

Measurement Based Reconfigurations in Optical Ring Metro Networks

Original

Measurement Based Reconfigurations in Optical Ring Metro Networks / Bianco, A., J. M., F., G., G., Neri, F., C., P.. - In: JOURNAL OF LIGHTWAVE TECHNOLOGY. - ISSN 0733-8724. - STAMPA. - 23:10(2005), pp. 3156-3166. [10.1109/JLT.2005.856158]

Availability:

This version is available at: 11583/1403706 since:

Publisher:

IEEE

Published

DOI:10.1109/JLT.2005.856158

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Measurement-Based Reconfiguration in Optical Ring Metro Networks

Andrea Bianco, *Member, IEEE*, Jorge M. Finochietto, *Student Member, IEEE*, Gianluca Giarratana, Fabio Neri, *Member, IEEE*, and Chiara Piglione, *Student Member, IEEE*

Abstract—Single-hop wavelength division multiplexing (WDM) optical ring networks operating in packet mode are one of the most promising architectures for the design of innovative metropolitan network (metro) architectures. They permit a cost-effective design, with a good combination of optical and electronic technologies, while supporting features like restoration and reconfiguration that are essential in any metro scenario. In this article, we address the tunability requirements that lead to an effective resource usage and permit reconfiguration in optical WDM metros. We introduce reconfiguration algorithms that, on the basis of traffic measurements, adapt the network configuration to traffic demands to optimize performance. Using a specific network architecture as a reference case, the paper aims at the broader goal of showing which are the advantages fostered by innovative network designs exploiting the features of optical technologies.

Index Terms—All-optical networks, dynamic network configuration, logical topology design, wavelength division multiplexing (WDM) packet rings.

I. INTRODUCTION

AS INTERNET usage continues its growth, carriers continue to see a steady increase of packet data traffic in their metropolitan networks (metros). Today's network solutions are mostly based on circuit-switched synchronous optical network/synchronous digital hierarchy (SONET/SDH) rings that are not efficient in carrying data traffic, due to their inherent asymmetry, and bursty and self-similar behavior. Several evolutions of legacy SONET/SDH to packet-switched technologies are currently being proposed. For example, the IEEE 802.17 resilient packet ring (RPR) standard aims at solving problems from which SONET/SDH networks suffer in supporting packet data by optimizing bandwidth sharing. However, as higher rates need to be supported, both SONET/SDH and RPR node costs increase, since all incoming/outgoing and in-transit traffic always needs to be processed electronically. Similar scaling problems arise in metro infrastructures based upon switched gigabit Ethernet, with additional concerns related to fair resource allocation and quality of service (QoS)

control. Basically, in current solutions, network scalability is limited because nodes must switch/process the full network bandwidth.

Due to advances in optical technology [1], new packet-switched networks can be devised that can cost-effectively sustain larger bandwidths. Metros seem to be one of the best arenas for an early penetration of these technologies. On the one hand, high-capacity requirements can be satisfied by exploiting fiber bandwidth by means of wavelength division multiplexing (WDM), without requiring node interfaces to access and electronically process the full network bandwidth. On the other hand, packet traffic can be handled by temporally sharing WDM channels, either by dynamically setting up lightpaths between nodes willing to communicate, or by exploiting statistical packet multiplexing in static channels.

In this context, single-hop optical ring networks operating in packet mode are considered a promising architecture for future metros [2]–[7]. The ring topology has been extensively proposed in the literature because of its simplicity and since it easily satisfies restoration requirements. Besides, the single-hop approach avoids complex switching in the optical domain and thus permits a cost-effective balance of optics and electronics. In these networks, nodes are equipped with few (typically one) transceivers, and each transceiver operates at the data rate of a single WDM channel. Paths between nodes are created by dynamically sharing, on a packet-by-packet basis, WDM channels, without requiring nodes to process the full network bandwidth. However, tunability at transceivers is required to exploit the fiber bandwidth by temporally allocating all-optical single-hop bandwidth between nodes in all available channels. Fast-tunable transceivers today are not yet fully mature, as noted in [8] and [9]; however, significant progresses were done, and these components are expected to be soon commercially available. Arrays of fixed-tuned transceivers can be a temporary feasible substitute for the tunable transceiver, particularly when the number of WDM channels is not very large.

Due to the cost of tunability at transceivers, media access protocols that require packet-by-packet tunability only at one end of the all-optical path (i.e., either only at the transmitter, or only at the receiver) have been studied to save the cost of the still quite expensive tunable devices. Usually, these protocols assume a fast-tunable transmitter and a fixed receiver, permanently tuned to a WDM channel [10], [11]. When a node needs to send a packet, it simply tunes its transmitter to the receiver's destination wavelength. This implies that transmitter tuning times must be negligible with respect to the packet duration to obtain a good efficiency. Simple distributed access protocols

Manuscript received December 15, 2004; revised July 12, 2005. This paper was partially funded through the Italian FIRB project Algorithms for Dynamic Optical Networks based on Internet Solutions (ADONIS), the Italian PRIN project Wdm Optical Network DEMonstrator over Rings (WONDER), and the Sixth Framework Programme (FP6) European Network of Excellence e-Photon/ONe.

The authors are with the Dipartimento di Elettronica, Politecnico di Torino Corso Duca degli Abruzzi 24 10129 Torino, Italy (e-mail: bianco@mail.tlc.polito.it; finochietto@mail.tlc.polito.it; giarratana@mail.tlc.polito.it; neri@mail.tlc.polito.it; piglione@mail.tlc.polito.it).

Digital Object Identifier 10.1109/JLT.2005.856158

can be designed for these tunable-transmitter/fixed-receiver architectures.

If the number of nodes is larger than the number of WDM channels, a decision problem arises concerning the allocation of the different receivers to WDM channels to equalize the traffic among the available channels. If fixed receivers are considered, any allocation is permanent and cannot be updated in response to long-term changes in the traffic pattern, which are typical in metros. Therefore, it may be worthwhile to reallocate, i.e., tune to different wavelengths, receivers, to dynamically keep the network in an optimal operation point. One elegant way of achieving this result is to introduce slow (hence inexpensive) tunability in receivers. This tunability does not need to be fast, since it must not track packet-by-packet variations, but longer term variations of the traffic pattern. Low-cost devices available today (e.g., mechanical or thermo-optic filters) can be suitable to implement this slow receiver-tunability feature.

