Optimal Design of Logical Topologies in Wavelength-Routed Optical Networks with Multicast Traffic

M.Mellia[1], A.Nucci[1], A.Grosso[2], E.Leonardi[1], M.Ajmone Marsan[1]
[1] – Dipartimento di Elettronica - Politecnico di Torino - Italy
[2] – Dipartimento di Automatica ed Informatica - Politecnico di Torino - Italy

Abstract— **In this paper we discuss the optimal design of logical topologies in wavelength-routed WDM networks supporting unicast and multicast transfer of IP datagrams. We first explain the key aspects of the problem, emphasizing the fact that in IP networks the routing algorithms are an input to the optimization problem, not an optimization target. We then provide a mixed integer linear programming formulation of the optimization problem, which however leads to unacceptably high complexity for networks of non-trivially small size. We then propose both greedy and metaheuristic approaches for the sub-optimal design of logical topologies with acceptable complexity. Finally, we derive lower bounds that allow the assessment of the performance of the proposed algorithms. Some numerical results indicate that the proposed metaheuristics largely outperform the greedy approaches, and are able to obtain very good logical topologies.**

Keywords—**Logical Topology Design; WDM Networks; Multicast**

I. Introduction

In Wavelength-Routed (WR) optical networks, which employ Wavelength Division Multiplexing (WDM), high-capacity (electronic) routers are connected through semi-permanent optical pipes called "lightpaths" that may extend over several physical links. Lightpaths, thus, can be seen as chains of physical channels through which packets are moved from a router to another toward their destinations. At intermediate nodes, incoming channels belonging to in-transit lightpaths are transparently coupled to outgoing channels through a passive wavelength router that does not process in-transit information. Instead, incoming channels belonging to terminating lightpaths are converted to the electronic domain, so that packets can be extracted and processed, and possibly retransmitted on outgoing lightpaths after electronic IP routing. In a WR network, a "logical topology", whose vertices are the IP routers and whose edges are the lightpaths, is overlayed to the "physical topology", made of optical fibers and optical cross-connects. In principle, the "logical topology" configuration is independent from the physical topology; however, a number of constraints exists: i) the establishment of each lightpath requires the reservation of physical resources (i.e., a WDM channel on the optical fibers along the path); ii) the number of lightpaths departing and terminating at each node is limited by the number of transmitters/receivers and by the processing capability of electronics; iii) the maximum allowable length of lightpaths may be limited by transmission degradation.

In order to best exploit the capacity of a WDM infrastructure, a crucial task thus is the identification of the best feasible logical topology for the transport of a given traffic pattern. In recent years, this problem was extensively studied in the case of purely unicast traffic. It was shown that the identification of the optimal logical topology is computationally intractable for large size networks [1], [2], and several heuristic approaches were proposed for the identification of suboptimal solutions (see [3] for a detailed overview).

A significant fraction of the Internet traffic is expected to come from native multicast applications, a few years from now. As a consequence, the development of logical topology design procedures that can account for multicast traffic flows is necessary.

Since multicast traffic flows are characterized by many destinations, replication (branching) of multicast packets somewhere in the network is necessary. Two different approaches are possible in WR networks: replication in the optical domain, or replication in the electronic domain. The first approach was considered in [4], where the concept of lightpath was extended to a "Light-Tree", which is a transparent one-to-many pipe in which injected packets are passively replicated and delivered to the light-tree end-points. This approach relies on the availability of multicast capabilities in the physical layer, which seem difficult to obtain. In addition, the performance of light-tree based solutions, with multiple packet flows sharing one light-tree, may be severely limited by undesirable replications of packets.

To the best of our knowledge, in this paper we investigate for the first time solutions in which packet replication is performed in the electronic domain. These solutions impact neither the optical architecture of nodes, nor the IP protocol, and in addition appear more flexible and efficient than optical domain replication approaches.

Moreover, we must consider the fact that in the IP over WDM context, the routing strategy is assigned 'a priori'; it cannot be considered a variable of the problem. This key observation led us to the optimization of the topology design, assuming that the IP routing algorithm is fixed.

II. Problem Formulation

The problem of the optimal Logical Topology Design (LTD) for the transport of unicast and multicast traffic can be stated as follows:

Given: i) an existing physical topology, comprising nodes equipped with a limited integer number of tunable transmitters and receivers, connected by optical fibers that support a limited number of wavelengths; ii) a description of the average traffic exchanged by sources and sets of destinations; iii) a multi-hop routing strategy, defined both for unicast and multicast flows; find the logical topology that minimizes a target function or cost.

While different functions can be chosen as optimization target, we selected the maximum congestion level in the network, defined as the maximum traffic flow on lightpaths, since this metrics normally drives important network performance indices, such as loss probability and delay (see [3] for a discussion). Some of the LTD algorithms proposed in this paper can be easily adapted to a different target function.

We do not consider in the optimization the constraint on the maximum number of wavelengths that each fiber can carry, since new transmission equipments allow the use of very high numbers of wavelengths (up to 128), and in addition the use of wavelength converters can be instrumental for the relaxation of this constraint.

