Traffic Grooming in Star Networks

Bensong Chen, Rudra Dutta, George N Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534

***Abstract*— We study the traffic grooming problem on WDM networks with the physical topology of a star. In star networks, several nodes are connected to a single hub directly through a bi-directional optical fiber, but are not connected to each other. Previous studies concentrated on the objective of minimizing the total amount of electronic switching. However, in order to lower the network cost, we consider the objective of minimizing the number of line terminating equipment (LTE), which is the dominant cost among optical devices. We first present new complexity results for this problem. The goal of minimizing LTE cost can be embodied in two alternative objective functions, the total number of LTE or the number of LTE at the node with the greatest such number (Min-Max). We provide heuristic algorithms for both objectives, and we present numerical results to demonstrate the effectiveness of our algorithms. One of our main contributions is to show that for star networks, it is usually possible to obtain near-optimal values for both objectives simultaneously, and our heuristics can obtain such solutions. The most consistently useful solutions for both objectives can be obtained by addressing the Min-Max objective as primary.**

I. Introduction

Wavelength division multiplexing (WDM) technology has the potential to satisfy the ever-increasing bandwidth needs of network users on a sustained basis. WDM is the process of transmitting data simultaneously at multiple carrier wavelengths over an optical fiber cable. The wavelengths can be kept sufficiently far apart that they do not interfere with each other. Thus, a single strand of fiber can be thought of as a collection of high capacity *virtual fibers*. Today, WDM systems are widely deployed in long-haul networks, and have a major presence in metro-area networks as well.

A *lightpath* is defined as a clear channel (or *wavelength* in this case) in which the signal remains in optical form throughout the physical path between two end nodes. The set of lightpaths defines a *logical topology*, which can be designed to optimize some performance measure for a given set of traffic demands. The logical topology design problem has been studied extensively in the literature. Typically, the traffic demands have been expressed in terms of whole lightpaths, while the metric of interest has been the number of wavelengths, the congestion (maximum traffic flowing over any link), or a combination of the two. In WDM networks, nodes are equipped with *optical cross-connects* (OXCs), devices which can optically switch wavelengths, thus making it possible to establish *lightpath* connections between pairs of network nodes. Since each network node needs to terminate data destined to the site, and initiate traffic from it to other network nodes, devices are needed to add/drop signals to/from the lightpaths, and switch them to other channels, if necessary. Since user data are expressed in electronic signals, some kind of transformation between optical and electronic signals needs to be done. The OEO (opto-electro-optical) transformation is done by *line terminating equipment* (LTE) at each network node. Digital cross-connects (DXC) can further be used to switch the electronic signals for rearranging data multiplexed onto the wavelengths.

With the deployment of commercial WDM systems, it has become apparent that the cost of network components, especially LTE, is one of the dominant costs in building optical networks, and is a more meaningful metric to optimize than, say, the number of wavelengths (since with current technology, more than a hundred wavelengths can be multiplexed in one fiber). Furthermore, with currently available optical technology, the data rate of each wavelength is on the order of 2.5-10 Gbps, while 40 Gbps rates are becoming commercially available. In order to utilize bandwidth more effectively, new models of optical networks allow several independent traffic streams to *share* the capacity of a lightpath. These observations give rise to the concept of *traffic grooming*, a variant of logical topology design, which is concerned with the development of techniques for combining lower speed components onto wavelengths in order to minimize network cost.

In this paper, we consider the traffic grooming problem in the star topology. Star networks arise in the interconnection of LANs or MANs with a wide area backbone. Cable TV networks and passive optical networks (PONs) are based on a tree topology, which can be decomposed into stars as well. Also, consider a relatively small optical WDM network with a general topology. If we require that only one of the network nodes has switching ability, the virtual topology formed by a traffic grooming solution will be in the form of a star network. Although direct lightpaths that "pass through" the hub node may not actually pass the OXC at the hub, the virtual topology (by ignoring the physical links) will look just like that in the star topology. Therefore, if we use the star virtual topology as building block, it may be possible to solve larger and more general network topologies by employing a decomposition method.

Typically, the goal of existing traffic grooming studies is to minimize the LTE cost, given a set of traffic demands. To this end, the problem is formulated so that its objective reflects the amount of LTE either directly (e.g., by counting the number of lightpaths established) or indirectly (e.g., by considering the amount of electronic routing required, as opposed to optical routing). However, most studies concentrate on some *aggregate* representation of the LTE cost. That is, the

This work was supported in part by NSF grant ANI-0322107.

objective to be minimized is usually expressed as the sum, over all network nodes, of the LTE cost at each individual node. While a metric that accounts for the network-wide LTE cost is important, minimizing the total cost in the network without imposing any bound on the cost of individual nodes may result in a solution in which some nodes (such as the hub node in the star topology) end up with a (very) large amount of LTE while some others with only a small amount of LTE. Such a solution may have a number of undesirable properties. First, a node that requires a large amount of LTE may be too expensive or even impractical to deploy (e.g., due to high interconnection costs, high power consumption, or space requirements). Second, the resulting network can be highly heterogeneous in terms of the capabilities of individual nodes, making it difficult to operate and manage. Third, and more important, a solution minimizing the total LTE cost can be extremely sensitive to the assumptions regarding the traffic pattern, as previous studies [10] have demonstrated. Specifically, a solution that is optimal for a given set of traffic demands may be far away from optimal for another such set that is only slightly different from the first. Since LTE involve expensive hardware devices that are difficult to move from one node to another on demand, an approach that attempts to minimize total LTE cost may not be appropriate for dimensioning a network unless the network operator has a clear picture of traffic demands far into the future *and* these traffic demands are unlikely to change substantially over the life of the network. Recently, we studied the problem of minimizing the maximum nodal LTE cost in a ring network [2]. We also note that a similar approach of minimizing the maximum nodal cost was taken in [4] in a different context, namely for routing and wavelength assignment in the presence of converters.

Because of its importance and difficulty, the traffic grooming problem has drawn a lot of attention in recent years. Some surveys on the general traffic grooming problem can be found in [11], [19], [6], [23]. Recent studies aimed at more general mesh topologies can be found in [9], [17], [22]. Additional studies with the objective of minimizing LTE cost can be found in [15], [16], [18], [13].

The new topic of dynamic grooming is emerging in some recent work, where the traffic is not completely static, and must be modeled by some means other than a single traffic matrix (e.g., by subwavelength call arrivals and departures). In [21], an efficient reconfiguration problem is studied using a novel graph representation of the problem. The grooming problem is then transformed into corresponding problems in graph theory. Note that the Min-Max objective we consider in this paper results in balanced equipment capability at each network node. Thus it is eminently suitable for the emergent area of dynamic traffic grooming, because such a solution offers more adaptability toward changing traffic patterns.

