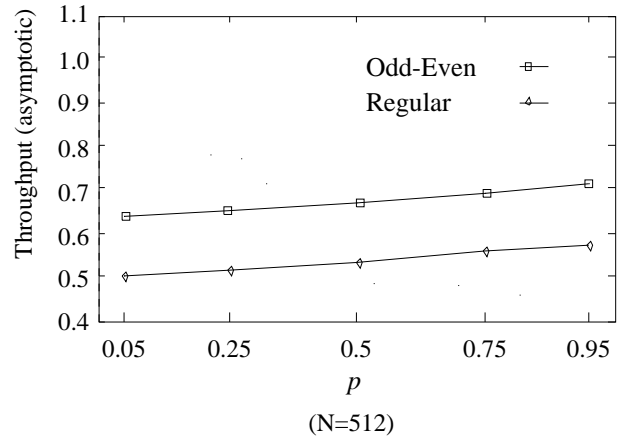


Fig. 7 Asymptotic throughput of the Odd-Even switch vs. Regular switch (IBP arrivals).

as p decreases (hence traffic is more bursty).

3.3 Non-uniformly distributed destinations

So far we have assumed uniformity for the output port addresses. In this section we evaluate the switch's performance under the assumption of a non-uniform, non-bursty traffic pattern. The output port mapping for each incoming cell is determined by a binomial distribution [8], where for $i = 2, \dots, N - 1$:

$$\begin{aligned} a_{i,j} &= Pr[\text{a cell arriving at the } i\text{-th input is} \\ &\quad \text{destined to the } j\text{-th output}] \\ &= \binom{N}{j} r_i^j (1 - r_i)^{N-j}, \quad j = 1, \dots, N \end{aligned} \quad (4)$$

and r_i is a probability associated with input i , $2 < i < N$. Note that for this binomial distribution the maximum probability occurs for $j = \lfloor Nr_i \rfloor$. Now, if we choose $N - i$ to be the output to which the highest percentage of traffic arriving on i goes (therefore it has the largest probability associated with input i), then we have

$$r_i = 1 - \frac{i}{N}, \quad i = 2, \dots, N - 1. \quad (5)$$

For inputs 1 and N the output address for $j = 1, \dots, N$ is given by a normalized Poisson-like distribution with rate r ,

$$a_{1,j} = \frac{\frac{r^{N-j}}{(N-j)!}}{\sum_{\forall j} \frac{r^{N-j}}{(N-j)!}}, \quad a_{N,j} = \frac{\frac{r^j}{j!}}{\sum_{\forall j} \frac{r^j}{j!}} \quad (6)$$

One of the advantages of using a binomial distribution for inputs $i = 2, \dots, N - 1$, is that we still retain the balanced traffic effect for the Odd and Even input queues in each input port. However, this observation does not apply for inputs 1 and N , where, depending on the value of r , there is some bias introduced against the even outputs, but as N gets very large this bias becomes negligible. Another reason for using this kind of binomial distribution is it gives us a substantial number of "hot spots"[†] (instead of one or two), which is more realistic as the number of switch ports increases (and thus the number of potential connections among hosts in an ATM network).

In our simulation runs, we chose a value of $r=0.4$; however r has no significant relevance on estimating the throughput of the system. Using the same input-to-output address assignment functions (4)-(6) for both the Odd-Even and the Regular switch, Figure 8 shows the throughput of each switch respectively.

4. The Odd-Even and the Look-ahead model

Lastly, we compare the Look-ahead scheme [1] (an input "window" policy for dealing with the HOL prob-

[†]Destinations that collect most of the incoming traffic.




Fig. 8 Throughput of the Odd-Even switch vs. Regular switch (non-uniform traffic).

lem) to the Odd-Even switch. In the Look-ahead strategy (also known as the bypass discipline) the arbitration and cell selection process employs w separate contention resolution rounds. During the first one all the HOL cells at the input queues contend for an output port. Those queues that lost the contention during that first round can participate in a new arbitration where the cells right behind the HOL cells (that were blocked) compete for outputs that are still idle (i.e. they have not received any cells so far in the current time slot). In the same fashion a contention cycle that is comprised of a total of w rounds, allows, during the i -th round, the i -th (deep in the queue) cells (from those input queues that have not sent a cell) to contend for an unclaimed




Fig. 9 Throughput of the Odd-Even switch vs. Look-ahead($w=2$) and Regular switches (IBP arrivals, $p=0.5$).

output. It is clear that only one cell is allowed to be switched from a single input and only one cell can be served by an output port during a time slot. For $w=1$,

we have the regular pure input-buffered switch.

In our comparison we consider a Look-ahead scheme with a window of $w = 2$. We feel that a window of two makes it comparable to the Odd-Even model where contention resolution also consists of two rounds (one for polling the odd and one for polling the even queues). Also, note that in both strategies only one cell can be cleared from a single input port (there is not any speed-up feature).

