Analysis and Optimization of Transmission Schedules for Single-Hop WDM Networks

George N. Rouskas, *Member, IEEE,* and Mostafa H. Ammar, *Member, IEEE*

***Abstract*— We consider single-hop lightwave networks with stations interconnected using wave division multiplexing. The stations are equipped with tunable transmitters and/or receivers. A predefined, wavelength-time oriented schedule specifies the slots and the wavelengths on which communication between any two pairs of stations is allowed to take place. We define a wide variety of schedules and develop a general framework for analyzing their throughput performance for any number of available wavelengths, any tunability characteristics, and general (potentially nonuniform) traffic patterns. We then consider the optimization of schedules given the traffic requirements and present optimization heuristics that give near-optimal results. We also investigate how the number of available wavelengths (channels) affects the system throughput, and we develop techniques to efficiently share the available channels among the network stations. As a result, we obtain systems that are easy to scale while having very good performance.**

I. Introduction

WAVE DIVISION MULTIPLEXING (WDM) is emerging as a promising technology for the next generation of multiuser high-speed communication networks. WDM divides the low-loss spectrum of the optical fiber into nonoverlapping channels, each operating at a data rate accessible by the transmitting stations. The multiple channels introduce transmission concurrency and provide a means to overcome the so called electronic bottleneck [13]. As a result, WDM networks have the potential of delivering an aggregate throughput that can grow with the number of wavelengths deployed, and can be in the order of Tb/s.

In single-hop WDM networks, packet transmissions are possible only when there is a direct communication path established between the source and the destination. Single hop systems require tunable transceivers with a large tuning range $\times$ tuning speed product in order to fully utilize the capabilities of WDM. The state of the art in tunable laser and filter technology is discussed in [4]. Although commercial products are not yet available, recent advances in lightwave technology are promising.

Research in the area of single-hop systems has focused on the problem of how to efficiently allocate the bandwidth among the network stations (see [17] for an overview of the various proposed schemes). In general, the schemes differ in whether they require coordination between the transmitter and receiver prior to the actual packet transmission, or not. In the former case, one [12], [16], [5], [6] or more [14] control channels are reserved for the arbitration of the transmission requests. Schemes with no pretransmission coordination employ either random access protocols [10] or a wavelength-time assignment of the optical bandwidth [7].

The wavelength-time assignment technique is the extension of TDMA over a multichannel environment. A schedule specifies the slots within each frame and the channel on which packet transmissions are permitted between any source-destination pair. Our work deals with analysis and optimization of a wide variety of schedules for any number of wavelengths, any transceiver tunability characteristics, and general (potentially nonuniform) traffic patterns. Similar optimization problems, although in a different context, are addressed in [15], [1], [2]. In [8], [7], [3] several schedules are studied, and models are developed to analyze their performance. These works, in contrast to ours, do not deal with schedule optimization, and the analysis is restricted to uniform traffic. In [9] a method to schedule transmissions to minimize the tuning time is presented, without considering the effects on throughput.

This paper is organized as follows. In Section II we describe our system model, and in Section III we obtain throughput expressions for the various schedules. Section IV investigates the problem of obtaining an optimal schedule and presents several heuristics which yield very good results. In Section V a more general heuristic is developed. Section VI presents some numerical examples, and Section VII contains some concluding remarks.

II. System Model

We consider a network of N stations, each equipped with one receiver and one transmitter, interconnected through an optical broadcast medium that can support C wavelengths, $\lambda_1, \lambda_2, \cdots, \lambda_C$. In general, $C \le N$.[1] Depending on the tunability characteristics (tunable or fixed transmitters/receivers) we refer to the three resulting systems as TT-FR, FT-TR, or TT-TR. If the receivers (transmitters) are fixed, wavelength $\lambda(i) \in \{\lambda_1, \cdots, \lambda_C\}$ is assigned to the receiver (transmitter) of station $i, i = 1, \cdots, N$. The tunable transmitters (receivers),

Manuscript received July 6, 1993; revised November 15, 1993 and June 29, 1994; approved by IEEE/ACM Transactions on Networking Editor R. Ramaswami. This paper was presented in part at the IEEE INFOCOM '93 Conference, San Francisco, CA, March 30, 1993–April 1, 1993.

G. N. Rouskas was with the College of Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA. He is now with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695 USA.

M. H. Ammar is with the College of Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA.

IEEE Log Number 9409113.

[1] With one transceiver per station, at most N transmissions are possible at any given time, thus we do not consider the case $C > N$.

1063–6692/95$04.00 © 1995 IEEE

on the other hand, consist of lasers (filters) tunable over all wavelengths $\lambda_c, c = 1, \cdots, C$.

The network operates in a slotted mode, with a slot time equal to the packet transmission time plus the tuning time. All stations are synchronized to the slot boundaries [19]. We define σ_i as the probability that a new packet arrives at station i during a slot time, p_{ij} as the probability that a packet arriving at station i is destined to station j, and $\sum_j p_{ij} = 1$. Each station has $N-1$ single-packet buffers, one for packets to each possible destination; packets arriving to a full buffer are lost. This is an extension of the single-channel network model developed in [15]. The case of L-packet buffers, $L \geq 1$, is treated in [18].

It should be emphasized that this is a model of the media access control (MAC) layer. Packets that cannot be buffered at the MAC layer are not really lost, but are typically buffered at a higher layer. Thus, $\sigma_i p_{ij}$ characterizes the arrival process from the higher layer to the single-packet buffer for station j when that buffer is empty.

Time slots are grouped in frames of M slots. Within a frame, a_{ij} slots are assigned for packet transmissions between the source-destination pair (i, j). A schedule indicates, for all i and j, which slots during a frame can be used for transmissions from i to j, and can be described by the variables $\delta_{ij}^{(t)}, t = 1, 2, \cdots, M$, defined as

$$\delta_{ij}^{(t)} = \begin{cases} 1, & \text{if station } i \text{ has permission to} \\ & \text{transmit to station } j \text{ in slot } t \\ 0, & \text{otherwise} \end{cases} \quad . \quad (1)$$

Obviously, $a_{ij} = \sum_{t=1}^{M} \delta_{ij}^{(t)}$. To ensure fairness, all schedules we consider have the property that, if the traffic originating at i and terminating at j is nonzero, then at least one slot per frame is assigned for transmissions from i to j. Formally,

$$\forall i, j : \text{if } \sigma_i p_{ij} > 0 \text{ then } a_{ij} \geq 1 (\text{Fairness Condition}) \ . \quad (2)$$

For $k = 1, 2, \cdots, a_{ij}$, we let $d_{ij}^{(k)}$ denote the distance, in slots, between the beginning of the kth slot that i has permission to transmit to j, and the beginning of the next such slot, in the same or the next frame (see Fig. 1). We also define $w_{ij}^{(t)} \in \{\lambda_1, \cdots, \lambda_C\}$ as the wavelength on which i will transmit a packet to j in slot t. For a TT-FR (FT-TR) system we have that $w_{ij}^{(t)} = \lambda(j) \forall i, t$ ($w_{ij}^{(t)} = \lambda(i) \forall j, t$). Whenever $C < N$, in a TT-FR (FT-TR) system a number of receivers (transmitters) have to be assigned the same wavelength, and share a single channel $\lambda_c, c = 1, \cdots, C$. We let R_c and X_c, subsets of $\{1, \cdots, N\}$, denote the set of receivers and transmitters, respectively, sharing channel λ_c

$$R_c = \{j \mid \lambda(j) = \lambda_c\} \ \forall c; \ X_c = \{i \mid \lambda(i) = \lambda_c\} \ \forall c \ . \quad (3)$$

A. Transmission Modes

We define the *transmitting set*, $I_i^{(t)}, t = 1, 2, \cdots, M$, of station i in the tth slot in a frame as the set of stations to which i is permitted to transmit. The *receiving set*, $J_j^{(t)}, t = 1, 2, \cdots, M$, of station j in the tth slot in a frame is the set of stations that have permission to transmit to j. In terms of $\delta_{ij}^{(t)}$,

$$I_i^{(t)} = \{j \mid \delta_{ij}^{(t)} = 1\} \ \forall i, t; \ J_j^{(t)} = \{i \mid \delta_{ij}^{(t)} = 1\} \ \forall j, t \ . \quad (4)$$

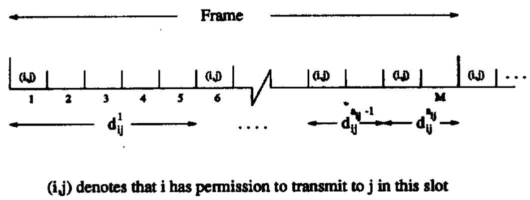

Fig. 1. Definition of $d_{ij}^{(k)}$ for $k = 1, \cdots, a_{ij}$.