In the next section we shall refer to the LTD optimization for purely unicast traffic using the acronym ULTD and to the LTD optimization for unicast and multicast traffic using the acronym MLTD, while we shall use the acronym LTD to refer to the general case.

A. MILP Formulation

Several different Mixed Integer Linear Programming (MILP) formulations of the ULTD problem appeared in the recent literature [1], [5], [6]. In this paper we extend the LTD formulation to the case in which multicast traffic is present. Note that the MLTD formulation is not a straightforward extension of the ULTD formulation, because the branching capabilities change the flow conservation equations.

This work was supported by the Italian Ministry for University and Scientific Research under the IPPO (IP Over Optics) project.

Let T be the set of traffic requests, whose elements t_k indicate the average traffic associated with connection k from source s_k to the destinations in set D_k. If $|D_k| > 1$ then we have a multicast connection. Let d_k^l be the l-th destination in set D_k. We use the binary variables $b_{ij} \in \{0, 1\}$ to indicate whether a lightpath originating from node i and terminating in node j (i.e., lightpath $i \rightarrow j$) belongs to the logical topology ($b_{ij} = 1$ if the lightpath is included in the logical topology, $b_{ij} = 0$ otherwise). The real variables r_{ij}^{kl} indicate the fraction of traffic t_k, flowing on lightpath $i \rightarrow j$, that reaches destination d_k^l. As already stated, multicast packets flowing on lightpath $i \rightarrow j$ can reach several destinations because of branching, thus $\sum_l r_{ij}^{kl}$ in general is larger than r_{ij}^k, the fraction of traffic t_k flowing on lightpath $i \rightarrow j$. We denote with f_{ij} the total traffic flowing on lightpath $i \rightarrow j$. Finally, let $f_{max} = \max_{ij}(f_{ij})$ be the maximum amount of traffic flowing on any lightpath in the logical topology. Given N, the number of nodes in the network, and given δ_O^i and δ_I^i, that represent, respectively, the numbers of transmitters and receivers available at node i (i.e., δ_I^i and δ_O^i are the maximum in/out degrees of node i in the physical topology), the LTD problem can be formulated as follows:

$$\min f_{max} \tag{1}$$

under the following constraints:

- Flow conservation at each node

$$\sum_j r_{ij}^{kl} - \sum_j r_{ji}^{kl} = \begin{cases} 1 & \text{if } i = s_k \\ -1 & \text{if } i = d_k^l \\ 0 & \text{otherwise} \end{cases} \qquad \forall k, l \tag{2}$$

- Total flow on lightpaths for each connection

$$r_{ij}^k \geq r_{ij}^{kl} \qquad \forall i, j, k, l \tag{3}$$

$$r_{ij}^k \leq b_{ij} \qquad \forall i, j, k \tag{4}$$

- Flow on links

$$f_{ij} = \sum_k r_{ij}^k t^k \tag{5}$$

$$f_{ij} \leq f_{max} \qquad \forall i, j \tag{6}$$

- Degree constraints

$$\sum_j b_{ij} \leq \delta_O^i \qquad \forall i; \tag{7}$$

$$\sum_i b_{ij} \leq \delta_I^j \qquad \forall j \tag{8}$$

- Variables range constraints

$$f_{ij} \geq 0 \qquad \forall i, j \tag{9}$$

$$0 \leq r_{ij}^k \leq 1 \qquad \forall i, j, k \tag{10}$$

$$0 \leq r_{ij}^{kl} \leq 1 \qquad \forall i, j, k, l \tag{11}$$

$$f_{max} \geq 0 \tag{12}$$

$$b_{ij} \in \{0, 1\} \qquad \forall i, j \tag{13}$$

The above formulation of the topology design problem rests on three assumptions: i) the traffic between any two nodes can be arbitrarily partitioned and routed on multiple paths; ii) no more than one lightpath is allowed between any two nodes; iii) multicast flow branching can be performed in every electronic node .

The first assumption can be released by imposing that r_{ij}^{kl} and r_{ij}^k are binary variables (i.e., $r_{ij}^{kl} \in \{0, 1\}$ and $r_{ij}^k \in \{0, 1\}$).

Modifying the problem formulation to allow parallel lightpaths requires more significant changes, since the space of the decision variables must be expanded, replacing the b_{ij} with new variables b_{ijp}, r_{ij}^{kl} with r_{ijk}^{klp}, r_{ij}^k with r_{ij}^{kp}, and f_{ij} with f_{ij}^p, being $p = 0, 1, \ldots, \min(\delta_I^j, \delta_O^i)$. The structure of equations (1), ..., (13), instead, remains the same.

The MLTD problem is NP-hard, since it falls in the class of general MILP problems, and thus is numerically intractable, even for networks with a moderate number of nodes. Moreover, the MLTD problem represents a generalization of the well-known NP-hard ULTD problem, in the sense that it includes ULTD as particular instance.

