On-line Competitive Algorithms for Call Admission in Optical Networks *

Baruch Awerbuch† Yossi Azar‡ Amos Fiat§

Stefano Leonardi¶ Adi Rosén||

October 8, 1996

Abstract

In the present work we study the on-line call admission problem in optical networks. We present a general technique to deal with the problem of call admission and wavelength selection by reducing this problem to the problem of only call admission. We then give randomized algorithms with logarithmic competitive ratios for specific topologies in two models of optical networks, and consider the case of full-duplex communication as well.

1 Introduction

A new generation of networks, known as *all-optical networks*, allows data transmission rates of the order of gigabits per second, several order of magnitude higher than current networks. Wavelength division multiplexing (WDM)

*A preliminary version of this paper appeared in *Proc. of the 4'th Annual European Symposium on Algorithms*, 1996.

†Johns Hopkins University & Laboratory for Computer Science, MIT Email: baruch@cs.jhu.edu.

‡Department of Computer Science, Tel-Aviv University Email:azar@math.tau.ac.il.

§Department of Computer Science, Tel-Aviv University Email:fiat@math.tau.ac.il.

¶Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza" Email: leon@dis.uniroma1.it.

||Department of Computer Science, University of Toronto Email: adiro@cs.toronto.edu.

allows to use the bandwidth available by partitioning it in many "channels", each at a different optical wavelength. A key issue in all-optical networks is to maintain the signal in optical form and at the same wavelength during the entire transmission, since the overhead cost of conversion to electric form or between different wavelengths would be prohibitive. The signal is then converted into optical form when transmitted and converted from optical form when received.

Such optical networks consist of routing nodes connected by fiber optic links, each of which can support several wavelengths and can carry the signal in only one direction. Sometimes, a transmitter or a receiver are connected to a specific routing node. A routing node is able to route the signal coming through a certain port at a certain wavelength, to one or more outgoing ports. Two signals coming on the same input port on different wavelengths can be routed to different output ports ("generalized switches"). However, it is impossible to route to a common output port more than one signal using the same wavelength. In general, two signals can share a link only if they are transmitted on different wavelengths.

We distinguish between *switchless* networks and *reconfigurable* networks. In switchless networks the routing pattern at each routing node is fixed: a signal entering a node at a certain wavelength is directed over all the outgoing edges on which this wavelength is available. In reconfigurable networks the routing pattern at each node can be changed: a signal entering a node at a certain wavelength can be directed to any subset of the outgoing edges on which this wavelength is available. If we consider switchless networks, the network associated with each wavelength is a directed acyclic graph in order to prevent interference of a transmission with its own signal. For reconfigurable networks, the network associated with a wavelength need not be acyclic since the signal can be stopped at any routing node.

The on-line call admission problem in optical networks is that of deciding in an on-line manner (i.e., upon the arrival of a request for a connection), if to accept the request or not. In case a request is accepted, the algorithm has to assign it to a wavelength (and a route in the reconfigurable case), in a way that the new signal does not interfere with any route of any on-going call that uses the same wavelength.

Previous related work. Optical routing has mainly been studied in the off-line case. Lower bounds on the number of necessary wavelengths for permutation routing in general reconfigurable networks have been devised in [BH92, BH93, PS93]. Hypercubes have been considered by Pankay [P92], who also proves an $\Omega(\log n)$ lower bound on the number of wavelengths necessary for routing a permutation in a bounded degree network, where n is the number of nodes in the network. Aumann and Rabani [AR94] show that in bounded degree networks any permutation can be routed with $O(\log^2 n/\beta^2)$ wavelengths, where β is the edge expansion of the network. Finally, the existence of a good permutation routing algorithm for general networks with generalized switches has been obtained in [RU94].

For switchless networks, Barry and Humblet [BH92, BH93] give lower bounds on the number of necessary wavelengths necessary for a set of connections, while an almost matching upper bound is presented in [ABC$^+$94]. The connection between packet routing and optical routing is also discussed in [ABC$^+$94].

Approximation algorithms for the off-line problem of minimizing the number of wavelengths necessary to schedule an arbitrary set of connections in reconfigurable networks, the *path coloring* problem, have been proposed for specific topologies. Raghavan and Upfal [RU94] give constant approximation algorithms for trees and trees of rings. Mihalis, Kaklamnis and Rao [MKR95] give similar results in the directed case. Kleinberg and Tardos [KT95] give an $O(\log n)$ approximation algorithm for meshes and a special class of planar graphs, where n is the number of nodes in the network. Rabani [R96] improved the result for meshes to $O(\text{poly}(\log\log n))$.

The path coloring problem has also been studied in its on-line version. Algorithms with an $O(\log n)$ competitive ratio have been devised by Bartal and Leonardi [BL96] for trees, trees of rings and meshes. The authors also present an $\Omega(\frac{\log n}{\log\log n})$ lower bound for trees and an $\Omega(\log n)$ lower bound for meshes. For general networks, it has been shown that even randomized on-line algorithms cannot approximate the optimal solution with a polylogarithmic competitive ratio [BFL96].