Sets $I_i^{(t)}$ and $J_j^{(t)}$, together with the tunability characteristics, specify whether *collisions, destination conflicts*, or both are possible under a given schedule. Collisions occur when two or more transmitters access the same channel during a slot, i.e., when there exist a slot t, transmitters i and $i' \neq i$, and receivers j and j', such that $j \in I_i^{(t)}, j' \in I_{i'}^{(t)}$, and $w_{ij}^{(t)} = w_{i'j'}^{(t)}$. On the other hand, systems with tunable receivers may experience destination conflicts if multiple stations are permitted to transmit to the same destination on different channels, or when there exist a slot t, transmitters i and $i' \neq i$, and a receiver j, such that $i, i' \in J_j^{(t)}$, and $w_{ij}^{(t)} \neq w_{i'j}^{(t)}$.

There is, however, a certain class of schedules that allow neither collisions nor destination conflicts. These schedules give exactly C permissions to different source-destination pairs, one per channel, in each slot; in the case of fixed transmitters we can write

$$\sum_{i=1}^{N}\sum_{j=1}^{N} \delta_{ij}^{(t)} = \sum_{i=1}^{N} | I_i^{(t)} | = C, \ \forall t$$
$$\sum_{i \in X_c} | I_i^{(t)} | = 1, \forall t, c; | J_j^{(t)} | \leq 1, \ \forall j, t \quad (5)$$

denoting that exactly one of the transmitters that share channel λ_c is given permission to transmit in a given slot, and each receiver can receive a packet from at most one source. A similar expression can be derived in the case of fixed receivers.

We now distinguish two *transmission modes*:

- **one-to-one** transmission mode, if no collisions or destination conflicts are possible;
- **many-to-many** transmission mode, otherwise.

The many-to-many mode is the most general mode as the transmitting and receiving sets can be any subset of $\{1, \cdots, N\}$. Thus, any schedule that can be described by giving appropriate values to variables $\delta_{ij}^{(t)}$ and $w_{ij}^{(t)}$, is a many-to-many schedule. The one-to-one mode can be considered as a special case with constraints (5) imposed on the transmitting/receiving sets.

Fig. 2 demonstrates the two different transmission modes for a FT-TR network with $N = 4$ stations and $C = 2$ wavelengths. Channel λ_1 is shared by the fixed transmitters of stations 1 and 3 ($X_1 = \{1, 3\}$), while channel λ_2 is shared by the transmitters of stations 2 and 4 ($X_2 = \{2, 4\}$). A one-to-one schedule is such that in Fig. 2(a), only one of the transmitters in $X_c, c = 1, 2$, has permission to transmit in any given slot, and in Fig. 2(b) each station may receive from at most one source in a given slot. If any of these conditions is violated, the resulting schedule is many-to-many.

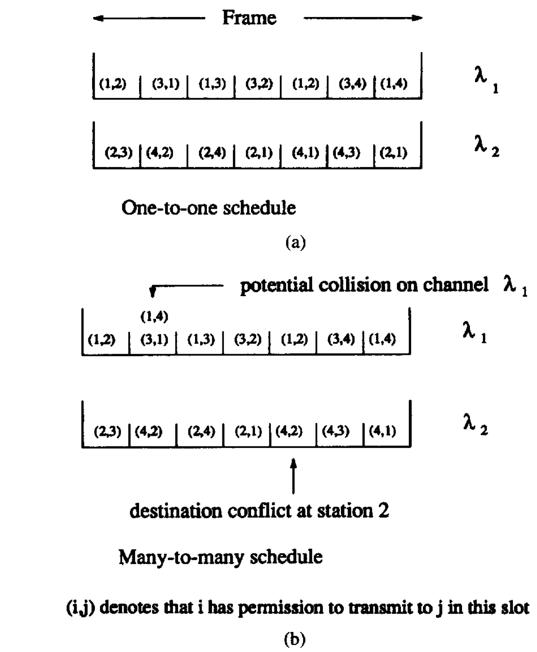

Fig. 2. (a) One-to-one and (b) many-to-many schedules for a FT-TR network with $N = 4$ stations and $C = 2$ wavelengths, $X_1 = \{1,3\}, X_2 = \{2,4\}$.

B. Selection Policies for Many-to-Many Schedules

With many-to-many schedules, whenever $\mid I_i^{(t)} \mid > 1$ for some i, t, transmitter i has the freedom to select the destination to which it will transmit a packet in slot t. Similarly, if $\mid J_j^{(t)} \mid > 1$ for some j, t *and* the receivers are tunable, receiver j may select one of many possible sources to which it can tune in slot t. The criteria used to select one of the stations in the transmitting and receiving sets define a *selection policy*. In order to be suitable for a high-speed environment a policy should be simple, very fast to execute, and should not require coordination between stations as that would cause severe degradation of the overall throughput. Two policies that share these characteristics are now described.

1) Random Policy (RP): At the beginning of a slot, t, transmitter i randomly and uniformly selects a destination $j \in I_i^{(t)}$ and transmits a packet, if it has one for j, on wavelength $w_{ij}^{(t)}$. Similarly, if the receivers are tunable, receiver j' will select a station $i' \in J_{j'}^{(t)}$ and will tune to wavelength $w_{i'j'}^{(t)}$ waiting for a transmission from i'.

2) Round-Robin Policy (RRP): In slot t, i considers one destination in $I_i^{(t)}$ for a packet transmission. By remembering the destination it considered in the previous frame, i considers all stations in $I_i^{(t)}$ in a round-robin order. A tunable receiver operates in a similar fashion.

When both a transmitter i and a receiver j implement a selection policy for a slot t, a necessary condition for a successful packet transmission from i to j in this slot is that both the policy at i select j for transmission *and* the policy at j select wavelength $w_{ij}^{(t)}$ to listen to.[2] In particular, the decisions of the round-robin policy at transmitter i (resp., receiver j) repeat after $\mid I_i^{(t)} \mid$ (resp., $\mid J_j^{(t)} \mid$) frames. It is easy to see that the policies at i and j will be synchronized (i.e., the policy at i will select j, and the policy at j will select i) once every least common multiple $(\mid I_i^{(t)} \mid, \mid J_j^{(t)} \mid)$ frames. For this reason, whenever both the transmitters and receivers implement the round-robin policy, a transmitter i should transmit to a receiver j only in slots in which their policies are synchronized.[3]

In situations where the policies above are implemented by a transmitter, further improvement is possible: when making a selection (random or round-robin), the transmitter ignores destinations for which it has no packets in its buffers. Note that such improvement is not possible in the selection policies implemented at the receivers. This is because a receiver has no way of knowing whether a particular source will have a packet for it or not. Therefore, when both the transmitters and receivers implement the round-robin policy, we do not propose implementing the improved version at the transmitters, as that might cause the policies at the transmitters and receivers to become completely unsynchronized.

In order to simplify our model we do not consider such improvement in the transmitter selection policies in the analysis to follow. Simulation results showing the degree of improvement possible in the case of the random policy are discussed in Section VI-C.

III. Throughput Analysis

A schedule can be characterized by three parameters:

- the transceiver tunability characteristics (TT-FR, FT-TR, or TT-TR),
- the transmission mode (one-to-one or many-to-many), and
- the policies used by the transmitters and receivers (random or round-robin).

In the following section we derive expressions for the throughput of the most general schedule, namely, a TT-TR system and a many-to-many transmission mode. We also show how to modify these expressions to apply to other tunability characteristics and the one-to-one mode. The throughput of a schedule S will be given by

$$T(P) = \sum_{i=1}^{N} \sum_{j=1}^{N} T_{ij}(P) \tag{6}$$

where $T_{ij}(P)$ is the throughput of the source-destination pair (i, j), i.e., the number of successful packet transmissions per slot between i and j, and P is the policy used by the transmitting and receiving stations.

A. Systems with Tunable Transmitters–Tunable Receivers

1) Many-to-Many Schedules with Random Policy: Consider the source-destination pair (i, j), and let $t_{ij}^{(1)}, \cdots, t_{ij}^{(a_{ij})}$, be the a_{ij} slots within a frame in which j is in the transmitting set of i. Without loss of generality, let $1 \leq t_{ij}^{(1)} < \cdots < t_{ij}^{(a_{ij})} \leq M$.

[2] This condition is not sufficient as some other packet transmission on the same wavelength will result in a collision.

