CSE Journal Articles

Computer Science and Engineering, Department of

12-1998

Optimizing Amplifier Placements in a Multiwavelength Optical LAN/MAN: The Unequally Powered Wavelengths Case

Byrav Ramamurthy
University of Nebraska-Lincoln, bramamurthy2@unl.edu

Jason Iness
IEEE

Biswanath Mukherjee
IEEE

Optimizing Amplifier Placements in a Multiwavelength Optical LAN/MAN: The Unequally Powered Wavelengths Case

Byrav Ramamurthy, *Student Member, IEEE*, Jason Iness, and Biswanath Mukherjee, *Member, IEEE*

Abstract—**Optical networks based on passive-star couplers and employing WDM have been proposed for deployment in local and metropolitan areas. These networks suffer from splitting, coupling, and attenuation losses. Since there is an upper bound on transmitter power and a lower bound on receiver sensitivity, optical amplifiers are usually required to compensate for the power losses mentioned above. Due to the high cost of amplifiers, it is desirable to minimize their total number in the network. However, an optical amplifier has constraints on the maximum gain and the maximum output power it can supply; thus, optical amplifier placement becomes a challenging problem. In fact, the general problem of minimizing the total amplifier count is a mixed-integer *nonlinear* problem. Previous studies have attacked the amplifier-placement problem by adding the "artificial" constraint that all wavelengths, which are present at a particular point in a fiber, be at the same power level. This constraint simplifies the problem into a solvable mixed-integer *linear* program. Unfortunately, this artificial constraint can miss feasible solutions that have a lower amplifier count but do not have the equally powered wavelengths constraint. In this paper, we present a method to solve the minimum-amplifier-placement problem, while avoiding the equally powered wavelength constraint. We demonstrate that, by allowing signals to operate at different power levels, our method can reduce the number of amplifiers required.**

Index Terms— **Amplifier placement, local area network/metropolitan area network, linear/nonlinear programming, optical network, optimization, passive star, WDM.**

I. INTRODUCTION

A. Network Environment

THE focus of this paper is on a class of the next-generation optical local area networks or metropolitan area networks (LANs/MANs) which span distances from fewer than a kilometer to a few tens of kilometers and which provide loop-free communication paths between all source–destination pairs.[1] A large-distance version of such a network is depicted in Fig. 1, and it consists of $N = 63$ stations and $M = 4$ passive optical star couplers ("stars"), such that each star is connected to other stars and/or stations via two unidirectional fiber links. The passive-star coupler provides a broadcast facility, but it must also be of the "nonreflective" type (to be elaborated below) in order to prevent loops in the network.

Our study will consider the case where each station in the network has a fixed-wavelength transmitter and is set to operate on its own unique wavelength channel. Each station either has a tunable receiver or a receiver array in order to receive signals from all of the other stations. The objective is to ensure that a station's transmission can be received by every other station after being subject to losses and gains as the signal traverses through different parts of the network. The network consists of optical stars that are nonreflective. A nonreflective star consists of pairs of inputs and outputs, and each output carries all of the wavelengths that were incident on all of the inputs except for the wavelengths that were carried on its own paired input (see Fig. 2 for an example). Such stars have been employed in the Level-0 All-Optical Network (AON) [2]. Nonreflective stars are needed in order to avoid interference due to loops ("echoes") in the network. A star in the network with k input fibers and k output fibers operates such that the power on each wavelength on an input fiber is divided evenly among the other $k - 1$ output fibers. This is referred to as the *splitting loss* at a star. Note that the splitting loss can be different for different-sized stars in the network.

As the sample network in Fig. 1 shows, these networks can be deployed as part of a MAN. We require that each signal (wavelength) be received at all of the other receivers at a power level greater than a station's receiver sensitivity level, denoted by p_{sen}. However, apart from the splitting loss due to the stars mentioned above, there is signal attenuation on the fibers given by the parameter α dB/km. Even though attenuation losses for fiber are relatively low (approximately 0.2 dB/km loss[2]) compared to other transmission media, larger networks (MANs) and networks with numerous splitting/coupling losses will require amplification to allow a transmitted signal to reach the receivers at a detectable level. The constraints on the

[1] Such networks have been referred to in the literature as access networks, passive optical networks (PON's) [22], etc.

[2] The 0.2-dB loss per kilometer of fiber is close to the absolute minimum due to the fundamental limits of Rayleigh scattering loss and infrared material loss.
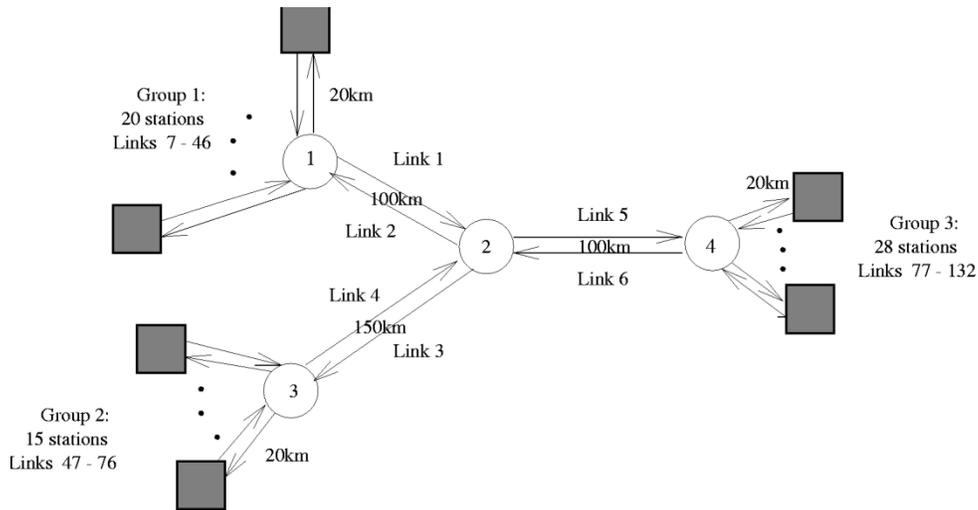
Fig. 1.   Example of a passive-star-based optical metropolitan area network (slightly modified version of the one used in [14]).
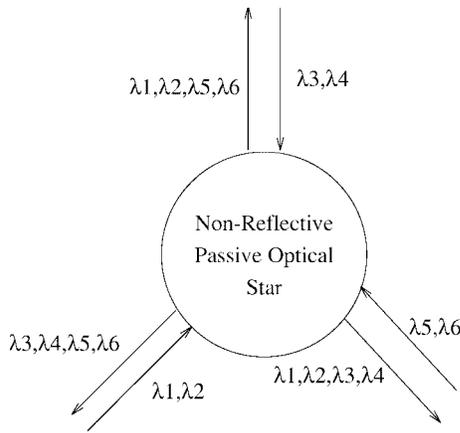

Fig. 2.   Example of a nonreflective star.

system are shown in Table I, along with typical values for each parameter. $P_{\mathrm{NONLIN,max}}$ defines the power level, in a fiber, above which a signal encounters significant nonlinear effects. However, the total power at any point in the network is usually bounded by a lower value $P_{\max}$, which is the maximum output power of an amplifier and a transmitter. P_{sat} is the internal saturation power of the optical amplifier. $G_{\max}$ is the maximum small-signal gain of the optical amplifier. These parameter values[3] (last column of Table I) will be used in our illustrative numerical examples in Section III.

