TDM-BASED WDM ACCESS PROTOCOLS: A COMPARISON OF RESERVATION AND PRE-ALLOCATION STRATEGIES FOR A PHOTONIC STAR-COUPLED CONFIGURATION

PATRICK W. DOWD[1] AND KALYANI BOGINENI
Department of Electrical and Computer Engineering
State University of New York at Buffalo
Buffalo, NY 14260

Abstract

This paper introduces a *collisionless* pre-allocation protocol and compares it to a *collisionless* reservation protocol for WDM star-coupled photonic networks. A detailed simulation analysis is developed to study the behavior of the protocols with varying system characteristics. The *pre-allocation* protocol is designed for a network with the WDM channels assigned to the nodes for data packet reception. Access to the channels is based on time multiplexing, eliminating data channel collisions and destination conflicts. The *reservation* protocol employs one of the WDM channels as a control channel to reserve access for data packet transmission on the remaining data channels. Time-division multiplexing is employed on the control channel enabling all active nodes the opportunity to transmit once every control cycle. Any available data channel can be used for transmission. Channel collisions and destination conflicts are avoided through the use of tables tracking the status of data channels and destination nodes maintained at every node. This approach is shown to significantly reduce the long synchronization delays typical of systems employing time-division multiplexing on data channels. The impact to the performance due to the switching latencies of the optical devices is decreased through overlapping mechanisms. The performance is evaluated in terms of network throughput and average packet delay for variations in the number of nodes and data channels, data packet length and switching latency.

Index Terms: media access protocol, performance analysis, wavelength division multiplexing, photonic network architecture.

1 INTRODUCTION

This paper introduces a collisionless media access protocol for a wavelength division multiplexed (WDM) star-coupled system with channels pre-allocated to the nodes for data packet reception. The performance of the proposed protocol is compared to a collisionless, control channel-based reservation protocol [1, 2]. The performance analysis, studied in terms of average packet delay and network throughput, is based on extensive discrete-event simulation.

Wavelength multiplexing eases the speed mismatch between optics and electronics by partitioning the enormous optical bandwidth into multiple, more manageable, channels that each operate at the data-rate limited by the interface electronics [3]. WDM creates multiple channels where each channel may operate in a multi-access mode. A multiple access environment can be achieved through a variety of optical channel topologies [4]. Star-coupled networks have an optical power budget advantage over optical bus-based systems so a larger system size can be supported before amplification is required. Furthermore, high fault-tolerance is achieved with a star-coupled interconnection due to its passive nature and complete unity distance connectivity [5, 6]. This high connectivity is achieved with low system complexity through the multiple access nature of the system.

[1]This work was supported by the National Science Foundation under Grant CCR-9010774 and ECS-9112435; and utilized the computing facilities at the National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, through NSF Grant TRA-910321.

Media access protocols developed for photonic star-coupled WDM networks may be broadly classified into *reservation* and *pre-allocation* strategies [7, 8]. *Reservation techniques* designate a wavelength channel as the *control channel* that is used to reserve access on the remaining channels for data packet transmission [9, 10, 11, 2, 12, 13, 14, 1]. The control channel is used to transmit control information and reserve access on the data channels. Media access protocols are required to provide arbitration on both the data and control channels. *Pre-allocation techniques* pre-assign the channels to the nodes, where each node has a *home channel* that it uses either for all data packet transmissions or all data packet receptions [7, 15]. This eliminates the requirement that a node possess both a tunable transmitter *and* a tunable receiver. Pre-allocation eliminates the need for a control channel so all channels are used for data transmission [16, 7].

Random and static access pre-allocation protocols were proposed in [7, 15, 17, 18]. Pre-allocation may be achieved by either specifying the channel a node will use to transmit (requires a tunable receiver and a fixed transmitter) or receive (requires a tunable transmitter and a fixed receiver). A source node tunes its transmitter to the home channel of the destination node and transmits according to the media access protocol with a system where channels are pre-allocated for data packet reception. A home channel may be shared with other nodes if the number of nodes (M) exceeds the number of channels (C). Any node in the system can determine the home channel of any other node in a decentralized fashion with knowledge of the destination node number and the total number of nodes and channels in the system [7, 18].

The reservation based protocols proposed for star-coupled systems in [9, 16] are based on random access schemes. The channels are shared by all nodes in [9] on a contention basis, with random access schemes employed on both control and data channels. The destination node is informed through the control channel by the source node of its intention to transmit, the data channel to expect the data packet, and its size. Packet collisions could occur during both control and data packet transmission. Improvements to these protocols were proposed to avoid collisions on the control channel in [13], reduce data channel collisions in [19] and avoid receiver conflicts in [14].

Reservation protocols may provide flexibility in the use of data channels and are particularly attractive over pre-allocation protocols when there are far fewer channels than nodes. However, reservation protocols are often more complex than pre-allocation approaches since the transfer is based on two stages: reservation and transmission [7]. Depending on the implemented protocol, collisions may occur during control and data packet transmission. The performance of reservation and pre-allocation random access protocols were compared in [7] where it was shown that pre-allocation protocols achieve a significant improvement in performance with reduced system complexity.

A hybrid protocol, denoted as DT-WDMA, was proposed in [10] that has data channels pre-allocated for data packet transmission and a time-multiplexed control channel used to inform a destination node of the intentions of a source node. Each node has two fixed (or slow tunable) transmitters and two receivers. One transmitter is parked on the control channel and the other is parked to transmit on its home channel. One receiver is fixed (or slow tunable) and continually monitors the control channel while the other receiver is (fast) tunable. A node tunes its fast receiver to the home channel of the source node after receiving a control packet that identifies it as the destination.

A protocol that also time multiplexes the control channel, denoted as TDMA-C, was proposed in [1]. However, the data channels with TDMA-C are not owned by any particular node. Time multiplexing the control channel provides each node a chance to transmit each control cycle. This protocol achieves collisionless data packet transmission without requiring the allocation of *home* channels. Packet collisions due to destination node or data channel contention are eliminated through status tables maintained at each

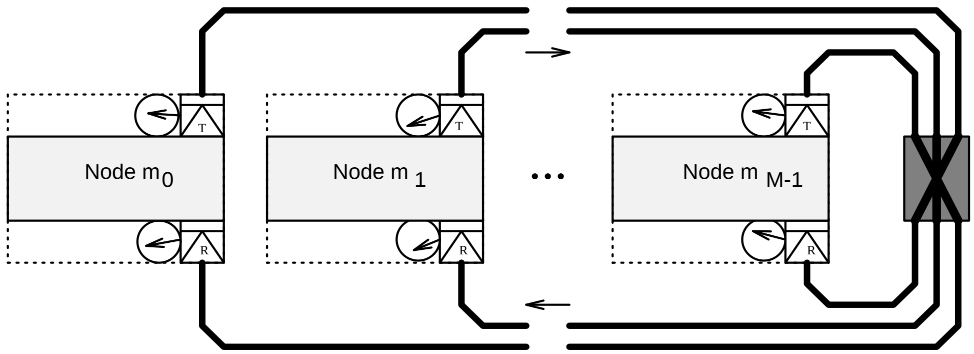

Figure 1: Network architecture: star-coupled configuration shown with tunable transmitters and receivers to achieve WDM with one fixed receiver at each node for I-TDMA*; one fixed and one tunable receiver at each node for TDMA-C.

node that track the availability of destination nodes and data channels. This protocol decouples the system size and the number of data channels. Variable sized packets are supported without loss of utilization with small packets. TDMA-C employs a single tunable transmitter for both data and control packets, one fixed receiver to monitor the control channel and one tunable receiver to receive data packets on any channel per node. The source node informs the destination node of the data channel that will be used for data packet transmission and the size of the data packet through a control packet. As shown in [2], TDMA-C achieves a significant improvement in performance over DT-WDMA. Additional advantages of TDMA-C over DT-WDMA are the support of variable packet sizes, elimination of destination conflicts requiring the execution of an arbitration algorithm as in [10], scalability since the number of channels and nodes were no longer coupled, and simpler implementation.