If slow tunability is present at receivers, the impact of reconfiguring network receivers must be taken into proper account [12] when solving the problem of allocating receivers to match traffic conditions. In fact, retuning a receiver implies introducing a period of service disruption during which nodes cannot transmit to that receiver. As a consequence, the reconfiguration must not only bring the network to an optimal operation point, but also minimize service disruption.

The scope of this paper is to introduce reconfiguration algorithms in a single-hop optical ring network with slow receiver-reconfiguration capabilities. Although a similar problem has been studied in [12], this paper looks at defining proper reconfiguration algorithms relying on traffic measurements to detect the traffic pattern, which is not assumed to be known. The main contribution is the introduction of a reconfiguration schema, which aims at keeping the network at an optimal operation point with minimum service disruption.

Before proceeding, we note that a possible alternative to the use of tunable devices is to equip nodes with W fixed receivers; this choice would simplify the access protocols, as packets can be transmitted in any free slot, independently of the destination, and potentially improve performance; the receiver-reconfiguration issue would vanish. This architecture however requires that receivers process in electronics the full network bandwidth, which is against the basic assumptions of our design, which aims at exploiting the available bandwidth in optics through WDM without requiring a speedup in electronic technology. Nodes would be more expensive, and the network less scalable. W receivers obviously permit larger output loads, and a performance increase when a high load is directed toward a specific node. However, this would require complex and expensive electronics in each node; moreover, on average, the W receivers would be significantly underutilized, since the overall network traffic is destined to multiple receivers whereas each receiver is equipped to receive the full network traffic. If more bandwidth is required for a specific node, it would always be possible to plug more than one single tunable-transceiver interface to support this requirement.

We would like to remark that the issue addressed in the paper has an interest beyond the application to the specific considered metro architecture. Indeed, the rate at which network resources




Fig. 1. Logical network model.

should be reallocated (which often translates into the speed at which transceivers and switches must tune in optical networks) cannot keep up with the continuous reduction of packet durations with increasing line rates. Packet-by-packet control and retuning introduces strong technological challenges, and consequently high costs, which may not be strictly necessary to provide acceptable levels of QoS to users. Indeed, service requirements do not scale with transmission speeds and packet durations, and most network control dynamics need to be matched to service requirements instead of to optical transmission rates.

The paper is organized as follows. In Section II, we describe the considered network architecture and the system model, providing motivations for the importance of network reconfigurability, i.e., adaptability to long-term traffic fluctuations. The problem is formalized in terms of mixed integer linear programming (MILP) in Section III. In Section IV, we introduce the basic reconfiguration mechanisms, assuming two different measurement schemas, named incoming- and in-transit-traffic measurements. Next, in Section V, we discuss simulation results to assess the properties of the proposed algorithms. Finally, we draw some conclusions in Section VI.

II. SYSTEM MODEL

We consider a specific WDM optical packet network, physically made of two counter-rotating rings. This architecture was proposed and is currently being studied and prototyped in the framework of the Italian national project called WONDER.¹ Each ring encompasses N nodes and conveys W wavelengths. The network is assumed to be synchronous and time slotted. During a time slot, at most one packet can be transmitted in one of the W available slots, one for each wavelength channel. Rings are used in a peculiar way: One ring is used for transmission only while the other one is used for reception only. To provide connectivity between the two rings, a folding point

¹<http://www.tlc-networks.polito.it/wonder/>

is needed, where transmission wavelengths are switched to the reception path, as sketched in Fig. 1. Transmitted packets travel towards the folding point in a first ring traversal, are switched to the reception path, and then received during a second ring traversal. If each node can become the folding point (i.e., if each node has a switching capability), then the network preserves the interesting restoration property of rings, as described in [13]. Although this architecture does not exploit wavelength spatial reuse, it avoids transmission impairment (e.g., noise recirculation) typical of ring topologies, while guaranteeing that all the network traffic accepted prior to the fault can be supported also after restoration (note that this may not be the case in ring networks with spatial reuse).

Nodes are equipped with a fast-tunable transmitter and exploit WDM to partition the traffic directed to disjoint subsets of destination nodes, each subset comprising the destinations whose receivers are currently tuned to the same wavelength. Nodes tune their transmitters to the receiver's destination wavelength, and establish a temporary single-hop connection lasting one time slot. Sharing of wavelength channels is therefore achieved according to a statistical time division multiple access (TDMA) scheme. Access decisions are based on a channel inspection capability (similar to the carrier sense functionality in Ethernet—see [13]), by which nodes know which wavelengths were not used by upstream nodes in each time slot. Priority is given to in-transit traffic, i.e., a multichannel empty-slot protocol is used.

The considered topology has intrinsic fairness problems, since the first nodes have a better access chance. Fairness-control issues are beyond the scope of this paper. It was, however, shown that simple algorithms can ensure throughput fairness; for example, in [10], we effectively used variations of the Metaring fairness-control scheme in WDM rings. We assume in the rest of the paper that fairness issues can be successfully solved at the access-protocol level.

While fast tunability at the transmitters is required to provide full connectivity between nodes, slow tunability at receivers is needed to follow traffic changes in time, since nodes' receivers may have to be retuned to balance the offered load among the available WDM channels. As an example, a configuration where N/W nodes are tuned on each of the W available wavelength channels is optimal for uniform traffic distribution, since the load on all available channels is equalized, but it may be highly inefficient for a different traffic distribution, as, for example, in the presence of heavy-traffic servers. Thus, if the traffic distribution is unknown or time variable, the slow tunability becomes a must to allow a dynamic network configuration to be obtained.

We consider that a fast-tunable transmitter is capable of retuning itself in negligible time between two consecutive time slots, while a slow-tunable receiver is characterized by a tuning latency τ .

Fast tunability is not assumed at the receiver to limit both network costs and reconfiguration algorithm complexity. Due to the receiver tuning latency, all nodes willing to transmit to receivers that are currently involved in a retuning process must refrain from transmission for a period of time at least equal to the tuning latency τ . This is to avoid packet losses, since

retransmission costs may be very high in networks with large delay \times bandwidth products.