We remark here that the value of the parameter p_{sen} can be chosen (by the user) such that it remains much higher than the noise levels at the intermediate amplifiers and at the receiver. The value of p_{sen} can also be changed iteratively, after the placement of the amplifiers, in order to obtain the desired signal-to-noise ratio (SNR) [3] or the bit-error rate (BER) at each of the receivers, following the approach in [14] and [15]. Thus, in this study, we do not consider system factors such as amplifier's amplified stimulated emission (ASE) noise and crosstalk at the receivers, explicitly; these factors are assumed to be implicitly incorporated in the parameter p_{sen}.

---

[3] The insertion loss at an amplifier can be included indirectly by requiring a higher small-signal gain and hence is omitted.

B. Problem Definition

In the network setting described above, it is important to quantify the minimum number of amplifiers required to operate the network and to determine their exact placements in the network. In such a network, when signals on different wavelengths originating from different locations in the network arrive at an amplifier, their power levels could be very different. This phenomenon is known as the *near–far effect* and it results in inefficient utilization of the individual amplifier. The difference in power levels of the input wavelengths can significantly limit the amount of amplification available, since the higher powered wavelengths could saturate the amplifier and limit the gain seen by the lower powered wavelengths. Fig. 3 shows, at some location on a fiber link, a case where three wavelengths have different power levels and a case where the three wavelengths have the same power level. In Fig. 3(a), the total power is 4.5 μW, and in Fig. 3(b), it is 3 μW. Since the per-wavelength amplifier sensitivity is 1 μW (-30 dBm), in both cases an amplifier will be required before the signals suffer any more attenuation. However, since an amplifier has a limited total output power, the amount of achievable gain is greater when the total input power is less. This would allow the signals in Fig. 3(b) to receive a higher gain than the signals in Fig. 3(a). Also, allowing signals in the same fiber to be at different power levels changes the minimal-amplifier-placement problem from a mixed-integer linear program (MILP) [20] into a mixed-integer nonlinear program, as we shall show later in this paper.

Previous optical amplifier-placement schemes [14], [20] bypassed these problems by restricting all of the wavelengths at any given point in a fiber to be at the same power level. Unfortunately, requiring wavelengths to be at the same power level often forces the designer to add more amplifiers than the minimum necessary in order for the receivers to receive signals at or above the receiver sensitivity level. Since each optical amplifier costs around \$25 000, every attempt should be made to minimize their number in the network. It is also desirable to reduce the number of amplifiers used in the network based on noise, maintenance, and fault-tolerance considerations.

TABLE I
IMPORTANT PARAMETERS AND THEIR VALUES USED IN THE AMPLIFIER-PLACEMENT ALGORITHMS

| Parameter	Description	Range	Value used
p_{sen}	Minimum signal power at receiver and the amplifier sensitivity level	-30 dBm at 1 Gbps	-30 dBm
G_{max}	Maximum small-signal gain	≤ 25 dB - Multi-Quantum Well (MQW) Amplifier [18] ≤ 30 dB Erbium-Doped Fiber Amplifier (EDFA) [1]	20 dB
$P_{NONLIN,max}$	Maximum total power in fiber	10-50mW	10 mW
P_{max}	Maximum total output power of amp and transmitter		1 mW
P_{sat}	Internal saturation power of the amp		1.298 mW
α	Fiber attenuation		0.2 dB/km
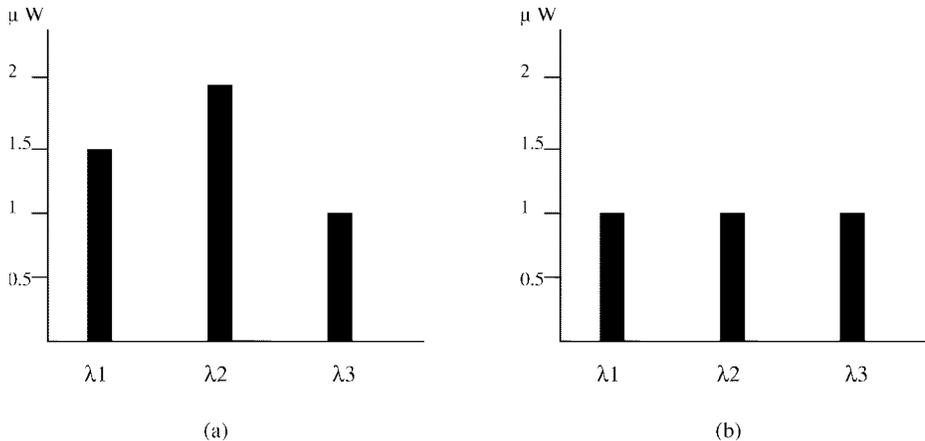
Fig. 3.   Two examples of powers on three wavelengths passing through a fiber.
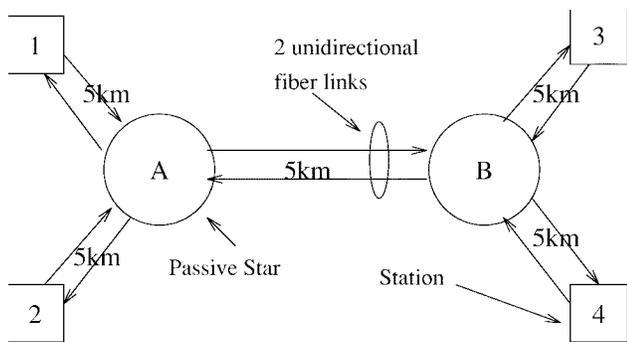

Fig. 4.   Simple two-star network that needs no amplifiers to operate.

