Bandwidth Reservation in Optical WDM/TDM Star Networks

N. Saberi and M.J. Coates
Department of Electrical and Computer Engineering, McGill University
Email: nsaber1@tsp.ece.mcgill.ca, coates@ece.mcgill.ca

***Abstract*— In this paper, we propose an off-line scheduling algorithm for an optical TDM/WDM star network composed of a non-blocking optical central switch and a set of edge switches, each capable of simultaneously transmitting (and receiving) at several different wavelengths. We implement a scheduling algorithm, which assigns a cost to each time slot destined to a particular destination and attempts to allocate each request a set of time slots with the lowest cost. This strategy provides low rejection probability for future requests. In order to reduce the signaling bandwidth and the computation complexity we require the scheduler to preserve the allocation of the existing connections by modifying the schedule for only the changes in the traffic request. For deallocating the terminating connections we propose two different techniques, with different performance and complexity. Then we enhance the performance of the simpler technique with a modification to the scheduling algorithm, which should be performed only for scheduling the first frame.**

I. Introduction

When the enormous switching capacity of optical switches is considered, (overlaid) star topologies become an attractive alternative to meshes. The resource allocation is focused at a single point in the network (the central, optical switch) and hence is a much simpler exercise [1]. In such a topology, assuming some buffering at the edge switches, dynamic time slot reservation can avoid contention issues, as arise for instance with burst-switching techniques [2], and still yield efficient utilization of the central switch in high-load scenarios. In this paper, we propose an off-line scheduling algorithm for an optical TDM/WDM star network composed of a non-blocking central switch (e.g., a Batcher-Banyan optical switch) and a set of edge switches, each capable of simultaneously transmitting (and receiving) at several different wavelengths. The algorithm develops upon the minimum cost search (MCS) scheduling algorithm for a star-coupled network with tunable transceivers proposed in [3].

II. A Minimum Cost Search Scheduling Algorithm

In this section, we develop a scheduling algorithm that receives a traffic demand matrix $\mathbf{D}$, where each entry (i, j) represents the requested number of time slots in the next frame for traffic from source node i to destination node j. The algorithm then assigns the wavelength-time slot pairs (t, w) comprising the frame to the source-destination pairs. The aim is to minimize the number of rejected time slot requests in each frame. In order to reduce signalling overhead and to reduce scheduling complexity, we require the algorithm to satisfy the *transparency* property [3], which requires that the scheduling is only modified for the new requests (the time slot allocation of persistent connections does not change from frame to frame).

We consider networks in which all of the edge switches can transmit and receive at the same set of W wavelengths and use a frame of length L time slots at each wavelength. Since individual transmitters can operate simultaneously, we can treat each wavelength independently. The network has point-to-point connections so the destinations access time slots separately (there is no broadcasting in the network), so we can consider a *destination frame* of length $L \times W$ for each destination.

We commence the development of the algorithm by assigning a *cost* for the allocation of a (i, j) source-destination pair to a wavelength-time slot pair (t, w). It is important to note that this cost is determined entirely by the extant scheduling. The cost function is:

$$C_{ij}(t, w) = N_{fs}(t, w) + \lambda K_{ij}(t, w), \qquad (1)$$

where $N_{fs}(t, w)$ is the number of free sources at this time slot, i.e., the number of sources not transmitting to any other destinations on wavelength w, λ is a small positive constant, and $K_{ij}(t, w) = \{0, 1, 2\}$ is the number of additional switching operations that the core switch must perform to accommodate the allocation. The motivation behind this cost function is simple. The first term represents the current flexibility of that time slot (the number of free sources for future allocation) and reflects the desirability of retaining flexibility by allocating demands to the most constrained slots where possible. The second term reflects the desirability of minimizing the power consumption of the optical switch, which is partially determined by the number of switching operations that it must perform each frame.