While in the case of ULTD it is possible to provide a formulation in which only aggregate flows appear [5], thus reducing complexity, disaggregate flows r_{ij}^{kl} must be considered in MLTD, in order to apply flow conservation equations at each node (2), leading to a significant increase of complexity in the formulation.

In addition, the MILP formulation naturally leads to a combined topology *and* routing optimization. Therefore, the routing strategy is a result of the optimization procedure, together with the logical topology configuration. Thus, MILP and heuristics based on MILP, such as those relying on continuous relaxation and rounding, do not permit the advance specification of the routing strategy to be considered in the optimization. In the IP over WDM context, however, the routing strategy is assigned 'a priori': it is not a variable of the problem, but a constraint. These facts motivated the development of heuristic approaches that permit the specification of the routing strategy as input of the MLTD problem.

III. Routing on the Logical Topology

Unicast IP routing algorithms have been standardized and are widely accepted [7], [8]. They can be described as minimum cost path and shortest path routing algorithms in which only static administrative costs can be considered. This observation suggested us to carry out the LTD optimization adopting a shortest path routing for unicast flows, i.e., a routing strategy that minimizes the number of hops (lightpaths) traversed by each unicast flow.

Many algorithms were instead proposed to route multicast IP traffic [9], [10], [11], some of which have been standardized, but are not yet widely adopted. The existing IP multicast routing algorithms can be roughly grouped into two classes: the first is based on distributed algorithms, where shortest paths are used to build the distribution tree. The second class is instead based on centralized algorithms, where a special node is elected, which has global visibility of the multicast tree, and runs a centralized algorithm to compute the best multicast tree.

In the context of WDM networks, considering routing of multicast flows while designing the logical topology leads to significant advantages only when the aggregate multicast traffic is a considerable percentage of the total traffic. The traffic scenarios foreseen for IP networks of the coming years envisage a significant amount of multicast traffic resulting from the distribution of music, video, large file caches, multimedia events. The number of users of such services can be large, so that centralized multicast routing algorithms can be advantageous. Moreover, since the WDM technology will be mainly deployed in the high-capacity network backbone, traffic scenarios containing single-rooted multicast trees can be expected to be quasi-static.

From a theoretical point of view, the definition of the optimal multicast routing (i.e., the routes that lead to the minimization of the capacity resources used for the transport of multicast flows), can be easily formalized as a minimal Steiner Tree problem; i.e., the problem of finding the minimal graph that connects the multicast traffic source to a specified set of destinations.However, the minimal Steiner Tree problem is NP-Hard. Thus, multicast routing algorithms proposed in the literature can be catalogued as heuristic procedures that obtain a minimal Steiner Tree at a limited computational cost.

These reasons suggested us to utilize as input of our MLTD optimization procedures a centralized multicast algorithm, but any multicast algorithm can be adopted in the optimization procedure. We selected the "*Selective Closest Terminal First*" ($SCTF$) algorithm proposed in [12], that, given the multicast traffic source, the group definition, and the network topology, finds a Steiner Tree with a good trade-off between performance and computational complexity.

IV. Logical Topology Design Heuristics

In this section we propose heuristic approaches for the MLTD problem; we first introduce two greedy algorithms, then we propose algorithms that exploit metaheuristic optimization algorithms developed in the operational research context.

A. Greedy Heuristics

Greedy heuristic algorithms for MLTD can be applied either to generate an initial solution for the metaheuristic-based algorithms described in the following section, or to quickly produce reasonable logical topology configurations when time constraints do not allow the application of computationally intensive algorithms.

A.1 Source Copy Multicast Algorithm ($SCOM$)

The first heuristics is named *Source COpy Multicast* ($SCOM$) algorithm; it relies on the assumption that multicast IP packets are routed by sending multiple copies from the multicast traffic source node. In this case, each multicast flow degenerates into a set of unicast flows connecting the source to each multicast destination. Any unicast topology design algorithm can be used to obtain the desired logical topology; we selected $MLTDA$, proposed in [3], that is aimed at the maximization of single-hop flows in the network [1].

A.2 Route & Remove ($R\&R$)

The second heuristic algorithm initially considers a fully-connected logical topology; an iterative removal from the logical topology of the least-loaded lightpaths is executed, until the degree constraints are satisfied. Note that the lightpath load is evaluated by routing both unicast and multicast flows according to a fixed routing strategy, thus fulfilling the routing constraints.

To describe this algorithm, a boolean variable is associated with each lightpath, which can assume the values Removable or Unremovable (a lightpath is removable if the topology obtained by removing it is still connected, Unremovable otherwise).

$R\&R$ *algorithm*

Step 0: Select the fully-connected logical topology and mark all lightpaths as Removable.

Step 1: Solve both the unicast and the multicast routing problems on the current topology and compute traffic flows on lightpaths.

Step 2: Find a set of lightpaths that can be removed from the graph by solving a minimal Weight Matching[2] (mWM). Only the edges that are marked as Removable can be chosen in the matching.

Step 3: Remove lightpaths in the logical topology, only if the resulting logical topology remains connected. Otherwise, mark the lightpath as Unremovable.