Previous related work on on-line call admission regards virtual circuit routing, the so called *call control* problem, also motivated by its application to ATM networks [AAP93, ABFR94, AGLR94, KT95]. In this case each call requires the establishment of a fixed connection on a path between the transmitter and the receiver at a certain transmission rate. The goal is that of maximizing the number of accepted calls without exceeding the maximum available bandwidth on any connection link.

Observe that if all the links have the same bandwidth and the transmission rate required by any call is the whole available bandwidth, then the call control problem is equal to the call admission problem in optical reconfigurable networks of one wavelength. Under this restriction the call control problem admits efficient randomized algorithms for specific topologies. These topologies are trees, meshes, hypercubes and a special class of planar graphs [ABFR94, AGLR94, KT95]. Such algorithms for the one-wavelength case will be used in the present paper as subroutines for algorithms for the general case of more wavelengths. On the other hand, an $\Omega(n^\epsilon)$ lower bound on the competitive ratio of randomized algorithms for the on-line call control problem in general networks, and hence for on-line call admission control in reconfigurable optical networks of one wavelength, has been proved in [BFL96].

Results of the present paper. The results of this paper regard the on-line call admission problem in optical networks. For each request for a call an on-line algorithm must immediately decide if to accept or reject it without knowledge of future requests. The benefit obtained by the algorithm is the cardinality of the set of accepted requests. We measure the performance of an on-line algorithm by its competitive ratio [ST85]. Let A be an algorithm for the on-line benefit problem and let OPT be the optimal (off-line) algorithm that knows the whole sequence of requests in advance and deals optimally with it. An algorithm A is ρ-competitive if for every sequence of requests σ $Benefit_A(\sigma) \geq \frac{1}{\rho} Benefit_{OPT}(\sigma)$, where $Benefit_A(\sigma)$ is the benefit obtained by algorithm A on requests sequence σ.

For randomized algorithms we consider the expectation (over the random choices of the algorithm) of the benefit of the on-line algorithm. We consider the *oblivious adversary* for randomized algorithms [BBK$^+$90], meaning that the adversary knows the on-line algorithm but cannot observe the results of the coin tosses on a particular sequence.

In this paper we first develop a general on-line technique for wavelength selection that allows us to reduce the problem of wavelength selection, admission and routing, to the problem of only admission and routing. That is, this technique allows to construct on-line algorithms for multiple wavelength networks, from on-line algorithms for a single-wavelength network, with a competitive ratio that remains almost the same.

Since the call admission problem in reconfigurable optical networks of a single wavelength is equivalent to the call control problem, the above reduction gives on-line call admission algorithms for any optical network topology, when there is such an algorithm for the call control problem in the same topology.

As to switchless networks, we develop on-line call admission algorithms for networks shaped as rooted forests. We do so by giving an algorithm for one wavelength, and applying the above discussed reduction. We get an on-line algorithm with a logarithmic competitive ratio. We complement this result by a lower bound that shows that for networks shaped as rooted forests, and having any number of wavelengths, any on-line randomized algorithm has a competitive ratio of $\Omega(\log n)$, where n is the number of nodes in the network.

We also consider the problem of establishing full-duplex communications between any two points of the network. Due to the nature of switchless optical networks, two different wavelengths must be used for the two directions. The introduction of dependencies between different wavelengths does not allow in general (see Section 6) to use our general technique in order to reduce the problem to the case of a single wavelength. Therefore we develop a specific on-line call admission algorithm for full-duplex connections in switchless optical networks. Here again, we do so for networks shaped as rooted forests. We obtain an on-line algorithm with a polylogarithmic competitive ratio.

In the following we summarize the results that we present in this paper:

- A general technique that enables to obtain a $(c+1)$-competitive algorithm for call admission and wavelength selection for the multi-wavelength case, from a c-competitive algorithm for the single-wavelength case.

- Using the above technique, we obtain on-line logarithmically-competitive randomized algorithms for call admission in reconfigurable optical networks on network topologies for which on-line algorithms with logarithmic competitive ratio for the call control problem are known.

- An on-line logarithmically-competitive randomized strategy for call admission in 1-wavelength switchless networks shaped as rooted directed forests. As a result we obtain logarithmic-competitive randomized algorithms for call admission and wavelength selection in switchless optical networks shaped as rooted directed forests.

- A logarithmic lower bound on the competitive ratio of randomized algorithms for on-line call admission in switchless optical networks of arbitrary number of wavelengths, and shaped as rooted directed forests.
- An $O(\log^2 n)$-competitive algorithm for full-duplex communication in switchless optical networks shaped as rooted directed forests and arbitrary number of wavelengths.