The minimum cost search algorithm we propose does not achieve optimal utilization, because it does not consider the global allocation problem, instead allocating requests sequentially on a single time slot basis. The algorithm operates by repeatedly visiting the (i, j) entries in the traffic demand matrix $\mathbf{D}$ in a round-robin fashion; at each visit, if the requested number of slots has not yet been assigned, the algorithm attempts to allocate a single wavelength-time slot pair to the (i, j) request. The round-robin allocation results in an approximately fair assignment of slots to each pair.

The scheduling of a single (i, j) time slot request is performed by first identifying the (i, w)-*eligible* time slots in the j-th destination frame. The (i, w)-eligible time slots are defined as the free time slots in the j-th destination frame during which source i is not transmitting to any other destination on wavelength w. The cost $C_{ij}(t, w)$ of each of these eligible time slots is evaluated, and the demand is assigned to the slot incurring minimum cost. In the case of ties, the demand is assigned to the earliest slot and the lowest wavelength (assuming wavelengths are ordered in some fashion).

Deallocation is implemented by a reverse procedure, in which we look for and release the most costly currently-allocated time slot. This “cost-based” deallocation procedure can be replaced by a “tail” deallocation technique, in which the last time slot is released (order is determined by position in the destination frame). The tail deallocation removes the need to re-evaluate costs, which can prove computationally expensive, and results in a greater consistency in the slots that are allocated to a pair. It does however, lead to a higher rate of request rejection. A modification to the initial scheduling can substantially improve performance of tail deallocation. For the first frame, each time slot is considered independently and the non-blocking allocations for that slot are identified (these are the allocations where every source is matched with a destination for that slot). These allocations are then extracted from the demand matrix. This procedure is much more computationally demanding than the MCS scheduling and is unsuitable for transparent operation, but it can be performed for the first frame. For all subsequent frames, tail deallocation and transparent MCS allocation are performed.

III. Simulation Results

A. Performance Comparison for Cost-based and Tail Deallocation techniques

We compare the performance of the scheduling algorithm for the two different deallocation schemes in terms of the utilization and the rejection percentage for a star network of 5 edge nodes. We explore the case where the demand matrix in each frame is determined as $D(i,j) \sim \lceil \mathcal{N}(4,1) \rceil$, for $i \neq j$, where $\mathcal{N}(\mu,\sigma)$ is the *Gaussian* distribution, and $D(i,i) = 0$. We consider a destination frame length of 14 time slots and $\lambda = 0$. Figure 1 shows the average utilization and the percentage of the rejection (over 25 trials) achieved by the scheduling algorithm for the two deallocation schemes for 1500 consecutive frames. It is clear that the cost-based deallocation scheme outperforms the tail deallocation scheme, although the latter approaches the former after many frames. Figure 2 shows the improvement achieved for tail deallocation when the modified scheduling is performed in the first frame.

B. The effect of λ on the performance of the scheduling algorithms

In order to reduce the amount of consumed power at the optical switch we involve the number of switching operations at the boundary of each time slot in the cost function described in formula 1. Changing λ provides different degrees of control on the power consumption. In other words, the greater the λ the smaller the number of switching operations per frame, the lower the power consumption. We compare the performance of our scheduling algorithms for two different values of λ (λ = 0 and λ = 1). In this experiment we consider the traffic matrix described in the previous section ($D(i,j) \sim \lceil \mathcal{N}(\mu,1) \rceil$, for $i \neq j$, and $D(i,i) = 0$) and we vary μ from 1 to 8 to evaluate the performance of our algorithms in both underloaded and overloaded regimes. As figure 3 shows, λ = 1 provides a lower power consumption by reducing the number of switching operations compared to the case that we consider λ = 0. However we have a very little degradation in the utilization and the rejection percentage of the two algorithms with λ = 1. It is worth of noting that our tail deallocation scheme shows a much better behavior in terms of the power consumption compared to our cost-based deallocation scheme. A simple explanation is that cost-based deallocation violates the arrangement of the allocations in each frame by deallocating the least costly time slots which can be placed anywhere within a frame.