The grooming problem in the star topology has also gained interest recently. The objective considered in relevant studies is to minimize the total amount of electronic switching (and, thus, the delay introduced by OEO transformation), which is related to, but different from, the objective we consider in this study. In [5], the authors first prove that the problem is equivalent to a Maximal Weighted Local Constraint Subgraph (MWLCS) problem, which is NP-Complete. They also describe a greedy algorithm that guarantees a solution whose cost is at most twice the optimal. In [1], the authors consider two versions of the problem: minimizing the electronic switching and maximizing the optical switching. Besides proving NP-Completeness, they also proved an important result regarding approximation of this problem. They have shown that an approximation algorithm for either version cannot act as an approximation for the other. They also provided approximation algorithms for both versions separately, by transforming the corresponding versions of the problem to existing NP-Complete problems. A polynomial-time optimization algorithm is also given for the special case in which only two wavelengths are available on each fiber. In earlier work [8] two of the authors studied the grooming problem in stars, among other elemental network topologies. For star networks, we gave complexity proofs, and by pruning the search tree, found a method that can give a series of upper and lower bounds. A greedy heuristic was provided to make improvement towards the objective at each iteration. The paper inspired some of the NP-Completeness results we present.

In this paper, we concentrate on LTE cost. This is an important measure of the actual cost of a network. We consider both the problems of minimizing this cost as a total over all the network nodes, as well as minimizing the LTE cost at the node where this cost is maximum (the Min-Max problem). We discuss the relation between the two, and present some intuitive insights in this regard, as well as provide heuristic algorithms for each objective that perform well in practice. We show that addressing the Min-Max objective as the primary one produces the most practically useful solutions, from the point of view of either objective.

The rest of the paper is organized as follows. In the next section we define the problem precisely. Section III presents our results on the computational complexity of the various star network grooming problems. In Section IV we present our heuristic algorithms. Section V presents numerical results to demonstrate the effectiveness of our approach, and Section VI concludes the paper.

II. Problem Definition

The traffic grooming problem is an extension of the well-known *Routing and Wavelength Assignment* (RWA) problem in wavelength routed optical networks. An optical network can be abstracted as a directed graph, with vertices representing network nodes (sites), connected with directed edges showing the optical fiber links. A traffic demand from node s to node d (denoted as $t^{(sd)}$) can be carried on a certain physical route of fiber links, and by certain wavelengths on each of the links. If we consider a set of such demands T, the RWA problem becomes: how to carry all the traffic demands from the respective sources to the destinations using the available wavelengths. Since data carried for each demand are distinct, this is a *multicommodity flow* problem. If two demands share the same fiber link, they must be carried on different wavelengths. The goal of the RWA problem is to satisfy all traffic demands in T, while minimizing the number of wavelengths used in the whole network. Generally, the RWA problem assumes no use of wavelength converters in the network, that is, each traffic demand is carried on the same wavelength throughout the routing path.

Previous studies show that wavelength assignment to minimize the number of wavelengths can be solved in polynomial time in paths and stars, but it is NP-hard [7] in tree topologies. Not surprisingly, the RWA problem is NP-hard in general network topologies [14] as well. However, if the topology is a star, the problem is equivalent to the minimum edge coloring problem in bipartite graphs, which is solvable in polynomial time [20].

A. The Star Traffic Grooming Problem

Current optical technologies allow very high bandwidths for each wavelength channel in a fiber link. Since individual $t^{(sd)}$ values are likely to be much lower, effective utilization of the wavelength channel bandwidth requires multiplexing lower-rate traffic streams onto a single wavelength using time-division multiplexing. If we were to require that only traffic belonging to the same source/destination pair be multiplexed onto the same lightpath, the grooming problem would be equivalent to the RWA problem. However, this constraint means that we have to set up direct lightpaths for each source/destination pair, which is generally impractical due to wavelength constraints or optical device constraints, and results in low utilization of the available bandwidth.

For this reason, each node in the optical network needs to perform both optical and electronic switching. Each node is designed to let some lightpaths pass through by optically switching them, and to terminate/originate other lightpaths. Some traffic may be switched electronically onto new lightpaths to be carried to its destination. This electronic switching function, referred to as *grooming*, allows for better use of wavelength capacity, reduces wavelength requirements, and enhances *virtual connectivity*. As a trade-off, expensive electro-optic devices (e.g., line terminating equipment) and electronic switches (digital cross-connects) need to be placed at the network nodes. The traffic grooming problem is thus defined as that of balancing the advantages and costs, and obtaining a tradeoff that minimizes some measure of the cost of network equipment.

We define a positive integer C as the capacity of one wavelength, expressed as units of some basic transmission rate (such as OC-3). The capacity C has also been variously called the *grooming factor*, or *granularity*. Let W be the number of wavelengths that each fiber can carry concurrently. A traffic demand matrix $T = [t^{(sd)}]$ can be defined, where integer $t^{(sd)}$ denotes the number of basic transmission units to be carried from node s to node d. (We allow the traffic demands to be greater than the capacity of a lightpath, i.e., it is possible that $t^{(sd)} > C$ for some s, d.) Given the traffic matrix, the traffic grooming problem involves the following conceptual

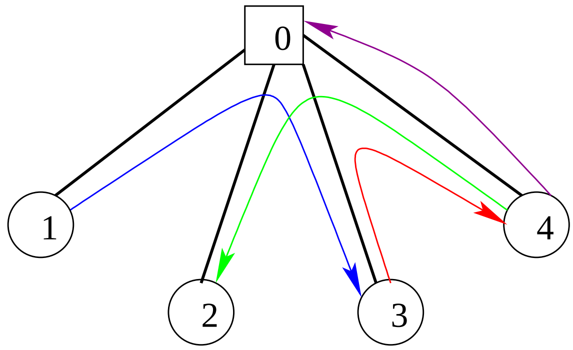

Fig. 1. A 5-node star with 4 lightpaths

subproblems (SPs): (1) *logical topology SP:* find a set R of lightpaths that forms a virtual topology, (2) *lightpath routing and wavelength assignment SP:* solve the RWA problem on R, and (3) *traffic routing SP:* route each traffic stream through the lightpaths in R.

The first and third subproblems together constitute the grooming aspect of the problem. Also, the number W of wavelengths per fiber link is taken into consideration as a constraint rather than as a parameter to be minimized. Note that this is only a conceptual decomposition that helps in understanding and talking about the problem, and we do not suggest that the grooming problem can be generally solved by solving these subproblems in order.