We focus in this paper on the control of slow-tunable devices, to adapt network configuration to slowly varying traffic distributions. The reconfiguration mechanism proposed is conceptually centralized, i.e., it runs in a given master node, which is responsible for reconfiguration decisions, as well as for collecting information used to drive the reconfiguration process. The master node collects information on the network status, as channel loads and node transmission needs, by proper traffic measurements, described later in more detail. If the new network status requires or suggests that receiver reconfiguration may be useful, a proper reconfiguration algorithm is run and a new network configuration, i.e., a new assignment of receivers to wavelengths, is computed. The master node broadcasts signaling messages containing the new configuration to slave nodes. To account, in a simple way, for the propagation time of signaling messages, we assume the worst case situation, in which a receiver becomes unavailable due to a retuning for twice the round trip time (RTT, equal to one ring-traversal latency) to allow for signaling-message propagation, plus the tuning latency to allow the slow-tunable device to tune to the proper wavelength. More precisely, the reconfiguration process goes through three steps: first, the master node disables all transmissions towards nodes that must be reconfigured; after two RTTs (to permit in-flight packets to reach the receiver after traversing the transmission and the reception paths), nodes retune their receivers, and finally, after the tuning latency, transmissions towards destinations involved in the reconfiguration process are reenabled. Note that, under light-load conditions, these periods of receiver unavailability only affect transmission delays, while under high-load conditions they may also cause buffer overflows and packet losses due to network-capacity reductions.

In Section IV, we define both the measurement techniques adopted to monitor network status, and the algorithms used to determine, first, the need of reconfiguration, and then, the new network configuration.

III. PROBLEM FORMULATION

In this section, we formally model the reconfiguration problem. We focus on a given measured traffic pattern, and formalize the problem of finding an optimal network configuration, i.e., an optimal allocation of receivers to wavelengths. The optimality is related to network performance: Thus, balancing the traffic load on the available wavelengths by properly assigning receivers to wavelengths is an obvious optimization goal, which translates into minimizing the maximum load \mathcal{L}_{\max} across all WDM channels. However, it is important to observe that reconfiguration decisions must take into account not only traffic balance among WDM channels but also service disruption due to temporary receiver unavailability. Whereas traffic balance depends on the aggregate traffic of each receiver and on the assignment of receivers to wavelengths, service disruption is related to network RTT and to receiver tuning latency. The algorithms used to determine whether a reconfiguration is needed or not should carefully consider these two aspects.

We will formalize the reconfiguration problem via a two-objective problem, taking care first of the wavelength's loads and then of reconfiguration costs. Indeed, if we assume that the traffic pattern is fully known, i.e., we assume to know also the newly detected traffic-pattern duration, we could compute exactly the benefit of each network configuration in terms of bandwidth utilization, taking into account both the bandwidth gain due to the new receiver configuration and the bandwidth loss due to RTTs and tuning latencies. This would lead to an optimization problem with a single objective function that considers overall network bandwidth only. However, this is clearly not possible in practice; moreover, sometimes it can be interesting to study the problem with cost functions induced by management issues, and not directly related to performance. For these reasons, we keep costs and loads as two separate objectives, rather than merging them in a linear combination and moving to a scalar problem.

The output of the measurement process is stored on the traffic matrix $\mathbf{T} = [t_{ij}]$, where t_{ij} is the traffic rate from node i to node j normalized to the channel bandwidth, and $t_j = \sum_{i=1}^N t_{ij}$, the aggregate receiver bandwidth of node j .

The cost of tuning receiver j to wavelength k is denoted by c_{jk} while the cost of a network reconfiguration is the sum of the costs of all retuned receivers. This cost can refer to the total number of receivers to be retuned, or to the amount of bandwidth lost, or to other different metrics. Although the reconfiguration algorithms we present can be adapted to different cost definitions, in the remainder of the paper, we define the reconfiguration cost as

$$c_{jk} = \begin{cases} 0, & \text{iff node } j \text{ is currently on wavelength } k \\ 1, & \text{otherwise.} \end{cases}$$

The MILP problem formulation is obtained by introducing a set of control variables x_{jk} that considers the potential allocation of the receivers after the reconfiguration:

$$x_{jk} = \begin{cases} 1, & \text{iff node } j \text{ will receive on wavelength } k \\ 0, & \text{otherwise.} \end{cases}$$

Thus, the mathematical model becomes a two-objective minimization function:

$$\text{Minimize} \quad \left[\mathcal{L}_{\max}; \sum_{j=1}^N \sum_{k=1}^W c_{jk} x_{jk} \right] \quad (1)$$

subject to the following constraints:

$$\mathcal{L}_{\max} \geq \mathcal{L}(k) = \sum_{j=1}^N t_j x_{jk} \quad \forall k, 1 \leq k \leq W \quad (2)$$

$$\sum_{k=1}^W x_{jk} = 1 \quad \forall j, 1 \leq j \leq N. \quad (3)$$

Equation (2) guarantees that no wavelength has a load $\mathcal{L}(k)$ larger than \mathcal{L}_{\max} , while (3) ensures that each receiver is allocated only to one wavelength, since nodes are equipped with a single receiver. The problem is solved in a lexicographic

fashion, i.e., finding out first the minimum \mathcal{L}_{\max} for the current traffic pattern, and then looking for the least disruptive allocation of receivers that equals \mathcal{L}_{\max} . In particular, the problem of finding a well-balanced allocation of receivers to wavelength channels is equivalent to the well-known problem of scheduling jobs on identical parallel machines [14], where the receiver's aggregate traffic t_j represents the job's duration, and wavelengths represent machines. This is an NP-hard problem; hence, our problem, being a generalization, is also NP-hard.

IV. RECONFIGURATION MECHANISM

In this section, we describe heuristic approaches to solve the reconfiguration problem. We classify reconfiguration mechanisms on the basis of the traffic measurement approach used to collect information on network status. Two different approaches are defined to collect measured data.

- 1) *In-transit-traffic measurement scheme*—The master node observes the amount of traffic transmitted on all WDM channels.
- 2) *Incoming-traffic measurement scheme*—all nodes measure the amount of data arriving for transmission to the metro and send this information to the master node.

The first scheme is simpler to implement and requires less signaling. However, only the traffic that was able to access the network can be observed and measured; thus, no overload situation can be measured. In the second one, the real traffic offered to the network can be detected, thereby leading to a more reliable estimate of the traffic pattern, at the price of signaling overhead to convey this information to the master node. In particular, the estimate is accurate also in overloaded conditions.

Since reconfiguration mechanisms are triggered by measurements, a key point is the estimation of the current traffic pattern. We assume long-term variations of the traffic with respect to packet dynamics; thus, measurements can be done periodically over a measurement window long enough to estimate steady-state traffic conditions. If traffic measurements are done by each node considering incoming traffic addressed to the metro, then a straightforward mean of the measured samples is accurate enough to characterize traffic conditions, provided that the measurement-window size is properly set. However, if measurements are done looking at in-transit traffic, then measured samples can be affected by transient phenomena. In fact, nodes could buffer packets either if the network configuration is not matched to the traffic pattern (i.e., prior to reconfiguration) or if a node is retuning. As a consequence, after network reconfiguration, the traffic sent on the network represents not only current incoming-traffic conditions but also transient ones. Therefore, an exponentially weighted mean is used to mitigate the effect of transient measured values. More details on the measurement process are provided later.