2 The Model

2.1 Switchless networks

We model a *switchless* all-optical network as follows. There is a set of *wavelengths*, sometime called *colors*. For each color λ, we are given a directed acyclic graph $G_\lambda = (V, E)$, $|V| = n$. The graphs of the different wavelengths need not be identical. All the graphs share the same set of vertices V. The same pair of vertices may be connected by edges of different graphs but these edges are considered different since they have different colors. All the vertices of the graph consist of routing nodes. In addition to those there are transmitters and receivers. Each transmitter or receiver is connected to one specific routing node and may tune to some color.

A request for a call consists of a pair: a transmitter and a receiver, (s, t). A node s can establish a connection to receiver t using a certain color if in the graph associated with this color, there is a directed path from s to t. In this case *all* the nodes reachable from s through a directed path receive the signal and cannot get any additional information on this color. Calls (s_1, t_1) and (s_2, t_2) are called *conflicting* at color λ if for this color either there is a directed path from s_1 to t_2 or there is a directed path from s_2 to t_1. If two calls conflict at certain color λ they cannot concurrently use that color.

We are given a sequence of requests for establishing connections. For each such request our call admission and color selection algorithm should either reject the call or accept it. In the latter case it has to assign the call to one of the colors in a consistent way with any previous call assigned to the same color, i.e, the call can be assigned to color λ only if there is no other conflicting call already assigned to color λ. The benefit accrued by the algorithm is the number of calls that it accepts.

We specifically consider networks with the additional property that the graph associated with each color is such that the in-degree of each node is at most one. In other words, the graphs G_λ are directed forests, with a designated root for each tree, and all edges directed away from the root. These networks have the nice property that two calls $(s_1, t_1), (s_2, t_2)$ conflict on some color if and only if s_1 is a descendant or an ancestor of s_2 in the forest of this color. In the following we will present results for on-line call admission in switchless optical networks with this kind of topology.

2.2 Reconfigurable networks

We model a *reconfigurable* all-optical network as follows. Each color λ is associated with an undirected graph $G_\lambda = (V, E)$, $|V| = n$. The different graphs need not be identical. All the graphs share the same set of vertices V.

A request for a call is a pair of vertices (s, t) such that a receiver is associated to s and a transmitter to t. A node s can establish a connection to receiver t using a certain color if s is connected to t in the graph associated with this color. To establish the call a route from s to t should be reserved; two calls scheduled on the same color cannot share any edge.

3 Reduction from many wavelengths to one

The method that we present here applies not only to routing in optical networks, but to any benefit problem, where the benefit is gained by accommodating "entities" in any of several (not necessarily identical) "bins". That is, given a set of entities, an algorithm has to maximize the benefit gained by accepting entities. To accept an entity the algorithm has to accommodate the entity in one of several "bins". Within each bin there may be restrictions as to the set of entities that can be accommodated in it concurrently. However, we assume *total independence* between the different bins: If a set S of entities can be accepted in a given bin B when all other bins are empty, then the same set S can be accepted in B, together with any sets accepted in any of the other bins.

Given an on-line (deterministic or randomized) algorithm A for one bin, we proceed as follows to build an on-line algorithm A' for many bins. Let

B_i, $1 \le i \le k$, be the set of bins. We run a set of k copies of A, one for each bin. We use a "first fit" type algorithm. When a request arrives, we first give it as input to A_1. If it is accepted by A_1, the handling of the request is terminated. Otherwise we present the request to A_2 and so on, until either the request is accepted or the set of k copies of A is exhausted.

Theorem 1 *If A is ρ competitive then A' is $\rho + 1$ competitive.*

Proof. We use the following notations:

- R - the sequence of entities presented to A'.
- O - the set of entities accepted by the optimal algorithm.
- O_i - the set of entities accepted by the optimal algorithm into bin number i.
- T_i - the set of entities accepted by A' into bin number i.
- $B(S)$ - the sum of benefits associated with the entities in the set S.

For clarity of presentation, we first prove the theorem for the case that A is deterministic.

By the definition of A' the sequence of entities presented to A_i is the subsequence of R obtained from R by eliminating $\cup_{j<i} T_j$. Denote this subsequence by R_i. Clearly,

$$O_i \setminus \cup_{j<i} T_j \subseteq R_i \ .$$

Therefore, the optimal gain that can be obtained from R_i is at least

$$B(O_i \setminus (\cup_{j<i} T_j)).$$

Since A_i is ρ competitive it will gain a benefit

$$B(T_i) \ge \frac{1}{\rho} \cdot B(O_i \setminus (\cup_{j<i} T_j)) = \frac{1}{\rho} B(O_i) - \frac{1}{\rho} B((\cup_{j<i} T_j) \cap O_i) \ .$$

It follows that

$$\begin{aligned}\sum_{i=1}^{k} B(T_i) &\geq \sum_{i=1}^{k} \frac{1}{\rho} B(O_i) - \sum_{i=1}^{k} \frac{1}{\rho} B((\cup_{j<i} T_j) \cap O_i)] \\ &\geq (\sum_{i=1}^{k} \frac{1}{\rho} B(O_i)) - \frac{1}{\rho} B(\cup_{i\leq k} T_i) \\ &\geq \sum_{i=1}^{k} \frac{1}{\rho} B(O_i) - \frac{1}{\rho} \sum_{i=1}^{k} B(T_i) \ ,\end{aligned}$$

where the second inequality follows since the sets O_i are pairwise disjoint. We get that

$$(1+\rho) \sum_{i=1}^{k} B(T_i) \geq \sum_{i=1}^{k} B(O_i).$$

and then A' is $(\rho + 1)$-competitive.