The optimization goal we consider is to minimize the cost of LTE in the network. We consider two flavors of the objective which gives rise to two different problems of star network grooming:

Overall: Minimize the total number of LTE to be used in the design. This is equivalent to minimizing the number of lightpaths established in the system (since each lightpath requires one LTE at both ends). Note that this objective is also equivalent to minimizing the number of edges in the logical topology formed by lightpaths. As we mentioned before, this cost function is more appropriate when custom-designed equipment may be placed at different nodes, and the traffic pattern is expected to be quite static.

Min-Max: Minimize the number of LTE to be used at the node at which the maximum such number is used, or equivalently minimize the *maximum number of lightpaths originating from or terminating at any node.* This is the same as minimizing the *maximum nodal degree* in the logical topology. The nodal degree is defined as the maximum of the in- and out-degrees of a node. This cost function is more appropriate when we are constrained to use the same equipment at each node, or the traffic pattern is expected to vary even though the overall traffic load may remain the same.

We restrict our study to the physical topology of a star. Figure 1 shows a star network with 5 nodes. We always label the central hub node with 0, followed by non-hub nodes. We assume that every physical link is bi-directional, and no traffic demand is allowed to traverse the same physical link more than once. This assumption can ensure efficient use of wavelength capacity. Thus, only the hub node can be an intermediate stop for a traffic component, non-hub nodes are not allowed to switch traffic.

Under these assumptions, there can be only two types of lightpaths in the logical topology. The first type consists of single-hop lightpaths which either originate at a non-hub node and terminate at the hub node, or vice versa. The second type consists of two-hop lightpaths that originate and terminate at non-hub nodes, and are switched optically at the hub node. Such a two-hop lightpath from node s to node d can only carry the traffic component $t^{(sd)}$ (in whole or part), and any remaining bandwidth on it cannot be used to carry any other traffic (due to the restriction that non-hub nodes switch no traffic).

In star networks, wavelength assignment can always be done in polynomial time, as long as other wavelength constraints are not violated [12]. However, we prove in the next section that the grooming problem as a whole remains intractable, whether bifurcated routing is allowed or not.

III. Complexity Results for Stars

We consider two types of traffic routing, based on whether bifurcation of traffic is allowed or not allowed. When bifurcation is allowed, any traffic component $t^{(sd)}$ can be split into various parts which may follow different logical routings from source to destination; however, the splitting has to be into parts that are each integer multiples of the base rate. When bifurcation is not allowed, such split-routing of traffic is not permitted; that is, for any source-destination pair (s, d) such that $t^{(sd)} \leq C$, we require that all $t^{(sd)}$ traffic units be carried on the *same* sequence of lightpaths from source s to destination d. On the other hand, if $t^{(sd)} > C$, it is not possible to carry all the traffic on the same lightpath. In this case, we allow the traffic demand to be split into $\lfloor \frac{t^{(sd)}}{C} \rfloor$ subcomponents of magnitude C and at most one subcomponent of magnitude less than C, and the no-bifurcation requirement applies to each subcomponent independently.

While many grooming problems have been known to be NP-Complete, the computational complexity of traffic grooming in star networks (a particularly simple topology) under the objective functions mentioned above has remained an open question so far. We have obtained proofs which settle these questions. For bifurcated routing, we have the following result:

Theorem 3.1: The decision version of the grooming problem in star networks with the Min-Max objective (bifurcated routing of traffic allowed) is NP-complete.

The proof is by reduction from the constrained PARTITION problem. By creating a star network with two source nodes and $2n$ destination nodes (n is related to the PARTITION parameters), we can set up the traffic demand matrix so that exactly half of the residual demands need to be carried from each source to n of the destinations, and that the sum of the demands are equal to a certain value. The detailed proof can be found in [3].

A similar construction allows us to prove the same result for the Overall objective as well. Then we can claim the following corollary; the detailed proof can be found in [3].

Corollary 3.1: The decision version of the grooming problem in star networks with the Overall objective (bifurcated routing of traffic allowed) is NP-complete.

Similarly, for the non-bifurcated case, we have:

Theorem 3.2: The decision version of the grooming problem in star networks with the Min-Max objective (bifurcated routing of traffic not allowed) is NP-complete.

The proof is by reduction from the BIN-PACKING problem. For each instance of the BIN-PACKING problem, we can create a star network with one source node and n destination nodes. The bin size is the capacity of one wavelength. By setting up the traffic demands, we can force the solution to partition the n residual demands into the available wavelengths while not violating the wavelength capacity constraints. The detailed proof can be found in [3].

Again, the problem restriction allows us to extend the result to the Overall objective as well. Further, because of the construction in the proof, we can claim the following aggregated result for the non-bifurcated case in the star topology:

Corollary 3.2: The decision version of the grooming problem in star networks with the Min-Max or the Overall objective (bifurcated routing of traffic not allowed) is NP-complete, and remains so even when a candidate logical topology is provided as part of the problem instance.

IV. Heuristic Grooming Algorithm for Stars

Clearly, for practical purposes a polynomial-time algorithm with good performance is needed for star networks of large size, preferably one that provides near-optimal grooming solutions. In this section, we present polynomial-time algorithms for star networks with either the Min-Max or the Overall objective. We have focused only on the case where bifurcated routing is allowed, since this is a more realistic design scenario; bifurcated routing is more flexible, allows better utilization of network resources and may be preferred in some situations because of security concerns.

A. Star Grooming for the Min-Max Objective

As we have mentioned before, wavelength assignment in stars can be performed in polynomial time, therefore we will concentrate only on the virtual topology and routing problem in our heuristic algorithm.

The detailed algorithm is described in Figure 2. The concept of *reduction* of a traffic matrix (in Line 1 of the algorithm) is to reduce the matrix T so that all elements are less than the capacity C of a single wavelength, by assigning a whole lightpath to traffic between a given source-destination pair that can fill it up completely. The available wavelengths on the links of the path segment from the source to the destination node are also decremented by the number of lightpaths thus assigned. Since breaking such lightpaths would increase the amount of LTE at some intermediate nodes of the path, this procedure does not preclude us from reaching an optimal solution, nor does it make the problem inherently easier or more difficult. We continue using the same notation for the

Min-Max Traffic Grooming Algorithm for Star Networks
Input: A star network with N non-hub nodes and a hub 0, W wavelengths, capacity C of each wavelength, and traffic matrix $T = [t^{(sd)}]$.
Output: The number of lightpaths b_{ij} from node i to node j of the star, and traffic routing quantities $t_{ij}^{(sd)}$; **or** failure if no feasible solution exists.

procedure **StarMinmaxGrooming**

1. Reduce the traffic matrix by assigning direct lightpaths to traffic components greater or equal to C
2. Use 1-hop lightpaths to carry the remaining two-hop traffic
3. Check feasibility. If infeasible, exit with failure
4. Initialize $b_{ij}, t_{ij}^{(sd)}$, record indegree I_j, outdegree O_j and remaining capacity r_{ij} on lightpath (i, j) in the current topology
5. $u \leftarrow$max degree of the non-hub nodes
6. **while** *max hub degree* $> u$ **do**
7.   Sort all the 2-hop residual $t^{(sd)}$ in non-increasing order
8.   **for each** of the sorted $t^{(sd)}$ **do**
9.   **if** carrying $t^{(sd)}$ directly does not increase u **then**
            re-route the traffic on direct lightpaths
            update all variables accordingly
10.   **endfor**
11.   **if** $u < W$ and *max hub degree*$> u$ **then** u++
12.   **else** break; **endif**
13. **endwhile**
14. Use the polynomial-time WLA algorithm in [12] to assign wavelengths to the lightpaths

end **procedure**

Fig. 2. Traffic Grooming Algorithm for Star Networks with the Min-Max Objective

traffic matrix and traffic components, but in what follows they stand for the same quantities after the reduction process.