A. Incoming-Traffic Measurement Scheme: The 3-Step Algorithm

In this scheme, nodes measure the traffic groomed from the local area and willing to travel over the metro. Each node

estimates one row of the traffic matrix \mathbf{T} , which is communicated to the master node to keep an updated estimation of the whole traffic matrix. The estimation is done periodically over a measurement window of duration T_m .

The reconfiguration scheme we propose for this case is named 3-step algorithm; it aims at a full network reconfiguration at the end of each measurement window. In the first step, named load balancing (LB), the problem of balancing receiver loads is solved without accounting for service disruption. Once a solution has been found, Step 2, named wavelength assignment (WA), and Step 3, named receiver swapping (RS), try to improve the LB solution by jointly keeping channels balanced and avoiding unnecessary reconfigurations.

1) *Load Balancing (LB)*: As previously stated, the problem of finding a well-balanced allocation of receivers to wavelength channels can be mapped to the problem of scheduling jobs on identical parallel machines [14]. Although the problem falls in the class of NP-hard problems, approximation algorithms that limit the distance from the optimal solution do exist. The longest processing time (LPT) algorithm is one of these, which guarantees that any solution is at most $4/3$ greater than the optimal one [15], thus, providing an upper bound on the distance of our algorithm from the optimal solution.

LPT works on W bins, representing wavelengths to which receivers must be allocated. Receiver aggregate bandwidths, obtained by matrix \mathbf{T} , are loaded in initially empty bins following these steps:

- 1) Sort nodes by decreasing $t_j \forall j = 1, \dots, N$.
- 2) Allocate largest t_j to least loaded bin.
- 3) If unassigned receivers do exist, go to 2).

The LPT algorithm is run each time a new traffic matrix is detected at the master node to find out if a new configuration is needed. Whether to schedule a reconfiguration or not depends on how much the new allocation improves network performance. A threshold is defined to decide whether it is worthwhile to reconfigure the network by comparing the old and the new allocations: The threshold is applied to $\sum_{k=1}^W \min(1, \mathcal{L}(k))$, the maximum overall traffic that can be handled by the network. Once LB has found a solution that improves over the previous one more than the threshold, the next step is to associate each of the loaded bins with wavelengths to minimize the number of nodes that must be reconfigured.

2) *Wavelength Assignment (WA)*: The next step is to associate the bins filled by LB with a proper wavelength, so as to minimize the number of nodes that should retune their receiver. The WA problem can be seen as a bipartite matching problem [16]. A bipartite graph has two sets of nodes: edges may not connect nodes in the same set. A matching is a subset of edges with the constraint that, at most, one edge in the matching can be connected to each node. If a weight is associated to each edge, the weight of a matching is the sum of the weights associated with the selected edges. A matching with weight w^* has maximum weight if no other matching exists with weight larger than w^* .

In Fig. 2, numbers close to the left nodes represent the receivers that must be allocated to the corresponding bin while numbers close to the right nodes represent the receivers cur-




Fig. 2. Bipartite graph for WA.

rently allocated to the corresponding wavelength. Each edge from bin i to wavelength j has a weight w_{ij} equal to the number of allocated receivers on bin i currently tuned to wavelength j . By running a maximum weight matching (MWM) algorithm, we obtain the WA that minimizes the number of nodes involved in the reconfiguration, thus minimizing the cost function. As a result, we minimize the number of reconfigurations required to obtain the load balancing on wavelength channels as determined by LB in the previous step.

3) *Receiver Swapping (RS)*: Even if WA tries to minimize the number of receiver reconfigurations, there may exist some unneeded reconfigurations. To understand why, let us assume that the two receivers i and j are similarly or equally loaded. Since no notion of previous wavelength allocation is used when determining bin assignments via LB, and since WA cannot modify receiver allocations to bins, it may happen that receivers i and j are assigned to wavelengths in such a way that both need a reconfiguration, whereas it would be possible, by exchanging their bin assignment, to avoid their reconfiguration. Thus, RS is a simple local search algorithm that swaps almost equally loaded receivers to avoid unnecessary reconfigurations.

The outcome of the 3-step algorithm is a new configuration with better load balancing properties than the previous one, and with low service disruption. Although the proposed heuristic may not lead to the optimal solution with minimum service disruption, it clearly eliminates unwanted reconfigurations that otherwise would cause larger blackout periods.

B. In-Transit-Traffic Measurement Scheme: The First-Fit Algorithm

Instead of measuring incoming traffic at each node, the in-transit measurement approach for traffic estimation measures in-band traffic. One way of implementing this scheme is to equip the master node with a receiver on each wavelength, so that measurements can be centralized in this node, thereby eliminating the exchange of signaling messages at the cost of increased hardware complexity. However, under high-load conditions, measurements may not estimate accurately the traffic matrix. Indeed, the master node is not able to observe all arrived traffic, due to contentions among nodes willing to access an overloaded channel.

When this scheme is adopted, measurements are exponentially weighted to smooth transient values that do not represent steady-state conditions. As a consequence, the master node divides the measurement window of duration T_m into K subwindows, each one weighted by w_k . If m_{jk} is the mean

value of bandwidth directed to receiver j in the k th subwindow, the resulting aggregate traffic bandwidth on this receiver is computed as

$$t_j = \frac{\sum_{k=1}^K m_{jk} w_k}{\sum_{k=1}^K w_k} \quad (4)$$

where $w_k = e^{k\tau}$, with τ being the slow receiver tuning latency, since larger tuning latencies imply longer transient periods.

If there is no congestion on the channels, the 3-step algorithm described in Section IV-A could be used since the traffic matrix can be correctly estimated. However, if congestion is present, the algorithm frequently exhibits convergence problems due to uncertainties in the estimation of the measured bandwidth. Moreover, running a fairly complex algorithm when the measured data are potentially affected by errors may not be the best solution.

As a consequence, we propose a simpler algorithm that aims at considering both LB and service disruption simultaneously. Instead of taking full reconfiguration decisions at once, we try to improve the network configuration by successive partial modifications, each modification being possibly triggered at the end of each measurement window. Although this approach may not lead to the optimal solutions, it reconfigures the network in a conservative way, so as to see how the new configuration accommodates traffic, later considering another partial reconfiguration that takes into account the effect of the previous modification.