[3] In order to implement this, a transmitter only needs to know the value of $\mid J_j^{(t)} \mid$ which can be immediately derived from the schedule, and the first frame in which the policy at j will select i.

As frames repeat over time, the lth slot in which j is in the transmitting set of i, $l = 1, 2, 3, \cdots$, is one of $t_{ij}^{(1)}, \cdots, t_{ij}^{(a_{ij})}$, within some frame m, and we can write $l = ma_{ij} + k$, where m, $m = 0, 1, 2, \cdots$, is the frame, and $1 \le k \le a_{ij}$. We now define

$$h_{ij}(l) = [(l-1) \bmod a_{ij}] + 1 \quad l = 1, 2, 3, \cdots \tag{7}$$

$$f_{ij}(l) = t_{ij}^{(h_{ij}(l))} \quad l = 1, 2, 3, \cdots . \tag{8}$$

Function $f_{ij}(l)$ maps the lth slot onto one of $t_{ij}^{(1)}, \cdots, t_{ij}^{(a_{ij})}$. Obviously,

$$h_{ij}(ma_{ij} + k)m = h_{ij}(k)$$

and

$$f_{ij}(ma_{ij} + k) = f_{ij}(k) = 0, 1, 2, \cdots, k = 1, \cdots, a_{ij}. \tag{9}$$

We observe the system at the instants just before the beginning of slots in which j is in the transmitting set of i. Consider the lth such slot. We define $q_{ij}^{(l)}$ as the probability that i has a packet for j at the beginning of the lth slot. $q_{ij}^{(l)}$ is equal to 1 if i had a packet for j at the beginning of the $(l-1)$th slot and j was not selected by the random policy at that slot. Otherwise, $q_{ij}^{(l)}$ is equal to the probability that at least one packet for j arrived at i during the $d_{ij}^{(h_{ij}(l-1))}$ slots between the $(l-1)$th and the lth slots. The probability that j is selected by the random policy in the lth slot is $\frac{1}{|I_i^{(f_{ij}(l))}|}$.

We can now express $q_{ij}^{(l)}$ as

$$\begin{aligned} q_{ij}^{(l)} &= q_{ij}^{(l-1)}\left(1 - \frac{1}{|I_i^{(f_{ij}(l-1))}|}\right) \\ &+ \left[1 - q_{ij}^{(l-1)}\left(1 - \frac{1}{|I_i^{(f_{ij}(l-1))}|}\right)\right] \\ &\times \left[1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(h_{ij}(l-1))}}\right] \\ &l = 2, 3, \cdots \end{aligned} \tag{10}$$

$$q_{ij}^{(1)} = 0 \text{ (Initial Condition)}. \tag{11}$$

The initial condition (11) is obtained by assuming that the frame starts at a slot in which i can transmit to j. After some algebraic manipulation of (10), we get

$$q_{ij}^{(l)} = A_{ij}^{(l-1)} q_{ij}^{(l-1)} + B_{ij}^{(l-1)}, l = 2, 3, 4, \cdots \tag{12}$$

where

$$A_{ij}^{(l)} = \left(1 - \frac{1}{|I_i^{(f_{ij}(l))}|}\right)(1 - B_{ij}^{(l)}) \tag{13}$$

$$B_{ij}^{(l)} = 1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(h_{ij}(l))}}. \tag{14}$$

Because of (9) we have that $B_{ij}^{(ma_{ij}+k)} = B_{ij}^{(k)}$ and $A_{ij}^{(ma_{ij}+k)} = A_{ij}^{(k)}, k = 1, \cdots, a_{ij}, m = 0, 1, \cdots$. It is now easy to show using induction that

$$q_{ij}^{(ma_{ij}+1)} = q_{ij}^{(a_{ij}+1)} \sum_{n=0}^{m-1} D_{ij}^n, D_{ij} = \prod_{k=1}^{a_{ij}} A_{ij}^{(k)}, m = 0, 1, \cdots . \tag{15}$$

The sum in (15) will always converge as $m \to \infty$ since $0 \le D_{ij} < 1$. Define $r_{ij}^{(k)}$ as the limiting probability that i has a packet for j in the kth slot within a frame that j is in the transmitting set of i,

$$r_{ij}^{(k)} = \lim_{m \to \infty} q_{ij}^{(ma_{ij}+k)} \quad k = 1, \cdots, a_{ij}. \tag{16}$$

We can now obtain $r_{ij}^{(k)}$ using the following recursion (see (15) and (12)):

$$r_{ij}^{(1)} = \frac{1}{1 - D_{ij}} q_{ij}^{(a_{ij}+1)} = \frac{1}{1 - D_{ij}} \sum_{n=1}^{a_{ij}} B_{ij}^{(n)} \prod_{r=n+1}^{a_{ij}} A_{ij}^{(r)} \tag{17}$$

$$r_{ij}^{(k)} = A_{ij}^{(k-1)} r_{ij}^{(k-1)} + B_{ij}^{(k-1)} \quad k = 2, \cdots, a_{ij}. \tag{18}$$

In order to compute the throughput of pair (i, j) we note that a packet will be successfully transmitted from i to j in the kth slot in a frame if (a) i has a packet for j at the beginning of the slot, (b) the random policy at i picks j to transmit to, and the random policy at j picks wavelength $w_{ij}^{(k)}$ to listen to, and (c) if another station, i', is also allowed to transmit in the same slot and on the same wavelength, it does not do so. As $m \to \infty$,

$$\begin{aligned} T_{ij}(RP) = \frac{1}{M} \sum_{k=1}^{a_{ij}} \Bigg\{ & \frac{r_{ij}^{(k)}}{|I_i^{(f_{ij}(k))}||J_j^{(f_{ij}(k))}|} \\ & \prod_{i' \ne i, j', w_{ij}^{(f_{ij}(k))} = w_{i'j'}^{(f_{ij}(k))}} \left(1 - \frac{r_{i'j'}^{(k)}}{|I_{i'}^{(f_{ij}(k))}|}\right) \Bigg\}. \end{aligned} \tag{19}$$

Expressions for all other schedules can be derived from (19) by imposing appropriate restrictions. For example, if the receivers are not tunable (TT-FR system), the factor $\frac{1}{|J_j^{(f_{ij}(k))}|}$ goes away and we have $w_{ij}^{(f_{ij}(k))} = \lambda(j), w_{i'j'}^{(f_{ij}(k))} = \lambda(j')$. Also, in the special case of one-to-one schedules, $|I_i^{(f_{ij}(k))}| = |J_j^{(f_{ij}(k))}| = 1$ and $A_{ij}^{(k)} = 0$ for $k = 1, \cdots, a_{ij}$. Then, $r_{ij}^{(k)} = B_{ij}^{(k-1)} = 1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(k-1)}}$, there is no $i' \ne i$ such that $w_{ij}^{(f_{ij}(k))} = w_{i'j'}^{(f_{ij}(k))}$, for some j', and we get the following expression, valid for any tunability characteristics.

$$T_{ij} = \frac{1}{M} \sum_{k=1}^{a_{ij}} 1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(k)}} \text{ (One - to - one schedules)}. \tag{20}$$

2) Many-to-Many Schedules with Round-Robin Policy: We again consider the pair (i, j). Let $z_i^{(t)} \in I_i^{(t)} (Z_j^{(t)} \in J_j^{(t)})$ be the decision of the round-robin policy at transmitter i (resp., receiver j) in slot $t, t = 1, \cdots, M$, i.e., the station i (resp., j) will transmit (resp., listen) to. The decisions of the round-robin policy repeat, for all i, j, after $F = \prod_{i,t} I_i^{(t)} \prod_{j,t} J_j^{(t)}$

frames; in other words, $z_i^{(t)} = z_i^{(t+MF)} \forall i, t$, and $Z_j^{(t)} = Z_j^{(t+MF)} \forall j, t$. We can now restrict our attention to a window of F consecutive frames. Let b_{ij} be the number of slots t within a window such that $z_i^{(t)} = j$ and $Z_j^{(t)} = i$. Using the same arguments leading to (19) we get

$$T_{ij}(\text{RRP}) = \frac{1}{MF} \sum_{k=1}^{b_{ij}} \left\{ \left[1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(k)}}\right] \prod_{i', i' \neq i, j', w_{i'j'}^{(k)} = w_{ij}^{(k)}} (1 - \sigma_{i'} p_{i'j'})^{d_{i'j'}^{(k)}} \right\}. \quad (21)$$

IV. Optimization of One-to-One Schedules

We now turn our attention to the problem of obtaining schedules that maximize the system throughput. As a first step, one has to determine the transmission mode and, if applicable, the type of policy to use. Let $Q_{ij} = 1 - (1 - \sigma_i p_{ij})^M$ be the probability that at least one packet with destination j will arrive at station i during a number of slots equal to the frame length, M. When Q_{ij} is close to one for all source-destination pairs (i, j), one-to-one schedules are favored over many-to-many schedules, as no packets are wasted due to collisions or destination conflicts. But in one-to-one schedules, slots are assigned for the exclusive use of a certain source-destination pair. If the above condition is not true, some slots may be unused for most of the time. Our approach is to first consider determining an optimal one-to-one schedule. Next we consider how we may obtain many-to-many schedules with good performance in situations where Q_{ij} is very low for some pairs of stations.