The following step, checking feasibility, is straightforward: determine if the all-electronic solution violates any wavelength constraint. Note that after getting the all-electronic solution, we have reached the lower bounds for the degrees at non-hub nodes. Since we are considering the Min-Max objective, if the hub degree is even smaller than one of the non-hub degrees, we have already reached the optimal; otherwise, we need to lower the hub degree without increasing our Min-Max objective. To this end, we use a greedy heuristic which, at each step picks the largest traffic component that has not been considered yet, and attempts to optically route it, if doing so will reduce the hub degree without increasing the maximum degree. Following this procedure, we may reach a point where the maximum degree is at one or more non-hub nodes, possibly jointly with the hub node, in which case the algorithm terminates. Otherwise the hub node still has a degree greater than that of any non-hub node. If it would not violate the wavelength limit at any non-hub node, we then decrease our Min-Max target by one and repeat the procedure.

The complexity of the algorithm is straightforward to obtain: the reduction and initial feasible solution takes time $O(N^2W)$; the *while* loop between Steps 6-13 is executed no more than W times, since a feasible solution requires $u \leq W$; and there are no more than $N(N-1)$ traffic elements to consider within the loop. Thus, the overall complexity of the algorithm is $O(N^2W)$.

B. Star Grooming for the Overall Objective

One of the main goals of this study was to understand the relationship between the Overall and the Min-Max objectives. Accordingly, we were interested in knowing whether the Min-Max heuristic presented above could be used with little or no modification to perform well for the Overall objective as well. We found that a different approach is needed to address the Overall objective, though there is some commonality. In this section, we present our development of the algorithm for and insight into the Overall objective; later in Section V, we present data regarding the performance of these algorithms not only for the objective each was designed for, but also for the other objective.

First, we argue that the Overall results given by our Min-Max heuristic in Figure 2 can hardly be improved without violating the Min-Max objective. This is because we aimed at the Min-Max objective alone in the algorithm. As we will see in Section V, the Min-Max algorithm performs very close to the optimal: for stars with 10 non-hub nodes, only 1 out of the 50 cases in our experiments results in a non-optimal solution. However, this is accomplished at the expense of neglecting the Overall objective. Let us consider the following approaches to improve the value of the Overall objective.

(1) After the Min-Max heuristic terminates, continue rerouting two-hop traffic onto direct lightpaths. This might improve the nodal degree at the hub, but cannot be done without violating the existing Min-Max solution, because at the last iteration of the *while* loop at Step 6-13, we have rerouted all such traffic that will not violate the current objective u in non-increasing order (refer to the *if* statement at Step 9).

(2) Break some direct lightpaths back into two at the hub node. This can only be done by examining the last iteration of the *while* loop step by step, considering the candidate traffic in nondecreasing order (because smaller traffic are more likely to fit into remaining capacities). Moreover, we need to ensure that (a) the hub degree cannot increase beyond u, and (b) the non-hub degree cannot increase beyond u. Note that at the previous iteration, we have reached the limit that the non-hub degrees cannot increase beyond $u-1$, so there is little room left for improvement.

Finding a good combination of the above two approaches leads to combinatorial explosion. Therefore, if we want to balance the two objectives, new approaches are needed. To find a good algorithm, we use a Integer Linear Programming (ILP) formulation for the star grooming problem. A more general (and more straightforward) ILP can be found in [11], but here we give a simplified version that has only binary (0-1) variables. The formulation allows us to relax the constraints to allow for traffic demands between each source/destination pair that exceed the wavelength capacity C.

In this context, we need to distinguish between the problem parameters as given by the instance, and after some preprocessing. First we perform the traffic matrix reduction (as described above in Section IV-A), then we set up the minimum number of single hop lightpaths to (and from) the hub node from (and to) each non-hub node which are required to carry traffic components that terminate (and originate) at the hub (that is, traffic components of the form $t^{(s0)}$ and $t^{(0d)}$). We define the following notation for quantities after this preprocessing:

$t_r^{(sd)}$: the traffic demand $(0 \leq t_r^{(sd)} < C)$ from s to d, $s \neq d \neq 0$. We further define $t_{out}^s = \sum_d t_r^{(sd)}, t_{in}^d = \sum_s t_r^{(sd)}$.

r_{out}^s: the remaining capacity left on the (possibly) underutilized lightpath $(s,0)$. r_{in}^d is defined similarly.

w_{out}^s: the full wavelengths available on link $(s,0)$. w_{in}^d is defined similarly.

We need to find $x^{(sd)} \in \{0,1\}$, in which 0 denotes electronic routing of remaining demand $t_r^{(sd)}$, while 1 denotes optical routing (setting up a two-hop lightpath dedicated to it).

Therefore, our goal to minimize the total number of lightpaths can be expressed as follows; note that we do not count the lightpaths set up during reduction, because they are necessary and we have no choice but to keep them.

Minimize:

$$\sum_s \left\lceil \frac{t_{out}^s - r_{out}^s - \sum_d t_r^{(sd)} x^{(sd)}}{C} \right\rceil + \sum_d \left\lceil \frac{t_{in}^d - r_{in}^d - \sum_s t_r^{(sd)} x^{(sd)}}{C} \right\rceil + \sum_{s,d} x^{(sd)} \quad (1)$$

Subject to (link capacity constraints):

$$t_{out}^s - r_{out}^s - \sum_d t_r^{(sd)} x^{(sd)} + C \sum_d x^{(sd)} \leq C w_{out}^s, \forall s \quad (2)$$

$$t_{in}^d - r_{in}^d - \sum_s t_r^{(sd)} x^{(sd)} + C \sum_s x^{(sd)} \leq C w_{in}^d, \forall d \quad (3)$$

Note that, except for $x^{(sd)} \in \{0,1\}$, all other values can be calculated from the traffic matrix very easily, and really count as parameters. Therefore, the total number of variables is at most $N(N-1)$. Each variable is binary, thus the solution space has $2^{N(N-1)}$ combinations (or less if some traffic components are zero).