This completes the proof for the case that A is deterministic.

If the algorithm A is randomized the same analysis goes through as well, with some additional technicalities.

Let r_i be the sequence of random choices taken by algorithm A_i. For any l, the set A_l is determined by the sequences r_j, $j \leq l$.

Since A_i is ρ competitive we know (by the same arguments as in the deterministic case) that

$$\forall \{r_j\}_{j=1}^{i-1}, \quad E_{r_i}[B(T_i) \mid T_1 \ldots T_{i-1}] \geq \frac{1}{\rho} B(O_i) - \frac{1}{\rho} B((\cup_{j<i} T_j) \cap O_i).$$

Therefore,

$$\begin{aligned}E_{r_j, 1\leq j\leq k}[B(T_i)] &= E_{r_j, 1\leq j<i}[E_{r_i}[B(T_i) \mid T_1 \ldots T_{i-1}]] \\ &\geq E_{r_j, 1\leq j<i}[\frac{1}{\rho} B(O_i) - \frac{1}{\rho} B((\cup_{j<i} T_j) \cap O_i)] \\ &= \frac{1}{\rho} B(O_i) - \frac{1}{\rho} E_{r_j, 1\leq j<i}[B((\cup_{j<i} T_j) \cap O_i)] \ .\end{aligned}$$

Since for any j the random choices of $A_l, l > j$, have no effect on T_j, we can also write

$$E_{r_j, 1\leq j\leq k}[B(T_i)] \geq \frac{1}{\rho} B(O_i) - \frac{1}{\rho} E_{r_j, 1\leq j\leq k}[B((\cup_{j<i} T_j) \cap O_i)] \ .$$

To evaluate $E_{r_j,1\leq j\leq k}[\sum_{i=1}^{k} B(T_i)]$, we sum the above expression for $1 \leq i \leq k$:

$$\begin{aligned}
E_{r_j,1\leq j\leq k}[\sum_{i=1}^{k} B(T_i)] &= \sum_{i=1}^{k} E_{r_j,1\leq j\leq k}[B(T_i)] \\
&\geq \sum_{i=1}^{k} \frac{1}{\rho} B(O_i) - \sum_{i=1}^{k} \frac{1}{\rho} E_{r_j,1\leq j\leq k}[B((\cup_{j<i} T_j) \cap O_i)] \\
&= \sum_{i=1}^{k} \frac{1}{\rho} B(O_i) - \frac{1}{\rho} E_{r_j,1\leq j\leq k}[\sum_{i=1}^{k} B((\cup_{j<i} T_j) \cap O_i)] \\
&\geq \sum_{i=1}^{k} \frac{1}{\rho} B(O_i) - \frac{1}{\rho} E_{r_j,1\leq j\leq k}[\sum_{i=1}^{k} B(T_i)],
\end{aligned}$$

where the last inequality follows from the same arguments as in the deterministic case.

It follows that the (randomized) algorithm A' is $(\rho+1)$-competitive. □

3.1 Applications to optical networks

The general technique of the previous section can be applied to optical networks both for the switchless and the reconfigurable case. Every wavelength λ with the corresponding network G_λ corresponds to a bin, and we can then use for every wavelength a base algorithm that deals with the problem of call admission (without wavelength selection). Observe that for different wavelengths the corresponding networks are not necessarily identical. However, we have to use for every network G_λ a base algorithm that is suited for its topology T_λ.

G_λ, $\lambda \in \Lambda$, has topology T_λ. Λ, if every network Assume that Then,

Corollary 2

a $(\rho+1)$-competitive (deterministic or randomized) algorithm A' for a ρ-competitive (deterministic or randomized) call admission algorithm for topology T_λ. call admission in optical networks with the set of wavelengths for every wavelength $\lambda \in \Lambda$, A_λ is is suited for topology T_λ. there exists

From the above corollary and from the results of [ABFR94, AGLR94, KT95] we obtain the following for reconfigurable optical networks.

Corollary 3 *There exists an $O(\log n)$-competitive randomized algorithm for the on-line call admission problem in reconfigurable optical networks if for every wavelength λ the corresponding network G_λ has topology isomorphic to a tree, a mesh or a nearly Eulerian planar graph.*

In the next section we give competitive algorithms for call admission in switchless optical networks of one wavelength, thus obtaining competitive algorithms for switchless optical networks as well.