Before proceeding to the general case, we will first consider the special problem of $C = N$, i.e., a number of available wavelengths equal to the number of stations. The solution method developed will provide the intuition for approaching the general problem with $C < N$, while some of the heuristics will be directly applicable to the general case.

A. Schedules for $C = N$

Our goal is to determine a one-to-one schedule S such that the overall throughput, as given by expressions (6) and (20), is maximized. Expression (20) is valid for all three types of systems. Without loss of generality, in the following discussion we only consider TT-FR and FT-TR systems, in which transmissions to, or from, a certain station take place on the same channel. The optimization problem can be formulated as

$$\mathcal{P}_1 : \max_{\delta_{ij}^{(t)}, M} \; T = \frac{1}{M} \sum_{i=1}^{N} \sum_{j=1}^{N} \sum_{k=1}^{a_{ij}} 1 - (1 - \sigma_i p_{ij})^{d_{ij}^{(k)}} \quad (22)$$

subject to

$$a_{ij} = \sum_{t=1}^{M} \delta_{ij}^{(t)} \; \forall i, j \quad (23)$$

$$d_{ij}^{(k)} = \begin{cases} f_{ij}(k+1) - f_{ij}(k), & k = 1, \cdots, a_{ij} - 1 \\ M + f_{ij}(1) - f_{ij}(k), & k = a_{ij} \end{cases} \quad (24)$$

$f_{ij}(k)$ as in (8)

$$\sum_{k=1}^{a_{ij}} d_{ij}^{(k)} = M \; \forall i, j \quad (25)$$

$$\sum_{i=1}^{N} a_{ij} = M \; \forall j; \; \sum_{j=1}^{N} a_{ij} = M \; \forall i \quad (26)$$

$$\sum_{i=1}^{N} \delta_{ij}^{(t)} = 1 \; \forall j, t; \; \sum_{j=1}^{N} \delta_{ij}^{(t)} = 1 \; \forall i, t \quad (27)$$

$$\delta_{ij}^{(t)} = 0, 1 \; \forall i, j, t, \; M \text{ integer.} \quad (28)$$

Constraints (26) specify that each station is permitted to receive and transmit in exactly M slots, while constraints (27), a special case of (5) with $|X_c| = 1, c = 1, \cdots, N$, guarantee that the resulting schedule is a one-to-one schedule.

As formulated, $\mathcal{P}_1$ is a hard allocation problem. We now present a heuristic to obtain near-optimal one-to-one schedules. The heuristic is based on a decomposition of $\mathcal{P}_1$ into three manageable subproblems.

1) Obtain a_{ij} that satisfy (26). This is described in Appendix A and is based on a decomposition of the problem using theory developed in [15], [1], [2].
2) Construct a schedule by considering each channel independently of others. This may result in allocations that violate either of the conditions in (27); considering channels independently may cause a transmitter (resp., receiver) in a TT-FR (resp., FT-TR) system to be assigned to transmit to (resp., receive from) two or more receivers (resp., transmitters) in the same slot.
3) Rearrange the schedule resulting from 2) above to remove the violations, converting it into a one-to-one schedule.

The following steps describe the heuristic for a TT-FR system (a FT-TR system is handled in a similar way). Our approach is based on the golden ratio policy developed in [15], where only frame lengths equal to the Fibonacci numbers are considered. This policy places the permissions of a source-destination pair in approximately equal distances within a frame (see Appendix A).

Heuristic 1 (Optimized One-to-One Schedules for TT-FR Systems and $C = N$):

1) Select the smallest Fibonacci number $M \geq N - 1$ and obtain a_{ij} as in Appendix A. If the a_{ij} do not satisfy fairness condition (2), repeat with the next Fibonacci number.
2) Consider receiver $j = 1$ and use the golden ratio policy [15] to allocate $a_{i1}, i = 1, \cdots, N$, slots in a frame for transmissions on channel $\lambda(j)$. Repeat for $j = 2, \cdots, N$.
3) Run algorithm REARRANGE, described below, to convert the schedule produced by Step 2 to a one-to-one schedule, and compute its throughput from (6), (20).

Algorithm REARRANGE

```
for d = 1 to N (* consider one destination at a time *)
begin
   list ← empty;
   for t = 1 to M (* consider all transmissions to d *)
     if s is such that δ_sd^(t) = 1 and it violates Σ_{j=1}^{d} δ_sj^(t) ≤ 1 then
        δ_sd^(t) ← 0; add s to the list;
   while the list is nonempty
   begin
     s ← the first item of the list; violation ← false; l ← d
     repeat
       find a t1 such that δ_sj^(t1) = 0 ∀j = 1,...,d
          δ_sl^(t1) ← 1;
       if there exists a m such that δ_ml^(t1) = 1 then
       begin
         find t2 such that δ_il^(t2) = 0 ∀i = 1,...,N; δ_ml^(t2) ← 1; δ_ml^(t1) ← 0;
         if there is a n < d such that δ_mn^(t2) = 1 then
                 violation ← true; s ← m;  l ← n;
       end
     until violation = false
   end
end
```

Fig. 3. Algorithm to construct a one-to-one schedule for a TT-FR system. A slight variation of the algorithm can be used for a FT-TR system.

4) Repeat Steps 1 through 3 for the next Fibonacci number, up to an upper limit, M_{max}. Select the frame length, and schedule, that yields the largest throughput.

The algorithm REARRANGE used in the heuristic for TT-FR systems is shown in Fig. 3. The algorithm operates as follows. Suppose that in the initial schedule transmitter i has been given permission to transmit to two receivers in the same slot, t. Then, because of the condition $\sum_{j=1}^{N} a_{ij} = M$, there must be a slot t' in which i is not assigned to transmit to any receiver. The algorithm will move one of i's permissions from slot t to slot t', and will move another transmitter's permission from t' to t. By repeating the process, a one-to-one schedule is constructed. The proof of correctness of the algorithm is given in Appendix B, where it is also shown that its worst case complexity is $O(N^2M^2)$. Note that the correctness of the algorithm implies that if all a_{ij} satisfy (26) then a one-to-one schedule always exists.

B. Schedules for $C < N$

Systems with a number of wavelengths, C, equal to the number of stations, N, do not scale well, for two reasons. First, the number of wavelengths that can be supported within a fiber is limited by the optical technology. Second, there is a tradeoff between tuning range and tuning speed in state of the art agile optical transceivers; even when an adequate number of wavelengths is available, the transceivers that can tune over the entire range of wavelengths may not meet the speed requirements for single-hop networks.

We now concentrate on the problem of obtaining near-optimal one-to-one schedules when $C < N$. As in Section IV-A, we only consider TT-FR and FT-TR systems, in which a single channel may be shared by a number of receivers or transmitters. The new optimization problems are harder than $\mathcal{P}_1$, as the maximization is over all partitions of $\{1, \cdots, N\}$ into sets R_c or X_c, the sets of receivers and transmitters, respectively, sharing channel λ_c. In [3] it is assumed that $N = mC$, and that m receivers (transmitters) share each channel. Although this assumption is acceptable under the uniform traffic conditions considered there, in the general case one would like to select the receivers or transmitters that share each channel so that the utilization of all channels be kept at almost the same level. In the following, we first assume that the sets are given and present heuristics to obtain near-optimal one-to-one schedules, and then discuss how to construct these sets, for any values of N and C.

1) Systems with Fixed Transmitters–Tunable Receivers: Given sets X_c of transmitters sharing channel $\lambda_c, c = 1, \cdots, C$, our optimization procedure is described by Heuristic 2. Note that we have decomposed the optimization problem into two subproblems: (a) how the bandwidth of channel λ_c should be allocated to stations in X_c, and (b) how the bandwidth each station in X_c receives should be allocated for transmissions to each destination.

