

Traffic Regulation and Control in the Manhattan Street Network

O.Tayan, D.Harle

Department of Electronic and Electrical Engineering,
University Of Strathclyde, Glasgow, G1 1XW, U.K.
Tel: (+44) 141 548 2082 Fax: (+44) 141 552 4968
Email: otayan@eee.strath.ac.uk, d.harle@eee.strath.ac.uk

Keywords: Medium Access Control, veto protocol, Manhattan Street Network

Abstract

The Manhattan Street Network (MSN) with the *Clockwork* Routing scheme has been proposed as a particularly suitable candidate for use in optical networks. However, without any medium access control (MAC), the success of the MSN with *Clockwork* Routing in optical networks would be inhibited by the destination-release protocol, since the protocol provides no guarantee of fair access or bounded traffic delay requirements of access nodes. Hence, new traffic regulation schemes providing medium access control are required to ensure fair access between network nodes and satisfy stringent quality of service (QoS) requirements of applications. This paper introduces a new scheme for medium access that supports fairness between competing access nodes while limiting the delay in the network. A ring model based on the "return-to-trail" property of the MSN is used to realise the logical behaviour of the MSN and connection services are regulated using a distributed view of the network loading.

1 Introduction

The MSN is a self-routing virtual topology proposed for local and metropolitan area network applications. An MSN is characterised as a two-connected regular mesh with the nodes connected as rows and columns. Each node consists of a 2x2 cross-bar switch that connects incoming links to outgoing links. The *Clockwork* Routing Scheme has been proposed as a candidate for implementing optical packet switched networks. The *Clockwork* Routing Scheme is a time-slotted system that enhances the MSN by including a simple routing mechanism employed at intermediate nodes that prevents optical layer contentions, requires no re-sequencing of packets at the destination node and has comparable throughput to conventional routing schemes [1].

This paper focusses on the use of rings to model the logical operation of the MSN, allowing traffic regulation and performance to be monitored at the ring level. In [2], it was shown that an $n \times n$ MSN may be decomposed into n^2 independent rings by exploiting the return-to-trail property. The MSN with *Clockwork* Routing was viewed as a collection of independent rings, each operating as a

slotted system while enforcing destination-release slot-reuse of the container at the destination node. A key advantage over single token passing protocols is evident in its support in allowing several nodes transmit simultaneously using different slots, one on each ring of the MSN. However, the drawback of using a destination-release based protocol was acknowledged: immediate reuse of the slot container at destination nodes may deprive other nodes of service. Such schemes are clearly restrictive in real time applications, whereby stringent and bounded delays must be guaranteed as part of the service. Therefore, traffic regulation schemes providing medium access control (MAC) have become an essential part of the research. In [2], the ring model was considered for purposes of deriving the stability criteria of each of the peripheral queues in the MSN and obtaining performance analyses of the ATM Ring medium access protocol. It was noted, however, that the ATMR scheme was initially proposed for bidirectional rings. Hence, it would be beneficial to study and compare alternative MAC protocols suitable for unidirectional rings. This study uses the ring approach to simulate and analyse an alternative MAC protocol; a protocol that regulates connection services to provide fair medium access by using a distributed view of the loading on the network.

The remainder of this paper is organised as follows. Section 2 explains the unfair medium access scheme associated with the destination-release based protocol. Section 3 explains how an $n \times n$ MSN with *Clockwork* Routing may be decomposed into n^2 independent rings. The behaviour of the proposed protocol in a single constituent ring is the focus of study in this paper. An alternative to the ATMR protocol (in respect of providing medium access control) is introduced in section 4, outlining the assumptions and objectives and explaining the methodology/implementation techniques used. Section 5 presents the results obtained for variations of the proposed model by altering design parameters. Assumed traffic metrics such as network load, are varied and the results are compared and discussed. Section 6 concludes the work.