Step 4: If all the in/out-degree constraints are satisfied then STOP, otherwise GOTO Step 1.

[1] The maximization of traffic flows that are routed in a single-hop fashion is equivalent to the well-known δ-Maximum Weight Matching (δ-MWM) problem on a bipartite graph when the in/out-degrees of all nodes are equal (i.e., $\delta_I^i = \delta_O^i = \delta\ \forall i$). The optimal solution of this problem can be found with $O(\delta N^3)$ operations [13]. Even if the in/out-degrees of nodes are different, a modification of the δ-MWM problem can be used to solve the ULTD problem.

[2] The minimal Weight Matching provides a matching with maximal size and minimal weight

The reported greedy heuristics entail an increasing degree of computational complexity. The first one is very simple; it transforms the MLTD problem in a ULTD problem, decomposing each multicast connection into several unicast connections. However, topology optimization is obtained ignoring the possibility of performing multicast replication inside the network (multicast replication is allowed only at the source interface). The second algorithm entails a significant larger degree of complexity, since the topology optimization considers the real multicast routing.

B. Metaheuristics - Tabu Search (TS)

The heuristics we propose next relies on the application of the Tabu Search (TS) methodology [14]. TS is based on a partial exploration of the space of admissible solutions. The exploration starts from an initial solution that is generally obtained with a greedy algorithm.

For each admissible solution, a class of neighbor solutions is defined. A neighbor solution is defined as a solution that can be obtained from the current solution by applying a *perturbation*. The set of all the admissible perturbations uniquely defines the *neighborhood* of the current solution.

At each iteration of the TS algorithm, all solutions in the neighborhood are evaluated, and the best is selected as the new current solution. A special rule, the *tabu list*, is introduced in order to prevent the algorithm to deterministically cycle among already visited solutions. The tabu list consists in a fixed-size list recording the last perturbations that were accepted. Until a perturbation is stored in the tabu list, it cannot be used to generate a new solution.

After a given number of iterations, the algorithm returns the best visited solution. Note that TS algorithms can be seen as an evolution of the classical local optimal solution search algorithms called Steepest Descent [14], however, thanks to the tabu list mechanism, TS can accept worse solutions than the current one, and thus it is not subject to local minima entrapments.

C. Metaheuristics - Simulated Annealing (SA)

Like Tabu Search, also Simulated Annealing (SA) [14] is based on a partial exploration of the space of admissible solutions, finalized to the discovery of good solutions. At each iteration of the algorithm, however, only one solution in the neighborhood is visited and evaluated. If the new solution performs better than the current one, it is accepted as the new current solution, otherwise it is accepted with probability p, and discarded with probability $1 - p$.

The probability of accepting a worse solution than the current one generally depends on the iteration count. Usually, the value of p is decreased several times during the exploration process. As a consequence, in the initial phases of SA the solution space exploration is dominated by randomness; while by decreasing the value of p, i.e., by decreasing the *algorithm temperature*, an increasing degree of determinism is introduced in the exploration. When p becomes negligible, the algorithm tends to behave as the First Improvement [14] algorithm.

D. Settings

Three fundamental aspects that must be defined in the metaheuristic algorithms just described concern: i) the choice of an initial solution; ii) the definition of the perturbation that generates the neighborhood; iii) the evaluation of the visited solutions.

As initial solution we select the result of the $R\&R$ heuristic.

The perturbation is defined performing a branch exchange algorithm; select two lightpaths, $n_1 \to n_2$ and $n_3 \to n_4$, and "exchange" their destinations, obtaining lightpaths $n_1 \to n_4$ and $n_3 \to n_2$.

This perturbation guarantees that degree constraints are not violated, thus generating a valid move. However, it does not guarantee that the new topology is connected, requiring an explicit test for connectivity.

Finally, the evaluation of new solutions is performed by running the routing algorithms for both unicast and multicast traffic, and computing the maximum flow on lightpaths.

V. Complexity

Let F be the number of multicast flows that must be transported in the network. Let $\Delta = \max_i(\delta_O^i, \delta_I^i)$ denote the maximum in/out degree of the logical topology.

The evaluation of the maximum congestion level of lightpaths requires the information about the unicast and multicast traffic routing. The computation of the unicast routing along shortest paths consumes $O(N^2 \log(N\Delta))$ operations [15]. Once shortest paths are known, the computation of each Steiner Tree according to algorithm $SCTF$ requires $O(N^3)$ operations in the worst case, since, at each iteration, the algorithm requires at most $O(N^2)$ operations to choose the closest source/destination pair, choice that must be performed for all the nodes in the destination set, that is $O(N)$. In conclusion, the computation of the unicast and multicast routing requires $O(N^2(\log(N\Delta) + FN))$ operations.

We can now evaluate the complexity of the greedy heuristics. $SCOM$ has complexity $O(FN) + O(\Delta N^3)$, since $O(FN)$ operations are required to transform the multicast requests into unicast requests, and $O(\Delta N^3)$ operations are needed to run the $MLTDA$ algorithm. $R\&R$ has complexity $O(N^3(\log(N\Delta) + FN) + N^4)$, since at most $O(N)$ iterations are needed to complete the algorithm, each time routing the traffic and performing a mWM algorithm.