Heuristic 2 (Optimized One-to-One Schedules for FT-TR Systems and $C < N$):

1) For each channel λ_c obtain the percentage of time, x_{ic} that $i \in X_c$ should transmit on λ_c. This can be done by solving problem $\mathcal{P}_3$ (see Appendix A). Allocate a_{ic} slots to station i so that $\lfloor Mx_{ic} \rfloor \le a_{ic} \le \lceil Mx_{ic} \rceil$ and $\sum_{i \in X_c} a_{ic} = M$.
2) For each station i obtain x_{ij}, the percentage of time that i should transmit to j. This can be done by solving problem $\mathcal{P}_4$ (see Appendix A). Allocate a_{ij} slots to pair (i, j), such that $\lfloor a_{ic}x_{ij} \rfloor \le a_{ij} \le \lceil a_{ic}x_{ij} \rceil$ and $\sum_j a_{ij} = a_{ic}$.[4]
3) Use the golden ratio policy [15] to place the a_{ij} slots, $i \in X_c, j = 1, \cdots, N$, on λ_c.

2) Systems with Tunable Transmitters–Fixed Receivers: When receivers are fixed, we can use a technique similar to Heuristic 2 to allocate slots for transmissions on a certain channel. Observe, though, that in a TT-FR system, a station transmitting on channel λ_c can reach any of the stations in R_c. Define s_{ic} as the probability that the destination of a new packet generated at i has its receiver tuned to channel λ_c

$$s_{ic} = \sum_{j \in R_c} p_{ij} \quad c = 1, \cdots, C \tag{29}$$

If we define a "destination" as a set of stations, R_c, rather than a single station, each station can have C, instead of $N - 1$, packet buffers, one for the receivers in R_c. By using s_{ic} instead of p_{ij}, all the results of Section IV-A are now directly applicable. The new model keeps information about sets of receivers, rather than individual ones, and, thus, is an approximation of our original model. Nevertheless, it allows us to use Heuristic 1 with the only difference that C, rather than N "destinations" have to be considered.

3) Construction of Sets X_c and R_c: We now address the problem of how to select the transmitters or receivers that share a certain channel so that some load balancing is achieved. We now formulate this problem in a more general context, as follows. We are given a set of N elements, and v_i is the weight

[4] At the end of Step 2, a check should be made to ensure that there is no j such that $\sum_i a_{ij} > M$. If there is one, the excess slots should be allocated to other source-destination pairs.

associated with element e_i. Our objective is to partition this set into $C, C < N$, subsets $Y_c, c = 1, \cdots, C$, such that the following quantity is minimized

$$\max_{1 \le k, l \le C} \mid \sum_{e_i \in Y_k} v_i - \sum_{e_j \in Y_l} v_j \mid . \tag{30}$$

For $C = 2$, the problem reduces to partitioning the set of elements into two subsets, Y_1 and Y_2 such that $\mid \sum_{e_i \in Y_1} v_i - \sum_{e_i \in Y_2} v_i \mid$ is minimized. But the problem of whether there exists a partition of the set of N elements into two subsets Y_1 and Y_2 such that $\sum_{e_i \in Y_1} v_i = \sum_{e_i \in Y_2} v_i$, is NP-complete [11, p. 223], and thus the minimization problem for $C = 2$ is expected to be hard. We now propose the following greedy procedure to construct the C subsets.

1) Sort elements e_i in decreasing order of v_i. Initialize $Y_c = \{e_c\}, c = 1, \cdots, C$, and $k \leftarrow C + 1$. Note that sets Y_c are also sorted in decreasing order of $\sum_{e_i \in Y_c} v_i$.
2) Set $Y_C \leftarrow Y_C \bigcup \{e_k\}$ and $k \leftarrow k + 1$. Sort $Y_c, c = 1, \cdots, C$, in decreasing order of $\sum_{i \in Y_c} v_i$. Repeat Step 2 while $k \le N$.

For obtaining the sets of transmitters, X_c, we used $v_i = \sigma_i$, the probability that a new packet will arrive at station i in a slot. For the sets of receivers, R_c, we set $v_j = \sum_{i=1}^{N} \sigma_i p_{ij}$.

V. Optimization of Many-to-Many Schedules

Whenever $Q_{ij} = 1 - (1 - \sigma_i p_{ij})^M \ll 1$ for some pairs (i, j), one-to-one schedules may result in low throughput because of the requirement that at least one slot be assigned to any pair of stations, regardless of how low its traffic requirements are (see (2)). In these cases, we can take advantage of the low traffic by using more general, many-to-many schedules that may yield a better throughput. The idea is that, instead of assigning one slot per pair of stations, it might be better, in terms of throughput, to assign a single slot for transmissions from a station, i, to a group of stations, provided that the probability of a packet for any of the stations in the group arriving at i during a time period equal to the frame length, is very small. Heuristic 3, described below, is based on this idea, and assumes that the frame length is given.

Heuristic 3 (Optimized Many-to-Many Schedules for $C \le N$):

1) For transmitter $i = 1, \cdots, N$, let $G_i = \{j \mid Q_{ij} \le \Delta\}$. If $\mid G_i \mid \ge 1$, partition G_i in disjoint sets $g_i^{(1)}, \cdots, g_i^{(k_i)}$ such that $\sum_{j \in g_i^{(l)}} Q_{ij} < \epsilon < 1, l = 1, \cdots, k_i$.
2) Run the appropriate heuristic (Heuristic 1 or 2) for the given frame length, M, modified so that for any groups $g_i^{(l)}, l = 1, \cdots, k_i$, produced in Step 1, only one slot is assigned for transmissions from i to the stations in each group. The resulting schedule, S, is, in general, a many-to-many schedule. Compute its throughput by means of the appropriate expressions. Use S if its throughput is better than that of the one-to-one schedule with the same frame length.

Parameter Δ has a small positive value and controls which destinations will *not* be assigned a slot of their own for transmissions from i. Parameter $\epsilon (> \Delta)$ controls the number of stations that are grouped together. The smaller the values of Δ and ϵ, the smaller the number of stations in a group. By adjusting the values of the two parameters we can keep the probability of collisions and/or destination conflicts within a slot to an acceptable level. Also, by choosing a small Δ we can have $G_i = \emptyset \forall i$, producing a one-to-one schedule. Thus, Heuristic 3 is very general.

0	0.30	0.30	0	0.30	0	0	0
0.30	0	0	0.30	0	0.30	0	0
0.30	0	0	0.30	0	0	0.30	0
0	0.30	0.30	0	0	0	0	0.30
0.30	0	0	0	0	0.30	0.30	0
0	0.30	0	0	0.30	0	0	0.30
0	0	0.30	0	0.30	0	0	0.30
0	0	0	0.30	0	0.30	0.30	0

Fig. 4. Mesh type traffic matrix (Network 1).

0	0.20	0.35	0.25	0.03	0.02	0.01	0.04
0.20	0	0.25	0.40	0.04	0.01	0.03	0.02
0.30	0.25	0	0.35	0.02	0.02	0.01	0.03
0.15	0.40	0.30	0	0.02	0.01	0.02	0.02
0.02	0.04	0.03	0.03	0	0.30	0.35	0.20
0.01	0.03	0.04	0.03	0.30	0	0.25	0.25
0.02	0.01	0.01	0.03	0.40	0.15	0	0.30
0.03	0.03	0.02	0.02	0.25	0.35	0.30	0

Fig. 5. Disconnected type traffic matrix (Network 2).

VI. Numerical Results

A. Optimized One-to-One Schedules for $C = N$

We consider the mesh type, disconnected type, ring type and quasiuniform traffic matrices shown in Figs. 4, 5, 6, and 7, respectively. The mesh type configuration is considered as it often arises in parallel and distributed computations. Fig. 6 also shows the optimized one-to-one schedule of frame length $M = 21$ produced by Heuristic 1. For these matrices we computed the throughput of a simple cyclic schedule, the throughput of the optimized schedule produced by Heuristic 1, and an upper bound on the throughput of the optimal solution

0	0.60	0.10	0.05	0.05	0.01	0.03	0.05
0.05	0	0.50	0.02	0.10	0.05	0.04	0.10
0.10	0.05	0	0.50	0.10	0.02	0.05	0.05
0.05	0.10	0.03	0	0.60	0.05	0.10	0.02
0.04	0.02	0.05	0.03	0	0.70	0.05	0.10
0.10	0.15	0.03	0.05	0.04	0	0.45	0.10
0.04	0.05	0.05	0.01	0.10	0.05	0	0.65
0.55	0.10	0.01	0.10	0.05	0.10	0.04	0

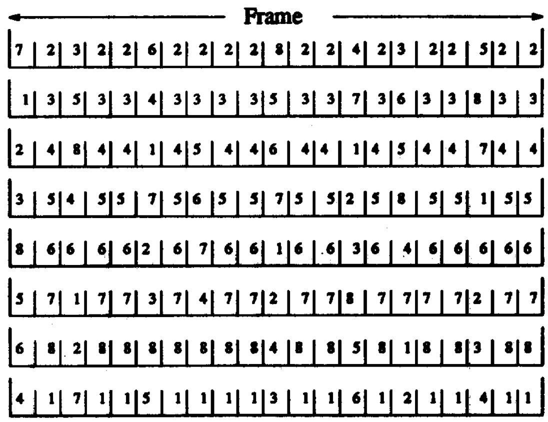

Fig. 6. Ring type traffic matrix (Network 3) and optimized one-to-one schedule for $M = 21$.