If we relax the ceiling operations in the objective, we can get a new formulation that inspires a greedy algorithm, as we show below. The new constraints are (link capacity constraints):

$$\sum_d (C - t_r^{(sd)}) x^{(sd)} \leq C w_{out}^s - t_{out}^s + r_{out}^s, \forall s \quad (4)$$

$$\sum_s (C - t_r^{(sd)}) x^{(sd)} \leq C w_{in}^d - t_{in}^d + r_{in}^d, \forall d \quad (5)$$

The goal is now to minimize:

$$\sum_{s,d} (C - 2t_r^{(sd)}) x^{(sd)} + \left(\sum_s (t_{out}^s - r_{out}^s) + \sum_d (t_{in}^d - r_{in}^d) \right) \quad (6)$$

where the latter part is a constant. The ILP formulation resembles the *Multiconstraint 0-1 Knapsack Problem* (MKP), also called Multi-Dimensional 0-1 Knapsack Problem (MDKP). However, it has special forms that characterize the star grooming problem, so better approaches can be taken to get near-optimal solutions.

We can now get insight into the development of a good heuristic by using the 0-1 ILP formulation for the problem. From the goal, we know that we should try to route all traffic demands that are greater than $C/2$ optically, so that the values $(C - 2t_r^{(sd)}) x^{(sd)}$ are negative, and will decrease the goal. However, we should guard against the constraints as well. Therefore, our intention is to greedily route the traffic demands that are more than $C/2$ optically, while making sure not to violate the constraints.

Note that minimizing the quantity $\sum_{s,d} (C - 2t_r^{(sd)}) x^{(sd)}$ is the same as maximizing $\sum_{s,d} (2t_r^{(sd)} - C) x^{(sd)}$. Hence, for the relaxed objective, at each iteration, we are closer to our goal by $2t_r^{(sd)} - C$ by routing $t_r^{(sd)}$ directly onto one lightpath. However, for the original objective with the ceiling operations, the performance of each iteration will depend on the remaining capacities on the two corresponding lightpaths $(s,0)$ and $(0,d)$. For instance, even if a $t_r^{(sd)} < C/2$, when both lightpaths have remaining capacity $\geq C - t_r^{(sd)}$, routing the traffic directly on a single lightpath will decrease the total degree by 2; accordingly, some traffic greater than $C/2$ may actually introduce a new lightpath without eliminating either of the two one-hop lightpaths, adding penalty to the objective. Therefore, we can make potential improvement by trying to route more traffic directly which does not exceed $C/2$, and if the results are better, we should accept the better solution.

This observation leads to the essentially greedy algorithm presented in Figure 3. The algorithm is similar to the one presented in [8], and that in [5] as well, although those algorithms were designed to minimize an electronic routing (not nodal degree) objective. Note the important difference that in our algorithm, while the order of considering traffic elements is essentially greedy, we continue until all traffic elements have been considered, and then pick the best one. This means that in case a succession of greedy steps produce first an increase, followed by a larger decrease, we shall be able to pick the best case; further, if several steps produce the same objective value, we shall be able to pick any of them. The importance of this ability will become clear in the discussion of Section V-C.

The idea of the algorithm comes from the fact that the optimal overall solution must lie between the all-electronic and the all-optical solutions. Furthermore, intuitively, carrying larger traffic demands on two-hop lightpaths will make better use of the wavelength capacity. The disadvantage of the all-electronic solution is excessive hub degrees, while the drawback of the all-optical solution is the many underutilized lightpaths that could have been avoided by grooming. With our approach, we are likely to find near-optimal results in-between

Traffic Grooming Algorithm for Star Networks to Minimize The Overall Number of Lightpaths

Input: A star network with N non-hub nodes and a hub 0, W wavelengths, capacity C of each wavelength, and traffic matrix $T = [t^{(sd)}]$.

Output: The overall number of lightpaths $\sum b_{ij}$ from node i to node j of the star, and traffic routing quantities $t_{ij}^{(sd)}$; **or** failure if no feasible solution exists.

procedure **StarOverallGrooming**

1. Reduce the traffic matrix, update corresponding variables, and record residual traffic $[t_r^{(sd)}]$
2. Use 1-hop lightpaths to carry all $t_r^{(sd)}$ electronically
3. Check feasibility. If infeasible, exit with failure
4. Set $i = 1$; record the current overall degree as u_0
5. Sort all the nonzero two-hop residual $t_r^{(sd)}$ in non-increasing order, labeled as $l_1, l_2, \ldots, l_m, m \le N^2$
6. **while** $l_i > 0$ **do**
7.   Reroute l_i directly on a new 2-hop lightpath if doing so does not violate the constraints
8.   Record the resulting overall degree as u_i; $i \leftarrow i + 1$
9. **endwhile**
10. Find the smallest of $u_0, u_1, \ldots, u_m$, and use the corresponding virtual topology as the solution
11. Use the polynomial-time WLA algorithm in [12] to assign wavelengths to the lightpaths

end **procedure**

Fig. 3. Traffic Grooming Algorithm for Star Networks with the Overall Objective

the two extremes.

The complexity of the algorithm is again straightforward to obtain. Up to N^2 iterations can be made in Step 4, each in $O(1)$ time. Finding the smallest value in a list of N^2 values takes time $O(\log N)$. Thus, the overall complexity for the greedy algorithm is $O(N^2)$.

V. Numerical Results

In this section, we present our experimental results with stars for the two algorithms we discussed in the previous section. First we present individual results for each algorithm, then consider joint issues.

For any set of results, we use 50 problem instances. The traffic matrix $T = [t^{(sd)}]$ of each of the instances is generated by drawing $N(N-1)$ random numbers (rounded to the nearest integer) from a Gaussian distribution with a given mean t and standard deviation σ that depend on the traffic pattern. We consider two traffic patterns, a *random* and a *quasi-uniform* pattern.

We consider the random traffic pattern because this is often the most challenging and general of traffic scenarios, when there is no particular structure to the traffic matrix that can be exploited by an algorithm. To generate a traffic matrix of this pattern, we first determine the mean value t of the Gaussian distribution according to the desired link load L, and we let the standard deviation be 150% of the mean t. Consequently, the traffic elements $t^{(sd)}$ take values in a wide range around the mean, and the loads of individual links also vary widely. With such a high standard deviation, the random number generator may return a negative value for some traffic elements; in this case, we set the corresponding $t^{(sd)}$ values to zero. Also, if a traffic matrix generated in this manner is infeasible (i.e., the load on some link exceeds the value WC), then we discard it and we generate a new matrix for the corresponding problem instance. Note that with the high variation, the link load L may not be balanced, and we can only guarantee a range of link loads around the target load for the instances.