4 The base algorithm for one-wavelength switchless networks

In this section we give an optimal call admission scheme for one-wavelength switchless optical networks with topology of rooted directed forests.

The network is a directed forest $D = (V, E)$. We consider every tree T of the forest separately. Consider a call from s to t. This call can possibly be scheduled in the network D if there exists a subtree rooted at s containing t and hence a directed path from s to t.

We apply the "classify and randomly select" paradigm [ABFR94, LT94]:

- Preprocess the graph and produce a partition of the vertices of the graph into several classes.
- Select a class, uniformly at random among the above classes, and consider only calls with source vertex in the selected class.

The classification algorithm that we use defines inductively tree T_i, with $T_0 = T$, as follows. Let S_i be the set of vertices of T_i that either are leaves or have a unique path to a leaf. Assign the vertices of S_i to class i. The tree T_{i+1} is obtained from T_i by removing the vertices in S_i. The procedure is continued until all the vertices are assigned to some class.

Lemma 4 *Every node of the graph D is assigned to some class i, such that $0 \leq i \leq \lceil \log n \rceil$.*

Proof. Vertices in distinct trees of the forest are classified independently. Thus we restrict our attention to a single rooted directed tree.

A directed tree can be partitioned into a set of disjoint directed paths in the following way: every path of the partition starts with either the root or a vertex such that its parent has out-degree bigger than 1. Every path ends with a vertex being a leaf or having out-degree bigger than 1. These paths can also degenerate into a single vertex. We call this partition the *path partition* of the tree. If we contract every path of the path partition of a tree T_i to a single vertex, then the resulting tree, that we call the contracted tree T_i^c, has internal vertices with out-degree at least 2.

Let $P(T_i)$ be the number of vertices in T_i^c. We prove that for any $i \geq 0$, either T_{i+1} is empty or $P(T_{i+1}) \leq P(T_{i)}/2$. Since $P(T_0) \leq n$ we have that the number of classes is at most $1 + \lceil \log n \rceil$.

Consider the tree T_i and its contracted tree T_i^c. All the vertices of T_i being leaves or having a unique path to a leaf are assigned to class i and removed to create T_{i+1}. These are vertices of paths that are represented by leaves in T_i^c. Since in T_i^c all internal vertices have out-degree at least 2, removing the leaves of T_i^c halves the number of vertices of T_i^c. Observe now that the number of vertices in T_{i+1}^c is at most the number of non-leaf vertices of T_i^c. This is because every path in the path partition of T_{i+1} is either a whole path of the path partition of T_i, or is obtained by the concatenation of a number of paths of the path partition of T_i. Therefore, $P(T_{i+1}) \leq P(T_i)/2$, and the claim is proved. □

In the following we denote by $i = 0, \ldots, L$, for $L \leq \lceil \log n \rceil$, the number of classes of the partition of D.

Recall that two calls (s_1, t_1) and (s_2, t_2) are consistent if there is no directed path from s_1 to t_2 and there is no directed path from s_2 to t_1. On a rooted forest, if each of the calls (s_1, t_1) and (s_2, t_2) is feasible by its own, (that is, t_1 is a descendent of s_1 and t_2 is a descendent of s_2), then the above condition is that s_1 is neither an ancestor nor a descendant of s_2. Note that vertices of class i form a set of vertex-disjoint paths where any vertex in a path is neither an ancestor nor a descendant of any vertex in another path. Therefore, two calls originating at vertices of class i that belong to different paths are consistent.

When a call (s, t) arrives, if there is no path from s to t we ignore the call. Note that no algorithm can accept such call. If there is a path from s to t the call is assigned to the class of node s. The randomized on-line algorithm selects uniformly at random one of the classes and considers only calls that belong to this class. To these calls it applies the *greedy* algorithm that accepts a new call (s,t) if and only if it is consistent with all previously accepted calls in the same class. I.e., if there is no call (s',t'), previously accepted, such that either s is a descendent of s' or s' is a descendent of s.

In the following we analyze this algorithm. Let $\mathcal{OPT}_i$ be the benefit of the optimal off-line solution, if the sequence of requests is restricted to the calls of class i, and let ALG_i be the benefit of the greedy algorithm applied to the calls of class i.