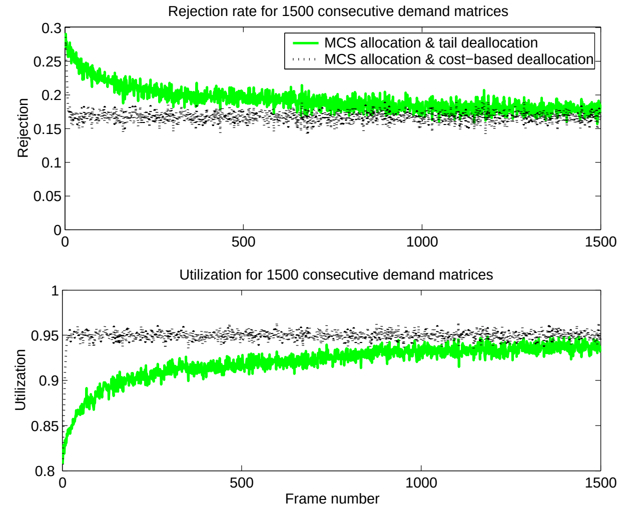

Fig. 1. Comparison of achieved utilization and rejection percentage of the MCS scheduling algorithm when cost-based and tail deallocation schemes are used.

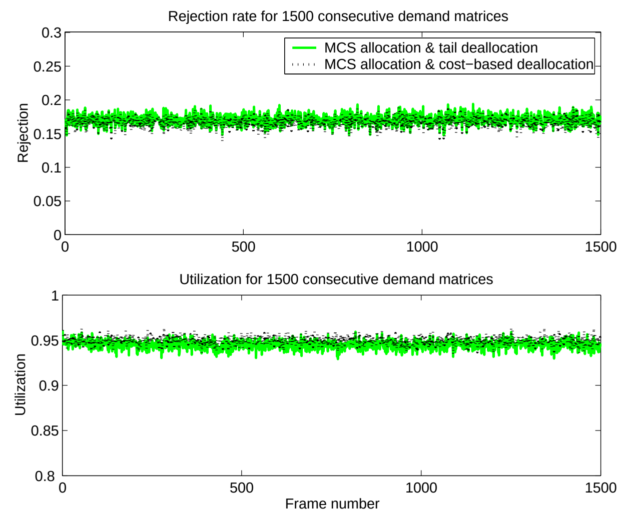

Fig. 2. Comparison of achieved utilization and rejection percentage of the MCS scheduling algorithm when cost-based and tail deallocation schemes are used with modification to the first frame allocation, in this figure the utilization and rejection percentage are virtually indistinguishable for the two deallocation schemes.