We also experiment on a quasi-uniform traffic pattern with high loads, because such a pattern is quite realistic. Traffic matrices are generated by the same procedure as above, but the Gaussian distribution parameter is changed to make the standard deviation 10% of the mean. In this case, when all links have high loads of traffic, the wavelength constraints will not allow for an all-optical solution, and the performance of the algorithms will change accordingly. We revisit this issue in Section V-D.

A. Results for the Min-Max Objective

Using the industry standard ILP solver CPLEX with the ILP formulations, we can get some good results for the random pattern with the Min-Max objective. Using an Sun UltraSparc system, we were able to get optimal results for stars with up to $N = 16$ non-hub nodes.

We collected data for 50 cases with $N = 16, 24$. We also used the all-electronic solution as the scale to evaluate our solutions and the optimal. The concept of "grooming effectiveness" is used in the evaluation, which is defined as the objective divided by the all-electronic result (at the hub node for this case); note that lower is better for this measure. We have found that our heuristic performs well in general. For $N = 16$, only 4 of the 50 cases give non-optimal result; for $N = 24$, the number of non-optimal cases rises to 18.

Figure 4 shows results for stars with 24 non-hub nodes. We find that all non-optimal results are very close to the optimal. We also find that the grooming effectiveness is smaller (better) when the star size grows: the typical grooming effectiveness was around 40% for 10-node stars (not shown here) and is around 25% for 24-node stars as seen in Figure 4, even though the load was less for the former. This result suggests that grooming techniques are more helpful in larger networks, even if we restrict the network to the simple topology of stars.

B. Results for the Overall Objective

The Overall algorithm, after applying all the constraints, also works quite well for stars with 10 non-hub nodes. When the star size grows to 16, most of the cases we generate take more than a few hours to solve with CPLEX, and we were unable to obtain bulk results for the optimal in these cases. The plot in Figure 5 shows results for 10-node stars. The typical optimal grooming effectiveness is just under 60%, which translates to an average of 446 lightpaths for each case, and our algorithm obtains solutions with only 2 to 4 additional lightpaths. The average difference is 2.96, which is less than one percent from the optimal values.

Having seen that our algorithms perform well for the objectives they were designed for, we investigate the performance of each for the objective it was *not* designed for in the next section. In Section V-D we investigate the behavior of the algorithm of Figure 3 when the limit on the number of available wavelengths in each fiber is ignored, again giving us insight into the relationship between the two objectives.

C. Cross-Objective Performance results

From the nature of strong symmetry in stars, we have the intuition that for the star network topology, the Min-Max and the Overall objective are closely related to each other. That is, by obtaining the Min-Max solution of the problem, the resulting overall degrees are also close to the Overall objective, and vice versa. Figure 6 shows how our solutions for the Min-Max objective perform with respect to the Overall objective. The results are for 50 instances of 10-node stars.

Figure 6 merits careful consideration. There are three separate curves, and they all compare the result obtained for some particular objective by a given algorithm versus the optimal value of *that objective*, for individual problem instances. In all cases, the optimal value of the two objectives were obtained by solving the ILPs exactly. Thus the vertical axis represents the percentage by which the output of a particular algorithm exceeds the optimal value. For example, considering the first three points on the dotted blue line with data points marked by asterisks ('*'), we see that the maximum degree at a node in the solution obtained by solving the ILP for the Overall objective was 12% more than the minimum possible maximum degree at a node (which was obtained by solving the ILP for the Min-Max objective) for the first problem instance; for the second instance, the maximum degree in the optimal solution of the Overall ILP was the same as the optimal maximum degree, and for the third instance it was 10% greater.

The first line (with plus '+' signs) plots the value of the overall degree in the optimal solution to the Min-Max ILP and the second line (with cross 'X' signs) plots the value of the overall degree in the solution returned by the Min-Max heuristic algorithm of Figure 2. The first obvious observation is that the Min-Max heuristic algorithm has very good performance with respect to the Overall objective, exceeding the optimal by only 3-4%. It is interesting to notice that with respect to the Overall objective, the Min-Max heuristic algorithm gives better results than the Min-Max optimal solutions given by CPLEX. This is not surprising because CPLEX simply returns the first optimal solution from its branch-and-bound search of the solution space, and the results may have many high-degree nodes due to the search sequence. For the same reason, we also note the important fact that the *variability* is much less for the Min-Max heuristic, thus the Min-Max heuristic minimizes the overall degree much more *consistently*.

Now we do a reverse comparison. The third line in Figure 6 (with asterisk '*' signs) plots the value of the maximum degree at a node in the optimal solution to the Overall ILP. The figure shows that while we try to minimize the overall degree, the maximum degree from the solution is also not far from the optimal Min-Max objective in star networks. However, it is quite inconsistent: several of the values are close to optimal while several others are between 10% and 15% higher. We can explain this result by recalling that often a reduction in the overall degree can be obtained at the cost of increasing the degree of the hub node and reducing the degrees at the non-hub nodes. In the next section, we further investigate the behavior of the Min-Max objective as the overall degree is sought to be minimized.

D. Evolution of the Two Objectives

In our heuristic algorithm for the Overall objective, a simple greedy approach is used without considering the Min-Max objective. In this section, we study the behavior of the two objectives as the algorithm proceeds from the all-electronic to the all optical solutions, and present the results in Figures 7-9. We record the value for both objectives at each iteration, and find the trends from all-electronic to all-optical solutions to study the trend on both objectives. If our goal is to find good solutions that are not far away from either of the two objectives, we can analyze the two figures showing the steps from the Overall algorithm, and find a point close to the *trough* of both.

Initially, to allow us to focus on this evolution, we ignore the constraint of the wavelength limit on a link. A star of size $N = 10$ means that the number of outgoing lightpaths from a node s has a variation of at most 9 between the all-electronic and all all-optical solutions. Therefore, the wavelength limit constraint does not come into play for instances with low traffic load and uniform pattern. For the Random traffic pattern, statistically, a few cases may happen to reach the limits, but the chances are low as long as the the link loads are low. On the other hand, when many links have very high loads concurrently (that is when the load is uniformly high), the wavelength limit realistically constrains feasible solutions.

From our experiments, we find that the Min-Max objective will generally go down and then later up during the transformation from all-electronic to all-optical solutions. The Overall objective has the same general trend, but may "thrash" up and down less smoothly. This justifies our strategy of picking the best value from the results of all iterations, rather than stopping when the Overall objective goes up for the first time.