1) *First-Fit Algorithm*: The key idea of the First-Fit algorithm, which is an extension of the algorithm presented in [17], is to schedule only partial reconfigurations that improve traffic balance among channels. In fact, each partial reconfiguration aims at retuning the least loaded receiver from the most loaded wavelength M to the least loaded one m . While it is straightforward to find out m , M cannot be easily determined if there is congestion in more than one channel. Indeed, it is difficult to distinguish a channel close to congestion from one strongly congested since both would look alike considering in-transit measurements (i.e., both measured loads are approximately equal to 1). Thus, there is a risk that a node is retuned from a channel not congested (but close to congestion) to the least loaded channel without any overall improvement. In fact, if the channel was not suffering congestion, the performance in the next measurement window is not improved. The First-Fit algorithm exploits this idea to keep track of the channels that are actually not congested. For this purpose, a list \mathcal{W} of potentially congested channels is used; while m is picked up among all channels, M is selected from the list \mathcal{W} . To evaluate the performance improvement of a previously scheduled reconfiguration, we look at the mean measured channel load defined as

$$\bar{\mathcal{L}} = \frac{\sum_{k=1}^W \mathcal{L}(k)}{W}.$$

Every time a reconfiguration has to be done, we save the current mean load as $\bar{\mathcal{L}}_{\text{old}}$ and the values of wavelengths M and m . After a measurement window, we evaluate the new mean channel load as $\bar{\mathcal{L}}_{\text{new}}$ from the in-transit measurements.

If $\bar{\mathcal{L}}_{\text{new}}$ is greater than $\bar{\mathcal{L}}_{\text{old}}$, then the previous reconfiguration helped; otherwise, an unneeded retuning was done and channels M and m must be removed from the list \mathcal{W} , since we know that they are not congested channels.

More formally, the First-Fit algorithm runs through the following steps.

- 1) Evaluate $\bar{\mathcal{L}}_{\text{new}}$.
- 2) If $\bar{\mathcal{L}}_{\text{new}} > \bar{\mathcal{L}}_{\text{old}}$, goto 4.
- 3) Update $\mathcal{W} = \mathcal{W} \setminus \{M, m\}$.
- 4) Find

$$\mathcal{L}(M) : M = \arg \max_{k \in \mathcal{W}} \mathcal{L}(k)$$

$$\mathcal{L}(m) : m = \arg \min_{1 \leq k \leq W} \mathcal{L}(k).$$

- 5) Select smallest t_j on wavelength M such that

$$\mathcal{L}(M) + \epsilon > \mathcal{L}(m) + t_j.$$

- 6) If t_j does exist, then
 retune node j to wavelength m and $\bar{\mathcal{L}}_{\text{old}} = \bar{\mathcal{L}}_{\text{new}}$;
 else, $\bar{\mathcal{L}}_{\text{old}} = 0$ and $\mathcal{W} = \{1, \dots, W\}$.

Initially, the list \mathcal{W} contains all channels; if, after step 3), the list becomes empty, then the list is replenished. Once M and m have been selected after step 4), we still need to find out the smallest t_j in M that does not load m beyond the current load on M ; otherwise, we are increasing traffic unbalance. When comparing the current maximum load $\mathcal{L}(M)$ with the possible load on m , a small value ϵ is used to account for inaccuracies in the estimation of channel loads. Finally, a retuning is scheduled at step 6) if there exists a t_j that satisfies the previous described condition; otherwise, the algorithm parameters are reset to the initial values.

V. RESULTS AND ANALYSIS

In this section, we present performance results obtained by simulation when considering a network with $W = 4$ wavelengths and a total of $N = 16$ nodes, where the distance between two adjacent nodes is about 27 km, i.e., 90 μs ; thus, the ring RTT is 1.45 ms. Slots last 1 μs , corresponding to a packet size of about 1250 bytes at 10 Gbit/s. The slow-tunable receiver has a tuning latency $\tau = 10$ ms. Nodes adopt a virtual output queuing (VOQ) architecture to avoid head of the line (HoL) blocking of packets waiting for access; thus, each node keeps a separate first-in first-out (FIFO) queue for each destination node, with a queue size of 32 000 fixed-size packets.

The duration of the measurement window is set to $T_m = 50$ ms. In the case of incoming-traffic measurements, the threshold to determine whether the new allocation is worth the reconfiguration cost is set to 5% of the previous allocation. For the in-transit-traffic case, each measurement window is divided into $K = 5$ intervals, exponentially smoothed as previously described.

We look at transient scenarios, i.e., algorithm behaviors when the traffic pattern changes, according to a predefined

scheduling, from an initial traffic pattern to a final one. We assume a linear transition from the initial traffic pattern to the final one, i.e., the transition occurs in a given number S of subsequent steps of uniform duration T_S . At a given step, the intermediate traffic matrix (representing the current traffic pattern) is obtained from a linear combination of the initial traffic matrix and the final traffic matrix, weighted by $S - i/S$ and i/S , respectively, where i is an integer ranging from 0 to $S - 1$, which increases at each step. In our simulations, we set $S = 10$ and $T_S = 100$ ms, i.e., 100 000 slot times; thus, the whole transition process lasts $S \times T_S = 1$ s, starting at simulation time 1 s and ending at simulation time 2 s. Although a 1 s traffic variation cannot be considered a long-term one, we use it only to describe the properties of our algorithms; moreover, it can represent a limit case of a fast-changing pattern that our algorithms are able to cope with. In practice, we envision that T_S should be much larger than 100 ms, since we are addressing long-term traffic changes; as a consequence, T_m was chosen to be smaller than T_S , to promptly evaluate changes in the traffic pattern. Values of T_m larger than T_S are not realistic, and would clearly induce worse performance due to the inability to accurately detect traffic-matrix modifications.

