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

Baruch Awerbuch† Yossi Azar‡ Amos Fiat§ Stefano Leonardi¶
Adi Rosén||

Abstract

We study the on-line call admission problem in optical networks. We present a general technique that allows us to reduce the problem of call admission and wavelength selection to the call admission problem. We then give randomized algorithms with logarithmic competitive ratios for specific topologies in switchless and reconfigurable optical networks. We conclude by considering full duplex communications.

1 Introduction

A new generation of networks, known as *all-optical networks* [Gr92], allows data transmission-rates of the order of gigabits per second, several order of magnitude faster than current networks. Wavelength division multiplexing (WDM) partitions the available bandwidth into many "channels", each of which is assigned a different 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. This is because the technology to perform optical conversion between different wavelengths is currently very expensive. An alternative, to convert the signal from optical to electrical form and then

*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. Research supported in part by the Israel Science Foundation and by the United States-Israel Binational Science Foundation (BSF). Email:azar@math.tau.ac.il.

§Department of Computer Science, Tel-Aviv University Research supported in part by the Israel Science Foundation and by the United States-Israel Binational Science Foundation (BSF). Email:fiat@math.tau.ac.il.

¶Dipartimento di Informatica Sistemistica, Università di Roma "La Sapienza", via Salaria 113, 00198-Roma, Italia. This work was partly supported by EU ESPRIT Long term Research Project ALCOM-IT under contract n. 20244, and by Italian Ministry of Scientific Research Project 40% "Algoritmi, Modelli di Calcolo e Strutture Informative". Email: leon@dis.uniroma1.it

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

back to optical form on a different wavelength, could be prohibitively expensive. Thus, we deal with the model where the the signal is entirely optical except for the transmission and receiving ends.

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. Some of the routing nodes may have transmitters or receivers attached to them. We consider routing nodes of the following type: the routing node maps incoming wavelengths on a given arc to outgoing arcs. However, if two incoming signals on the same wavelength (from two different incoming arcs) are mapped to the same outgoing arc then this creates interference. 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 to all the outgoing arcs 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 arcs on which this wavelength is available. Clearly if one is to establish a connection between two nodes then is is never necessary to direct an incoming signal to more than one outgoing arc.

If we consider switchless networks, the network associated with each wavelength is a directed acyclic graph. Any cycle implies that a transmission interferes with its own signal. For reconfigurable networks, the network associated with a wavelength need not be acyclic since cycles can be broken at a routing node by appropriately routing the signal to an appropriate subset of the outgoing arcs.

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 connection request), if to accept the request or not. In case a request is accepted, the algorithm has to assign it to an available wavelength. For reconfigurable networks the algorithm also has to decide upon an appropriate routing. For both switchless and reconfigurable networks, the algorithm must ensure that the new signal does not create interference with any existing call (such interference is only possible for calls using the same wavelength).

Previous related work. Optical routing has been primarily studied in the context of off-line decision making. Lower bounds on the number of wavelengths necessary for permutation routing in general reconfigurable networks have been devised in [BH92, BH93, PS93]. Hypercubes have been considered by Pankaj [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 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 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. Mihail, 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; Gerstel, Sasaki and Ramaswami consider the problem on rings and trees, and give $O(\log n)$-competitive algorithms [GSR96]. Bartal and Leonardi also present an $\Omega(\frac{\log n}{\log \log n})$ lower bound for trees and an $\Omega(\log n)$ lower bound for meshes. For the line network the problem is equivalent to the well-studied problem of on-line coloring of interval graphs; Kierstead and Trotter give a 3-competitive algorithm for this problem [KT81]. On the other hand it has been shown that there exist networks for which 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, LMPR97]. 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.

One can observe that if all the links have the same bandwidth and every call requires the entire link bandwidth then the call control problem is equivalent to the call admission problem in optical reconfigurable networks of one wavelength. Under this restriction the call control problem admits good 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 our work as subroutines for algorithms for the multiple-wavelengths case. 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 has been proved in [BFL96]; this is also a lower bound for on-line call admission in reconfigurable optical networks of one wavelength.

Results of the present paper. This paper deals with the on-line call admission problem in optical networks. Upon every call request the on-line algorithm must immediately decide if to accept or reject it without knowledge of future requests. The benefit for an algorithm is the number of calls accepted.

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 $\text{Benefit}_A(\sigma) \geq \frac{1}{\rho}\text{Benefit}_{OPT}(\sigma)$, where $\text{Benefit}_A(\sigma)$ 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. In the context of randomized algorithms, we only consider the *oblivious adversary* [BBK+90], *i.e.*, the adversary knows the algorithm but cannot observe the results of the coin tosses on a particular sequence. A randomized algorithm A is ρ-competitive if for every sequence of requests σ, $E[\text{Benefit}_A(\sigma)] \geq \frac{1}{\rho}\text{Benefit}_{OPT}(\sigma)$, where the expectation is over the random choices of algorithm A.