Our study was motivated by the network in Fig. 4. For reasonable network parameters, this network can operate without using any amplifiers. However, if the power levels for all wavelengths must be equal on any given link, as required by the MILP approach in [20], then an amplifier (on one of the links between stars A and B) will have to be added to the network. This is because, if we fix the output power of star A to be some value x, then the signals from stations 3 and 4 must reach star B with an output power higher than x. Without an amplifier, signals from stations 1 and 2 reach star B at a power less than x, which means that wavelengths on the link from star B to station 3 (and similarly on the link from star B to station 4) will have unequal powers. Therefore, requiring equal power on all wavelengths adds an unnecessary amplifier to this

network. As we shall soon observe, allowing wavelengths to be at unequal powers eliminates the need for any amplifiers in this network.

In this paper, we propose a scheme that minimizes the number of amplifiers for the network setting described in [14] without the restriction that wavelengths in the same fiber be at the same power level. The method works as follows.

1) Determine whether or not it is possible to design the network taking into consideration the limitations of the devices (e.g., the power budget of the amplifiers).
2) Generate a set of constraints to accurately describe the problem setting, which turns out to be a nonlinear program.
3) Pass the set of constraints to a nonlinear solver, such as C code for Feasible Sequential Quadratic Programming (*CFSQP*) [19], in order to solve for the minimum number of amplifiers needed for the entire network
4) Determine the exact placements of the optical amplifiers.

Numerical examples will show that this network-wide optimization method without the equal-power constraint often results in solutions that require fewer amplifiers than the solutions in [14], [20].

C. Amplifier Gain Model

Currently, we employ a simplified model for the gain of a generic optical amplifier. The simplifying assumptions are that the amplifier has a flat gain over the wavelengths being
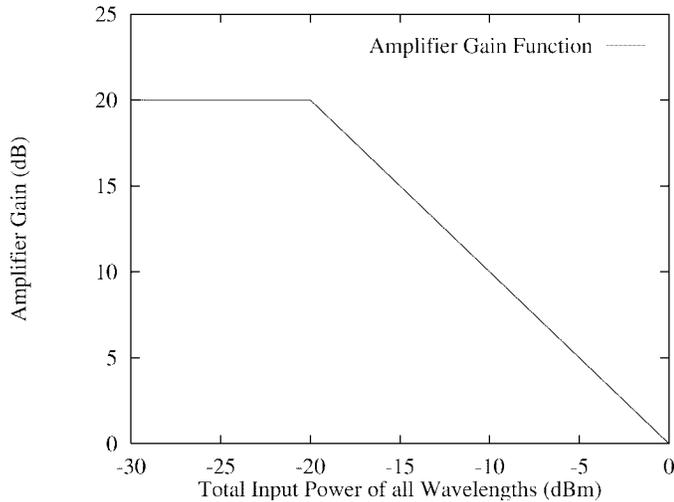
Fig. 5. Original amplifier gain model approximations used in previous studies [14].

Fig. 6. More accurate amplifier gain model used in this study.

amplified and that the amplifier gain is homogeneously broadened.[4] A flat gain can be achieved through various techniques such as: 1) notch filters [23]; 2) different pump laser powers [7]; 3) Mach–Zehnder filters [12]; and 4) demultiplexers and attenuators [5]. However, assuming that optical amplifiers are homogeneous is an approximation. For each specific amplifier (Erbium-doped fiber amplifier (EDFA), semiconductor optical amplifier (SOA), etc.), we need to develop a gain model depending on its degree of homogeneity in order to accurately solve the amplifier-placement problem.

Based on the above assumptions, the gain model for our amplifiers is given by (from [21])

$$\frac{P_{\text{in}}}{P_{\text{sat}}} = \frac{1}{G-1} \ln\left(\frac{G_0}{G}\right) \qquad (1)$$

where P_{in} is the total input power (across all wavelengths) to the amplifier in milliwatts, P_{sat} is the internal saturation power in milliwatts, G is the actual gain achieved (in absolute scale, *not* decibels), and G_0 is the small-signal gain (which is the gain achievable for small values of input power when the amplifier does not saturate, again in absolute scale). Since the formula for G is not an explicit formulation, we use an iterative method to solve for the value of G. Our amplifier model has been designed into our solution as a generic gain module that can be easily replaced when a more-accurate model for a specific amplifier is used.

Previous studies [14] used the gain model in Fig. 5. In this model, it is assumed that the full small-signal gain of the amplifier is realizable until the point at which the amplifier output becomes power limited. At this point, the amplifier is assumed to enter saturation and the gain starts to drop. This "point" of saturation[5] occurs in the example of Fig. 5 at a total input power of −20 dBm (0.01 mW). At lower input powers, the amplifier is assumed to be able to supply the
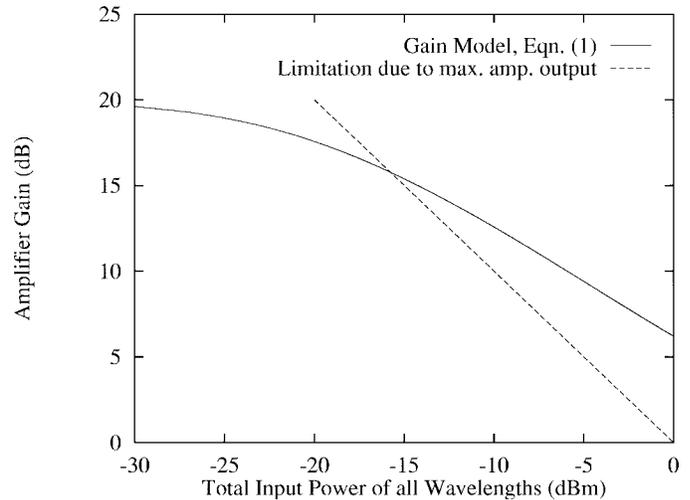
full small-signal gain of $G_{\text{max}} = 20$ dB. The more-accurate model (1), which is used in this paper (and also in [20]), is plotted in Fig. 6 and shows how saturation does not happen at a specific point, but is really a continuous effect. In fact, we note that, even for small input powers, the amplifier is not able to supply the full small-signal gain of $G_{\text{max}} = 20$ dB. The numerical differences between the models are not huge, but are significant enough so that a network designer may have thought a design was feasible (based on the model in Fig. 5), when, in fact, it may not satisfy the design specifications (based on the more-accurate model in Fig. 6). Notice, also, that there is a limit on the total available output power (P_{max}) from the amplifier. This limit is shown as the dashed line in Fig. 6. Hence, the gain curve used in this paper follows the curved line for low input powers and the straight dashed line in Fig. 6 for higher powers.