The initial traffic pattern is a uniform traffic pattern; since the same number of nodes is initially assigned to each wavelength, the whole capacity of the network is equally shared by all nodes. Thus, the element t_{ij} , $1 \leq i, j \leq N$, of the uniform traffic matrix is given by

$$t_{ij} = \lambda_{in} \frac{W}{N} \frac{1}{N-1}$$

where λ_{in} represents the normalized input load. The final traffic pattern is named “two server”; in this scenario, nodes are partitioned into two separated subsets: servers \mathcal{S} and clients \mathcal{C} . The two nodes belonging to \mathcal{S} transmit at a high rate, equal to the capacity of one wavelength per node, with equal probability to the other $N - 2$ nodes belonging to \mathcal{C} . The remaining network capacity is shared by client nodes, which transmit only to the servers with equal probability. In other words, t_{ij} , $1 \leq i, j \leq N$, is given by

$$t_{ij} = \lambda_{in} \begin{cases} 0, & \text{iff } i \in \mathcal{S} \wedge j \in \mathcal{S} \\ \frac{1}{N-2}, & \text{iff } i \in \mathcal{S} \wedge j \in \mathcal{C} \\ \frac{W-2}{N-2} \frac{1}{2}, & \text{iff } i \in \mathcal{C} \wedge j \in \mathcal{S} \\ 0, & \text{iff } i \in \mathcal{C} \wedge j \in \mathcal{C} \end{cases} .$$

In the sequel, we analyze algorithm performance when varying the traffic pattern according to the previously described scheduling, starting from the initial uniform traffic and finally leading to the two-server traffic pattern. In all subsequent figures, with the exception of Fig. 3, we plot the instantaneous throughput on the left and the normalized cumulative throughput, i.e., the total amount of data sent normalized to the input load, on the right. In the left figures, each traffic-pattern transition is highlighted at the top of the plot by a cross point when the transition occurs.




Fig. 3. Influence of tuning latency on the throughput at 100% load.

A. Incoming-Traffic Measurement

We first analyze the behavior when the incoming-traffic measurement scheme is considered. Several algorithms are considered and compared: fixed receiver (“Fixed RX”), theoretical tunable receiver with LB only [“Tun. Rx (LB)”], and tunable receiver with the full 3-step algorithm [“Tun. Rx (3-step)”]. Moreover, although the First-Fit algorithm is more naturally matched to the in-transit-traffic measurement case, we also report its curve as a reference [“Tun. Rx (First-Fit)”]. The fixed receiver refers to the case when no reconfiguration takes place, with the configuration being matched to the initial uniform traffic pattern. The theoretical case refers to an ideal case when the new network configuration is instantaneously optimally matched to the current real traffic pattern. This curve is not obtained by simulation; it simply refers to the maximum achievable throughput given the traffic matrix under consideration. Therefore, this case does not consider buffering effects; as a consequence, the curve may provide worse performance than other algorithms in simulation situations where buffering may temporarily increase throughput. Note that this is not equivalent to a simulation under an idealized scheme, i.e., where the network configuration is immediately adapted to traffic changes without waiting for the measurement algorithm to detect the new traffic pattern, and both RTT and tuning latencies are neglected. The three other cases refer to the full 3-step algorithm, to the same algorithm when using the LB step only, and to the First-Fit algorithm; in the case of LB only, bin i , $1 \leq i \leq W$, is directly associated with wavelength i .

In Fig. 4(a), the instantaneous network throughput (averaged over all wavelengths) is plotted when the network load is 100%, i.e., $\lambda_{in} = 1$. The theoretical curve shows that some of the intermediate traffic matrices, obtained as a linear combination of the uniform and of the two-server scenarios, cannot be completely scheduled, i.e., they are not admissible (the instantaneous throughput becomes less than 1). The case of fixed receivers highlights the performance degradation when traffic changes occur and the network configuration becomes increasingly unsatisfactory with respect to the current traffic pattern. When slow tunability is present at the receivers,




Fig. 4. Instantaneous and normalized cumulative throughput with incoming-traffic measurement scheme; $\lambda_{in} = 1$.

reconfigurations are scheduled to keep the configuration matched to traffic variations. Although these reconfigurations introduce blackout periods causing throughput falls, they put the network close to an optimal operational point with respect to the new traffic pattern. As shown on Fig. 4(a), reconfiguration algorithms produce fairly different results; in particular, the 3-step algorithm clearly outperforms the algorithm using LB only. The performance difference between the algorithms can be better appreciated in Fig. 4(b), where we report the normalized cumulative throughput. The 3-step algorithm performs best, thanks to the ability of scheduling reconfigurations that minimize the number of retuned receivers, whereas the fixed solution is clearly unacceptable. The First-Fit algorithm is able to reach the optimal allocation, but with increased delay, which implies a lower normalized cumulative throughput.

Coming back to Fig. 4(a), observe that only three main network reconfigurations take place at around 1.35, 1.75, and 2.05 s. This means that even if the algorithms detect all ten traffic modifications, only three times the reconfiguration process provides a significantly better channel equalization than the previous configuration, thus suggesting that the reconfiguration may be worthwhile. Reconfigurations are clearly detected

by the significant instantaneous throughput drop caused by the inability of transmitting toward receivers involved in the reconfiguration process, and are delayed, with respect to traffic-matrix changes, roughly by one measurement window. Observe that the second reconfiguration is delayed more than one measurement window with respect to the traffic-matrix change. The reason is that the measurement process, which starts synchronously but runs independently from the traffic-matrix reconfiguration process, is suspended when nodes are retuning, to avoid measurement errors. As such, after the first reconfiguration, the measurement process is shifted, with respect to traffic-matrix dynamics, and the second reconfiguration is delayed by one measurement window plus one tuning time with respect to traffic reconfiguration. As a result of the time shift between the measurement and the traffic reconfiguration processes, the third reconfiguration is slightly anticipated, since the shifted measurement process is able to detect the traffic change in the measurement window including the time instant in which the traffic change occurs. The First-Fit algorithm has longer reconfiguration times, and also less evident drops in the instantaneous throughput, due to the single receiver-reconfiguration constraint. It is worth noticing that sometimes the throughputs of both the 3-step and LB algorithms are higher than the theoretical case; this phenomenon is due to the buffering process at nodes. After a traffic-matrix change, additional traffic, with respect to the newly generated one, is offered temporarily, due to packets stored in buffers; thus, nodes may end up transmitting more than it would be possible with the current packet generation process alone (as it is for the theoretical-case curve). Finally, after the third reconfiguration, traffic becomes steady according to the two-server scenario and the network configuration becomes optimal for both algorithms.

In Fig. 5, the same scenario is simulated with an input load $\lambda_{in} = 0.9$. In this case, all traffic matrices are admissible, and the buffered data during reconfiguration time can be transmitted once the new configuration has been set up, using the excess bandwidth available. This phenomenon is clearly visible when the instantaneous throughput goes above the input load after each reconfiguration. In this scenario, differences among algorithms are highlighted not only during the blackout periods but also immediately afterwards. All slow reconfiguration algorithms allow nodes to transmit all generated packets, as shown in Fig. 5(b). However, the First-Fit and the 3-step algorithms help the network to recover faster.