Lemma 5 *For any sequence of calls and any class i, $ALG_i = \mathcal{OPT}_i$.*

Proof. Consider the vertices of a class i. These vertices can be partitioned into a number of mutually vertex-disjoint directed paths of the tree, such that there is no directed path connecting two vertices on different paths. Therefore, at most one call with source vertex on any given such path can be accepted (by any algorithm), and calls accepted on different paths are mutually consistent. The lemma follows by observing that both the optimal solution and the greedy solution use the same number of such distinct paths. □

Theorem 6 *The algorithm for one color is $O(\log n)$-competitive.*

Proof. Let $\mathcal{OPT}$ be the number of calls accepted by the optimal off-line algorithm. We obtain $\mathcal{OPT} \leq \sum_{i=0}^{L} \mathcal{OPT}_i$. The on-line algorithm selects uniformly at random one class amongst the $L+1$ classes and obtains the optimal benefit for that class. Hence, the expected benefit of the on-line algorithm is:

$$E(ALG) = \sum_{i=0}^{L} \frac{1}{L+1} \mathcal{OPT}_i \geq \frac{1}{L+1} \mathcal{OPT}.$$

Since $L \leq \lceil \log n \rceil$, the algorithm is $O(\log n)$-competitive. □

Using the above algorithm and the reduction of Section 3 we obtain the following result for switchless optical networks with any number of wavelengths.

Corollary 7 *There exists an $O(\log n)$-competitive randomized algorithm for the on-line call admission problem in switchless optical networks if for every wavelength λ, the graph corresponding to this wavelength, D_λ, is a directed forest.*

5 Lower bound for switchless networks

In this section we prove that the $O(\log n)$-competitive randomized algorithm for on-line call admission in switchless optical networks shaped as rooted directed forests is optimal up to a constant factor. Note that this lower bound applies to any number of wavelengths present in the network.

Theorem 8 *Any randomized algorithm for on-line call admission in switchless optical networks shaped as rooted directed forests has competitive ratio* $\Omega(\log\ n)$.

Proof. Let Λ be the set of wavelengths. Let $n = 2^d - 1$ for some integer value $d > 1$. The network D_λ associated with every color $\lambda \in \Lambda$ is a complete binary rooted directed tree of n vertices and d levels. The levels are indexed from 0 (the root) to $d-1$ (the leaves). In this proof every call is denoted only by its source vertex, while the destination vertex is an arbitrary vertex in the subtree rooted at the source vertex.

We show that for any randomized on-line algorithm, there is a sequence of requests in the network defined above, such that the ratio between the benefit of the optimal off-line solution and the expected on-line benefit has lower bound $\Omega(\log n)$.

For this purpose we consider a sequence of $d-1$ sets of calls. Set j, $j = 0, \ldots, d-2$, has $|\Lambda|2^j$ calls: for each vertex of level j, the set includes $|\Lambda|$ calls originating at this node.

Now fix the randomized on-line algorithm. Let x_j be a random variable representing the number of calls of set j accepted by the on-line algorithm, if presents with sets 0 to j. Observe that for any j, the random variable x_j has the same distribution regardless of whether any set l, $l > j$ is presented.

Let us consider the number of vertices of level $d-2$ in every subtree rooted at an accepted call. Every accepted call of level j covers 2^{d-2-j} vertices of level $d-2$. For any outcome of the coin tosses the overall number of covered vertices cannot exceed the total number of such vertices which is $|\Lambda|2^{d-2}$. We therefore get that for any outcome of the coin tosses $\sum_{j=0}^{d-2} x_j 2^{d-2-j} \leq |\Lambda|2^{d-2}$. Since the above holds for any outcome of the coin tosses we have also that $E[\sum_{j=0}^{d-2} x_j 2^{d-2-j}] \leq |\Lambda|2^{d-2}$. It follows from the linearity of expectation that $\sum_{j=0}^{d-2} E[x_j] 2^{d-2-j} \leq |\Lambda|2^{d-2}$, and therefore $\sum_{j=0}^{d-2} E[x_j] 2^{-j} \leq |\Lambda|$.

Let $\frac{1}{\rho_j}$ be the ratio between the expected benefit of the on-line algorithm and the benefit of the off-line algorithm after set number j is presented. At this point the optimal algorithm has benefit $2^j|\Lambda|$. Hence, we have that

$$\frac{1}{\rho_j} = \frac{2^{-j}}{|\Lambda|} \sum_{l=0}^{j} E[x_l] \ .$$

Summing over all j we get

$$\sum_{j=0}^{d-2} \frac{1}{\rho_j} = \sum_{j=0}^{d-2} \frac{2^{-j}}{|\Lambda|} \sum_{l=0}^{j} E[x_l] = \sum_{l=0}^{d-2} \frac{E[x_l]}{|\Lambda|} \sum_{j=l}^{d-2} 2^{-j} \leq \frac{2}{|\Lambda|} \sum_{l=0}^{d-2} E[x_l] 2^{-l} \leq 2.$$

Therefore, there exists at least one index k such that $1/\rho_k \leq \frac{2}{d-1}$. If the algorithm is presented with the sequence composed of the sets $0 \leq j \leq k$, then the ratio between the expected benefit of the on-line algorithm and the benefit of the off-line algorithm is at most $\frac{2}{d-1}$. Since $d = \log(n+1)$ the lemma follows. □

6 Full-duplex communication

In full-duplex communication every call requires to set up a bidirectional communication between a node u and a node v. Due to the nature of switchless networks, the two directions of the communication have to be scheduled on different wavelengths. Thus, in general, it is not possible to use our technique of Section 3 in this case, since the sets of accepted calls in each wavelength are not independent [1]. Therefore, we give a specific algorithm to handle this case, for networks of topology of a rooted directed forest.