II. SOLUTION APPROACH

Given a network as in Fig. 1, we would like to minimize the number of amplifiers used in the network without violating the device capabilities and constraints. Throughout this paper, we assume that the stars are connected together in the form of a tree and that all neighbors have two unidirectional links connecting each other. A mathematical formulation of the problem is provided in Section II-A. Unfortunately, the resulting mixed-integer nonlinear optimization problem is extremely difficult to solve. Hence, we carefully avoid the integral constraints by modifying the formulation, specifically the objective function, and solve the resulting nonlinear optimization problem. The description of the solution strategy is provided in Section II-B. The output from the solver is fed to an Amplifier-Placement Module which outputs the exact positions and gains of the amplifiers. The functionality of the Amplifier-Placement Module is described in Section II-C.

A. Formulation

In this section, the amplifier-placement problem is formulated as a mixed-integer nonlinear (constrained) optimization

---

[4] By homogeneous broadening, we mean that a single high-powered wavelength, which saturates the amplifier, can bring down the gain available for all of the wavelengths uniformly.

[5] The "true" point of saturation occurs when the amplifier gain is reduced by 3 dB from its maximum [1].

problem. First, the notation used in the formulation is introduced, and then the constraints and objective functions are described.

1) Device Parameters:

- p_{sen} = Minimum power required on a wavelength for detection in decibels referred to 1 mW, i.e., in dBm. This represents both the receiver sensitivity level and the amplifier sensitivity level, which have been assumed to be equal.
- $P_{\max}$ = Maximum power available from an amplifier in milliwatts = Maximum power of a transmitter in milliwatts.

It is not necessary that the maximum amplifier output and transmitter powers be identical. For simplicity, we have assumed them to be equal.

- $G_{\max}$ = Maximum (small-signal) amplifier gain in decibels.
- α = Signal attenuation in decibels per kilometer.

2) Problem Variables: This section introduces the variables used in the problem formulation. Note that, among the variables representing the power levels, those beginning with lowercase ($p_l^{\text{min,beg}}, p_{x,i}, p_i^{\text{xmit}}$) are measured in decibels referred to 1 mW and those with uppercase ($P_l^{\text{beg}}, P_l^{\text{min}}$) in milliwatts. Also, the variables in lowercase represent the per-wavelength power levels, whereas the ones in uppercase represent the aggregate power over all the wavelengths on the respective link.

- N = number of access stations in the network = number of wavelengths in the network.
- M = number of stars in the network.
- L = number of links in the network = $2 \times (N + M - 1)$.

Note that stars are identified by the indices $1, 2, \cdots, M$ and stations by the indices $M+1, M+2, \cdots, M+N$. As we shall soon observe, this provides notational convenience when we refer to the source/destination of a link, irrespective of whether it is a station or a star. Also, the wavelengths in the network are identified by the indices $M+1, M+2, \cdots, M+N$ of the source stations. We associate the following parameters with each link $l, 1 \le l \le L$:

s_l	source of link $l, 1 \le s_l \le (M+N)$;
d_l	destination of link $l, 1 \le d_l \le (M+N)$;
Λ_l	set of powered wavelengths carried by link l;
n_l	number of amplifiers on link l;
L_l	length of link l in kilometers;
SG_l	actual total <u>S</u>upplied <u>G</u>ain on link l in decibels;
$p_l^{\text{min,beg}}$	power level of the least-powered wavelength arriving at link l, in decibels referred to 1 mW;
P_l^{beg}	total power at the beginning of link l, in milliwatts;
P_l^{min}	total power on link l when all signals are $\ge p_{\text{sen}}$ and at least one wavelength is equal to p_{sen}, in milliwatts;
$gmax_l$	maximum gain available from an amplifier on link l, in decibels.

Consider the star $i, 1 \le i \le M$.

D_i in-degree of star i = out-degree of star i.

$p_{x,i}$ power of wavelength x at the output of star i, in decibels referred to 1 mW.

Consider the station $i, (M+1) \le i \le (M+N)$.

- p_i^{xmit} is the transmitted power of wavelength i at station i, in decibels referred to 1 mW.

3) Useful Functions: The following functions allow conversion between the milliwatt (regular) and decibels referred to 1 mW (log) scales:

$$ToDB(\xi) = 10 \cdot \log_{10}(\xi)$$
$$ToMW(\xi) = 10^{\xi/10}.$$

They are used to express the constraints conveniently in the appropriate scale.

4) Basic and Nonbasic Variables: Given a network, the values of the topology-specific variables $N, M, L, s_l, d_l, \Lambda_l, L_l$, and D_i are fixed, irrespective of the amplifier-placement algorithm chosen. The only basic variables used in the formulation are p_i^{xmit}, SG_l, and n_l. Note that the variables $P_l^{\text{beg}}, p_l^{\text{min,beg}}, P_l^{\text{min}}, p_{x,i}$, and $gmax_l$ are nonbasic variables and can be expressed in terms of the basic variables as follows.

For link l, the source of which is a star, i.e., $1 \le s_l \le M$, we have

$$p_l^{\text{min,beg}} = \min_{x \in \Lambda_l} p_{x,s_l} \qquad (2)$$

and we also have

$$P_l^{\text{beg}} = \sum_{x \in \Lambda_l} ToMW(p_{x,s_l}). \qquad (3)$$

For link l, the source of which is a station, i.e., $(M+1) \le s_l \le (M+N)$, we have

$$p_l^{\text{min,beg}} = p_{s_l}^{\text{xmit}} \qquad (4)$$

and we also have

$$P_l^{\text{beg}} = ToMW(p_{s_l}^{\text{xmit}}). \qquad (5)$$

For any link l, the total power drops to its minimum level when at least one of the wavelengths is equal to the sensitivity level (p_{sen}). Hence, on link l, starting with an aggregate power level P_l^{beg}, when the weakest signal is at a power level $p_l^{\text{min,beg}}$, after appropriate scale changes, we have

$$P_l^{\text{min}} = ToMW(ToDB(P_l^{\text{beg}}) - (p_l^{\text{min,beg}} - p_{\text{sen}})). \qquad (6)$$