B. In-Transit-Traffic Measurement

Let us focus on the in-transit measurement scheme, plotted in Fig. 6 for the same traffic pattern described above, with $\lambda_{in} = 1$. Recall that in this scheme, the measurement information is obtained by looking at wavelength channels only, thus being influenced by the nodes' ability to successfully access wavelength channels. We report results for the fixed receiver scheme as a "worst case" reference and for two reconfiguration algorithms: the 3-step algorithm, defined for the incoming-traffic measurement scheme, and the simpler First-Fit algorithm described in Section IV-B-1. As expected, throughput performance decreases as compared to the previous scheme. Due to





Fig. 5. Instantaneous and normalized cumulative throughput with incoming-traffic measurement scheme; $\lambda_{in} = 0.9$.

the large amount of buffered data, reconfiguration instants are not as evident as in the previous case, and can be identified by a sharp throughput increase or decrease.

Uncertainties on traffic estimates delay reconfiguration decisions and even force wrong ones, especially in the case of the 3-step algorithm. In fact, Fig. (6a) shows that this approach does not even finally set the network to its optimal operational point, while the First-Fit approach does it.

As mentioned in Section IV-B, full reconfiguration decisions, as those taken in the 3-step algorithm, when based on uncertain data, may be dangerous, since they can temporally improve network throughput, but, being not matched to real traffic, could sometimes produce more harm than benefits. In fact, as seen in Fig. 6(a), initially, both algorithms perform similarly and schedule the same reconfigurations (i.e., retune only one node at each reconfiguration). However, during the last three pattern changes, the 3-step algorithm cannot find a well-matched configuration since measurements on congested wavelengths become very inaccurate. Indeed, performance degrades to the point that it is not possible to find the optimal operational point once the traffic pattern is in steady state. However, it is worth noticing that the First-Fit may not always find a better





Fig. 6. Instantaneous and normalized cumulative throughput with in-transit measurement scheme; $\lambda_{in} = 1$.

operational point in other traffic scenarios, since it schedules partial reconfigurations only. Finally, the situation becomes really difficult only when the network is in overload. Indeed, Fig. 7 shows that when the network load $\lambda_{in} = 0.9$, both schemas perform well and no throughput loss is observed.

C. Effects of Tuning Latency

We wish to discuss the effect of tuning-latency values on network throughput. The receiver tuning latency is a key variable whose duration influences the hardware architecture, thus, network cost, as well as network performance. In Fig. 3, we present the results of several simulations with different receiver tuning latencies and report the mean network throughput. The simulation scenario is the same as above, and the instantaneous throughput is averaged over a time window starting at 1 s of simulation time and ending at 3 s of simulation time.

At a glance, the incoming-traffic measurement scheme performs better than the in-transit one. While both approaches based on incoming-traffic measurements show a stronger throughput decrease with increasing tuning latencies, both




Fig. 7. Instantaneous and normalized cumulative throughput with in-transit-traffic measurement scheme; $\lambda_{in} = 0.9$.

algorithms based on in-transit measurements are less sensible to the tuning latency. Indeed, the tuning latency has a dominant effect in the incoming-traffic measurement case given to the ability of computing a reconfiguration well matched to the traffic pattern. Conversely, in the in-transit measurement case, the dominant effect is the difficulty in obtaining a good network configuration, thus leading to larger throughput losses, in a manner less dependent from tuning latencies.

VI. CONCLUSION

The work in this paper was motivated by the idea that slow-tunable receivers, in addition to fast-tunable transmitters, are suitable for packet WDM metros, and are a fundamental asset that permit the following of long-term traffic variations while preserving low network complexity. Indeed, receivers do not have to be reallocated to WDM channels on a packet-by-packet basis, but only to react to user aggregate demands. As a consequence, since groomed demands change slowly with respect to packet dynamics, slow tunability is enough to pursue this goal. The increased cost due to slow-tunable devices can

be considered to be marginal with respect to traditional fixed devices.

Likewise, it is worth remarking that several switching and network-configuration functions do not need to scale with the packet duration, but should instead be matched to constraints related to the quality of the offered service, hence, can act on a less stringent time scale.

For a specific WDM packet-ring metro architecture, we discussed reconfiguration schemes that permit the reallocation of receivers to wavelengths according to long-term variations in the traffic pattern. The proposed reconfiguration schemes address the reconfiguration problem with two different measurement-based approaches. While the incoming-traffic measurement uses a complex but more accurate measurement mechanism (at the cost of requiring periodic signaling messages), the in-transit measurement approach implements a simpler scheme that, however, may become unreliable under highly loaded conditions. In this second scenario, we showed that it is wiser to run simpler algorithms that try to improve network performance by a trial-and-error procedure, rather than using an optimization approach that implies a complete redefinition of receiver allocation.

The proposed algorithms show good ability in tracking traffic variations by receiver reconfigurations, accurately balancing performance and reconfiguration costs. As expected, incoming-traffic measurement schemes perform better than in-transit ones, even if the difference is significant only in overloaded conditions. While results were presented only in a specific dynamic traffic setup, they well represent general trends.

Looking beyond the details of the proposed algorithms, the considered network architecture exhibits very interesting features that well exploits optical technologies.

REFERENCES

- [1] S. Yao, B. Mukherjee, and S. Dixit, "Advances in photonic packet switching: An overview," *IEEE Commun. Mag.*, vol. 38, no. 2, pp. 84–94, Feb. 2000.
- [2] J. Cai, A. Fumagalli, and I. Chlamtac, "The Multitoken Interarrival Time (MTIT) access protocol for supporting variable size packets over WDM ring network," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2094–2104, Oct. 2000.
- [3] Y. Cai, R. M. Fortenberry, and R. S. Tucker, "Demonstration of photonic packet-switched ring network with optically transparent nodes," *IEEE Photon. Technol. Lett.*, vol. 6, no. 9, pp. 1139–1141, 1994.
- [4] J. Fransson, M. Johansson, M. Roughan, L. Andrew, and M. A. Summerfield, "Design of a medium access control protocol for a WDMA/TDMA photonic ring network," in *Proc. IEEE Global Telecommunications Conf. (GLOBECOM)*, Sydney, Australia, Nov. 1998, vol. 1, pp. 307–312.
- [5] C. S. Jelger and J. M. H. Elmirghani, "Photonic packet WDM ring networks architecture and performance," *IEEE Commun. Mag.*, vol. 40, no. 11, pp. 110–115, Nov. 2002.
- [6] A. Smiljanić, M. Boroditsky, and N. J. Frigo, "High-capacity packet-switched optical ring network," *IEEE Commun. Lett.*, vol. 6, no. 3, pp. 111–113, Mar. 2002.
- [7] L. Dittman *et al.*, "The European IST project DAVID: A viable approach towards optical packet switching," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1026–1040, Sep. 2003.
- [8] M. Ibsen *et al.*, "8- and 16-channel all-fiber DFB laser WDM transmitters with integrated pump redundancy," *IEEE Photon. Technol. Lett.*, vol. 11, no. 9, pp. 1114–1116, Sep. 1999.
- [9] E. Y. Chan *et al.*, "A 12-channel multimode fiber-optic 1.0625-Gb/s fiber channel receiver based on COTS devices and MCM-L/COB/BGA packaging," *IEEE Photon. Technol. Lett.*, vol. 12, no. 11, pp. 1549–1551, Nov. 2000.