Let us now evaluate the complexity of the algorithms based on metaheuristics. At each iteration, the evaluation of all solutions in the neighborhood is necessary; this requires $O(D^2N^4(\log(N\Delta) + FN))$ operations, since $O(D^2N^2)$ neighbors are evaluated (the actual number is upper bounded by $0.5N[D(N-D)]^2$), and the evaluation of each solution requires the solution of the routing problem. If the number of iterations is I, the resulting complexity is $O(ID^2N^4(\log(N\Delta)+FN))$. The complexity of the SA algorithm is similar. Let I denote the number iterations; the resulting complexity is $O(IN^2(\log(N\Delta) + FN))$.

VI. Lower Bounds on Congestion

We compute three bounds: the first bound does not consider the routing constraints imposed by shortest paths, i.e., it allows the splitting of source-destination traffic relations on many links; the second and third bounds instead explicitly consider the shortest path constraints, and are based on the well-known NP-hard problem of Bin-Packing. We derive a closed-form bound for that problem as our second bound, while we propose a Branch and Bound approach as our third bound.

A. Minimum Flow Tree with Multicast

Our first bound is an extension of the *Minimum Flow Tree* (MFT) bound [2], [5], that takes into account multicast traffic. We refer to this bound as *Minimum Flow Tree with Multicast* ($MFTM$). Given a network with N nodes, and maximum in/out degree Δ, the best admissible topology we can build considering a source node s is a tree of degree Δ. There will be then Δ destinations at one hop distance from s, Δ^2 at two hops, Δ^3 at three hops, etc.

For unicast traffic, in a network with E links, we have that

$$f_{\max} \geq \frac{1}{E}\sum_{ij} f_{ij} = \frac{1}{E}\sum_k t_k \sum_{ij} r_{ij}^k = \frac{\bar{H}}{E} \qquad (14)$$

Where $\bar{H}$ is the average number of hops between a source and a destination, computed using traffic loads as weights. Note that (14) computes the minimum aggregate flow that is injected in the network, and then splits it over all the links that are present in the network. In [5], the authors prove that considering such an idealized topology, and dispacing the destination so that the higher is the traffic they receive from s, the closer they are to s, minimizes $\bar{H}$. For the expression of $\bar{H}_{\min}$, the reader is invited to refer to [5].

Consider now one of the F multicast groups, say the $k-th$, whose traffic load is t_k, and whose number of destinations is $M_k = |D_k| > 1$; the minimum traffic that the network has to transport to deliver this multicast traffic flow is given by $T_k = M_k t_k$, since at least M_k links are involved in the distribution tree. Thus, the total minimum traffic that flows in the network due to multicast is $\sum_k T_k = \sum_k M_k t_k$. Then, reworking equation (14) to include also the multicast contribution from (VI-A), we finally have

$$f_{\max} \geq \frac{1}{E}\left[\bar{H} + \sum_k T_k\right] \qquad (15)$$

B. Integral Flow Bounds

Consider source node s. Be $\mathcal{T}(s)$ the set of all traffic relations t_k which have s as source. $|\mathcal{T}(s)| = N - 1 + G(s)$, where $G(s)$ is the number of multicast groups having s as root node. Each element in $\mathcal{T}(s)$ has an associated 'cost' t_k. Being δ_O^s the number of lightpaths departing from s, we have to find the disposition of $|\mathcal{T}(s)|$ objects in δ_O^s bins, such that the maximum cost (i.e., the maximum congestion level) among the bins is minimized. This is the well known *Bin-Packing* problem [16]. Let $C_{\min}(s)$ be the maximum minimum cost. Then:

Proposition 1: Let $C_{\min}(s)$ the minimum congestion level reached when disposing $|\mathcal{T}(s)|$ objects, each presenting a cost $t_k \geq 0$, in δ_O^s bins. Then:

$$f_{\max} \geq \max_s C_{\min}(s) \qquad (16)$$

Proof: Let s^* be the source node that originates the maximum minimum congestion level. In the evaluation of $C_{\min}(s^*)$ we consider only flows originating from s, neglecting the fact that other flows can be routed through one of the $\delta_O^{s^*}$ links. Thus, there can be other objects that have to be packed in the bins. Being $t_k > 0$, and being the cost function additive, then $C_{\min}(s^*)$ is an increasing monotonic cost function with the number of objects. ■

A similar reasoning allows us to consider destination d, and the set $\mathcal{T}'(d)$ of all traffic flows entering node d. Then $|\mathcal{T}'(d)| = N - 1 + G'(d)$, where $G'(d)$ represents the number of multicast groups sending traffic to d. Being δ_I^d the number of lightpaths ending in d, we have to find the disposition of $|\mathcal{T}'(d)|$ objects in δ_I^d bins, such that the maximum congestion level among the bins is minimized. Let $C'_{\min}(d)$ be the maximum minimum cost. Then

$$f_{\max} \geq C = \max\left(\max_s C_{\min}(s), \max_d C'_{\min}(d)\right) \qquad (17)$$

Unfortunately, the Bin-Packing problem is known to be NP-hard [16]. Thus, in order to derive bounds that can be computed with limited computational cost, we compute two simple lower bounds to its solution.