2 Using the 'Return-to-trail' Property to Model the MSN

This section exploits the 'return-to-trail' property to decompose an $n \times n$ MSN into n^2 constituent rings. In [2], it was shown that each ring is independent and can be used to model the logical behaviour of a *Clockwork*-Routed MSN. A 4×4 MSN is used as an example to explain how an $n \times n$ MSN with *Clockwork* Routing can be decomposed into 16 independent rings. Consider a single slot-container released from the horizontal output port of node 1 on timeslot 0. Any data on the currently selected horizontal buffer in node 1 (governed by the current timeslot) would consequently occupy the slot-container. The return-to-trail path of this slot-container is highlighted in Figure 1. After the first hop (timeslot 1), the slot-container visits node 2, allowing data selection from its horizontal buffers if the slot-container were unoccupied. On the second hop (timeslot 2), the slot-container visits node 3 where the horizontal queue associated with the

current timeslot is selected to transmit data. An observation from Figure 1 is that the slot-container returns to its original node after a total of eight timeslots ($2n$ hops). It must be noted that each ring represents node 1 using two different nodes, whereby the first node would be the initial and returning position of the slot-container and the second node would be the midpoint intermediary node in the return-to-trail path. Additionally, horizontal and vertical peripheral queues are associated with the respective nodes. Hence, using a slot-container to traverse the MSN in a return-to-trail path has demonstrated the clockwise rotation of the slot-container in a predefined ring ($0 \rightarrow n-1$). Similarly, the decomposition obtained by considering node 1 at timeslot x for $\{0 \leq x \leq n-1\}$ can be extended for each MSN node by considering each node at timeslot x . In essence, this implies that an $n \times n$ MSN consists of n^2 independent rings each of which maintains the logical operation of the network. This paper focusses on a single ring to model the behaviour of the MSN.

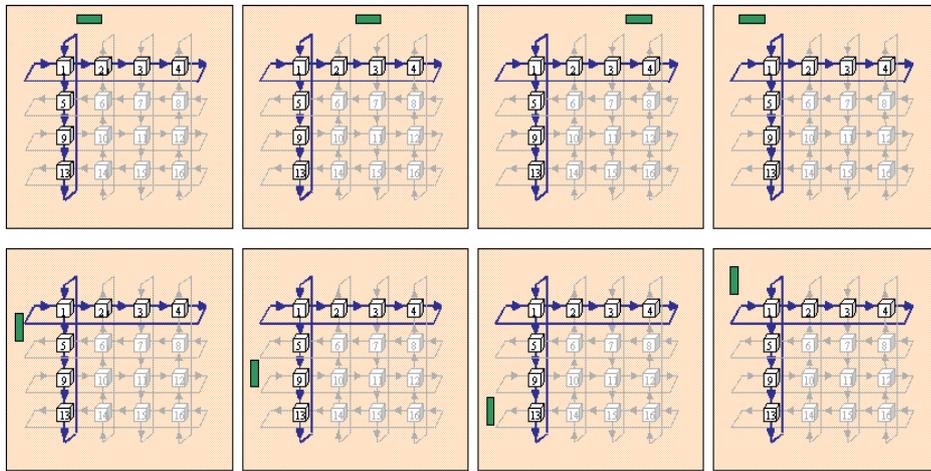


Fig. 1. Highlighting one possible sequence of a 'return-to-trail' path

3 Medium Access Unfairness in a Time-Slotted MSN

The ring based approach offers numerous advantages; providing a higher bandwidth utilisation than single-slotted systems and allowing for increased traffic to be offered to the network. A major drawback of the ring based approach, however, is that it provides no means of ensuring fair use of resources by access nodes. Traffic matrix parameters and positioning of nodes on the ring are influential in provisioning access to the shared medium. In a highly loaded network, destination-release of the slot-container may prevent some nodes from ever accessing the slot-container, therefore depriving these nodes of service [2]. For

instance, Figure 2 shows two nodes (nodes 1 and 5) in isolation: neither of these nodes can ever access a free slot-container when upstream stations are continuously transmitting to downstream stations of the respective nodes. Consider when node 2 (upstream from node 5) continuously transmits to node 9 (downstream from node 5), where the slot-container is released. In this scenario, node 5 only ever sees in-transit packets and never any free slot-container and hence is unable to transmit data.