The pre-allocation protocol (I-TDMA*) being introduced and investigated employs time-division multiplexing on all the channels in the system. Each node has a chance to transmit on any channel once per cycle. The network architecture has one tunable transmitter and one fixed receiver that is parked on its home channel. I-TDMA* is similar to the I-TDMA protocol [18], but eliminates the severe head-of-line effects [20] that significantly hindered the performance of I-TDMA due to a single transmitter queue [18, 21]. I-TDMA* employs C transmitter queues at every node, one for each channel.

Section 2 presents the photonic network architecture under consideration. Refer to [2] for background information on optical devices and network architectures. Section 3 defines the access protocols. The behavior of the protocols is then examined in Section 5 in terms of the performance metrics and varying system parameters. The performance comparison of the two different classes of protocols enables comparison in terms of the hardware required and the performance improvement levels attainable. The analysis is done through extensive discrete-event simulation.

2 Network Architecture

Fig. 1 illustrates the network architecture used in this paper. A *fixed* transmitter or receiver is defined as a device that cannot alter its operating wavelength, *slow tunable* devices require a time greater than the nominal data packet transmission time to switch between channels, while *fast tunable* devices require a time less than the nominal data packet transmission time to switch between channels. Slow tunable devices are typically tuned to a specific wavelength at system start-up. The device will remain parked at this wavelength and is not switched during normal operation.

M:	number of interconnected nodes
C:	number of data channels in the network
m_i:	address of node i; $1 \leq i \leq M$
c_0:	control channel in TDMA-C
c_i:	data channel i, $1 \leq i \leq (C-1)$ for TDMA-C
	data channel i, $0 \leq i \leq (C-1)$ for I-TDMA*
α:	switching latency
L:	ratio of length of data packet to length of control packet
λ:	packet generation rate at each node
Γ:	network throughput in packets per unit time
t:	average delay per packet
U:	utilization of data channels

Table 1: Summary of notation

Both systems consider M nodes which are numbered as $\{m_0, m_1, \ldots, m_{M-1}\}$, and C wavelength channels numbered $\{c_0, c_1, \ldots, c_{C-1}\}$. Table 1 summarizes the notation definitions.

2.1 I-TDMA*

The objective of pre-allocation is to avoid the requirement that each node possess both a tunable transmitter and a tunable receiver. This approach does not require a control channel and all available channels are used for data packet transmission.

Channels may be pre-allocated either for transmission or reception. Channels pre-allocated for packet transmission require each node to possess a tunable receiver and a fixed or slow tunable transmitter. The destination node must tune its receiver to the home channel of the source node to receive a packet. DT-WDMA is such a protocol, however, it uses a control channel to inform the destination node of a source nodes' intention to transmit.

The I-TDMA* protocol introduced in this paper is based on channels pre-allocated for packet reception where each node receives on its home channel. I-TDMA* avoids collisions by time multiplexing access to the channels. Time is slotted to data packet lengths on each channel, and every node in the system has a chance to transmit to a destination node on each channel every cycle. Each node in the system has a tunable transmitter and a fixed or slow tunable receiver. A source node tunes its transmitter to the home channel of the destination node and transmits according to the access protocol. A node receives and processes all traffic along its home channel. Optical self-routing, where a node only receives data that is destined to it, is achieved when $M = C$ since a home channel is not shared. A home channel is shared and partial self-routing is achieved when $M > C$ [7]. Global tables mapping home channel allocation are not needed. A source node can determine the home channel of a destination node through the destination node number, the number of nodes in the system M and the number of channels C.

Each node has a receiver which is tuned to its home channel. A source node can compute the home channel of the destination node through a simple computation based on the channel allocation policy. Node m_i is assigned c_i as its home channel based on the allocation policy, where $c_i \in \{0, 1, 2, \ldots, C-1\}$ and $0 \leq i \leq M-1$. Two possible channel allocation schemes for a given station m_i are:

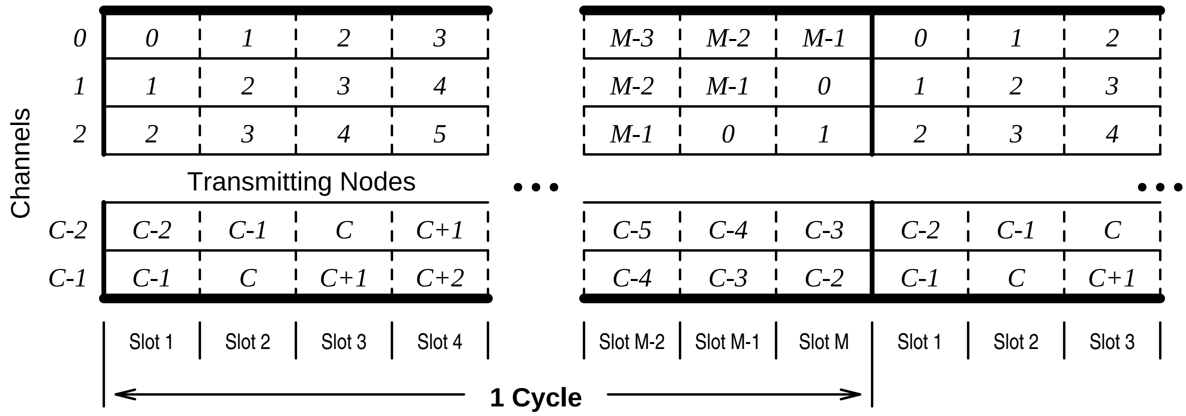

Figure 2: Allocation map for I-TDMA*. Channel/Transmitting Node space-time diagram when $M > C$ and cycle length is M slots.

Interleaved allocation: $c_i = m_i \bmod C$

Neighbor allocation: $c_i = \left\lfloor \frac{m_i}{\lceil M/C \rceil} \right\rfloor$

C transmitter queues are maintained at each node to buffer packets directed to each channel. This reduces the *head-of-line* effect seen in I-TDMA due to the single transmitter queue [18]. Access is slotted to data packet boundaries but time is described in terms of control slots as described in the following section.

2.2 TDMA-C

A node may transmit or receive on any data channel as well as the control channel. The control channel is defined as c_0; and c_i, $1 \le i \le (C-1)$, denotes a data channel. Each node has a tunable transmitter capable of tuning to any channel. Concurrent data packet transmission and data packet reception is supported. The receiver subsystem consists of two receivers. The first receiver (R_0) continually monitors the control channel to receive all transmitted control packets. The second receiver (R_1) is tunable and is used to receive data packets along any of the $C-1$ data channels. The control packet has four integer fields s, d, i and L: s, $1 \le s \le M$, identifies the source node address m_s; d, $1 \le d \le M$, identifies the destination node address m_d; i, $1 \le i \le C-1$, identifies channel c_i as the selected data channel, and L indicates the data packet length.

Time is normalized to the *control slot*, the time required for the transmission of a control packet, and taken to be one unit of time. Data packets are taken to be a positive integer L times the length of a control packet.

3 Description of Protocols

This section describes the two TDM/WDM protocols being considered in this paper. I-TDMA* employs time-division multiplexing on all the WDM channels whereas the TDMA-C protocol employs time-division multiplexing only on its control channel.

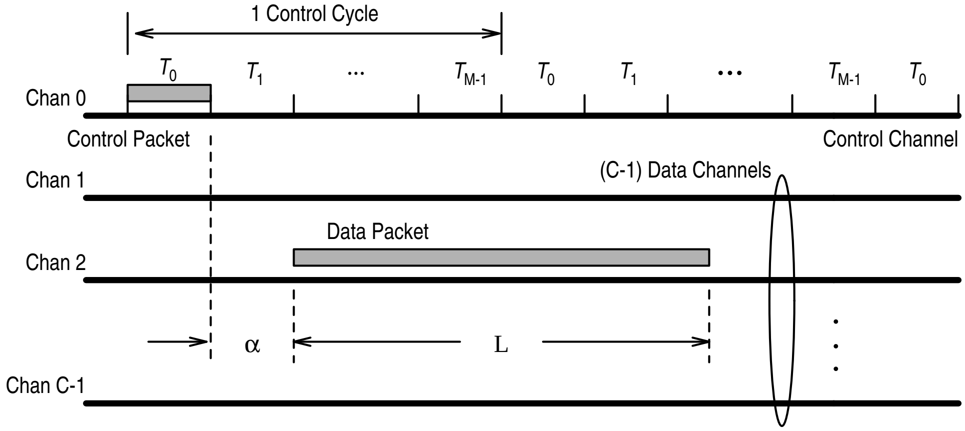

Figure 3: Time-Space diagram illustrating access to the control channel and transmission on the data channels.