The equation above is best explained with an example. Consider a link containing three wavelengths λ_1, λ_2, and λ_3. Suppose the power levels on these wavelengths at the beginning of the link were 2 μW (-26.99 dBm), 3 μW (-25.23 dBm), and 5 μW (-23.01 dBm), respectively. Now, the weakest signal is on wavelength λ_1, and from (2), we have $p_l^{\text{min,beg}} = -26.99$ dBm. Also, from (3), we have $P_l^{\text{beg}} = 2 \mu W + 3 \mu W + 5 \mu W = 10 \mu W$. Now, with a link attenuation (α) of 0.2 dB/km, and a sensitivity level (p_{sen}) of -30 dB (1 mW), this group of wavelengths can travel $(p_l^{\text{min,beg}} - p_{\text{sen}})/\alpha = (-26.99 + 30)/0.2 = 15.05$ km before the power of wavelength λ_1 drops below p_{sen}. At this point, the powers on

the three wavelengths are 1 μW, (-30 dBm), 1.5 μW, (-28.24 dBm), and 2.5 μW, (-26.02 dBm), respectively. Hence, the aggregate "minimum" power ($P_l^{\min}$) is $1 + 1.5 + 2.5\,\mu$W $= 5\,\mu$W. This value can be derived from the above equation, since

$$ToMW(ToDB(10\,\mu\text{W}) - (-26.99 + 30))$$
$$= ToMW(-20 - 3.01) = 5\,\mu\text{W}.$$

For links from stations to stars, i.e., $(M + 1) \leq s_l \leq (M + N)$ and $1 \leq d_l \leq M$, we have

$$p_{s_l,d_l} = p_{s_l}^{\text{xmit}} + SG_l - \alpha \cdot L_l - ToDB(D_{d_l} - 1). \qquad (7)$$

For links between stars, i.e., $1 \leq s_l, d_l \leq M$, we have

$$\forall_{x \in \Lambda_l}\ p_{x,d_l} = p_{x,s_l} + SG_l - \alpha \cdot L_l$$
$$- ToDB(D_{d_l} - 1). \qquad (8)$$

For any link l,

$$gmax_l = G(P_l^{\min}, G_{\max}, P_{\text{sat}}). \qquad (9)$$

We note that various amplifier gain models can be used to obtain this function G.

5) Constraints:

Inequalities: Consider the link $l, 1 \leq l \leq L$. The powers on each of the wavelengths at the beginning of the link l should be at least the sensitivity level, p_{sen}. This can be ensured by requiring that the *weakest* signal has a power level of at least p_{sen} as follows:

$$p_l^{\min,\text{beg}} \geq p_{\text{sen}}. \qquad (10)$$

The powers on each of the wavelengths at the end of each link l should be at least p_{sen}. This is to enable the receivers to detect the signals correctly. Thus,

$$p_l^{\min,\text{beg}} + SG_l - \alpha \cdot L_l \geq p_{\text{sen}}. \qquad (11)$$

The above inequalities [(10) and (11)] ensure that the signal powers remain at or above p_{sen} everywhere along the fiber links and throughout the network.

There are upper limits on the maximum power carried by all the signals in a link. This value $P_{\max}$ is the same for transmitters and amplifiers and, hence, at the beginning of link l, we have

$$P_l^{\text{beg}} \leq P_{\max}. \qquad (12)$$

Similarly, at the end of the link l, we have

$$ToDB(P_l^{\text{beg}}) + SG_l - \alpha \cdot L_l \leq ToDB(P_{\max}). \qquad (13)$$

Since we need to divide the total supplied gain SG_l among the n_l amplifiers on link l, we have

$$SG_l \leq gmax_l \cdot n_l. \qquad (14)$$

However, the gain SG_l should require no fewer than n_l amplifiers; thus,

$$SG_l > gmax_l \cdot (n_l - 1). \qquad (15)$$

Integrality Constraints: Consider the link $l, 1 \leq l \leq L$. The number of amplifiers, n_l, on any link l, is an integral value. Hence, we require that

$$n_l \quad \text{is an integer}. \qquad (16)$$

6) Objective function: Minimize

$$\sum_{l=1}^{L} n_l. \qquad (17)$$

7) Complexity: The only basic variables used in the formulation are p_i^{xmit}, SG_l, and n_l. The others can be computed either beforehand from the topology or at run time as a function of the basic variables. Hence, we have the following:

- number of variables $= 2 \cdot L + N$;
- number of integer constraints $= L$;
- number of nonlinear inequalities $= 6 \cdot L$.

8) Reasons for Nonlinearities: The approach presented in this paper differs from the one in [20] in that it allows the different wavelengths on a link to be at different power levels. Whereas the method in [20] needed to place amplifiers whenever all the wavelengths on the link were at their lowest power level, now the placement of the amplifier is constrained by the *weakest* signal on the link. Hence, on each link, we need to identify the wavelength coming in with the lowest power level ($p_l^{\min,\text{beg}}$). This introduces a nonlinear term in the formulation [(2)]. Moreover, the maximum gain ($gmax_l$) available at an amplifier on a link is dependent on the precise mix of the power levels on its incoming wavelengths. This computation cannot be performed off-line and results in nonlinear constraints [see (14) and (15)].

B. Solver Strategies

The mixed-integer nonlinear optimization problem resulting from Section II-A is an extremely difficult one to solve and is highly computation intensive. Surveys on techniques employed in solving such problems can be found in [9] and [10]. For such highly nonlinear, general, integer programming problems, branch-and-bound-based methods which are employed in integer linear programming problems do not work well [16]. In order to reduce the computation complexity, we choose to eliminate the integral constraints altogether. In our case, this can be done by removing the variables n_l from the formulation and, hence, the constraints in (14) and (15) disappear. A similar approach is described in [16], where the integrality of variables is expressed as an additional constraint to the original problem.

So, we define a *new objective function.*
Minimize

$$\sum_{l=1}^{L} SG_l / gmax_l \qquad (18)$$

which is close to the original one, since $n_l = \lceil SG_l / gmax_l \rceil$. The *starting point* of the problem space is especially important for this nonlinear search. We initialize the basic variables of the problem, namely, SG_l and p_i^{xmit} such that

$$SG_l = 0$$
$$p_i^{\text{xmit}} = ToDB(P_{\max})$$

i.e., the network is initialized to a state when all the transmitters are operating at their highest powers and all of the links have zero gain. However, we could also use the solution from [20] as a feasible starting point. Since the new objective function is not identical to the original one for the integral case, the solver might end up minimizing the function $SG_l/gmax_l$ and not the number of amplifiers in the network. To handle this situation, we adopt a *nonintrusive measurement* approach, where, at every feasible point along the search path to the optimum solution taken by the nonlinear program solver, we evaluate the original objective function and remember the point in the search space which resulted in the minimum value for the original objective function thus far.