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. *I.e.*, this technique allows one to construct on-line algorithms for multiple wavelength networks, from on-line algorithms for a single-wavelength network, without (much) change in the competitive ratio.

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

For switchless networks, we develop on-line call admission algorithms for a special type of network, the rooted forest network. On such networks we get an $O(\log n)$ competitive ratio, where n is the number of nodes in the network. We do so by giving an algorithm for one wavelength, and applying the reduction mentioned above. We complement this result by giving a lower bound that shows that for such networks, with an arbitrary number of wavelengths, the competitive ratio of any on-line randomized algorithm is bounded below by $\Omega(\log n)$.

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 different directions. The introduction of dependencies between different wavelengths does not allow us in general (see Section 6) 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 rooted forest networks. We obtain an on-line algorithm with a polylogarithmic competitive ratio.

In the following we summarize the results presented in this paper:

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

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

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

- 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, for rooted forest networks.

- An $O(\log^2 n)$-competitive algorithm for full-duplex communication in switchless optical rooted forest networks and an 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*, also referred to as *colors*. Each color λ is associated with a directed acyclic graph $G_\lambda = (V, E)$, $|V| = n$. The graphs associated with 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 arcs of different graphs but these arcs are considered different since they have different colors. All the vertices of the graph consist of routing nodes. In addition to these there are transmitters and receivers. Each transmitter or receiver is connected to one specific routing node and may tune in to a color of its choice.

A request for a call consists of a pair: a transmitter and a receiver, (s, t). A connection between s and t can be established by 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 in the graph receive the signal and cannot get any other signal on this color. Calls (s_1, t_1) and (s_2, t_2) are called *conflicting* at color λ if there is a directed path from s_1 to t_2 colored λ or there is a directed path from s_2 to t_1 colored λ. If two calls conflict at color λ then they cannot both be assigned color λ.

We are given a sequence of requests for connections. For each such request the call admission and color selection algorithm must either reject the call or accept it. In the latter case it has to assign the call to one of the colors so as to be consistent with any previous call assigned the same color, i.e, the call can be assigned to color λ only if there is no other conflicting call already assigned 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 arcs are directed away from the root. These networks have the property that two calls (s_1, t_1) and (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 rooted forest for this color. In the following we present results for on-line call admission in switchless optical networks of this topology, called rooted forest networks.

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), where s is the transmitter and t is the receiver. A node s can establish a connection to t using color λ if there is a path from s to t in the graph colored λ. To establish the call, a route from s to t has to be reserved; two calls assigned the same color cannot share an arc.

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 we have k bins, then for any i, if the set S_i can be accepted in bin i with sets $S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_k$, in the other bins, then the set S_i can be accepted to bin i also if the other bins contain arbitrary other sets $S'_1, \ldots, S'_{i-1}, S'_{i+1}, \ldots, S'_k$.

Given an on-line (deterministic or randomized) algorithm A for one bin, we proceed as follows to build an on-line algorithm A' for multiple bins. Let B_i, $1 \leq i \leq k$, 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 the request is accepted by A_1, it is accommodated in B_1 according to the decisions of A_1. Otherwise it is presented to A_2. In general, if the request is accepted by A_i it is accommodated in bin B_i, and otherwise it is presented to A_{i+1}. If the entity is presented to A_k and is rejected, then A' rejects the request.

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 placed by the optimal algorithm in bin B_i.

- T_i - the set of entities placed by A' in bin B_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 deterministic A.

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

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

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

This completes the proof for deterministic A.

For randomized A the same analysis applies with some modification for the random choices of the algorithms.

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

Since A_i is ρ-competitive we know (by the same argument as for deterministic A_i) 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,

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

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

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

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

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

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. □

The above analysis can be shown to be tight. *I.e.*, there are sets of bins and sequences such that when applying the first-fit strategy, in every bin one can accepts a $\frac{1}{\rho}$ fraction of the optimum benefit achievable from the entities presented to the bin, and the total benefit is (arbitrarily close to) a $\frac{1}{\rho+1}$ fraction of the total achievable benefit. Consider the following instance. For any $k \ge 0$ we have bins B_i, $0 \le i \le k$, and use entities of types Y_i, $0 \le i \le k$; in addition we have entities of type X. Each bin B_i can accommodate (any number of) entities of type Y_i or (any number of) entities of type X. However, it cannot accommodate an entity of type Y_i along with an entity of type X. Assume that for every bin we have a ρ-competitive algorithm, $\rho > 1$. [1] We present the following sequence of entities: Let $\mathrm{C} = 2^k$, and let $g_0 = \mathrm{C}$; define recursively $g_i = g_{i-1} - \lceil g_{i-1}/\rho \rceil$. Note that

$$C(1-\frac{1}{\rho})^i \ge g_i \ge C(1-\frac{1}{\rho})^i - \sum_{j=0}^{i-1}(1-\frac{1}{\rho})^j \ge C(1-\frac{1}{\rho})^i - \rho \ .$$

First, C entities of type X are presented. Then, for any i, $0 \le i \le k$, we present g_i entities of type Y_i.