3.1 I-TDMA*

I-TDMA* has evolved from I-TDMA described in [18, 22]. Each node now maintains C queues; one queue per channel. This eliminates the severe head-of-line problems of I-TDMA which hampered its ability to take advantage of an increase in the number of channels. This was due to a single queue of variable capacity per node which stored packets destined for all channels [18]. Packets to be transmitted on a channel in I-TDMA* are stored in the corresponding channel queue. This eliminates the head-of-line effect observed in I-TDMA and improves the network utilization.

Time is slotted where a slot is the time required for the tunable transmitter to tune to the required channel plus the data packet transmission time. Every node has a chance to transmit a data packet on every channel in a cycle. Determining the slot which is assigned to a particular source-destination pair is simple and decentralized and based on the home channel allocation policy defined in Section 2.1. Fig. 2 shows a channel/transmitting node allocation map when $M > C$. When $M = C$ the cycle length is $M - 1$ slots where each node has a slot reserved for it on each channel (other than its home channel) during each cycle assuming that a node will not be required to transmit to itself. Optical self-routing is achieved in this case where a node only receives traffic destined to it. Partial self-routing is achieved with $M > C$ and the cycle has a length of M slots since home channel transmission may be necessary since it may be shared. A node is assigned a total of C slots per cycle and remains idle for the remaining $M - C$ slots. However, as shown Fig. 2, the channels are fully allocated.

3.2 TDMA-C

Access to the control channel is based on a static cyclic slot allocation scheme. Each node is assigned one control slot per cycle, and all nodes have the opportunity to transmit a control and data packet during each cycle. The slot allocation is static and does not change with load.

A control cycle consists of M control slots as shown in Fig. 3. Every node has an assigned control slot it uses to reserve access on a data channel if backlogged. In Fig. 3, node m_0 transmits a control packet in control slot T_0. The transmitter then waits for α time slots before transmitting the data packet on the selected data channel. The delay α is defined as the *switching latency*. The switching latency is defined as $\alpha = \max\{t_s, t_r\}$, where t_s is the time required by the transmitter of the source to switch to the selected data channel wavelength, and t_r is the time required by the target node to receive and decode the control packet and switch its tunable receiver to the selected data channel. As described below, the dependence on the optical devices switching delay can be reduced by overlapping α.

Collisionless transmission is achieved by this protocol through the use of status tables. Each node maintains two tables: a table to track the status of the data channels to eliminate data channel collision, and a table to avoid destination conflict by tracking the status of the R_1 receiver at each node. This is why each node has receiver R_0 *parked* on the control channel: all transmitted control packets are received by all nodes (including the node that transmitted the control packet). R_0 updates the two status tables at the end of each control slot after receiving and decoding a control packet. If m_j transmits a control packet targeting m_i on data channel c_k, all nodes add $L + \alpha$ against entry i in their node status table and entry k in the channel status table. The entries indicate the number of time slots that the resources will be busy. All positive entries of each table are decremented at the end of every control time slot to update the remaining busy control slots.

The feasibility of TDMA-C protocol depends on the ability of the node to receive and process the information contained in the control packets at the speed of the optical network. A VLSI chip has been designed to be used in the receiver subsystem for maintaining the channel and node status tables [23]. The chip is designed to receive serial data at a speed of 2 Giga bits per second. The data is decoded, the entries in the tables are identified and updated. The chip was designed using *Octtools* on a MOSIS tiny chip frame and fully simulated with MUSA. The design of the chip incorporated the results of the performance analysis of [1] where it was shown that the optimum ratio of nodes to channels is 2:1. The chip was designed in a modular fashion, each implementing the status tables for 4 nodes and 2 channels, such that an array of chips could be used for larger systems.

A backlogged node checks its status tables at the beginning of its preallocated time slot. If the target node has a status table entry of less than or equal to 0, it is considered idle. If the target node is idle, the transmitter then checks for any available data channel. A data channel is considered idle if its status table entry is less than or equal to α. This achieves *overlap* of the switching latency α. The control packet is then formed with the source, target, selected data channel and packet length identifiers. If a node is not backlogged, its control slot remains idle during that cycle. In case the target is busy or an idle data channel is not available, the transmitter waits until the next cycle to attempt transmission.

The features of the TDMA-C protocol are:

- ♢ No collision on either the control channel or the data channels
- ♢ Arbitration is not required at target node
- ♢ Support of variable sized data packets
- ♢ The switching latency is overlapped to decrease the impact of the optical device switching characteristics
- ♢ Flexibility in using channels since any free channel can be used
- ♢ Easily adaptable to a change in the number of interconnected nodes and channels
- ♢ Complexity of implementation is reduced due to simple access arbitration of control channel along with its collisionless nature

4 Performance Models

Performance models enable prediction of the protocol behavior due to changes in system parameters such as M, C, L, α and λ. Table 1 summarizes the architecture parameters, and the notation used to denote the performance metrics.

The performance metrics of primary concern are network throughput and average packet delay. Throughput of the network is studied in terms of packets per time unit. The packet delay is defined as the time from when it is generated until it is received by the target node.

The model assumptions are:

1. All nodes are assumed to behave independently.
2. Each case considers a fixed length packet L times the control packet length.
3. Packet generation follows a poisson process with a rate λ packets per control slot per node.
4. A data packet can be transmitted on any idle data channel with equal probability for TDMA-C protocol.
5. Data packet synchronization is at control *slot* boundaries for TDMA-C and at data packet boundaries for I-TDMA*.
6. Uniform reference model: a packet generated at m_i is targeted to m_j with probability $\frac{1}{M-1}$ for $i \neq j$, $1 \leq i \leq M$ and $1 \leq j \leq M$; and with probability 0 when $i = j$.
7. Single transmitter queue of infinite capacity at each node for TDMA-C and C transmitter queues of infinite capacity at each node for I-TDMA*.

The simulators are based on a stochastic self-driven discrete event model. The simulators were written in the C programming language with SimPack. SimPack is a C based library of routines that provide discrete-event and random variate facilities [24, 25]. Steady state transaction times and throughput were measured. Simulation convergence was obtained through the replication/deletion method [26], with a 98% confidence in a less than 2% variation from the mean.

Average Packet Delay: The packet delay (t) is defined as the time taken from the instant a packet is generated at the node to the instant it is received at the target node. This includes the synchronization delay on the control channel for TDMA-C, the synchronization time on data channels for I-TDMA*, waiting time in the queue and the packet transmission time.

Data Channel Utilization: The percent of time the data channels are busy transmitting data packets defines the data channel utilization (U).

Network Throughput: Throughput is expressed as the number of packets successfully transmitted across the network per unit time. The throughput of the network in terms of packets per unit time, denoted as Γ, is $\Gamma = \frac{U}{L}$.

Data channels can be used according to their availability in the TDMA-C protocol. Any variation in C causes a large impact on the performance of the protocol. Two important conditions result from this: *data-channel starvation* and *control-channel limitation*. Data channel starvation occurs when data packet transmission is blocked due to the unavailability of a data channel when the target node is not blocked (when $C < M$). A system is defined as control limited when the throughput is bound due to destination blocking or insufficient access to the control channel. An increase in data channels improves performance characteristics for a protocol limited by data channel starvation, but has little or no effect on a control limited system. This issue is examined in greater detail in Section 5.4 where the impact of variations in C is studied.

The next section analyzes the protocol in terms of the parameters given in the above sections for variations in the system parameters.

5 Performance Analysis

This section analyzes the performance of the two protocols in terms of average packet delay and network throughput. The effect of varying the packet generation rate, data packet length, number of nodes and data channels, and switching latency is analyzed in the following sections. Note that time is normalized to control slot lengths.