Figures 7 and 8 present results for instances with *random* traffic pattern and low loads (in the range of 50%) and *uniform* traffic pattern with high load (95%) respectively. The horizontal axis represents the number of traffic elements which are routed optically instead of electronically, and the vertical axes represent the actual objective values. For the first instance, both objectives simultaneously reach their best values at around the 40^{th} iteration. However, for the instance of Figure 8, the Overall objective goes up well before the Min-Max objective reaches its best value.

Finally, Figure 9 shows the effect of the wavelength limit. It is the wavelength limit constrained version of Figure 8. (For the case of Figure 7, the constraints do not make a difference in the output; this is expected at low loads, as we mentioned above.) In this figure, if rerouting the traffic optically at a certain iteration would result in wavelength violation, we skip the element and leave the results for corresponding iteration blank. For instance, there is a big gap between Iterations 78 and 84, which means that the elements considered in those iterations will generate infeasible solution if we attempt to route them optically. We find that the effect of the constraints is a dampening of the objective curves, but the general behavior remains the same. On the whole, the removal of the constraints (at high traffic loads) produce effects similar to those produced by reducing the traffic load somewhat. We have obtained more numerical results which confirm these observations. Interestingly, the constraints do not in general affect the best value of the objectives by much.

VI. Conclusions

We considered the traffic grooming problem in WDM networks with the important physical topology of a star. We considered two objectives related to minimizing the number of Line Terminating Equipment. We proved that, for both objectives, the problem is NP-Complete. We proposed polynomial-time algorithms for both cases, and tested their performance. We also studied the relationship of the two objectives.

The results show that our algorithms are practically useful ones for this computationally hard problem. In star networks, it is usually possible to get good values for both objectives simultaneously; and our algorithms can be used for this purpose. We also show that our algorithms will produce solutions that are more robust to changes in traffic pattern and random variation in the traffic matrix than even an optimal solution for either objective, when such an optimal solution is produced by an algorithm that is unaware of this duality of objectives (e.g., by simply solving an ILP model). Finally, we have shown that our Min-Max heuristic consistently produces solutions that are close to the Min-Max optimal, *as well as close to the Overall optimal*. For star networks, unless there exist specific patterns in the traffic that can be exploited, this appears to be the best practical design approach. Using our greedy heuristic is an easy way to get an estimation of whether good values can be achieved simultaneously for a given problem instance.

Our algorithms can be extended to cope with more complicated network topologies with only one switching node. In this case, the virtual topology will be like a star. For larger-sized networks, decomposition approaches can be applied to break them into small pieces, and use the star approach to solve each piece hierarchically. This is part of our ongoing research in this area.

References

[1] Eric Angel, Evripidis Bampis, and Fanny Pascual. Traffic grooming in a passive star WDM network. *SIROCCO 2004*, June 2004.

[2] B. Chen, G. N. Rouskas, and R. Dutta. Traffic grooming in WDM rings with the Min-Max objective. In *Proceedings of Networking 2004*, volume LNCS 3042, pages 174–185, May 2003.

[3] Bensong Chen, Rudra Dutta, and George N. Rouskas. Traffic grooming in WDM star networks. Technical Report TR-2004-16, Dept. of Computer Sciences, NC State University, North Carolina, USA, May 2004.

[4] L-W. Chen and E. Modiano. Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters. In *Proceedings of IEEE INFOCOM*, 2003.

[5] Hongsik Choi, Nikhil Garg, and Hyeong-Ah Choi. Minimal delay traffic grooming for WDM star networks. *Proceedings of OPTICOMM 2003*, October 2003.

[6] T. Cinkler. Traffic and λ grooming. *IEEE Network*, 17(2):16–21, March/April 2003.

[7] S. Corteel, M. Valencia-Pabon, D. Gardy, D. Barth, and D. Denise. The permutation-path coloring problem on trees. *Theoretical Computer Science*, 297:119–143, 2003.

[8] R. Dutta, S. Huang, and G. N. Rouskas. On optimal traffic grooming in elemental network topologies. In *Proceeding of Opticomm 2003*, pages 13–24, October 2003.

[9] R. Dutta, S. Huang, and G. N. Rouskas. Traffic grooming in path, star, and tree networks: Complexity, bounds, and algorithms. In *Proceeding of ACM SIGMETRICS*, pages 298–299, June 2003.

[10] R. Dutta and G. N. Rouskas. On optimal traffic grooming in WDM rings. *IEEE Journal on Selected Areas in Communications*, 20(1):110–121, January 2002.

[11] R. Dutta and G. N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, November/December 2002.

[12] Rudra Dutta. *Virtual Topology Design for Traffic Grooming in WDM Networks*. PhD thesis, North Carolina State University, Raleigh, NC, August 2001.

[13] M. Esfandiari, S. Gloeckle, A. Zolfagheri, G. Clapp, J. Gannett, H. Kobrinski, V. Poudyal, and R. Skoog. Improved metro network design by grooming and aggregation of STS-1 demands into oc-192/oc-48 lightpaths. In *NFOEC 2003*, September 2003.

[14] E. Harder, S.-K. Lee, and H.-A. Choi. On wavelength assignment in WDM optical networks. In *Proceedings of the 4th International Conference on Massively Parallel Processing Using Optical Interconnections*, Geneva, Switzerland, April 1997.

[15] J.Q.Hu and B.Leida. Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks. Technical report, Manufacturing Engineering, Boston University, April 2002.

[16] Zenilton K.G. Patrocínio Jr. and Geraldo Robson Mateus. A lagrangian-based heuristic for traffic grooming in WDM optical networks. In *IEEE GLOBECOM 2003*, volume 5,1-5, pages 2767–2771, Dec 2003.

[17] V. R. Konda and T. Y. Chow. Algorithm for traffic grooming in optical networks to minimize the number of transceivers. In *IEEE Workshop on High Performance Switching and Routing*, pages 218–221, 2001.

[18] Chiewon Lee and E. K. Park. A genetic algorithm for traffic grooming in all-optical mesh networks. In *IEEE SMC 2002*, volume 7, 6-9, Oct 2002.

[19] E. Modiano and P. J. Lin. Traffic grooming in WDM networks. *IEEE Communications*, 39(7):124–129, Jul 2001.

[20] D. West. *Introduction to Graph Theory*. Prentice Hall, 1996.

[21] Hongyue Zhu, Hui Zang, Keyao Zhu, and Biswanath Mukherjee. Dynamic traffic grooming in WDM mesh networks using a novel graph model. *SPIE Optical Networks Magazine*, 4(3):65–75, May 2003.

[22] K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1):122–133, Jan 2002.

[23] K. Zhu and B. Mukherjee. A review of traffic grooming in WDM optical networks: Architectures and challenges. *Optical Networks*, 4(2):55–64, March/April 2003.