0	0.25	0.10	0.05	0.25	0.10	0.05	0.15
0.20	0	0.05	0.20	0.05	0.15	0.15	0.20
0.05	0.20	0	0.10	0.15	0.30	0.10	0.05
0.10	0.05	0.25	0	0.10	0.05	0.30	0.10
0.05	0.05	0.25	0.20	0	0.20	0.10	0.15
0.10	0.30	0.05	0.05	0.15	0	0.20	0.15
0.15	0.20	0.15	0.20	0.05	0.15	0	0.05
0.30	0.10	0.10	0.15	0.25	0.05	0.05	0

Fig. 7. Quasiuniform type traffic matrix (Network 4).

TABLE I
THROUGHPUT RESULTS FOR NETWORKS WITH $N = C = 8$ USING HEURISTIC 1

Net No	One-To-One Throughput Cyclic	One-To-One Throughput Heuristic	Increase from Cyclic	Upper Bound	% from Upper Bound
1	3.146	5.076	61.3%	5.256	3.4%
2	3.714	4.981	34.1%	5.330	6.5%
3	3.337	5.317	59.3%	5.568	4.5%
4	4.736	4.874	2.9%	5.270	7.5%

TABLE II
THROUGHPUT RESULTS FOR TT-FR NETWORKS WITH $N = 8$, $C = 2, 4$, USING HEURISTIC 1

Net	One-To-One Throughput I-TDMA*		One-To-One Throughput Optimized		Increase from I-TDMA*		Upper Bound		% from Upper Bound	
No	$C=2$	$C=4$	$C=2$	$C=4$	$C=2$	$C=4$	$C=2$	$C=4$	$C=2$	$C=4$
1	1.942	2.827	1.962	3.422	1.0%	21.0%	1.988	3.520	1.3%	3.0%
2	1.839	2.732	1.987	3.513	8.0%	28.6%	1.995	3.648	0.4%	3.7%
3	1.787	2.493	1.984	3.597	11.0%	44.3%	1.995	3.711	0.6%	3.1%
4	1.982	3.332	1.982	3.490	0.0%	4.7%	1.991	3.603	0.5%	3.1%

to problem $\mathcal{P}_1$.[5] The simple cyclic schedule is a one-to-one schedule of frame length $M = N - 1$, and is such that a station is given permission to transmit to each other station in exactly one slot per frame [8] (i.e., it is optimal for uniform traffic). The upper bound was computed as in Appendix A. The value of M_{max} in Heuristic 1 was 987.

Our results are summarized in Table I. Our heuristic produces schedules that are very close to the upper bound. Similar results have been obtained over a very wide range of traffic parameters. Our schedules also represent a significant improvement over the cyclic one, especially when the traffic is far from uniform (matrices 1, 2, and 3 in the table).

A. Optimized One-to-One Schedules for $C < N$

We again consider the 8-station networks of the previous section and investigate how the throughput performance is affected when we vary the number of available channels.

1) Tunable Transmitters–Fixed Receivers: The sets of receivers, R_c, that share channel $\lambda_c, c = 1, \cdots, C$, were obtained as in Section IV-B3). Each set was considered as a single "destination" (see Section IV-B2)) and we used Heuristic 1 to construct near-optimal one-to-one schedules. Our results are shown in Table II for 4 and 2 channels. The interleaved TDMA (I-TDMA*) schedule, of which the throughput is also shown, has N slots per frame, and is such that each source is allowed to transmit in exactly C slots during the frame, once to the receivers sharing each channel [3]. The upper bound in the table is an upper bound to the optimization problem *for these sets* R_c (recall that sets R_c are not fixed in the optimization problem, therefore we cannot get an overall upper bound).

The throughput of our schedules is significantly higher than that of the I-TDMA* schedule, especially for $C = 4$. Our schedules, also, are very close to the upper bound. When $C = 2$, the improvement over I-TDMA* is somewhat diminished. This can be explained by the fact that sets R_c are constructed so that the utilization of each channel is kept at approximately the same level. Since there are only two channels, by clustering the receivers, the traffic to each channel loses its original characteristics and becomes more "uniform".

2) Fixed Transmitters–Tunable Receivers: Table III shows, for $C = 4$ and $C = 2$, the throughput of the one-to-one

[5] For a fair comparison, we assume multiple packet buffers per station for the cyclic schedule, as well as for the I-TDMA* schedule in the next section, exactly like our optimized schedules.

TABLE III
THROUGHPUT RESULTS FOR FT-TR NETWORKS WITH $N = 8$, $C = 2, 4$, USING HEURISTIC 2

Net No	One-To-One Throughput: Optimized Schedule $C=4$	Optimized Schedule $C=2$	Cyclic Schedule $C=4$	Cyclic Schedule $C=2$	Increase from Cyclic Schedule $C=4$	Increase from Cyclic Schedule $C=2$
1	3.422	1.944	1.703	0.857	100.9%	126.8%
2	3.232	1.875	2.310	1.389	39.9%	35.0%
3	3.416	1.943	2.336	1.524	46.2%	27.5%
4	3.275	1.924	3.146	1.838	4.2%	4.7%

schedules, produced by Heuristic 2, for the same 8-station networks under the assumption that the transmitters are fixed. We also compare our results to a cyclic schedule which allocates exactly one transmission to the (i, j) pair in a frame. Again, our schedules perform much better than the cyclic one.

B. Optimized Many-to-Many Schedules

We now consider a 20-station FT-TR network with the following traffic parameters.

Network 5

$$\sigma_1 p_{12} = \sigma_1 p_{13} = \sigma_2 p_{21} = \sigma_2 p_{23} = \sigma_3 p_{31} = \sigma_3 p_{32} = 0.49$$

$$\sigma_i p_{ii} = 0 \; i = 1, \cdots, 20; \sigma_i p_{ij} = 10^{-5} \text{ for all other } i, j.$$

For these traffic parameters the upper bound in the throughput of one-to-one schedules is 2.223, and the throughput of a cyclic schedule is 0.320. In Table IV we show the throughput of the many-to-many schedules produced by Heuristic 3 for this network for various values of the Δ and ϵ parameters, for various frame lengths, and for the random policy. We also show the throughput of the one-to-one schedules produced by Heuristic 1 for the corresponding frame length.

As we can see, the throughput of the many-to-many schedules is always better than that of the one-to-one schedules with the same frame length, and, especially for smaller values of M, the improvement is significant. This is due to the fact that in a one-to-one schedule, slots assigned to pairs of stations for which $Q_{ij} \ll 1$ are mostly wasted. Heuristic 3 improves this throughput by assigning only one slot for transmissions to a group of stations for which there is low traffic (see Fig. 8 for an example). From the table, the best overall performance is obtained when $\Delta = 0.01$ and $\epsilon = 0.2$. By experimenting with different values for the two parameters, Heuristic 3 can produce a large variety of schedules from which we can choose the one with the best performance.

As we mentioned in Section II-B, a transmitter may implement modified versions of the selection policies, whereby it will ignore destinations for which it has no packets in its buffer (recall that this is *not* possible for the policies implemented at the receivers). We do not, however, expect that the modified policies will result in a significant increase in performance, for the following reasons. First, by construction (see Heuristic 3), the transmitting sets $I_i^{(t)}$, of source i, consist of destinations for which there is very low traffic and, thus, transmissions to these destinations have only a slight effect on overall throughput. Second, Heuristic 3 assigns only one slot per frame for transmissions to a group of stations; as the frame length

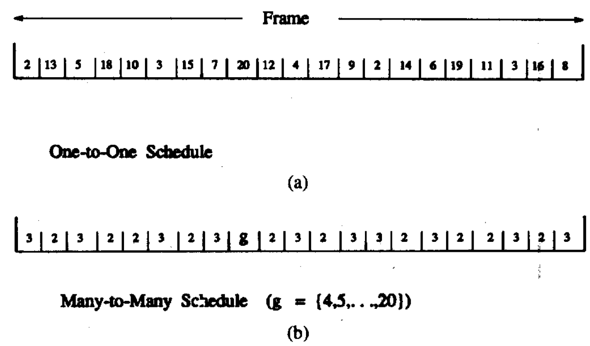

Fig. 8. Transmissions *from* station 1 under the one-to-one and many-to-many schedules ($\Delta = 0.01, \epsilon = 0.2$) of frame length $M = 21$ slots, for Network 5.