The ensuing heuristic search has the following interesting properties.

1) It contains significantly fewer variables and constraints. In fact, it has only:

   a) $L + N$ variables;
   b) $4L$ inequalities;
   c) *zero* integer constraints.

2) All the constraints and the objective function are easily differentiable. Hence, the gradients can be fed to the nonlinear program solver to aid it in its search for the optimum solution.

The nonlinear program solver, $CFSQP$ [19], which we have used for this study, consists of a set of C functions for the minimization of the maximum of a set of smooth objective functions, subject to general smooth constraints. If the initial guess provided by the user is infeasible for some inequality constraint or some linear equality constraint, $CFSQP$ first generates a feasible point for these constraints; subsequently, the successive iterates generated by $CFSQP$ all satisfy these constraints. Nonlinear equality constraints are turned into inequality constraints and the objective function is replaced by an exact penalty function which penalizes nonlinear equality constraint violations only. Given a feasible iterate x, the basic SQP direction d^0 is first computed by solving a standard quadratic program using a positive definite estimate H of the Hessian of the Lagrangian. d^0 is a direction of descent for the objective function; it is almost feasible in the sense that it is at worst tangent to the feasible set if there are nonlinear constraints and it is feasible otherwise. The user has the option of either requiring that the objective function (penalty function if nonlinear equality constraints are present) decrease at each iteration after feasibility for nonlinear inequality and linear constraints has been reached (monotone line search), or requiring a decrease within at most a few, say three, iterations (nonmonotone line search). The user must provide functions that define the objective function and constraint functions, and may either provide functions to compute the respective gradients or require that $CFSQP$ estimate them by forward finite differences. Additional details on the $CFSQP$ solver can be found in [13].

$CFSQP$ provides the user with some flexibility in the choice of algorithms and values for various parameters. We describe below some of our choices and characteristics unique to the amplifier placement problem at hand. In the formulation presented above, we do not have *any* nonlinear equality constraints and there is only one objective function [(18)]. This enables $CFSQP$ to employ the objective function directly (and not any penalty function) in its search. We require that $CFSQP$ use a nonmonotone line search [8], forcing a decrease of the objective function within at most three iterations. The gradients of the objective and some constraint functions are estimated by $CFSQP$ using forward finite differences. When there are no nonlinear equality constraints (as in our case), $CFSQP$ terminates when the norm for the Newton search direction d^0 falls below ϵ which is taken to be 10^{-8}.

While eliminating integer variables greatly simplifies the problem, there are, however, limitations to this approach, and they are discussed below.

1) *Local Minima*: The nonlinear program solver might terminate at a point corresponding to a local minimum for the objective function. This happens, for example, when the starting point corresponds to the Linear Program solution (see Table II and the examples in Figs. 1 and 11).

2) *Feasible Point Generation*: When the starting point is infeasible, subject to the constraints, the solver may not be able to locate a feasible point in the problem space. With $CFSQP$, this problem can be fixed by using a different quadratic programming solver to generate the feasible point. However, finding a feasible point becomes increasingly difficult as the number of network elements grows (i.e., more network elements means more variables).

3) *Integer Variables*: The nonlinear program solver ($CFSQP$), which we used in this study, is not well suited to handle integer variables. Hence, its results for this problem could be improved upon by using specialized mixed-integer nonlinear program solvers.

The output of the nonlinear program solver is fed to the Amplifier-Placement Module, which is described next.

C. Amplifier-Placement Module

The module uses the values of SG_l and p_i^{xmit} output by the nonlinear program solver to determine the exact location and gain of the amplifiers in the network. It operates on a link-by-link basis as follows. It computes the maximum value of the gain available from each amplifier on a link l ($gmax_l$) using (9) and, hence, the number of amplifiers (n_l) required on that link.[6] It also computes the power levels of the different wavelengths at the output of the stars ($p_{x,i}$). Several methods of splitting the gain (SG_l) among the n_l amplifiers on a link l are possible. We describe two methods below—the As Soon As Possible (ASAP) method and the As Late As Possible (ALAP) method.

The ASAP method for amplifier placement operates as follows. For all but the last amplifier on a link, this method places an amplifier on a link as soon as the input power is low

---

[6] The nonlinear program solver could possibly come up with a solution with negligible gains (SG_l) at certain links. We use $n_l = \lceil SG_l/gmax_l - \sigma \rceil$, where σ is a small number to handle this situation ($\sigma = 0.01$ in our numerical examples).

TABLE II
RELATIVE PERFORMANCE OF THE VARIOUS AMPLIFIER-PLACEMENT SCHEMES. A "*" IN COLUMN 4 INDICATES THAT THE NONLINEAR-PROGRAM SOLVER (NLP) COULD NOT PERFORM BETTER THAN THE LP SOLUTION, EVEN WHEN IT WAS GIVEN MULTIPLE FEASIBLE STARTING POINTS, INCLUDING THE SOLUTIONS FOUND IN [14] AND [20]. COLUMN 6 SHOWS THE TOTAL CPU TIME TAKEN BY THE NONLINEAR SOLVER RUNNING ON AN OTHERWISE-UNLOADED DEC 5000/240 TO SOLVE EACH PROBLEM

Network	Link-by-link method [14]	LP method [20]	NLP method (this work)	Absolute lower bound [11]	CPU time for NLP (this work)
Simple 2 star (Figure 4)	6	1	0	0	2s
Tree (Figure 9)	44	14	0	0	49s
MAN (Figure 10)	38	6	2	2	4h 6m 42s
Scaled-up MAN (Figure 11)	48	16	16*	14	10h 7m 42s
Scaled-down MAN (Figure 12)	38	4	0	0	1m 19s
Previous MAN (Figure 1)	79	77	77*	74	9m 50s
Denser MAN (Figure 13)	50	7	4	4	2h 5m 19s
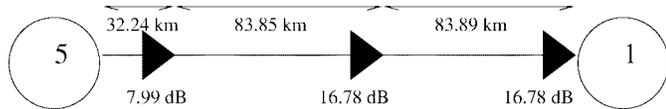
Fig. 7.   Amplifier placement using the ASAP method.
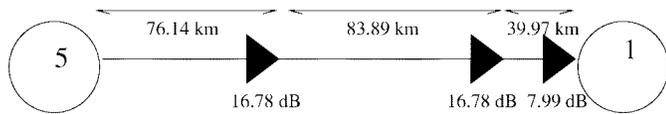
Fig. 8.   Amplifier placement using the ALAP method.