[1]For the case of $\rho = 1$ a simpler similar construction applies.

Note that an off-line algorithm can accept all the entities of type Y_i, $0 \leq i \leq k-1$, into bins B_i, $0 \leq i \leq k-1$, respectively, and all the entities of type X into bin B_k. This yields a solution of size

$$\begin{aligned}
OPT &= \mathrm{C} + \sum_{i=0}^{k-1} g_i \\
&\geq \mathrm{C} + \mathrm{C} \cdot \sum_{i=0}^{k-1} (1 - \frac{1}{\rho})^i - \sum_{i=0}^{k-1} \rho \\
&= \mathrm{C} \cdot (1 + \rho(1 - (1 - \frac{1}{\rho})^k))) - k\rho) \\
&\geq \mathrm{C} \cdot ((1 + \rho(1 - (1 - \frac{1}{\rho})^k)) - \frac{k\rho}{C}) \ .
\end{aligned}$$

On the other hand, the first-fit strategy uses a ρ-competitive algorithm for every bin. Considering the first phase of the sequence in which only entities of type X are presented, it is easy to see by induction on i that a possible situation is that g_i entities of type X are presented to the algorithm of bin B_i, and $\lceil g_i/\rho \rceil$ entities of type X are accepted into bin B_i. Now observe that the optimal solution achievable from the entities of type X and Y_i presented to the algorithm of bin B_i is of size g_i. Therefore, it is possible that no entity of type Y_i is accepted into bin B_i. We conclude that it is possible that the on-line solution will be of size

$$\begin{aligned}
ON &= \sum_{i=0}^{k} \lceil g_i/\rho \rceil \leq \sum_{i=0}^{k} (g_i/\rho + 1) \\
&\leq \mathrm{C} \cdot \frac{1}{\rho} \sum_{i=0}^{k} (1 - \frac{1}{\rho})^i + (k+1) \\
&= \mathrm{C} \cdot (1 - (1 - \frac{1}{\rho})^{k+1}) + (k+1) \\
&= C \cdot ((1 - (1 - \frac{1}{\rho})^{k+1}) + \frac{k+1}{C}) \ .
\end{aligned}$$

We thus have that the ratio between the optimal solution and the solution of the first-fit strategy can be $f(\rho, k) = \frac{1+\rho(1-(1-\frac{1}{\rho})^k)-\frac{k\rho}{C}}{1-(1-\frac{1}{\rho})^{k+1}+\frac{k+1}{C}}$. Recall that $C = 2^k$. For any fixed $\rho \geq 1$ $\lim_{k \to \infty} f(\rho, k) = \rho + 1$, which shows that our analysis is tight.

Note that the above construction uses bins of different types. In fact, if we restrict the instances to be those with identical bins then a better bound can be deduced from an analysis for a similar-in-spirit (off-line) algorithm [CFN77]. Their analysis can be modified to suit our setting and to show an upper bound of $\frac{e^{1/\rho}}{e^{1/\rho}-1}$ (instead of $\rho + 1$) on the competitive ratio. Their analysis can also be modified, in the same way as is done in our analysis above, to show the same bound for the randomized case as well.

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 cases. Every wavelength λ with network G_λ corresponds to a bin, and we can then use for every wavelength an algorithm that deals with the problem of call admission (without wavelength selection). Observe that for different wavelengths the corresponding networks are not necessarily identical. For every color λ we use an algorithm suitable for the appropriate topology (G_λ), where G_λ is the network consisting of edges of color λ.

The following follows from the general theorem above:

Corollary 2 *Assume that for every wavelength $\lambda \in \Lambda$, A_λ is a ρ-competitive deterministic (randomized) call admission algorithm for network G_λ. Then, there exists a $(\rho+1)$-competitive deterministic (randomized) algorithm A' for call admission in optical networks with the set of wavelengths Λ.*

From the above corollary and from the results of [ABFR94, AGLR94, KT95, LMPR97] 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 algorithm for one-wavelength switchless rooted forest optical networks.

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 that contains 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 trees 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, $P(T_{i+1}) \leq P(T_i)/2$. □

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 in advance 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 pairwise 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 pairwise 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. Clearly $\mathcal{OPT} \leq \sum_{i=0}^{L} \mathcal{OPT}_i$. 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 rooted forest networks is optimal up to a constant factor. Note that this lower bound applies to any number of wavelengths.