5.1 *Impact of Variations in Packet Size*

This section compares the performance of the two protocols in terms of packet delay and network throughput for variations in data packet length. This section illustrates the advantages of time-multiplexing the control channel rather than the data channels.

Figs. 4(a)-(c) compare the average packet delay and network throughput for the TDMA-C and I-TDMA* protocols with varying packet lengths $L \in \{16, 32, 64\}$, $\alpha = 0$, $M = 32$ and $C \in \{8, 16, 32\}$. The average packet delay is comprised of the synchronization delay, any queueing delays plus the packet transmission delay. For the I-TDMA* pre-allocation protocol, the synchronization delay is given by

$$\mathcal{T}_{\mathcal{P}} = \begin{cases} \lceil (M-1)L/2 \rceil & \text{if } C \leq M \\ \lceil (M-2)L/2 \rceil & \text{if } C = M \end{cases} \tag{1}$$

Synchronization of the TDMA-C reservation based protocol is on control slot boundaries so the synchronization delay is given by

$$\mathcal{T}_{\mathcal{R}} = \begin{cases} \lceil (M-1)/2 \rceil & \text{if } C \leq M \\ \lceil (M-2)/2 \rceil & \text{if } C = M \end{cases} \tag{2}$$

The synchronization delay is a large percentage of the average packet delay. Fig. 4(a) plots the delay-throughput graphs for $M = 32$, $C = 8$, $\alpha = 0$ and $L \in \{16, 32, 64\}$. The average packet delay for $L = 16$ and very low packet generation rates is about 32 for TDMA-C in comparison to about 250 for I-TDMA*. This illustrates the penalty of multiplexing based on data packet lengths which results in the

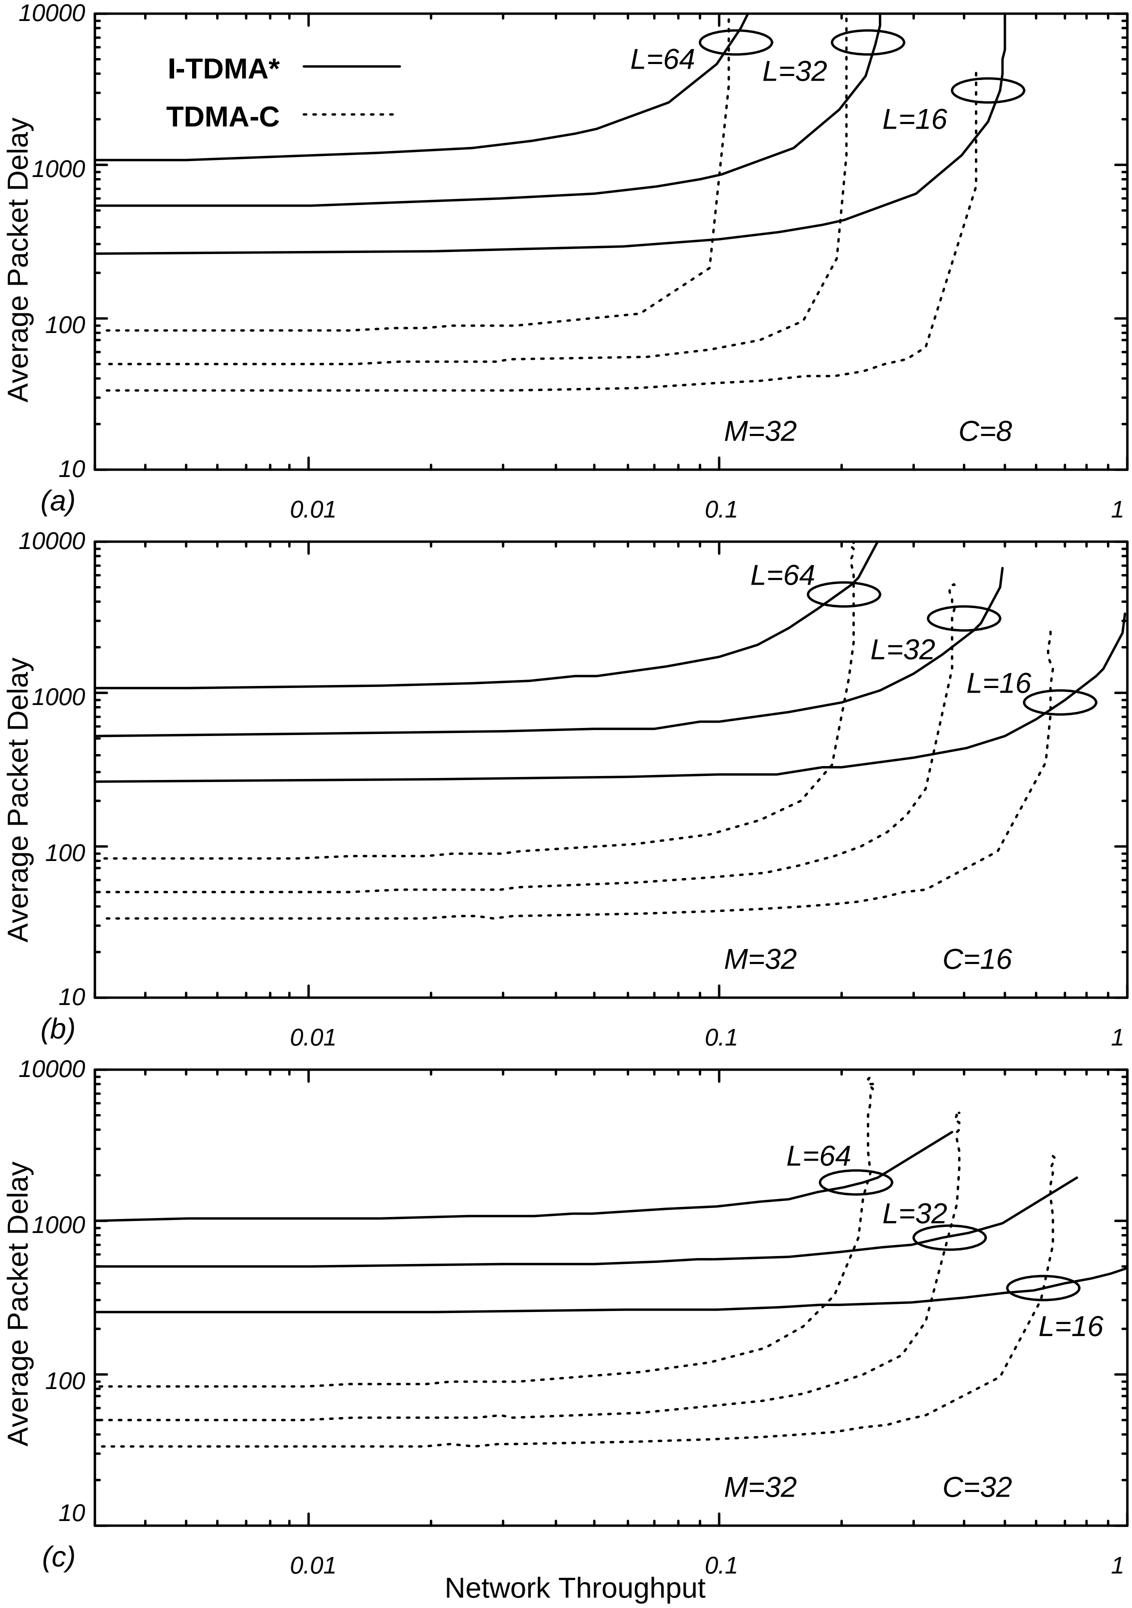

Figure 4: Comparison of I-TDMA* and TDMA-C protocols: Variation in C with $L \in \{16, 32, 64\}$, $\alpha = 0$ and $M = 32$. (a) $C = 8$ (b) $C = 16$ (c) $C = 32$.

large differences in the synchronization delays for the two protocols. The average delay of I-TDMA* is shown to be very sensitive to increases in packet size since the cycle length is directly proportional to data packet length.