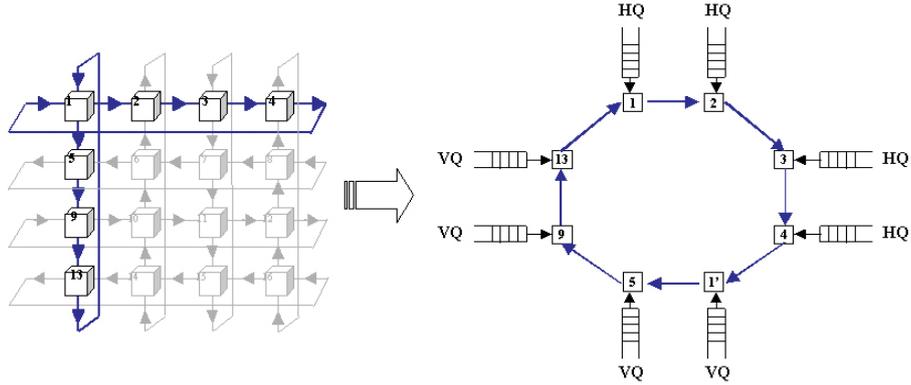


Fig. 2. Nodes deprived of service in the destination-release protocol

4 Medium Access Control for a Slotted-Ring Implementation of the MSN

4.1 Problem and Motivation

Sections 1 and 3 have stated the need for a medium access control scheme that ensures fair use of resources between access nodes. Without medium access control, the success of the MSN with *Clockwork* Routing in optical networks would be inhibited by the destination-release protocol since the protocol provides no guarantee of fair access or bounded traffic delay requirements of network nodes. Hence, new traffic regulation schemes providing medium access control are required to ensure fair medium access between network nodes and satisfy stringent quality of service (QoS) requirements of real time applications.

This section introduces a new scheme for medium access that supports fairness between competing access nodes. The MAC protocol uses a ring model to realise the logical behaviour of the MSN and connection services are regulated using a distributed view of the network loading.

4.2 Implementation and Methodology

Guaranteed QoS parameters in the MSN can be achieved by regulating traffic using a distributed view of the network loading, whereby all affected nodes

are able to support/oppose required connection services. Hence, the ring model implementation must allow intermediate nodes to decide whether they support or oppose each proposed connection. The protocol would then accept or reject the proposed connection by considering its own local perception of the loading on the network and the stance taken by each affected node. Therefore, distributed protocols may be employed to establish connection request services, thereby removing the need for a central management node. A thorough explanation of the operational concepts of a simple "Veto" protocol is now discussed.

The 2N bit Connection Request Structure Initially, the veto protocol was realised using a $2n$ bit connection request field. Each node in a $2n$ node ring is allocated a single bit for use in the request field. The bit corresponding to each node is set when the node requests a connection. Conversely, the bit is reset if the node has not requested connection as for example initialisation or when an existing request has been 'vetoed' by an intermediary node that lies along the path of which the request circulates towards the source node. The veto protocol implies the behaviour of a slotted ring system supported by a token that circulates through the ring. To avoid reducing network throughput by transmitting connection requests and data separately on the slot container, the protocol operates on a single (uniform) cell format that includes the connection request field and the veto signal(s) in the header structure of a payload cell. A generic cell structure used in each version of the protocol is illustrated in Figure 3 and Figure 4 showing the connection request field used in this simulation.

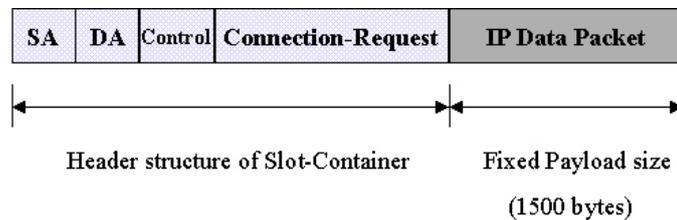


Fig. 3. Generic cell structure for each implementation of the veto protocol

The uniform cell format enables separate payload and connection request fields to be filled and emptied independently on each cell. Additionally, each traffic type must be stored separately before being transmitted on the slot container. Hence, each ring node must support an incoming connection request buffer and a payload data buffer.

The protocol initially circulates a connection-request cell around the ring (following a 'return-to-trail' path in the MSN), giving each intermediate node an opportunity to signal its opposition to the proposed connection. Any node wishing to block the proposed connection would alter the circulating connection

awaits its return after $2n$ timeslots. As each ring node receives the request, it must decide whether or not to accept the proposed connection. The decision to oppose (or accept) at each intermediate node can be based on any number of criteria including, instantaneous or average buffer occupancy or average service rate or the throughput into the node. If an intermediate node decides to oppose the connection, it resets a single bit in the connection request field before it is passed on towards the source node. Once the request has arrived at the source node, the source frees the request field for the following request and must decide whether or not to accept the connection. The decision of acceptance/rejection is based on the bit-value (within the request field) corresponding to the source node. If the corresponding bit-value was set, it would be deduced that no other node has vetoed this particular request. Hence, the data packet(s) associated with the request are accepted into data buffers awaiting transmission on a free payload slot within the circulating cell. A reset bit value, however, would indicate that either the node did not previously transmit a request or that the node sent a request which must have been vetoed by a ring node. The connection attempt would therefore be rejected and the associated data packets discarded. The protocol operates on the data fields independently of connection requests: when a node with an occupied data buffer sees a free data slot, it fills the empty slot with the packet at the head of the data queue. Once the cell has reached its destination, the data field becomes free again.