B.1 Fluidic Bin-Packing (FBP) Bound

By equally spitting all the objects among the bins, we obtain a simple lower bound on $f_{\max}$. At the same time, $f_{\max}$ is lower bounded by the largest object that must be accommodated in the bins. Thus:

$$C_{min}(s) \geq \max\left(\frac{1}{\delta_O^s}\sum_{t_k \in \mathcal{T}(s)} t_k, \max(t_k \in \mathcal{T}(s))\right) \qquad (18)$$

TABLE I

OPTIMIZATION RESULTS FOR THE DIFFERENT TRAFFIC SCENARIOS (A).

	μ %	BR	AR	WR	SCOM	R&R	TS	SA	MFTM	FPB	BBBP
Scenario (A)	5.21	248.60	326.01	476.19	271.22	235.35	165.84	173.10	100	73.50	84.33
	7.19	253.71	328.43	503.08	266.88	289.19	161.80	170.57	100	77.62	86.78
	8.63	255.80	331.53	509.66	238.42	265.28	165.34	171.34	100	74.61	78.47
	9.82	252.75	329.51	462.07	240.63	276.18	160.82	178.80	100	77.00	80.97
	18.28	246.65	330.13	587.37	242.32	252.63	162.22	171.85	100	76.99	87.88
Scenario (B)	2.12	132.97	210.55	374.01	118.50	115.54	100	100	23.50	100	100
	2.63	142.58	211.58	340.51	136.30	111.74	100	100	26.74	100	100
	4.25	153.87	219.89	400.22	112.83	115.18	100	100	30.54	100	100
	22.49	126.02	187.10	311.31	121.42	121.04	100	100	27.49	100	100
	31.86	129.23	191.19	300.73	113.00	113.54	100	100	29.29	100	100
Scenario (C)	22.37	192.67	290.82	518.38	248.28	230.00	122.10	128.32	55.08	87.61	100
	29.38	210.70	304.50	496.29	248.28	290.48	131.42	125.36	64.66	84.69	100
	32.66	173.68	252.51	415.12	233.44	235.05	109.34	102.13	47.11	100	100
	36.39	202.18	289.50	480.22	239.58	265.15	129.87	117.88	59.72	76.96	100
	53.43	198.79	295.36	519.33	239.85	291.94	132.84	121.02	64.37	81.39	100

The FBP bound is very simple; however, it allows splitting of flows through many paths, which is unrealistic, as we noted a number of times.

B.2 Branch and Bound Bin-Packing ($BBBP$) Bound

To obtain a more accurate bound, a Branch and Bound algorithm can be applied. This approach however requires more computational power than other bounds, and therefore may not be suitable to asses the quality of the solutions generated by the TS or SA algorithms.

VII. EXPERIMENTAL ANALYSIS

We report results for networks with either $N = 24$ or $N = 32$ nodes, and in/out degree 4 for logical topologies, i.e., $\delta_I^i = \delta_O^i = 4 \; \forall i$. We consider three traffic scenarios comprising $N_g = 6$ multicast connections. For each multicast connection, both the source node and the destination nodes are picked at random among all nodes in the network. The average number of destinations is fixed to $2/3$ of the total number of nodes. The three traffic scenarios can be described as follows:

(A) – Uniform, low variance – Each unicast and multicast traffic flow offers an average bandwidth which is randomly extracted from an exponential distribution with mean $\mu = 1$. This traffic scenario exhibits a significant degree of uniformity and regularity, since no hot-spots are present.

(B) – Uniform, high variance – Each unicast and multicast traffic flow offers an average bandwidth which is randomly extracted from a hyperexponential distribution with mean $\mu = 1$ and variance $\sigma^2 = 10$. This scenario exhibits a significantly smaller degree of uniformity and regularity than the previous one, since the bandwidth variance has been significantly increased.

(C) – Low-High Traffic – One fifth of the nodes are 'High-Traffic' (HT),while the remaining nodes are 'Low-Traffic' (LT). Traffic relations are subdivided into three classes: HT to HT, that offer a large average load, (exponentially distributed with $\mu = 10$); LT to HT and HT to LT (exponentially distributed with $\mu = 5$); finally, LT to LT traffic relations, (exponentially distributed with $\mu = 1$). The average bandwidth required by multicast connection is exponentially distributed with $\mu = 10$.

Although we did not perform a fine tuning of the metaheuristic settings, a small set of simulation was run in order to decide which parameter value to use. In the TS algorithm the tabu size was set to 12. The number of iterations was fixed to 60, giving a total number of evaluated topologies that is approximately 10^6.