For increased values of packet sizes, the delay increases for both protocols but I-TDMA is far more sensitive. For example, as L increases from 16 to 32 and from 32 to 64, the lightly loaded delay of I-TDMA* increases by 100% at each step. This is expected since the lightly loaded delay of I-TDMA* is approximately $\mathcal{T}_{I-TDMA*} = L[1 + (M-1)/2]$. TDMA-C has a lightly loaded delay of the form $\mathcal{T}_{TDMA-C} = [L + (M-1)/2]$ so is less sensitive to an increase in packet size, particularly for large systems. Fig. 4 shows that an increase in the number of channels has little impact on the delay of a lightly loaded system but significantly increases the maximum capacity of both protocols.

The maximum throughput for $M = 32$, $L = 16$ and $C = 8$ is 0.43 and 0.5 for TDMA-C and I-TDMA*, respectively. The maximum possible throughput (when data channel utilization is 100%) is C/L for I-TDMA*, and $(C-1)/L$ for TDMA-C since it must allocate one WDM channel as its control channel. Fig. 4 shows that both protocols can achieve their maximum channel capacities with varying packet lengths. However, TDMA-C supports varying packet sizes while I-TDMA* is slotted on data packet boundaries so support for dynamically varying packet sizes is not provided. I-TDMA* is able to achieve the maximum throughput attainable. TDMA-C has a slightly lower system capacity because the source nodes contend (in a collisionless fashion) for the data channels and destination receivers. As noted in [1], TDMA-C suffers head-of-line effects since transmission could be blocked due to receiver contention. The head-of-line effects with TDMA-C could be avoided by having C queues at each node, as in I-TDMA*, so transmission will not be blocked due to a single busy receiver. However, we have not done this since the performance degradation is not nearly as severe as with I-TDMA.

Figs. 4(b)-(c) plot the delay-throughput graphs as the number of channels increase from $C = 16$ to $C = 32$. These figures show that the protocols still maintain maximum throughput as the number of channels increase. However, there is little performance improvement forTDMA-C with $C > M/2$ due to the head-of-line effects caused by receiver contention. In fact, Fig. 4 shows that the performance of TDMA-C with $C = M/2$ is only slightly less than with $C = M$.

This section shows that as the data packet size increases, the average packet delay increases proportionately and the maximum network throughput decreases correspondingly. The maximum network throughput also depends on the number of channels in the network. For TDMA-C, the maximum throughput does not increase with increases in channels beyond $M/2$. This is examined in greater detail in Section 5.4. I-TDMA* attains higher maximum throughput because of the static allocation on the data channels and the elimination of head-of-line effects.

5.2 *Impact of Switching Latency*

This section studies the effect of varying switching latency (α) on the protocols and illustrates the dependence of the protocols on the availability of fast tunable devices. As technology improves, devices which support a large number of WDM channels with a fast switching time may become available. Such devices are being fabricated in research laboratories and are not expected to be commercially available for some time. However, devices (transmitters and receivers) able to support a limited number of WDM channels with microsecond to millisecond switching speeds are (or are expected soon) available.

The switching latency of the I-TDMA* protocol is the time required by the tunable transmitter to tune to

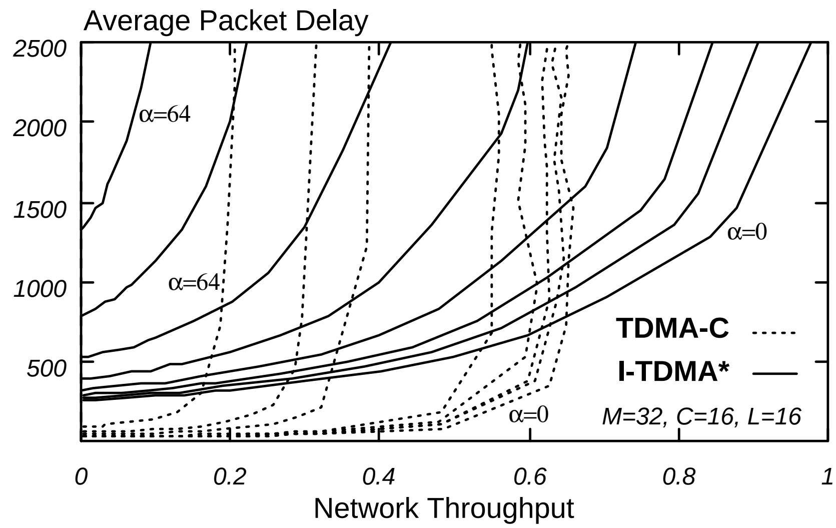

Figure 5: Impact of switching latency on I-TDMA* and TDMA-C protocols: $\alpha \in \{0, 1, 2, 4, 8, 16, 32, 64\}$, $M = 32$, $C = 16$ and $L = 16$.

---

the wavelength of the required home channel. The switching latency for TDMA-C, defined in Section 3.2, is also mainly composed of the switching latency of the optical devices. This section assumes that the switching latency has the same duration for both protocols.

There is a cycle synchronization delay with both TDM/WDM protocols. The cycle synchronization delay of the I-TDMA* protocol can be used to partially overlap the switching latencies in a lightly loaded system. As the system load increases, the probability that a packet arrives for transmission at an idle node decreases. This decreases the likelihood that α can be overlapped with the synchronization delay. Depending on the relative values of α and L, different strategies can be devised for I-TDMA* overlap [27]. This paper considers the simplest technique that is well suited to the situation when $\alpha < L$: extend the slot length to be of duration $L + \alpha$ rather than L. This strategy essentially sacrifices $\frac{\alpha}{L + \alpha}$ utilization, but this section shows that this is an effective technique for small α.

TDMA-C cannot use its synchronization delay to overlap the switching latency but achieves overlap through the mechanism described in Section 3.2. This overlap technique is particularly useful when $M > C$ and at higher traffic rates. This section examines the impact of switching latency on the delay and throughput characteristics of both the protocols. In particular we are interested in the case when data channel starvation may occur ($C < M$) for TDMA-C.

Fig. 5 shows the impact of α on the delay and throughput of TDMA-C and I-TDMA*. For a packet size of $L = 16$ slots, this figure varies the switching latency as $\alpha \in \{0, 1, 2, 4, 8, 16, 32, 64\}$ slots. This figure assumes that $M = 32$ and $C = M/2$. This graphs varies the switching latency from 0 to four times the nominal data packet transmission time, with smaller values ($\alpha < 16$) corresponding to fast tunable devices. The remainder of this section considers the magnitude of the variations in performance as the switching latency increases.

Average Delay: Fig. 5 illustrates an advantage of TDMA-C where the the impact of the delay is additive when the system is lightly loaded. For low packet generation rates, the average delay of TDMA-C is approximately $\frac{M-1}{2} + L + \alpha$. For example, the delay is approximately 32 when $\alpha = 0$ and holds the form of $32 + \alpha$ as α is increased. This can be verified from the figure where the delay is 33, 34, 36, 40, 49, 64 and 96 for $\alpha = 1, 2, 4, 8, 16, 32, 64$ respectively. The overlap mechanism cannot eliminate the impact of the switching latency to the delay, but it can reduce its impact to the maximum network throughput as considered below.

The impact of the switching latency on the delay is far more significant for I-TDMA* with the slot extension strategy than with TDMA-C. Fig. 5 illustrates the impact to delay for I-TDMA* when the slot length is extended to $L + \alpha$. The average cycle synchronization time with I-TDMA* is $\left\lceil \frac{M-1}{2} \right\rceil (L + \alpha)$, so the average delay for a lightly loaded system is $(M + 1)(L + \alpha)/2$. This can be verified from the figure where the delay is 264, 281, 297, 330, 397, 529, 794, 1323 time units for $\alpha = 0, 1, 2, 4, 8, 16, 32, 64$ respectively for $M = 32$, $C = 16$ and $L = 16$.