Theorem 8 *Any randomized algorithm for on-line call admission in switchless optical rooted forest networks 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 is $\Omega(\log n)$. For this purpose we consider a sequence of $d-1$ sets of calls. Set j, $j = 0, \ldots, d-2$, has $|\Lambda| \cdot 2^j$ 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 $|\Lambda| \cdot 2^{d-2}$. We therefore get that for any outcome of the coin tosses $\sum_{j=0}^{d-2} x_j 2^{d-2-j} \leq |\Lambda| 2^{d-2}$, and therefore $E[\sum_{j=0}^{d-2} x_j 2^{d-2-j}] \leq |\Lambda| 2^{d-2}$. 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 in time the optimal algorithm has benefit $2^j \cdot |\Lambda|$. 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. $\square$

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 for full-duplex communication the technique of Section 3, since the sets of accepted communications with different wavelengths are not independent. [2] We call such networks a *rooted forest network*. We give specific algorithms to deal with such networks.

In what follows we define our algorithm. For a graph D_λ the algorithm assigns each vertex to one of the classes as defined by the classification procedure of Section 4. Let L be the maximum index 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 every full-duplex communication request (u, v) we arbitrarily define the direction from u to v as forward and the direction from v to u 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 every color a direction on every wavelength, forward or backward, according to the choice made as described above. We define the class $C^H_{i,j}$ 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 $C^H_{i,j}$.

The algorithm applies a greedy strategy to accept calls from the selected class $C^H_{i,j}$: Call (u, v) is accepted with a forward color λ_f and a backward color λ_b if this assignment does not conflict with any call of the same class accepted previously with λ_f or λ_b.

Let $\mathcal{OPT}^H$ be the optimal benefit on the input sequence which is composed of the calls that can 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 $\mathcal{OPT}^H_{i,j}$ and $ALG^H_{i,j}$ be the benefit of the optimal solution and the benefit of the on-line solution, respectively, on the input sequence which is restricted to those calls (u, v) such that u is of class i and v is of class j, and can be accepted according to the directions specified by H. Let $E_H(\cdot)$ 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 by $C(\mathcal{OPT})$ the set of calls accepted in the optimal solution. We associate a $0-1$ random variable $X_c(H)$ with every call $c \in C(\mathcal{OPT})$. Let us assume that call c is accepted

[2] 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 arc 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 scheduling the opposite direction on the set $\{\Lambda'_i\}_{i=1}^k$ respectively.

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 there are $2^{|\Lambda|}$ different possible values of H. We observe that λ_1 and λ_2 appear as forward and backward colors, respectively, in $2^{|\Lambda|-2}$ possible values of H. Therefore, $X_c(H) = 1$ for $\frac{2^{\Lambda}}{4}$ different values of H. Hence, we derive:

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

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 for the calls in class $C^H_{i,j}$ can accept the call (u, v) with two different colors, and at most 2 other 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 $\frac{1}{(L+1)^2}$. 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 Symposium 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*, pages 32–40, 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*, pages 312-320, 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*, pages 412-423, 1994

[AR94] Y. Aumann, Y. Rabani. Improved bounds for all-optical routing. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 567-576, 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*, pages 379-386, 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*, pages 531-540, 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. In *Proc. of the 24th International Colloqium on Automata, Languages and Programming*, Springer Verlag LNCS 1256, pages 516–526, 1997.

[CFN77] G. Cornuejols. M. Fisher, G. Nemhauser. Location of Bank Accounts to Optimize Float. *Management Science* 23, pages 789–810, 1977.

[Gr92] P.E. Green. *Fiber-Optic Communication Networks*. Prentice Hall, 1992.

[GSR96] O. Gerstel, G.H. Sasaki, and R. Ramaswami. Dynamic Channel Assignment For WDM Optical Networks With Little Or No Wavelength Conversion. In *The 34th Allerton conf. on Communication, Control, and Computing*, 1996.

[KT81] H. A. Kierstead and W. T. Trotter. An Extremal Problem in Recursive Combinatorics. In *Congress. Numer.*, 33, pages 143–153, 1981.

[KT95] J. Kleinberg, E. Tardos. Disjoint paths in densely embedded graphs. In *Proceedings of the 36th IEEE Annual Symposium on Foundations of Computer Science*, pages 52–61, 1995.

[LMPR97] On-line Randomized Call Control Revisited. In *Proc. of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), pages 323-332, January 1998.

[LT94] R.J. Lipton, A. Tomkins. On-line interval scheduling. In *Proceedings of the 5th ACM-SIAM Symposium on Discrete Algorithms*, pages 302-311, 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. Pankaj. 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*, pages 400–409, 1996.

[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(2): 202-208, 1985.