In the SA algorithm, the probability of accepting a worse solutions is initially set to $p = 5 \cdot 10^{-3}$. Every 1000 iterations, p is decreased tenfold. The number of iterations was set to 10^6, so that the numbers of topologies visited by SA and TS are similar.

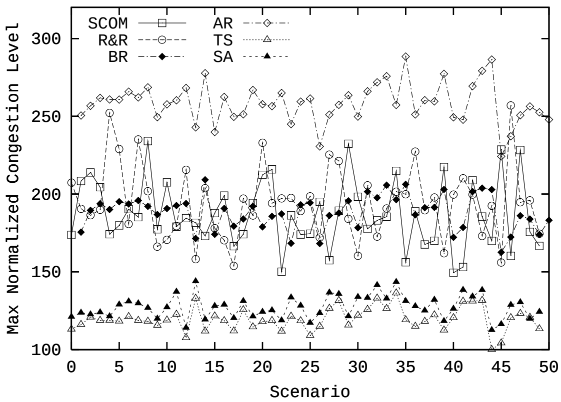

Fig. 1. Maximum congestion level normalized to the best solution, for 50 different traffic configurations

A. Uniform Traffic Scenarios

In this case we report in Fig. 1 the results obtained with the four proposed algorithms for 50 different traffic configurations loading a network with $N = 24$ nodes. Markers correspond to the maximum congestion level, normalized such that the minimum maximum congestion level has value 100, in the logical topologies generated by the proposed algorithms ($SCOM$, $R\&R$, TS, and SA). For each traffic configuration we generated at random 1000 topologies; markers labeled BR (Best Random) refer to the the lowest experienced maximum congestion level among all the random topologies, while markers labeled AR (Average Random) refer to the average congestion level.

Let μ denote the percentage of multicast traffic in the network. From Fig. 1 we see that both metaheuristic approaches provide much better results than greedy algorithms, with TS taking the edge over SA by few percentage points. It must be also noted that both $SCOM$ and $R\&R$ provide solutions similar to BR. This is due to the fact that both $SCOM$ and $R\&R$ rely only on the knowledge of the traffic T, but in this case, having all the traffic relations the same statistical characteristics, the information carried by T is not very detailed. However, also note that the solution obtained with greedy heuristics is better than AR. Being the fraction of multicast traffic not very high ($\mu \simeq 5\%$), $SCOM$ and $R\&R$ give performance in the same range.

To better understand the evolution of the TS and SA metaheuristics, Figs. 2 and 3 report the congestion level of the topologies analyzed by the TS and SA algorithms, respectively, for one of the 50 traffic configurations. The congestion values are normalized so that the best solution found (by TS) has a maximum congestion level

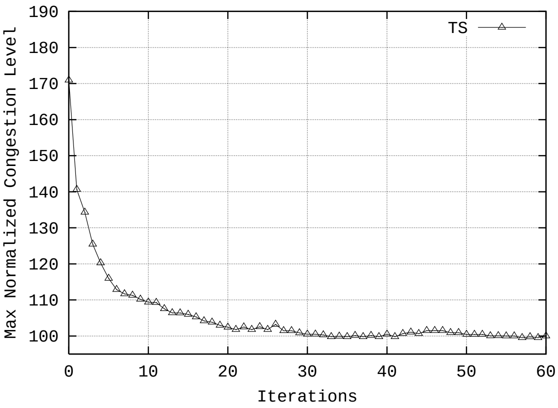

Fig. 2. Evolution of the TS algorithm.

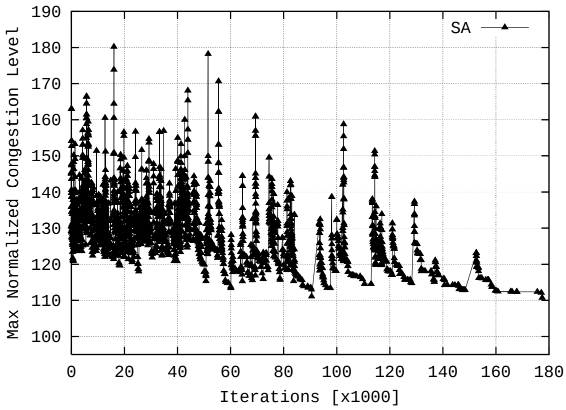

Fig. 3. Evolution of the SA algorithm.

equal to 100. Note that Fig. 2 reports for each iteration of TS the maximum congestion level of the best solution found in the neighborhood, while Fig. 3 reports the maximum congestion level of all the topologies selected by the SA algorithm, thus one for each iterations. The difference between the two approaches is clear from their evolution: TS performs a 'steepest descent' evolution until iteration 20; then the tabu list sometimes forces the algorithm to select a worse solutions, that eventually leads to a better final topology. Instead, SA exhibits a more random evolution, especially when p is high: the number of worse solutions selected by the algorithm is much higher than in TS, and the improvement toward a better solution is not as deterministic.

B. Different Traffic Scenarios

To further compare the different algorithms, we considered a network with $N = 32$ nodes, and generated 5 instances of traffic configuration for each traffic scenario, with different μ. We computed the maximum congestion level in the logical topology generated by the different algorithms, and also evaluated the three lower bounds derived in Section VI. Results for congestion levels are reported in Table I with a normalization against the strictest lower bound value, which is assigned a value 100.