This shows that the percent increase in delay with I-TDMA* is directly proportional to the ratio of α to L. For example, $\alpha < L$ is required with I-TDMA* for the percent increase in delay to be less than a 100% increase from its value at $\alpha = 0$. This constraint is relaxed with TDMA-C where $\alpha < \frac{M-1}{2} + L$ is required for less than a 100% increase delay from its delay at $\alpha = 0$.

Maximum Network Throughput: The impact of the switching latency on the network throughput can also be seen in Fig. 5 where the maximum network throughput is shown to decrease as α increases for both protocols. However, the impact on TDMA-C is less than with I-TDMA* for lower values of α. Fig. 5 illustrates how the overlap mechanism can reduce the impact to maximum throughput of TDMA-C with increases in switching latency.

With slot extension, the maximum network throughput of I-TDMA* is $\frac{C}{L + \alpha}$ so the percent decrease in maximum network throughput from $\alpha = 0$ is $\frac{\alpha}{L + \alpha} \times 100$. For example, the maximum network throughput decreases by 14% and 30% as α is increased from $\alpha = 1$ to $\alpha = 8$ for TDMA-C and I-TDMA*, respectively.

For larger values of α, the overlap mechanism is no longer able to mask the latency and the maximum throughput degrades with TDMA-C. Maximum network throughput decreases by 47% and 60% for TDMA-C and I-TDMA*, respectively, as α is increased from 16 to 64. For very large values of switching latencies, the network capacity is about the same for both protocols ($C/(L + \alpha)$ for I-TDMA* and $(C - 1)/(L + \alpha)$ for TDMA-C) since TDMA-C becomes data channel limited.

5.3 Impact of Variations in System Size

This section analyzes the effect of variations in the number of nodes and packet length on the delay and throughput when the number of data channels is scaled at $C = M/2$ with increased system size. Figs. 6(a)-(c) plot the average packet delay and network throughput for varying system size as $M \in \{8, 16, 32\}$, $\alpha = 0$, $C = M/2$ and $L \in \{16, 32, 64\}$.

Figs. 6 illustrates the limitation of the I-TDMA* in terms of scalability since the cycle is proportional to the number of nodes and packet size. The delay impact of increasing M is shown to be far more pronounced with I-TDMA*. As observed before, the synchronization delay increases as $(M - 1)L/2$ for I-TDMA*

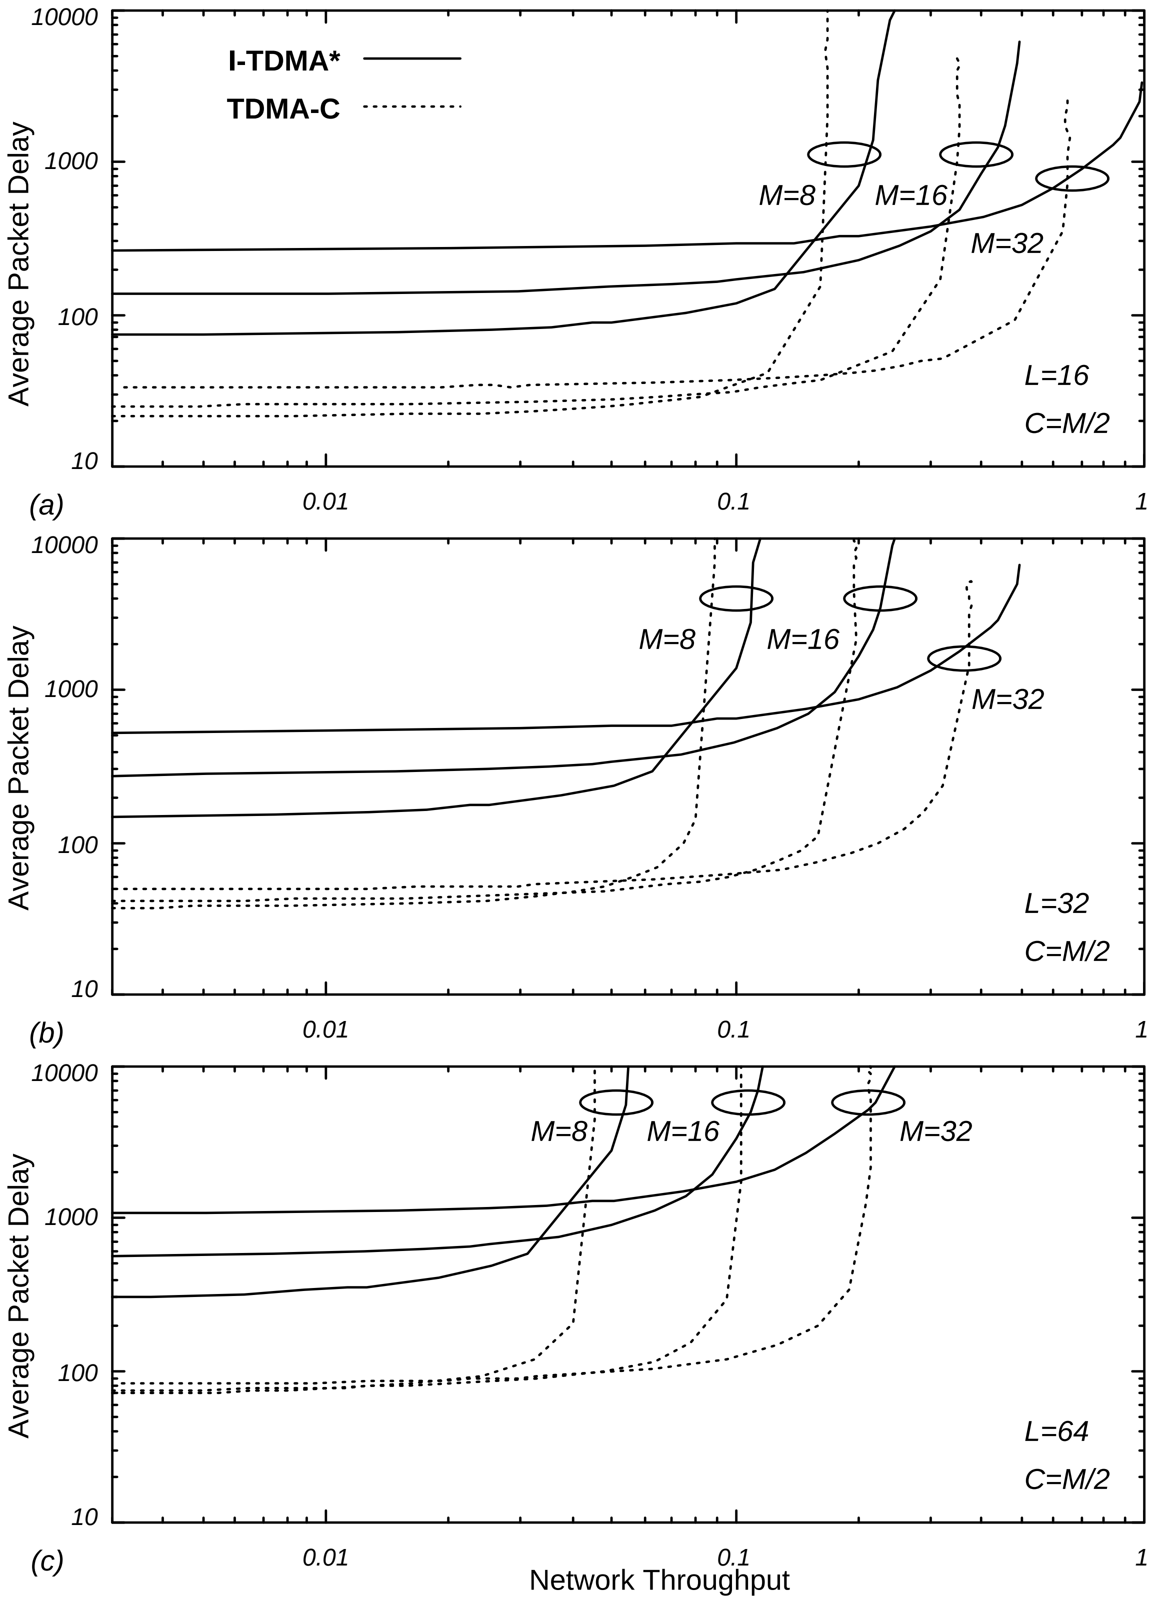

Figure 6: Comparison of I-TDMA* and TDMA-C protocols: Variation in L with $M \in \{8, 16, 32\}$, $\alpha = 0$ and $C = M/2$. (a) $L = 16$ (b) $L = 32$ (c) $L = 64$.

compared to $(M-1)/2$ for TDMA-C. This results in higher delays and lower capacity for increases in the system size with I-TDMA*. The delay impact on the TDMA-C is relatively insignificant in comparison to I-TDMA* for a lightly loaded system, especially for large packet sizes as seen from Fig. 6(c). As M increases, the cycle length increases thereby causing higher delays in the I-TDMA* protocol.