TABLE III
EXACT AMPLIFIER PLACEMENTS FOR THE NETWORK DEPICTED IN FIG. 10

Link (Star→Star)	LP (gain and distance from start of link)	NLP (gain and distance from start of link)
1→5	Gain 11.52 dB at 20.00 km	Gain 0.91 dB at 20.00 km
5→1	Gain 10.80 dB at 20.00 km	Gain 0.91 dB at 20.00 km
2→5	Gain 16.55 dB at 1.00 km	
5→2		
3→5	Gain 15.76 dB at 4.96 km	
5→3		
4→5	Gain 9.76 dB at 10.00 km	
5→4	Gain 9.01 dB at 10.00 km	

enough to allow the *maximum* gain, and for the last amplifier on a link, it places the amplifier as soon as the input power is low enough to allow the *remaining* gain. The ALAP method operates in a similar fashion, except that it attempts to place amplifiers as close as possible to the destination of the link. Both these methods split the total gain on the link among the amplifiers by operating all but one of them at their maximum possible gain.

The differences between the ASAP method and the ALAP method can be seen in Figs. 7 and 8, where the total gain of 41.55 dB is divided among three amplifiers. The link shown runs from star 5 to star 1 in the scaled-up MAN network (Fig. 11) and the gains shown are taken from the Linear-Program (LP) solution [20] for this network. Several other methods of splitting the gain, including equal distribution among the amplifiers on a link, are possible. The ALAP method was chosen in our study (see Table III). Further discussions on various approaches to gain splitting can be found in [17].

III. NUMERICAL EXAMPLES

The link-by-link method in [14] was designed to equalize the powers of the wavelengths in the network, as opposed to trying to minimize the number of amplifiers in the network. By forcing the powers of all wavelengths to be equal to p_{sen} at the beginning of most links (all links except those from stations to stars), the algorithm placed amplifiers simply by knowing how many wavelengths were on a link. If the number of wavelengths on a link is precomputed, this allowed the algorithm to operate on each link individually (locally) without knowing what was happening on other links. This led to a very simple amplifier-placement algorithm. Unfortunately, as was

shown in [20] and can also be seen in Table II, this approach does not minimize the number of amplifiers needed in the network. The transmitter powers can be adjusted to avoid placing amplifiers on the links which originate at a station. However, since signals on all other links start off with the minimum power (p_{sen} on each wavelength), we know that the algorithm will place an amplifier on every single link not originating at a station in the network. We note that there are $L - N$ such links in the network which originate at a star (recall that $L = $ number of links, $N = $ number of stations, and $M = $ number of stars); thus, we obtain the lower bound of $L - N = 2 \times (N + M - 1) - N = N + 2 \times (M - 1)$ on the number of amplifiers used by the method in [14]. This algorithm performs the poorest, in comparison to other placement schemes, on networks that have short links because the other algorithms can usually avoid placing an amplifier on a short link simply by exiting the originating star with enough power to traverse the short link. We show the results of this algorithm for various networks in column 2 of Table II.

The global method in [20] allowed wavelengths at the beginning of the links to be above the absolute minimum allowed, p_{sen}. However, the powers on all of the wavelengths at any given point in the network were required to be equal; this equally powered wavelengths constraint enabled the computation of the maximum gain ($gmax_l$) available on a link by knowing just the number of wavelengths on the link. The amplifier-placement problem can be formulated as a mixed-integer linear program and solved exactly. Consider a pair of adjacent stars in the network. Taking into account the attenuation loss along the links connecting the stars and the splitting losses at the stars, we require that there be at least one amplifier on either of these links. The lower bound on the number of amplifiers required using the LP method in [20] is, thus, $M - 1$, where M is the number of stars in the network.
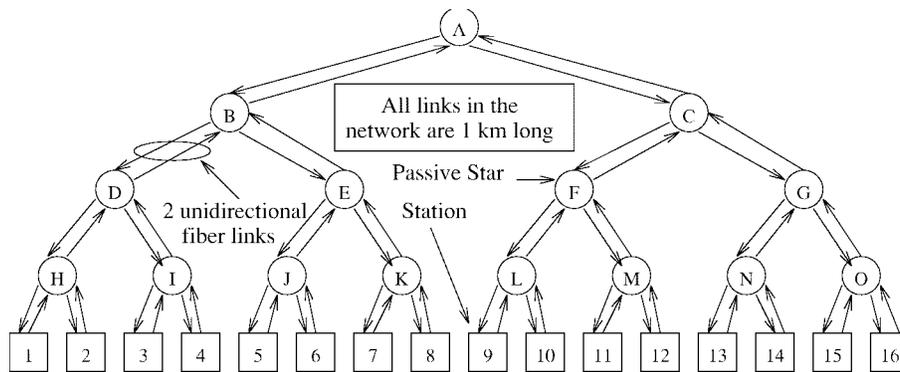
Fig. 9. Midsized tree-based network needing no amplifiers to function.

We show the results of this algorithm for various networks in column 3 of Table II (see [20] for details).

The method described in this paper (see Section II) is a global one, too; however, unlike the LP method in [20], it allows the wavelengths at any point in the network to operate at unequal powers.[7] The solution obtained to the amplifier-placement problem is not guaranteed to be the optimum because of the presence of local minima. We show the results of this algorithm for various networks in column 4 of Table II. The absolute lower bound was developed in [11] by first utilizing the number of wavelengths on each link and the physical constraints on the amplifiers to derive the maximum gain available from each amplifier on a given link. These values were then included in an LP-solvable solution to derive the lower bound on the number of amplifiers required in the network. We show the results of the lower bound computation for various networks in column 5 of Table II (see [11] for more details). Next, we compare the results of these three approaches to amplifier placement on certain sample networks (see Table II).

As mentioned earlier, the network in Fig. 4 motivated this study. While both the earlier approaches (the link-by-link method and the LP method) required a few amplifiers to operate the network, the NLP method described in this paper does not require any.

The network in Fig. 9 is the motivating network, described above, taken to the extreme. This network has many stars and yet it needs no amplifiers to function. Table II reveals that the new method was indeed able to come up with the solution of not needing any amplifiers. This is the type of network where the unequally powered wavelengths solution is clearly superior to the previous two amplifier-placement methods. Although it is arguable whether this network is realistic or not, we have presented it here in order to give the reader some insight as to the conditions in which the new method performs best.