TABLE IV
THROUGHPUT RESULTS FOR NETWORK 5 ($N = C = 20$) USING HEURISTIC 3

Frame Length	Throughput: One-to-One Schedules	Many-To-Many Schedules $\Delta = 0.01$ $\epsilon = 0.2$	$\Delta = 0.001$ $\epsilon = 0.2$	$\Delta = 0.01$ $\epsilon = 0.02$
21	0.567	1.843	1.843	1.843
34	1.265	1.990	1.990	1.990
55	1.694	2.022	2.022	2.022
89	1.853	2.083	2.083	2.083
144	1.974	2.089	1.974	2.051
233	2.050	2.107	2.050	2.063
377	2.085	2.118	2.085	2.091
610	2.106	2.123	2.106	2.107
987	2.118	2.128	2.118	2.117

increases, any gain in throughput within that slot decreases as a percentage of the overall throughput. Finally, although a transmitter will always pick a receiver (if any) for which it has a packet to transmit, if the receiver also implements a selection policy, it is not guaranteed that the transmitted packet will in fact be received.

We used simulation to assess the degree of improvement from implementing the modified policies at the transmitters. We considered the many-to-many schedules for Network 5, as well as for a 4-station, 2-wavelength FT-TR network (Network 6) with traffic parameters $\sigma_1 p_{12} = \sigma_2 p_{21} = \sigma_3 p_{34} = \sigma_4 p_{43} = 0.7$, $\sigma_i p_{ii} = 0, i = 1, \cdots, 4$, and $\sigma_i p_{ij} = 0.1$ for all other i, j.

The simulation results showed a negligible improvement for Network 5 (less than 0.5%), which diminished as the frame length increased. For Network 6, we used high values for parameters ϵ and Δ in Heuristic 3 (i.e., $\Delta = 0.41, \epsilon = 0.82$ for $M = 5$, and even higher for larger values of M) and the groups constructed were $g_1^{(1)} = g_2^{(1)} = \{3, 4\}$, and $g_3^{(1)} = g_4^{(1)} = \{1, 2\}$. The highest increase was observed for the smallest frame length, $M = 5$, when the throughput increased from 1.617 (random policy) to 1.690 (improved random policy), or 4.5%. Note, however, that this increase was possible only because we grouped destinations for which there is relatively high traffic (i.e., $g_1^{(1)} = \{3, 4\}$ although $Q_{13} = Q_{14} \geq 0.4 \forall M$) and, therefore, the improved random policy will almost always find a packet in the buffers to transmit. In general, we expect the increase to be much smaller.

VII. Concluding Remarks

In this paper we have considered single-hop lightwave networks employing WDM and a predefined, wavelength-time oriented schedule to coordinate packet transmissions. We have defined a wide variety of transmission schedules based on the transceiver tunability characteristics, the transmission mode, and the policy used to select one of possibly many stations to communicate with. For the random and round-robin policies we have presented a model for analyzing and optimizing the throughput performance of schedules for any number of wavelengths and general traffic patterns. Our results indicate that a significant improvement over previously proposed schemes is possible for the potentially nonuniform traffic patterns that we expect to encounter in practice.

Appendix

A. Heuristic To Obtain a_{ij}

As a first step, we will try to get an upper bound to problem $\mathcal{P}_1$. It can be shown [15] that

$$T \leq \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{a_{ij}}{M}\left[1-(1-\sigma_i p_{ij})^{\frac{M}{a_{ij}}}\right]. \tag{31}$$

Let $x_{ij} = a_{ij}/M$. x_{ij} indicates the percentage of time that station i should transmit to station j. (31) implies that the Mx_{ij} slots assigned for transmissions from i to j should be equally spaced, separated by a number of $1/x_{ij}$ slots. However, this is not possible in general. First, $1/x_{ij}$'s may not be integers, and second, even if they are, scheduling the transmissions between all sources and destinations in equally spaced slots may violate constraints (27). If we relax (27), maximizing the upper bound in (31) can be formulated as

$$\mathcal{P}_2 : \max_{x_{ij}} \; N - \sum_{i=1}^{N}\sum_{j=1}^{N} x_{ij}(1-\sigma_i p_{ij})^{\frac{1}{x_{ij}}} \tag{32}$$

subject to

$$\sum_{i=1}^{N} x_{ij} = 1 \; \forall j; \; \sum_{j=1}^{N} x_{ij} = 1 \; \forall i, \; x_{ij} \geq 0 \; \forall i, j. \tag{33}$$

The solution to $\mathcal{P}_2$ is independent of the frame length M. Given M, we can get a_{ij} from

$$\lfloor M x_{ij} \rfloor \leq a_{ij} \leq \lceil M x_{ij} \rceil \tag{34}$$

and so that constraints (26) are satisfied. Unfortunately, problem $\mathcal{P}_2$ does not yield an analytical solution. We now develop a heuristic for obtaining a_{ij}, by decomposing $\mathcal{P}_2$ into $2N$ easy to solve problems.

As a first step, we relax the N constraints $\sum_{j=1}^{N} x_{ij} = 1 \forall i$, of $\mathcal{P}_2$, to obtain a new problem, $\mathcal{P}_3$. Obviously, an optimal solution to $\mathcal{P}_3$ is an upper bound for the optimal solution to $\mathcal{P}_2$ and, consequently, an upper bound for $\mathcal{P}_1$. Furthermore, by removing the dependencies among j, $\mathcal{P}_3$ naturally decomposes into N subproblems, $\mathcal{P}_{3,j}, j = 1, \cdots, N$, whereby a given destination j appears only in one of the subproblems

$$\mathcal{P}_{3,j} : \max_{x_{ij}} \; 1 - \sum_{i=1}^{N} x_{ij}(1-\sigma_i p_{ij})^{1/x_{ij}} \tag{35}$$

$$\text{subject to} \; \sum_{i=1}^{N} x_{ij} = 1, \; x_{ij} \geq 0 \; \forall i. \tag{36}$$

Note that the sum of the objective functions of the N problems $\mathcal{P}_{3,j}$ for the different j, is equal to the objective functions of $\mathcal{P}_3$ (and $\mathcal{P}_2$). For a fixed j, $\mathcal{P}_{3,j}$ is the single channel problem in [15], and the solution is

$$x_{ij}^{(1)} = \frac{\ln(1-\sigma_i p_{ij})}{\sum_{m=1}^{N} \ln(1-\sigma_m p_{mj})}, i = 1, \cdots, N. \tag{37}$$

Using the same reasoning, the solution to another problem, $\mathcal{P}_4$, obtained by relaxing the N constraints $\sum_{i=1}^{N} x_{ij} \forall j$, of $\mathcal{P}_2$, provides an upper bound for $\mathcal{P}_1$. $\mathcal{P}_4$, in turn, can be decomposed into N subproblems, $\mathcal{P}_{4,i}, i = 1, \cdots, N$, which can be formulated in a way very similar to $\mathcal{P}_{3,j}$. For a given source, i, the solution is

$$x_{ij}^{(2)} = \frac{\ln(1-\sigma_i p_{ij})}{\sum_{n=1}^{N} \ln(1-\sigma_i p_{in})}, j = 1, \cdots, N. \tag{38}$$

If $x_{ij}^{\star}$ are a solution to both $\mathcal{P}_3$ and $\mathcal{P}_4$, for all i, j, they also constitute a solution to $\mathcal{P}_2$. By equating the right hand sides of (37) and (38), we can obtain the conditions under which a solution to either $\mathcal{P}_3$ or $\mathcal{P}_4$ will also solve $\mathcal{P}_2$. One special case is uniform traffic, i.e., $\sigma_i p_{ij} = \sigma_{i'} p_{i'j'} \forall i, j, i', j'$.

Given the frame length, M, we use the following steps to obtain the number of slots, a_{ij}.