In what follows we define our algorithm. For a graph D_λ the algorithm assigns each vertex to one of the classes defined by the classification procedure of Section 4 for one-way communication. Let L be the maximum index

[1]Note that in special cases we can use the algorithm of Section 4 also for full duplex communication. Assume we have a set of wavelengths $\{\Lambda_i\}_{i=1}^{n}$ such that the topology of each corresponding network is a rooted forest, and another set $\{\Lambda_i'\}_{i=1}^{k}$, where the topology of the network associated with Λ_i' is the topology of the network associated with Λ_i, with edge directions inverted. Then we can use the algorithm of Section 4 on the network of $\{\Lambda_i\}_{i=1}^{k}$, requesting one-directional communications, and schedule the opposite direction on the set $\{\Lambda_i'\}_{i=1}^{k}$ respectively.

of a class used in any graph D_λ. The classes are indexed from 0 to L with $L \leq \lceil \log n \rceil$. For duplex communication on switchless optical networks it is clearly necessary to use different wavelengths for the two communication directions. We therefore define arbitrarily one direction as *forward* and the other direction as *backward*. Every color is reserved with probability $\frac{1}{2}$ to forward communications (a *forward color* in the following), and with probability $\frac{1}{2}$ to backward communications (a *backward color* in the following). Let $|\Lambda|$ be the number of available colors. Let H be a $|\Lambda|$-bit vector defining for each color a direction, forward or backward, according to the choice made as described above. We define the class $C_{i,j}^H$ as the set of calls (u, v) such that

- There exists a forward color λ_f such that u has class i in D_{λ_f} and there is a path from u to v in D_{λ_f}.
- There exists a backward color λ_b such that v has class j in D_{λ_b} and there is a path from v to u in D_{λ_b}.

The algorithm chooses uniformly at random in $\{0, \ldots, L\}$ an integer i for the forward direction, and an integer j for the backward direction, and considers only calls that belong to class $C_{i,j}^H$.

If a call is considered according to the above rule, then the algorithm applies a greedy strategy for this call: Call (u, v) is accepted with a forward color λ_f and a backward color λ_b if the assignment does not conflict with any call of the same class accepted previously with λ_f or λ_b in one of the two directions.

Let $\mathcal{OPT}^H$ be the optimal benefit if any call (u, v) can only be accepted according to the directions specified by H, i.e., from u to v with a forward color, and from v to u with a backward color. Let $\mathcal{OPT}_{i,j}^H$ and $ALG_{i,j}^H$ be the benefit of the optimal solution and the benefit of the greedy algorithm, respectively, if any call (u, v) can only be accepted according to the directions specified by H, and with class i from u to v and with class j from v to u. Let $E_H(*)$ be the expected value over all the possible choices of H.

Lemma 9 *If $\mathcal{OPT}$ is the size of the optimal solution on a sequence of calls, then $E_H(\mathcal{OPT}^H) \geq \frac{1}{4}\mathcal{OPT}$.*

Proof. Let us denote with $C(\mathcal{OPT})$ the set of calls accepted in the optimal solution. We associate a $0-1$ random variable $X_c(H)$ with any call $c \in$

$C(\mathcal{OPT})$. Let us assume that call c is accepted by the optimal solution with some color λ_1 in its forward direction, and with some color λ_2 in its backward direction. $X_c(H) = 1$ if λ_1 is a forward color in H and λ_2 is a backward color in H, $X_c(H) = 0$ otherwise. Clearly, $\mathcal{OPT}^H \geq \sum_{c \in C(\mathcal{OPT})} X_c(H)$. If $|\Lambda|$ is the number of colors, then the set of different assignments for H is equal to 2^Λ. We now observe that λ_1 and λ_2 appear as forward and backward color respectively in $2^{\Lambda-2}$ assignments of H. Therefore, $X_c(H) = 1$ for $\frac{2^\Lambda}{4}$ values of H. Hence, we derive the claim of the lemma:

$$\begin{aligned} E_H(OPT^H) &\geq \tfrac{1}{2^{|\Lambda|}} \textstyle\sum_H \sum_{c \in C(\mathcal{OPT})} X_c(H) \\ &= \tfrac{1}{2^{|\Lambda|}} \textstyle\sum_{c \in C(\mathcal{OPT})} \sum_H X_c(H) \\ &= \tfrac{1}{4}\mathcal{OPT}\ . \end{aligned}$$

□

Lemma 10 *For any choice of H and any class $C^H_{i,j}$, $ALG^H_{i,j} \geq \frac{1}{3}\mathcal{OPT}^H_{i,j}$.*