Figure 3 illustrates how all three models, namely the Odd-Even, the Look-ahead and the regular one compare, assuming Bernoulli arrivals. We see that the throughputs of the Odd-Even and the Look-ahead are very close. However as traffic becomes more bursty (which means that cell arrivals are correlated) this is not any longer true. The Odd-Even performs much better, since a window of $w = 2$ for Look-ahead does not provide much of a benefit (see figure 9).

5. Conclusion

In this paper we have described an innovative crossbar, input-buffered ATM switch, the *Odd-Even* switch, which allows incoming cells to be stored in either of the two FIFOs placed in every input port, according to their destination, namely odd or even. Its performance has been quite extensively examined and compared it to the simple input-queueing switch under various traffic conditions. Furthermore, we compared the Odd-Even model to the Look-ahead scheme for which we assumed a window of $w = 2$, in order to make the latter comparable to the Odd-Even. In all cases, it has been demonstrated that the Odd-Even can achieve a noteworthy gain of more than 20% in throughput over the regular input-buffered switch. Other performance results (with respect to mean waiting times, cell drop ratios) clearly suggest a superior performance of the Odd-Even switch over the ordinary input-buffered switch. We recognize that the Odd-Even switch may require additional hardware and complexity (e.g. input controller, more complex arbitrator), however it is still simple in its implementation and effective in terms of performance.

We have expanded the idea of the Odd-Even switching scheme to include multiple (instead of just two) input FIFO queues connected to the same input port. This led us to introducing the Multiple Input-Queueing switch which we study in [4]. As an additional application of the Odd-Even switch we mention the use of this dual input-queueing strategy in the buffered-banyan multistage interconnection networks. This has resulted in us proposing a high-performance switching network architecture which we call the *Dual-Banyan* (DB) switch (cf. [5]).

References

[1] M. G. Hluchyj and M. J. Karol, "Queueing in High-

- Performance Packet Switching", *IEEE J. on Selected Areas in Commun.*, vol. 6, no. 9, pp. 1587-1597, Dec. 1988.
- [2] J. Y. Hui and E. Arthurs, "A Broadband Packet Switch for Integrated Transport", *IEEE J. on Selected Areas in Commun.*, vol. SAC-5, no. 8, pp. 1264-1273, Oct. 1987.
- [3] M. J. Karol, M. G. Hluchyj and S. P. Morgan, "Input versus Output Queueing on a Space-Division Packet Switch", *IEEE Trans. on Commun.*, vol. COM-35, no. 12, pp. 1347-1356, Dec. 1987.
- [4] C. Koliass and L. Kleinrock, "Throughput Analysis of Multiple Input-Queueing in ATM Switches", in *Broadband Communications*, pp. 382-393, L. Mason and A. Casaca (eds.), Chapman & Hall, United Kingdom, 1996.
- [5] C. Koliass and L. Kleinrock, "The Dual-Banyan (DB) Switch: A High-Performance Buffered-Banyan ATM Switch", in the proceedings of the *International Conference on Communications (ICC) '97*, pp. 770-776, Montreal, Canada, June 1997.
- [6] C. Koliass and L. Kleinrock, "An Analytically Simplified Performance Study of ATM Switches with Input-Queueing", submitted for review to the *IEEE Transactions on Communications*, March 1997.
- [7] R. Onvural, "Asynchronous Transfer Mode Networks: Performance Issues", Artech House, Norwood, MA, 1994.
- [8] S.Q. Li, "Nonuniform Traffic Analysis on a Nonblocking Space-Division Packet Switch", *IEEE Trans. on Commun.*, vol. 38, no. 7, pp. 1085-1096, July 1990.

Christos Koliass received his B.S. degree in mathematics from the University of Athens, Greece in 1989. In 1991 he was awarded a Fulbright Scholarship. Since then he has been with the Department of Computer Science at University of California, Los Angeles (UCLA) where he is pursuing his Ph.D. in the area of Computer Networks and Communications. He received the M.S. and Engineers degrees both in Computer Science from UCLA in 1993 and 1995 respectively. His current research interests are in the area of Asynchronous Transfer Mode (ATM) and in particular of ATM switching. Other interests include wireless ATM, queueing systems and applications, performance evaluation.

Leonard Kleinrock, a father of the Internet, received his B.S. degree in electrical engineering from the City College of New York in 1957 and the M.S.E.E. and Ph.D.E.E. degrees from M.I.T. in 1959 and 1963, respectively. He is currently a Professor in the Computer Science Department at UCLA. His research interests focus on nomadic computing and performance evaluation and design of many kinds of networks (e.g. packet switching networks, packet radio networks, local area networks, metropolitan area networks, broadband and gigabit networks) and of parallel, wireless and distributed systems. Dr. Kleinrock is a member of the National Academy of Engineering, an IEEE Fellow, a Guggenheim Fellow, recipient of the Swedish L.M. Ericsson Prize, of the 12th Marconi Award and of the UCLA Distinguished Teaching and Faculty Research Lecturer Awards.