The protocol achieves simultaneous operation of data and connection requests by operating on the data and request fields independently of each other within the circulating cell.

The motivation in this initial veto scheme had considered simplicity in implementation and the idea of a simple request field was to ensure simple calculations and signalling operations were required and realised from simple optical logic. The veto scheme aims to provide medium access control in the MSN deployed in ultrafast optical networks.

4.3 Multi-Queue Single Server Model

The veto scheme introduced in Section 4.2 can be described as a multi-queue single server (MQSS) model [3] in which a number of queues are served according to a certain discipline with the mean inter arrival rate λ_i at queue i .

The server visits each node sequentially, filling and freeing data and connection request slots independently as the cell circulates through the network. A cyclic-discipline governs the order by which connection-request queues are serviced by the server, whereas data packets are serviced using a stochastic mechanism¹. The MQSS model provides support for cyclic connection-request services and stochastic data service schemes by inserting/removing connection-requests at fixed (deterministic) time intervals ($2n$ hops) and inserting/removing data randomly from the data slot when data packets are transmitted and received respectively. A polling system with Markovian server routing had been considered

¹ i.e. only occupied buffers at each station are serviced

in the literature [2][4] to model a destination-release based protocol. Several previous studies, including [2] were concerned with stochastic polling mechanisms to model arbitrary (data) packet arrivals and packet transmissions from access nodes. However, the model presented in this study uses a circulating slot-container, guaranteeing that each node is visited after $2n$ hops. The stochastic nature of data flows are maintained in the model described in Section 4.2 since the data-slot within the circulating cell serves data stations according to the destination-release protocol. Once the data-slot has been filled by a packet from a data station, it remains occupied until it reaches its destination, at which point the data-slot is freed and available to the destination station. Meanwhile, connection requests are inserted and removed at the designated fields within the request-slot on each cycle of the slot-container (i.e after a period known as the *cycle* time). Interestingly, the model introduced in Section 4.2 seems to combine the features of the strictly cyclic order discipline commonly found in token-ring based networks and the random polling scheme used in stochastic data traffic modelling [2][4]. Furthermore, the MQSS is classified as a 1-limited service model since on each cycle, the slot-container can only service a single packet at the head of each queue.

Arrival Process In [3], the system utilisation ($p\theta$) was defined as the sum of the products between the arrival rate (λ_i) and service time (S_i) for each node i , as expressed in Equation 1.

$$p\theta = \sum_{i=1}^n \lambda_i \cdot S_i \quad \forall i \in [1, 2n] \quad (1)$$

An MQSS model operating in stability enforces the condition that $p\theta \leq 1$. The load on each node in the return-to-trail path was assumed to be uniform. A uniform load for each node in the return-to-trail path is equivalent to a non-uniform load on the MSN.²

4.4 Packet Generation Process

Each node generates packets using Bernoulli trials and each packet is assigned an arbitrary destination from a uniform distribution. The veto protocol utilises a connection-length parameter³, indicating the length of the call in constant size data cells (1500 bytes). A connection request is then generated for each call awaiting transmission and remains in the connection-request queue until arrival of the slot-container.

² A balanced MSN was shown to be equivalent to a heterogeneous loading on the nodes in the return-to-trail path [2].

³ Different simulations were considered for each connection-length in order to demonstrate the impact of call length on performance.

5 Results

This section provides a performance evaluation of the veto scheme when used with a *Clockwork* Routed MSN. The two figures of merit investigated in this section are the mean call drop probability (i.e. loss of discarded packets as a proportion of the offered load) and the mean call delay. The loss is an important network design parameter, as it indicates the point at which the network becomes overloaded and hence saturated. Mean call delay, on the other hand, is essential for monitoring packet delays and ensuring that they conform to the stringent requirements of intended applications. A number of variables including the call connection length, veto threshold and data buffer threshold were considered to obtain a detailed analysis of the veto protocol. Several assumptions were made in order to model the performance of the veto protocol. These assumptions include:

- An $n \times n$ MSN with Clockwork Routing may be decomposed into n^2 independent rings, each ring established by following the 'return-to-trail' path from each node at a chosen timeslot.
- Each ring is independent and equally loaded, hence a single ring was used to model the logical operation of the MSN.
- Uniform traffic generation was assumed at each node. In other words, the probability of success of Bernoulli trials was constant for all nodes. Additionally, source and destination nodes were chosen from a uniform distribution.
- All links between adjacent nodes were assumed to have 1 timeslot length. Moreover, it was assumed that the links were lossless and the propagation delay to traverse each link was considered negligible.
- The destination release protocol, allowing immediate reuse of the slot-container was considered. Essentially, this implies that a slot-container containing a packet destined for an arbitrary destination could be immediately reused at the destination if there was a packet at the head of the destination station.
- Processing times at the source and intermediate nodes were assumed negligible.

An analysis of the results obtained from a simulation must consider the degree of uncertainty introduced by random samples during simulation execution. This work uses a Student's t distribution to obtain a *small sample confidence interval* since the number of samples taken are typically below thirty.

A 95% confidence interval is assumed in all results presented in this section. All simulations have an initial warmup phase to eliminate any bias due to transient effects and obtain steady-state behaviour before recording statistics. Simulations were run for 10^7 trials and 10^3 warmup trials. A constant buffer threshold of 30 entries was used at each node. Consequently, exceeding the buffer threshold would result in a veto of the next connection request.

The experiments are based on a number of simulation runs of variable connection lengths in order to analyse the performance of a basic implementation of the veto protocol. Figures 5 and 6 show the probability of loss and the mean call delay, respectively, for a call length of 1, 5 and 20 packets in a 4x4 MSN.

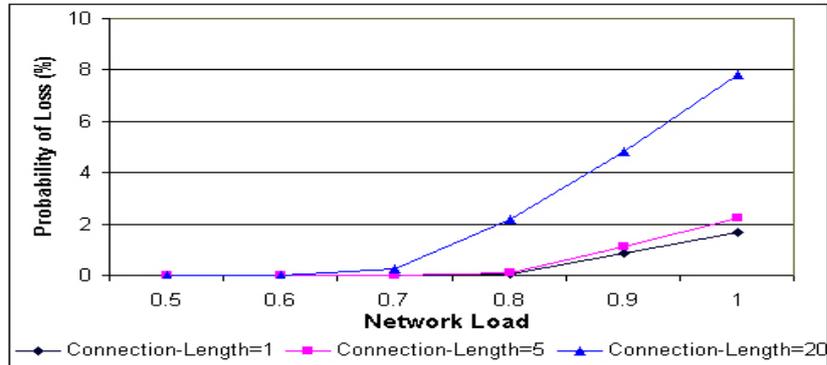


Fig. 5. Call Drop Probability for MAC scheme

Clearly, Figure 5 shows that the longest call duration results in the highest loss, since for each rejected call, 20 packets are dropped. It is expected that longer calls correspond to larger mean call delays, as shown in Figure 6 for a call length of 20. However, long calls require fewer connection requests, allowing for more efficient use of the bandwidth at higher loads. In contrast, call lengths of unity require a separate connection request for each packet, consequently occupying more time for connection-request passing and resulting in larger packet delays than longer calls at high loads. Figure 7 shows that at 0.5 load, the mean packet delay for short calls increases above the corresponding packet delay for longer calls, as a result of the additional time each node spends on connection request passing.

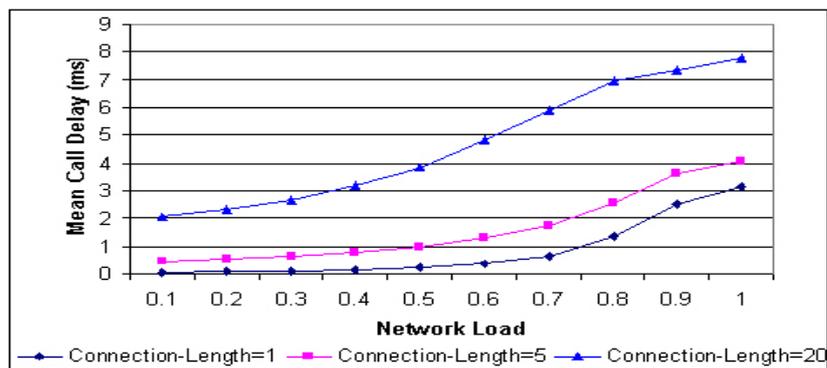


Fig. 6. Mean Call Delay for MAC scheme

Furthermore, it may be observed that for longer calls in particular, the mean packet delay levels off as the the load approaches unity, suggesting that the veto protocol limits the delay present in the network. The mean packet delay is obtained from the mean call delay and the connection-length, as defined in Equation 2.