- [10] M. A. Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri, "MAC protocols and fairness control in WDM multi-rings with tunable transmitters and fixed receivers," *J. Lightw. Technol.*, vol. 14, no. 6, pp. 1230–1244, Jun. 1996.
- [11] K. V. Shrikhande, I. M. White, D.-R. Wonglumsom, S. M. Gemelos, M. S. Rogge, Y. Fukushima, M. Avenarius, and L. G. Kazovsky, "HORNET: A packet-over-WDM multiple access metropolitan area ring network," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2004–2016, Oct. 2000.
- [12] I. Baldine and G. Rouskas, "Traffic adaptive WDM networks: A study of reconfigurations issues," *J. Lightw. Technol.*, vol. 19, no. 4, pp. 433–454, Apr. 2001.
- [13] A. Carena, V. De Feo, J. Finochietto, R. Gaudino, F. Neri, C. Piglione, and P. Poggiolini, "RingO: An experimental WDM optical packet network for metro applications," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 8, pp. 1561–1571, Oct. 2004.
- [14] M. Pinedo, *Scheduling: Theory, Algorithms, and Systems*. Upper Saddle River, NJ: Prentice-Hall, 2002.
- [15] R. L. Graham, "Bounds on multiprocessing timing anomalies," *SIAM J. Appl. Math.*, vol. 17, no. 2, pp. 416–429, 1969.
- [16] R. E. Tarjan, *Data Structures and Network Algorithms*. Philadelphia, PA: SIAM, Nov. 1993.
- [17] I. Alfouzan and A. Jayasumana, "Dynamic reconfiguration of wavelength-routed WDM networks," in *Proc. IEEE Local Computer Networks (LCN)*, Tampa, FL, Nov. 2001, pp. 477–485.



Andrea Bianco (M'98) was born in Turin, Italy, in 1962. He received the Dr.Ing. degree in electronics engineering and the Ph.D. degree in telecommunications engineering from Politecnico di Torino, Torino, Italy, in 1986 and 1993, respectively.

He is an Associate Professor in the Electronics Department, Politecnico di Torino. He has coauthored over 100 papers published in international journals and presented in leading international conferences in the area of telecommunication networks. His current research interests are in the fields of protocols for all-

optical networks and switch architectures for high-speed networks. He was involved in several European (ACTS-SONATA, IST-DAVID) and Italian projects (EURO, IPPO, WONDER) on optical networks and switch architectures.

Dr. Bianco was Technical Program Co-Chair of the High-Performance Switching and Routing (HPSR) 2003 Workshop. He has participated in the technical program committees of several conferences, including the IEEE INFOCOM 2000, the IEEE GLOBECOM 2005, the IFIP Working Conference on Optical Network Design and Modeling (ONDM) 2002, 2003, and 2004, and Networking 2002 and 2004.



Jorge M. Finochietto (S'99) was born in Buenos Aires, Argentina, in 1978. He received the Diploma in electronics engineering from the Universidad Nacional de Mar del Plata, Mar del Plata, Argentina, in 2000, and the Ph.D. degree in telecommunications from Politecnico di Torino, Torino, Italy in 2005.

He worked from 2000 to early 2002 at Techtel Telecomunicaciones, a South American network operator, in the areas of routing performance and asynchronous transfer mode (ATM) networks. In 2002, he joined the Dipartimento di Elettronica of Politecnico

di Torino as a Ph.D. student. He has been involved in the European IST-DAVID project, which investigated the potential of optical packet switching in metropolitan and backbone networks. He has also participated in the Italian research projects ADONIS and EURO on optical networks and switching architectures. Currently, he takes part in the European Network of Excellence e-Photon/ONe on optical networks and in the Italian project WONDER, which aims at building an optical network demonstrator for the metropolitan area.



Gianluca Giarratana was born in Carrara, Italy, on January 24, 1979. He received the degree in telecommunications engineering from Politecnico di Torino, Torino, Italy, in 2004.

He specialized in access protocols for reconfigurable optical networks; his final project was "Resource allocation in packet WDM ring networks." He is currently working in Accenture, a global business and technology consulting company.



Fabio Neri (M'98) received the Dr.Ing. and Ph.D. degrees in electrical engineering from Politecnico di Torino, Torino, Italy, in 1981 and 1987, respectively.

He is a Full Professor in the Electronics Department, Politecnico di Torino. His teaching includes graduate-level courses on computer communication networks and on the performance evaluation of telecommunication systems. He leads a research group on optical networks at Politecnico di Torino. He has coauthored over 100 papers published in international journals and presented in leading inter-

national conferences. His research interests are in the fields of performance evaluation of communication networks, high-speed and all-optical networks, packet switching architectures, discrete event simulation, and queueing theory.

Dr. Neri was General Co-Chair of the 2001 IEEE Local and Metropolitan Area Networks (IEEE LANMAN) Workshop and General Chair of the 2002 IFIP Working Conference on Optical Network Design and Modeling (ONDM).



Chiara Piglione (S'05) was born in Racconigi, Italy, in 1977. She received the degree in telecommunication engineering (*summa cum laude*) in July 2002. Since January 2003, she has been pursuing the Ph.D. degree at the Dipartimento di Elettronica, Politecnico di Torino, Torino, Italy. From August to May 2004, and from July to September 2005, she was Visitor Student at the Department of Electrical Engineering, Arizona State University, Tempe, AZ.

Between October and December 2002, she had a research contract with the Italian National Inter-university Consortium for Telecommunications (CNIT), Firenze, Italy. Her research interests include the study of multicast traffic, input-queued switches, and all-optical networks.