Although I-TDMA* is shown to suffer delay penalties as the system is expanded, Fig. 6 shows that the maximum throughput continues to scale and the maximum theoretical value is obtained. Fig. 6 shows that I-TDMA* consistently out-performs the TDMA-C protocol in terms of maximum throughput. For example, the maximum throughput for TDMA-C is about 33% lower than I-TDMA* at $M = 32$ and $L = 16$. As the packet size increases, the maximum throughput decreases and the difference between the maximum throughput of the two protocols also decreases. TDMA-C has a 26% and 16% lower maximum throughput than I-TDMA* for $L = 32$ and $L = 64$, respectively.

5.4 Impact of Variations in Data Channels

This section analyzes the effect on delay and throughput due to variations in the number of data channels. Any idle channel can be used for data transmission in the TDMA-C protocol so a change in the number of WDM channels results in a significant impact on the network throughput providing the system is data-channel starved. Fig. 7 plots the effect of varying the number of the channels in a system with a fixed number of nodes $M = 32$. Three cases of packet sizes are graphed, $L \in \{16, 32, 64\}$, and the number of channels is varied as $C \in \{M, M/2, M/4\}$.

As C is increased from $M/4$ to $M/2$, the average packet delay reduces and the network can handle higher packet generation rates. The effect of increasing C from $M/2$ to M is not significant with TDMA-C. This is because the system is *data channel starved* when $C = M/4$, and *control channel limited* as C is increased to M due to the head-of-line effect caused by destination blocking. Fig. 7 shows that there is little advantage of expanding the number of channels beyond $C = M/2$ for TDMA-C for all three cases of packet sizes. This effect is not seen in the I-TDMA* protocol because of the multiple queues at each node to avoid the destination blocking head-of-line effect. This shows that I-TDMA* can take full advantage of an increase in the number of channels. An increase in the number of channels does not decrease the average delay of a lightly loaded I-TDMA* system but consistently increases the maximum network throughput.

6 CONCLUSIONS

This paper introduced a time multiplexed collisionless WDM media access protocol for systems with channels pre-allocated to nodes for data packet reception. The network architecture has a tunable transmitter and a fixed receiver. Pre-allocation is a low-cost approach which can be built with *off-the-shelf* components [28]. The performance of the proposed protocol is compared to a collisionless control channel based reservation protocol. The network architecture has nodes with one tunable transmitter and two receivers. One of the receivers is fixed and used to sense the control channel, and the other receiver is tunable for the data packets. Collisionless transmission is achieved through the use of status tables, and arbitration is not required by the target node. The control channel is time multiplexed, providing maximum throughput and high stability of the network with heavy traffic. The number of interconnected nodes and data channels are independent with both protocols, reducing the problems associated with system expansion. The reservation-based protocol has flexibility in the allocation of data channels, and supports