1) Solve $\mathcal{P}_{3,j}$, for all j, and obtain $x_{ij}^{(1)}$. Solve $\mathcal{P}_{4,i}$, for all i, yielding $x_{ij}^{(2)}$.
2) From $x_{ij}^{(1)}$ ($x_{ij}^{(2)}$) obtain the $a_{ij}^{(1)}$ ($a_{ij}^{(2)}$) that satisfy (34) and constraints (26). For all i, j, set $a_{ij} = \min\{a_{ij}^{(1)}, a_{ij}^{(2)}\}$.
3) If the a_{ij} satisfy constraints (26), stop. Otherwise, consider all pairs (i, j) such that $\sigma_i p_{ij} \neq 0$, and add 1 to a_{ij}, if doing so does not violate these constraints. Repeat until the constraints are satisfied or until for all (i, j) with $\sigma_i p_{ij} \neq 0$, adding 1 to a_{ij} would violate them. If the latter is true, repeat adding 1 to all other a_{ij} until the constraints are satisfied.[6]

The upper bound shown in Tables I and II is obtained as the minimum of the optimal solutions to problems $\mathcal{P}_3$ and $\mathcal{P}_4$ (i.e., the sum of the objective functions at the optimal solutions to the N problems $\mathcal{P}_{3,j}$ and $\mathcal{P}_{4,i}$, respectively.)

[6] In this case, slots are assigned although they will never be used. However, we have found that the number of such slots is always very small and, in the vast majority of schedules, zero.

B. Proof of Correctness of Algorithm Rearrange

We now show that algorithm REARRANGE can be used to rearrange a schedule S satisfying constraints (25), (26) and the first of (27), so that the second of (27) is finally satisfied.

Proof. By induction on the number k of the destinations considered. For $k = 1$ it is obvious that the schedule produced is one-to-one. Suppose that for $k = d-1$ a one-to-one schedule is produced, and consider destination d. Since S satisfies the first of (27), there will be at most one s such that $\delta_{sd}^{(t)} = 1$, for some t. Let s and t be such that $\delta_{sd}^{(t)} = 1$, violating

$$\sum_{j=1}^{d} \delta_{sj}^{(t)} \leq 1. \tag{39}$$

Violation of (39) means that there exists a $d_1 < d$ such that $\delta_{sd_1}^{(t)} = 1$. Combined with the second of (26) we have $\sum_{j=1}^{d} \sum_{t=1}^{M} \delta_{sj}^{(t)} < M$, and there exists a slot t_1 such that s is not assigned to transmit to any destination in t_1. Thus, $\delta_{sj}^{t_1} = 0, j = 1, \cdots, d$, and s can transmit in slot t_1 without violating (39). If no source transmits to d in slot t_1, we make s transmit to d in that slot, and we are done. Otherwise, let m (again, it can be at most one) be the station assigned to transmit to d in slot t_1. By considering the first of (26) we get that $\sum_{i=1}^{N} \sum_{t=1}^{M} \delta_{id}^{(t)} < M$. Then, there exists a slot t_2 such that $\delta_{id}^{(t_2)} = 0, i = 1, \cdots, N$, which means that no station transmits to d in slot t_2. We can now make s transmit in t_1 and m transmit in t_2. By doing this, none of t_1, t_2 violates the first of (27). t_1 also satisfies (39), but t_2 may violate it, if m has been assigned to transmit to some station $n < d$ in this slot. In the latter case, we execute the **repeat** loop once more, for destination n. The loop will be executed at most d times, after which both t_1 and t_2 will satisfy the second of (27). □

By maintaining appropriate data structures the check for violation of (39) takes constant time. Also, the time needed for the execution of the **repeat** loop, is, in the worst case, proportional to the number of slots, M. Each time, d, the outer **for** loop is executed the list will have at most M items. For each item the **repeat** loop will be executed at most d times. Thus, the complexity of the **while** loop is proportional to dM^2. The outer loop is executed for $d = 1, \cdots, N$, and the worst case complexity of the algorithm is $O(N^2M^2)$.

REFERENCES

[1] M. H. Ammar and J. W. Wong, "The design of teletext broadcast cycles," *Perform. Eval.*, vol. 5, no. 4, Nov. 1985.

[2] ——, "On the optimality of cyclic transmission in teletext systems," *IEEE Trans. Commun.*, vol. COM-35, pp. 68–73, Jan 1987.

[3] K. Bogineni, K. M. Sivalingam, and P. W. Dowd, "Low-complexity multiple access protocols for wavelength-division multiplexed photonic networks," *IEEE J. Select. Areas Commun.*, vol. 11, pp. 590–604, 1993.

[4] C. A. Brackett, "Dense wavelength division multiplexing networks: Principles and applications," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 948–964, Aug. 1990.

[5] M.-S. Chen, N. R. Dono, and R. Ramaswami, "A media-access protocol for packet-switched wavelength division multiaccess metropolitan area networks," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 1048–1057, Aug. 1990.

[6] R. Chipalkatti, Z. Zhang, and A. S. Acampora, "Protocols for optical star-coupler network using WDM: Performance and complexity study," *IEEE J. Select. Areas Commun.*, vol. 11, pp. 579–589, 1993.

[7] I. Chlamtac and A. Gan, "Channel allocation protocols in frequency-time controlled high speed networks," *IEEE Trans. Commun.*, vol. 36, pp. 430–440, Apr. 1988.

[8] A. Ganz, "End-to-end protocols for WDM star networks," in *IFIP/WG6.1-WG6.4 Workshop on Protocols for High-Speed Networks*, May 1989, pp. 219–235.

[9] A. Ganz and Y. Gao, "Time-wavelength assignment algorithms for high performance WDM star based networks," *IEEE Trans. Commun.*, Apr. 1994.

[10] A. Ganz and Z. Koren, "WDM passive star—protocols and performance analysis," in *Proc. INFOCOM '91*, pp. 991–1000, Apr. 1991.

[11] M. R. Garey and D. S. Johnson, *Computers and Intractability.* New York: Freeman, 1979.

[12] 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," *J. Lightwave Tech.*, vol. 5, no. 12, pp. 1782–1793, Dec. 1987.

[13] P. S. Henry, "High-capacity lightwave local area networks," *IEEE Commun. Mag.*, pp. 20–26, Oct. 1989.

[14] P. A. Humblet, R. Ramaswami, and K. N. Sivarajan, "An efficient communication protocol for high-speed packet-switched multichannel networks," *IEEE J. Selected Areas Commun.*, vol. 11, pp. 568–578, 1993.

[15] A. Itai and Z. Rosberg, "A golden ratio control policy for a multiple-access channel," *IEEE Trans. Automat. Contr.*, vol. 29, pp. 712–718, Aug. 1984.

[16] N. Mehravari, "Performance and protocol improvements for very high-speed optical fiber local area networks using a passive star topology," *J. Lightwave Tech.*, vol. 8, no. 4, pp. 520–530, Apr. 1990.

[17] B. Mukherjee, "WDM-Based local lightwave networks Part I: Single-hop systems," *IEEE Networking Mag.*, pp. 12–27, May 1992.

[18] G. N. Rouskas, "Single-hop lightwave WDM networks and applications to distributed computing," Ph.D. dissertation, Georgia Inst. of Tech., Atlanta, GA, May 1994.

[19] G. Semaan and P. Humblet, "Timing and dispersion in WDM optical star networks," in *Proc. INFOCOM '93*, 1993, pp. 573–577.

George N. Rouskas (M'95) received the Diploma in electrical engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1989, and the M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology, Atlanta, GA, in 1991 and 1994, respectively.

He joined the Department of Computer Science, North Carolina State University, in August 1994, as an Assistant Professor. His research interests include lightwave network architectures, high-speed networks, and performance evaluation.

Dr. Rouskas was the recipient of the 1994 College of Computing Graduate Research Assistant Award. He is a member of the ACM and of the Technical Chamber of Greece.

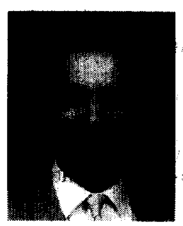

Mostafa H. Ammar (M'85) received the Ph.D. degree in electrical engineering from the University of Waterloo, Ont., Canada, in 1985. He also received the S.B. and S.M. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, in 1978 and 1980, respectively.

He is currently an Associate Professor in the College of Computing at Georgia Tech. He has been with Georgia Tech since 1985. From 1980 to 1982, he worked at Bell-Northern Research (BNR), first as a Member of Technical Staff and then as Manager of Data Network Planning. His research interests are in the areas of computer network architectures and protocols, distributed computing systems, and performance evaluation. He is the co-author of the textbook *Fundamentals of Telecommunication Networks* (New York: Wiley).

Dr. Ammar is the holder of a 1990-1991 Lilly Teaching Fellowship and received the 1993 Outstanding Faculty Research Award from the College of Computing. He is a member of the editorial board of *Computer Networks and ISDN Systems*. He is a member of the ACM and a member of the Association of Professional Engineers of the Province of Ontario, Canada.