$$\text{MeanPacketDelay} = \text{MeanCallDelay}/\text{ConnectionLength} \quad (2)$$

The maximum throughput obtained was 0.248, in comparison to 0.31 from an MSN model that had avoided connection requests in [2]. However, [2] had derived the maximum theoretical throughput as the inverse of the mean path length in the MSN. Hence, the reduction in throughput can be explained, since in [2], a shortest path routing discipline was used, yielding a mean path length of 2.94 compared to a mean of 4.0 for the ring considered in this study. The theoretical maximum throughput for the ring topology was calculated as 0.25, for a mean path length of 4. Furthermore, a fair comparison of the loss performance with the losses obtained in [2] is possible when a connection length of 1 is considered. Figure 5 shows a maximum drop probability of 1.6% in comparison to a theoretical loss probability 0.7. However, Figure 5 shows that the loss probability remains negligible until 0.9 load, where the loss approaches the theoretical maximum value. The additional loss may be explained, since at very high loads the veto scheme must discard some packets in order to prevent high demand nodes starving other nodes of service. Maximum loading causes queue occupancies at nodes to exceed the buffer threshold, hence arising in vetoes from all affected nodes.

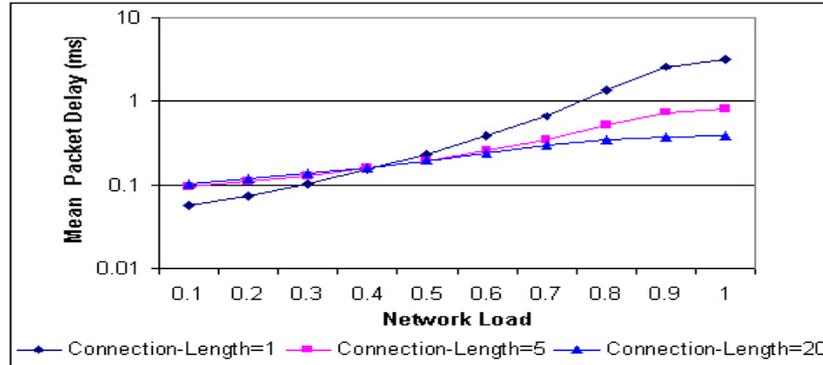


Fig. 7. Mean Packet Delay characteristics

6 Conclusions

This paper has investigated medium access control for the MSN with *Clockwork* Routing. In previous work, the MSN performance was inhibited by the destination-release protocol and clearly a new traffic regulation scheme was needed to prevent nodes being deprived of service. The significance of the "return-to-trail" property was evident, as it had enabled the MSN to be modelled as ring network with $2n$ nodes. A fair mechanism of controlling medium access and preventing node starvation, while limiting the delay in the network was achieved using the veto protocol. An analysis of the results had shown the significance of the veto protocol at high loads: nodes veto connections in order to keep their queue occupancy below the buffer threshold, thereby limiting the delay in the network with only the expense of a slight increase in loss as compared to the original model that lacks medium access control.

References

1. O.Tayan, J.Md.Desai, D.Harle, "A Practical Embedding Problem Exposing the Affect of Link Failures in the Manhattan Street Network", Proceedings of HET-NET's 2003, Ilkley, West Yorkshire, July 2003.
2. Chevalier, F., "Performance Evaluation Of The Clockwork Routing Scheme in Optical Packet-Switching Networks", PhD Thesis, 2000.
3. Akimaru H., Kawashima K., 'Teletraffic Theory and Applications', Published by Springer 1999.
4. Zafir-Vukotic M., "Waiting Time Estimates in Symmetric ATM- Oriented Rings with the Destination Release of Used Slots", IEEE/ACM Transactions on Networking, Volume 7, Number 2, April 1999, pp251-261.
5. Hock NG C., "Queueing Modelling Fundamentals", Published by John Wiley and Sons 1996.
6. K. Watkins, "Discrete Event Simulation in C", Published by McGraw Hill Book Company, 1993.
7. F. Chevalier, D.Cotter and D.Harle, "Routing and Fairness Issues Associated with a Self Routing Control Strategy in Regular Mesh Photonic Networks", International Teletraffic Congress ITC-16, Edinburgh, Scotland, June 1999.