From the tables we can see that when the traffic is uniform with low variance, as in scenario (A), $MFTM$ is the most stringent bound, while if the traffic variance is high, $BBBP$ and FBP provide better bounds, since they do not allow traffic splitting. Results for traffic scenario (A) confirm the relative performance between algorithms that was already observed in section VII-A, with TS providing the best performance, which is 60-65% higher than the tightest lower bound value. The greedy heuristics perform similarly to BR, but the solution they provide is not better than 2.5 times the lower bound.

When the traffic has a much higher variance, as in traffic scenario (B), the $SCOM$ and $R\&R$ heuristics are finally able to obtain good performance, only 10-20% higher than the tightest lower bound. In this case however both TS and SA are able to find the best logical topology, as seen by the fact that the maximum congestion level reached at the end of the optimization procedure is equal to the lower bound.

Finally, when the traffic differences are concentrated on few nodes, like in scenario (C), a metaheuristic approach is again necessary to obtain a good solution, since both $SCOM$ and $R\&R$ obtain performance similar to BR ($SCOM$ becomes again generally superior to $R\&R$ in this scenario). Note that TS is not able to produce the best solution any more, because it is trapped in some local minima, whereas SA, which exhibits a more random behavior, can more easily escape from such local minima.

VIII. CONCLUSIONS

In this paper we have presented and evaluated new greedy and metaheuristic algorithms for the optimal design of logical topologies in wavelength-routed WDM networks supporting multicast and unicast traffic, and we have provided a MILP formulation of the problem, which is however NP-hard.

Focusing on the case of IP networks, we assumed that the IP datagram routing algorithms (both unicast and multicast) are fixed, thus not a target of the optimization procedure.

Lower bounds and numerical results have been presented for the assessment of the performance of the proposed algorithms, indicating that the proposed metaheuristics largely outperform the greedy approaches, and are able to obtain very good logical topologies, quite close to the theoretical lower bounds.

REFERENCES

[1] D.Bienstock, O.Gueluek, "Computational Experience with a Difficult Mixed-Integer Multicommodity Flow Problem", *Mathematical Programming*, Vol.68, pp.213-235, 1995.

[2] B.Yener, T.Boult, "A Study of Upper and Lower Bounds on Minimum Congestion Routing in Lightwave Networks", *IEEE Infocom'94*, Toronto, Ontario, Canada, June 1994.

[3] E.Leonardi, M.Mellia, M.Ajmone Marsan, "Algorithms for the Topology Design in WDM All-Optical Networks", *Optical Networks Magazine, Premiere Issue*, Vol.1, n.1, pp.35–46, January 2000.

[4] L.H.Sahasrabuddhe, B.Mukherjee, "Light Trees: Optical Multicasting for Improved Performance in Wavelength Routed Networks", *IEEE Communications Magazine*, Vol.37, n.2, pp.67-73, February 1999.

[5] R.Ramaswami, K.N.Sivarajan, "Design of Logical Topologies for Wavelength-Routed Optical Networks", *IEEE Journal of Selected Areas in Communications*, Vol.14, n.6, pp.840-851, June 1996.

[6] B.Mukherjee, D.Banerjee, S.Ramamurthy, A.Mukherjee, "Some Principles for Designing a Wide-Area WDM Optical Network", *IEEE Transactions on Networking*, Vol.4, n.5, pp. 684-695, October 1996.

[7] J.Moy, "Open Shortest Path First Version 2", RFC 2178, July 1997.

[8] Y.Rekhter, T.Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[9] D.Waitzman, S.Deering, C.Partridge, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.

[10] S.Deering, D.Estrin, D.Farinacci, V.Jacobson, C.Liu, L.Wei, "The PIM Architecture for Wide Area Multicasting", *IEEE/ACM Transactions on Networking*, Vol.4, n.2, pp.153-162, April 1996.

[11] A.Ballardie, "Core Based Trees (CBT) Multicast Routing Architecture", RFC 2201, September 1997.

[12] S.Ramanathan, "Multicast Tree Generation in Networks with Asymmetric Links", *IEEE/ACM Transaction on Networking*, Vol.4, n.4, pp.558-568, August 1996.

[13] R.E.Tarjan, "Data Structures and Network Algorithms", *Society for Industrial and Applied Mathematics*, Pennsylvania, USA, November 1983.

[14] G.L.Nemhauser, A.H.G.Rinnoy Kan, M.J.Todd, "Optimization - Handbooks in Operations Research and Management Science", Vol. 1, *North- Holland*, 1989.

[15] T.H.Corman, C.E.Leiserson, R.L.Rivest, "Introduction to Algorithms", *MIT Press*, Cambridge, MA, USA, 1990.

[16] M.Garey, D.S.Johnson, "Computers and intractability: a guide to the theory of NP-completeness", *Freeman & C.*, New York, USA, 1978.