Proof. Assume that the on-line greedy algorithm accepts the full-duplex communication (u, v) with forward color λ_1 and backward color λ_2. The optimal off-line solution in class $C^H_{i,j}$ can accept the call (u, v) with two different colors, and in the worst case at most 2 others full-duplex communications *of the same class* that are conflicting with (u, v) and cannot be accepted by the on-line algorithm elsewhere. This is the case where one of these two communications is conflicting with the communication from u to v in λ_1, and the second one is conflicting with the communication from v to u in λ_2. □

Theorem 11 *The algorithm for full-duplex communication is $O(\log^2 n)$-competitive.*

Proof. For any H, class $C^H_{i,j}$ is selected with probability at least $\frac{1}{(L+1)^2}$. Since any possibility to accept a call according to H is considered in at least one class $C^H_{i,j}$, we derive the following upper bound on the optimal solution for a given choice of H: $\mathcal{OPT}^H \leq \sum_{i=0}^L \sum_{j=0}^L \mathcal{OPT}^H_{i,j}$. Therefore, by Lemma 9 and Lemma 10 the expected benefit of the on-line algorithm is:

$$\begin{aligned} E(ALG) &\geq E_H\left(\tfrac{1}{(L+1)^2} \textstyle\sum_{i=0}^L \sum_{j=0}^L ALG^H_{i,j}\right) \\ &\geq \tfrac{1}{3(L+1)^2} E_H(\mathcal{OPT}^H) \\ &\geq \tfrac{1}{12(L+1)^2}\mathcal{OPT}. \end{aligned}$$

Since $L \leq \lceil \log n \rceil$, the claim is proved. □

References

[ABC+94] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswani, B. Schieber, and M. Sudan. Efficient routing and scheduling algorithms for optical networks. In *Proceedings of the 5th Annual ACM-SIAM Symposuim on Discrete Algorithms*, pages 412–423, 1994.

[AAP93] B. Awerbuch, Y. Azar, and S. Plotkin. Throughput competitive on-line routing. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, November 1993.

[ABFR94] B. Awerbuch, Y. Bartal, A. Fiat, and A. Rosén. Competitive non-preemptive call control. In *Proceedings of 5th ACM-SIAM Symposium on Discrete Algorithms*, 1994.

[AGLR94] B. Awerbuch, R. Gawlick, F.T. Leighton, and Y. Rabani. On-line admission control and circuit routing for high performance computing and communication. In *Proceedings of the 35th IEEE Annual Symposium on Foundations of Computer Science*, 1994

[AR94] Y. Aumann, Y. Rabani. Improved bounds for all-optical routing. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, 1995.

[BBK+90] S. Ben-David, A. Borodin, R.M. Karp, G. Tardos, A. Widgerson. On the power of randomization in on-line algorithms. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, 1990.

[BFL96] Y. Bartal, A. Fiat, S. Leonardi. Lower bounds for on-line graph problems with application to on-line circuit and optical routing. In *Proceedings of the 28th ACM Symposium on Theory of Computing*, 1996.

[BH92] R.A. Barry and P.A. Humblet. Bounds on the number of wavelengths needed in wdm networks. In *LEOS'92 Summer Topical Mtg. Digest*, pages 114–127, 1992.

[BH93] R.A. Barry and P.A. Humblet. On the number of wavelengths and switches in all optical networks. *IEEE Trans. Comm.*, 1993.

[BL96] Y. Bartal and S. Leonardi. On-line routing in all-optical networks. Manuscript, 1996.

[KT95] J. Kleinberg, E. Tardos. Disjoint paths in densely embedded graphs. In *Proceedings of the 36th IEEE Annual Symposium on Foundations of Computer Science*, 1995.

[LT94] R.J. Lipton, A. Tomkins. On-line interval scheduling. In *Proceedings of the 5th ACM-Siam Symposium on Discrete Algorithms*, 1994.

[MKR95] M. Mihail, C. Kaklamanis, and S. Rao. Efficient Access to Optical Bandwidth. In *Proceedings of the 36th IEEE Annual Symposium on Foundations of Computer Science*, pages 548–557, 1995.

[P92] R.K. Pankay. Architectures for linear light-wave networks. PhD thesis, MIT, 1992.

[PS93] G.R. Pieris and G.H. Sasaki. A linear lightwave benes network. *IEEE/ACM Trans. on Networking*, 1993.

[R96] Y. Rabani. Path coloring on the mesh. In *Proceedings of the 37th IEEE Annual Symposium on Foundations of Computer Science*, 1996. To appear.

[RU94] P. Raghavan and U. Upfal. Efficient routing in all optical networks. In *Proceedings of the 26th Annual Symposium on Theory of Computing*, pages 133–143, 1994.

[ST85] D. Sleator, R.E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of ACM* 28, 1985.