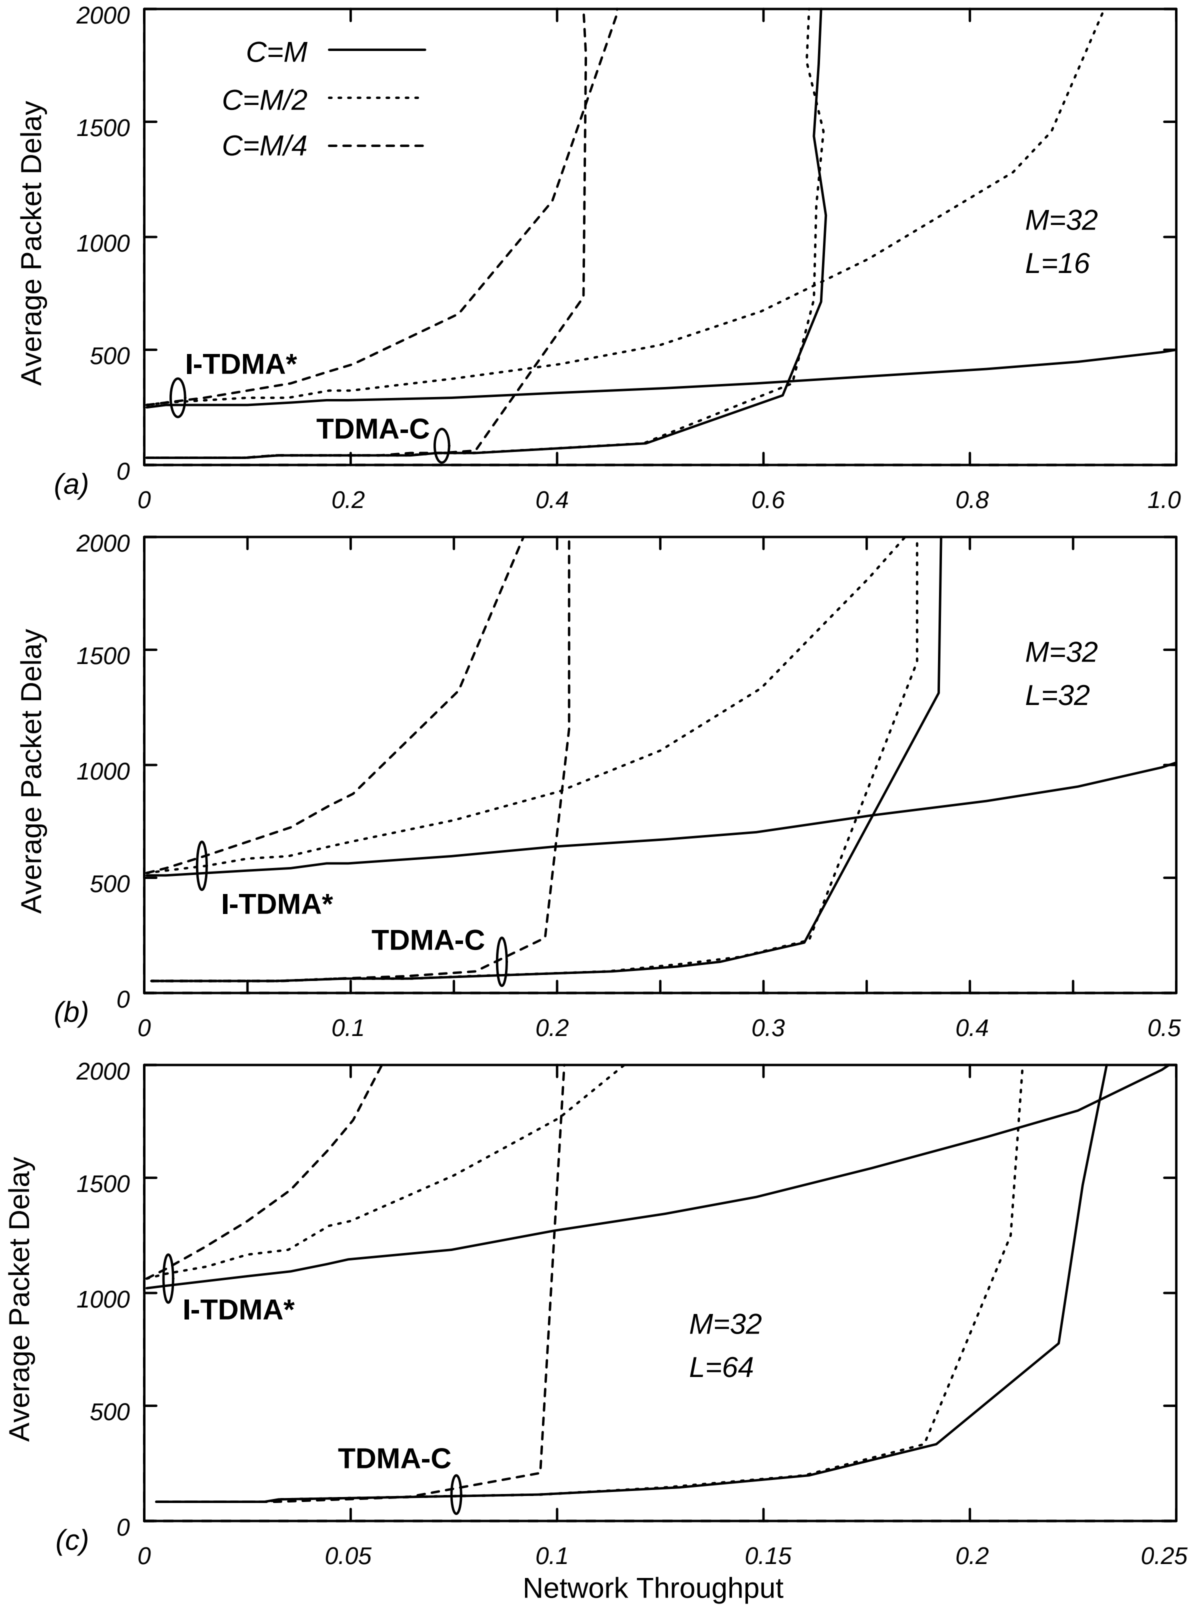

Figure 7: Comparison of I-TDMA* and TDMA-C protocols: Variation in L with $C \in \{M, M/2, M/4\}$, $\alpha = 0$ and $M = 32$. a) $L = 16$ (b) $L = 32$ (c) $L = 64$.

variable sized packets and is adaptable to changes in the number of nodes and channels. Packet delay and throughput of the network are studied with variations in the packet length, number of interconnected nodes and data channels in the network, and switching latency.

REFERENCES

[1] K. Bogineni and P. W. Dowd, "A collisionless multiple access protocol for a wavelength division multiplexed star-coupled configuration: Architecture and performance analysis," *IEEE Journal on Lightwave Technology*, (To Appear), 1992.

[2] K. Bogineni, M. Carrato, and P. Dowd, "Collisionless multiple access protocols for wavelength division multiplexed star-coupled photonic networks," *International Journal in Computer Simulation*, (To Appear), 1992.

[3] C. A. Brackett, "Dense wavelength division multiplexing networks: Principles and applications," *IEEE Journal on Selected Areas of Communications*, vol. 8, pp. 948–964, Aug. 1990.

[4] M. Fine and F. A. Tobagi, "Demand assignment multiple access schemes in broadcast bus local area networks," *IEEE Transactions on Computers*, vol. c-33, pp. 1130–1159, Dec. 1984.

[5] P. W. Dowd, "Optical interconnections for computer communications," Tech. Rep. TR01.A961, IBM Corporation, Apr. 1989.

[6] P. W. Dowd, "Optical bus and star-coupled parallel interconnection," in *Proc.* 4^{th} *International Parallel Processing Symposium*, (Los Angeles, CA), pp. 824–838, Apr. 1990.

[7] P. W. Dowd, "Random access protocols for high speed interprocessor communication based on a passive star topology," *IEEE Journal on Lightwave Technology*, vol. 9, pp. 799–808, June 1991.

[8] P. W. Dowd, "High performance interprocessor communication through optical wavelength division multiple access channels," in *Proc.* 18^{th} *International Symposium on Computer Architecture*, (Toronto, Canada), pp. 96–105, May 1991.

[9] I. M. I. Habbab, M. Kavehrad, and C. E. W. Sundberg, "Protocols for very high-speed optical fiber local area networks using a passive star topology," *IEEE Journal on Lightwave Technology*, vol. LT-5, pp. 1782–1793, Dec. 1987.

[10] M. S. Chen, N. R. Dono, and R. Ramaswami, "A media access control protocol for packet switched wavelength division multiaccess metropolitan area network," *IEEE Journal on Selected Areas of Communications*, vol. 8, pp. 1048–1057, Aug. 1990.

[11] H. B. Jeon and C. K. Un, "Contention-based reservation protocol in fiber optic local area network with passive star topology," *IEE Electronic Letters*, vol. 26, June 1990.

[12] H. Shi and M. Kavehrad, "Aloha/Slotted-CSMA protocol for a very high-speed optical fiber local area network using passive star topology," in *Proc. IEEE INFOCOM'91*, pp. 12D.4.1. – 12D.4.6., Mar. 1991.

[13] G. Sudhakar, N. Georganas, and M. Kavehrad, "Slotted Aloha and Reservation Aloha protocols for very high-speed optical fiber local area networks using passive star topology," *IEEE Journal on Lightwave Technology*, vol. 9, pp. 1411–1422, Oct. 1991.

[14] F. Jia and B. Mukherjee, "The receiver collision avoidance (RCA) protocol for a single-hop WDM local lightwave network," in *Proc. ICC'92*, (Chicago), pp. 6–10, June 1992.

[15] A. Ganz, "End-to-end protocols for WDM star networks," in *IFIP/WG6.1-WG6.4 Workshop on Protocols for High-Speed Networks*, (Zurich, Switzerland), May 1989.

[16] A. Ganz and I. Chlamtac, "Path allocation access control in fiber optic communication systems," *IEEE Transactions on Computers*, vol. 38, pp. 1372–1382, Oct. 1989.

[17] A. Ganz and Z. Koren, "WDM passive star-protocols and performance analysis," in *Proc. IEEE INFOCOM'91*, pp. 9A.2.1. – 9A.2.10., Mar. 1991.

[18] K. M. Sivalingam, K. Bogineni, and P. W. Dowd, "Pre-allocation media access control protocols for multiple access WDM photonic networks," in *Proc. ACM SIGCOMM'92*, (Baltimore, Maryland), Aug. 1992.

[19] N. Mehravari, "Performance and protocol improvement for very high speed optical fiber local area networks using a passive star topology," *IEEE Journal on Lightwave Technology*, vol. 8, pp. 520–530, Apr. 1990.

[20] M. Karol, M. Hluchyj, and S. Morgan, "Input verses output queueing on a space division packet switch," *IEEE Transactions on Communication*, vol. COM-35, pp. 1347–1356, Dec. 1987.

[21] K. Bogineni and P. W. Dowd, "Performance comparison of reservation and preallocation protocols for star-coupled WDM photonic networks," in *Proc. First International Symposium on High-Performance Distributed Computing (HPDC-1)*, (Syracuse, NY), Sept. 1992.

[22] K. Bogineni and P. W. Dowd, "Performance impact of switching latency on WDMA protocols," *SPIE Proceedings (High-Speed Fiber Networks and Channels)*, vol. 1784, Sept. 1992.

[23] P. W. Dowd, R. Sridhar, K. Bogineni, and A. Ranganathan, "Design of a VLSI chip for use in the receiver subsystem in an optically interconnected parallel computer," Tech. Rep. CES-91-0802, State University of New York at Buffalo, Department of Electrical and Computer Engineering, Dec. 1991.

[24] P. A. Fishwick, "Simpack: C-based simulation seed tool package, Version 2.0," tech. rep., Department of CIS, University of Florida, Gainesville FL, Apr. 1992.

[25] M. MacDougall, *Simulating Computer Systems*. The MIT Press, 1987.

[26] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*. McGraw Hill, 1991.

[27] K. M. Sivalingam, K. Bogineni, and P. Dowd, "Design and performance analysis of preallocation protocols for WDM photonic networks," *Proceedings SPIE (High-Speed Fiber Networks and Channels)*, vol. 1784, Sept. 1992.

[28] M. Girard, C. Husbands, and P. Dowd, "Media access protocols for spectral sliced WDMA testbed," *SPIE Proceedings (High-Speed Fiber Networks and Channels)*, vol. 1784, Sept. 1992.