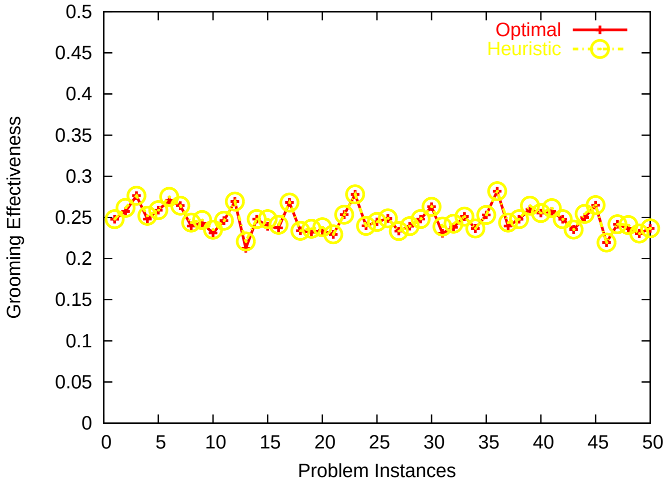

Fig. 4. Star, Random, MinMax, $N = 24$

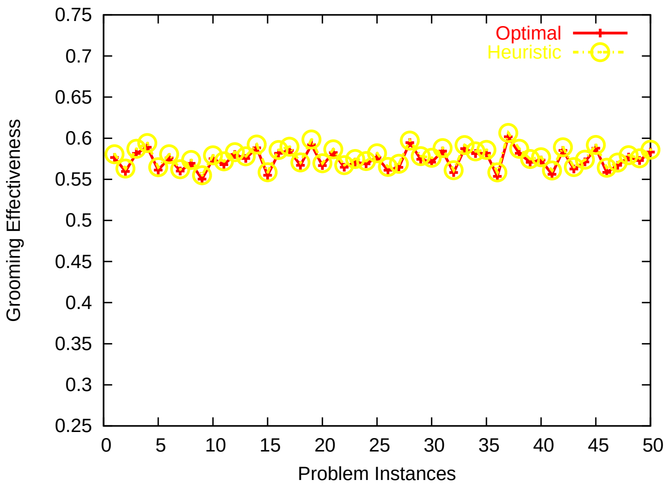

Fig. 5. Star, Random, Overall, $N = 10$

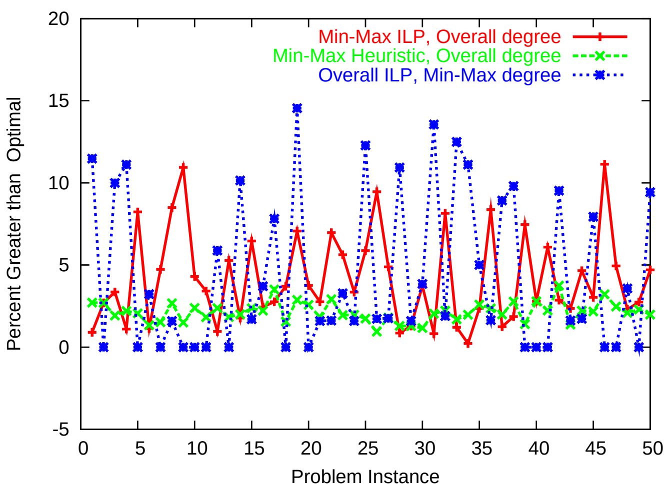

Fig. 6. Cross-Objective results, $N = 10$

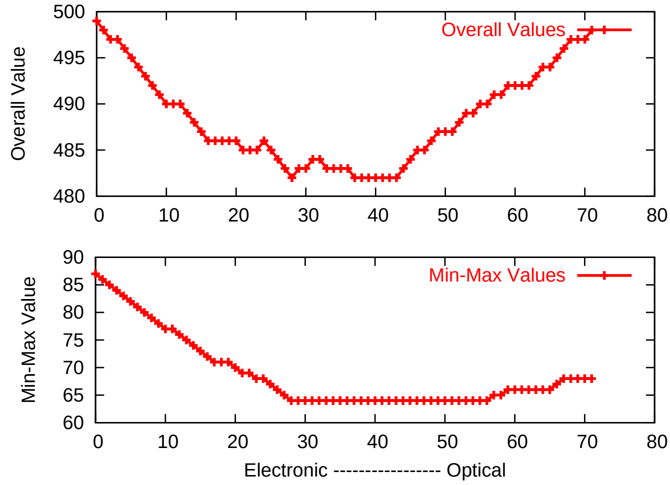

Fig. 7. Evolution of Objectives, Random Traffi c

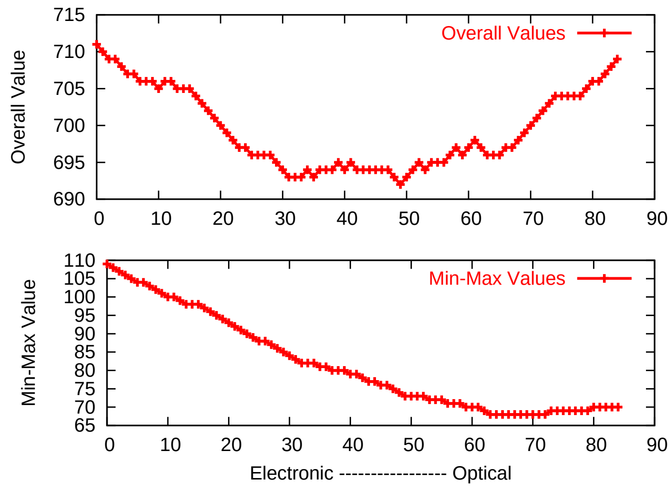

Fig. 8. Evolution, Uniform, High Load

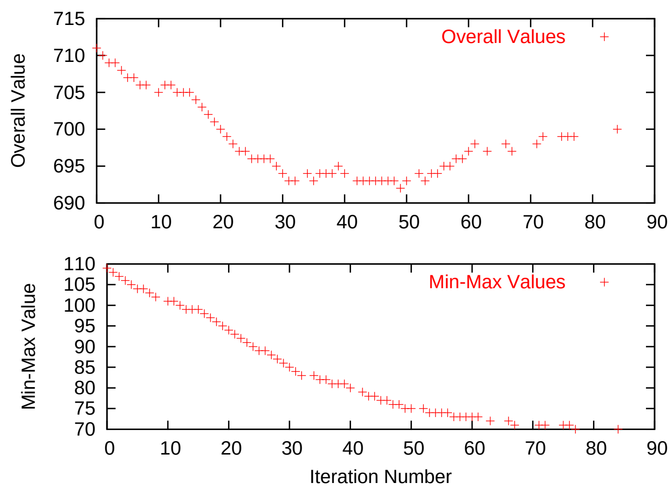

Fig. 9. Evolution, with Constraints