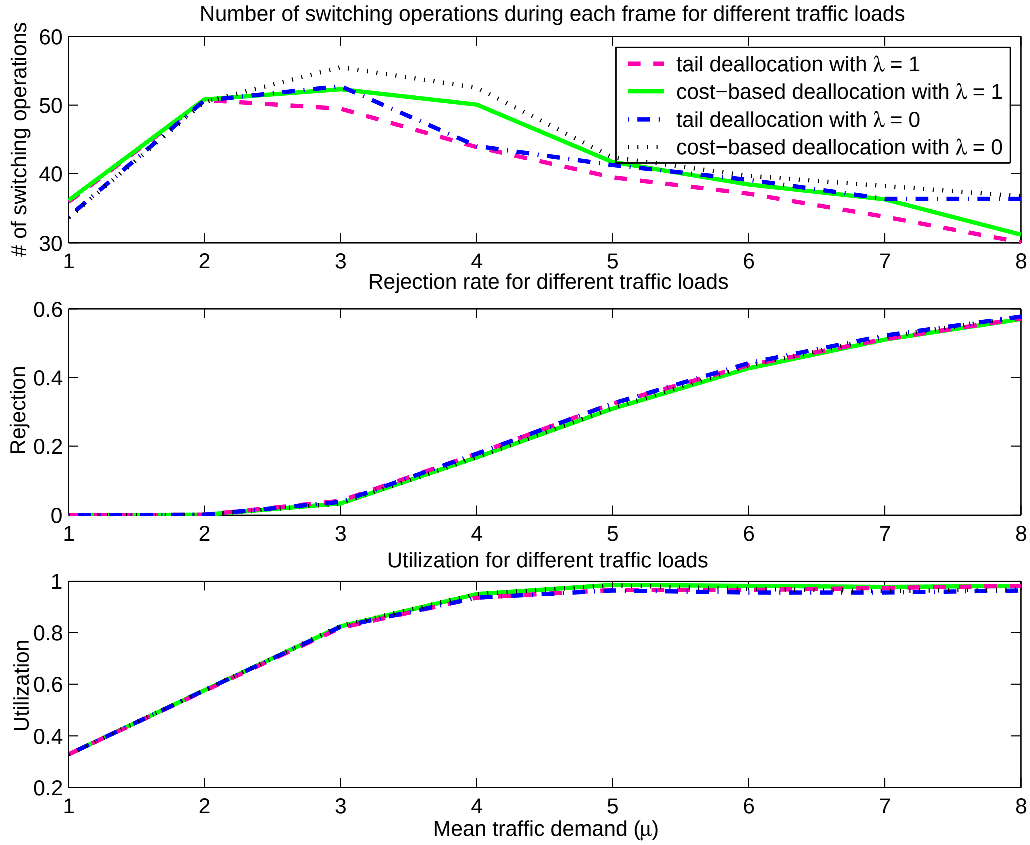

Fig. 3. Comparison of achieved utilization, rejection percentage, and number of switching operations of the MCS scheduling algorithm when cost-based and tail deallocation schemes are used with different values of λ, in both underloaded ($\mu < 3$) and overloaded ($\mu > 3$) regimes the number of switching operations are decreased with λ = 1, while the performance degradation is negligible.

While tail deallocation has a lower impact on the allocations arrangements, hence providing a more continuous distribution of the allocated time slots to each request.

IV. Conclusion

We have presented a scheduling algorithm for an optical TDM/WDM star network composed of a non-blocking central switch and a set of edge switches, each capable of simultaneously operating at several different wavelengths. The algorithm described in this article develops upon the so-called MCS algorithm, originally introduced for star-coupled networks with tunable transmitters [3] for the network described above. Our new algorithm applies the minimum cost search approach to every destination frame, as well as modifying the cost function to include a factor determined by the number of switching operations introduced by a scheduling.

Our proposed algorithm mitigates the delay due to computation time by implementing the transparency property in scheduling in that some modification to the schedule of the previous frame is sufficient to obtain the new schedule. Transparent scheduling is an acceptable solution for slowly varying traffic patterns; highly bursty traffic requires more sophisticated and accurate scheduling, which can be achieved using prediction-based methods. In these methods, prior to the arrival of the new requests, the traffic is predicted and a schedule is created based on the predictions [4]. This method can be used in conjunction with the minimum cost search algorithm to provide us with an algorithm with a very low computation time.

References

[1] B. Wen and K. Sivalingam, "Routing, wavelength and time-slot assignment in time division multiplexed wavelength-routed optical WDM networks," in *Proc. IEEE Infocom*, New York, NY, June 2002.

[2] M. Yoo, C. Qiao, and S. Dixit, "Optical burst switching for service differentiation in the next-generation optical internet," *IEEE Communication Mag.*, pp. 98–104, Feb. 2001.

[3] M.A. Marsan, A. Bianco, E. Leonardi, F. Neri, and A. Nucci, "Simple online scheduling algorithms for all-optical broadcast-and select networks," *IEEE European Trans. Telecommunications*, vol. 11, no. 1, pp. 109–116, Jan. 2000.

[4] E. Johnson, M. Mishra, and K. Sivalingam, "Scheduling in optical WDM networks using hidden Markov chain based traffic prediction," *J. Photonic Network Comms.*, vol. 3, no. 3, pp. 271–286, July 2001.