The network in Fig. 10 is meant to be a realistic design of a MAN. This network was designed in a semirandom fashion with some heuristics to guide the design. Table II
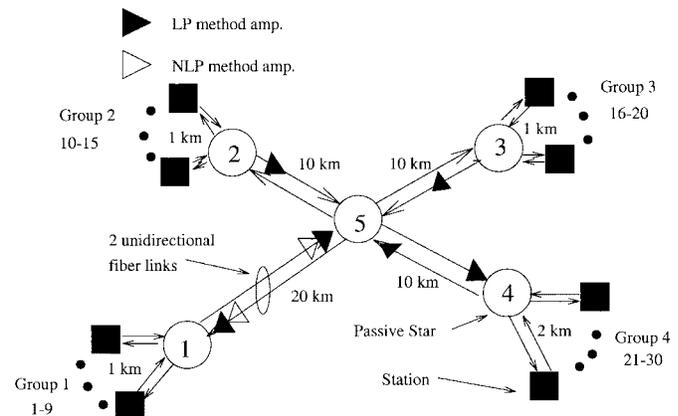

Fig. 10. A possible MAN network.

shows that the new method was able to find a solution which required fewer amplifiers than the methods in [14] and [20]. Fig. 10 also provides an insight into how the actual placements of amplifiers differ between the LP method and the NLP method. The triangles that are filled black are the locations at which the equally powered wavelengths method placed the six amplifiers it deemed necessary to operate the network. The empty, or filled-white, triangles are the locations where the unequally powered wavelengths method placed the two amplifiers it deemed necessary. The numerical information on exact gains and exact placements of the amplifiers can be seen in Table III. The power levels of the signals at the transmitters and receivers can be found in Table IV. Note that the equally powered wavelengths constraint results in more amplifiers and a higher overall gain in the network. Note also that the transmitters are unable to operate at their maximum power for the same reason. However, when wavelengths are allowed to operate at different power levels, we find that the NLP solution requires just the minimum overall gain to operate the network.

This network serves as the reference point for a study into the effects of scaling network distances up and scaling network distances down, which will be discussed below.

As previously noted in Section I, an amplifier becomes less efficient when multiple wavelengths passing through it are operated at different power levels. If a link were long

[7]Recall that, in these experiments, the NLP solver's starting point is chosen such that all transmitters are operating at the maximum power without any amplifiers in the network. From this possibly infeasible starting point, the solver reaches a feasible point for all the example networks, except for the previous MAN (Fig. 1).

TABLE IV
TRANSMITTER AND RECEIVER POWERS FOR THE NETWORK DEPICTED IN FIG. 10

| Stations	LP method		NLP method	
	Transmitter power	Receiver power	Transmitter power	Receiver power
1–9 (Group 1)	−13.02 dBm	−22.96 dBm	0.00 dBm	From G1: −9.94 dBm
				From G2: −27.58 dBm
				From G3: −26.79 dBm
				From G4: −30.00 dBm
10–15 (Group 2)	−21.82 dBm	−30.00 dBm	0.00 dBm	From G1: −27.58 dBm
				From G2: −8.18 dBm
				From G3: −23.94 dBm
				From G4: −27.15 dBm
16–20 (Group 3)	−21.82 dBm	−29.21 dBm	0.00 dBm	From G1: −26.79 dBm
				From G2: −23.94 dBm
				From G3: −7.39 dBm
				From G4: −26.36 dBm
21–30 (Group 4)	−12.61 dBm	−23.41 dBm	0.00 dBm	From G1: −29.80 dBm
				From G2: −27.15 dBm
				From G3: −26.36 dBm
				From G4: −10.80 dBm
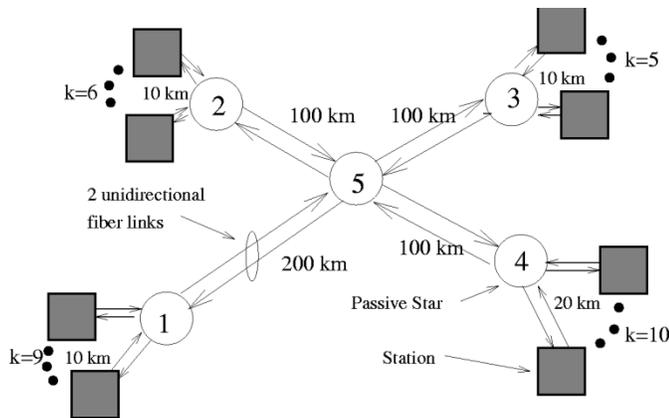
Fig. 11. A scaled-up version of the MAN network in Fig. 10.
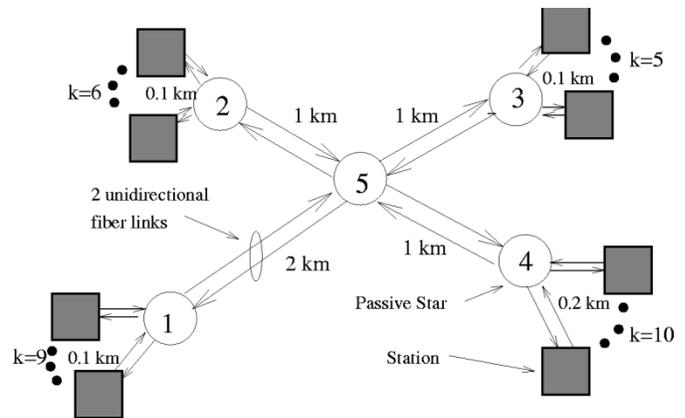

Fig. 12. A scaled-down version of the MAN network in Fig. 10.

enough, we would expect that this inefficiency would start to require the addition of more amplifiers. On the other hand, we would expect that, if links were short, then wavelengths at different power levels might not require the addition of more amplifiers and might allow us to potentially save even more amplifiers at critical points in the network. The network in Fig. 11 is meant to study the effects on the solution when we have links that span longer distances, and the network in Fig. 12 is meant to study the effects on the solution when a network has shorter links. Both of these networks are the same as the network in Fig. 10, except that the distances have been scaled up and down, respectively, by a factor of ten. As we see in Table II, the results seem to verify our earlier predictions. The new method is not able to find a better solution than the equally powered wavelengths solution for the larger network in Fig. 11, even when it was given multiple feasible starting points (including the solutions found in [14] and [20]). Our method's solution is not guaranteed to be the best because it could have become stuck at a local minimum. If our new method is stuck at a local minimum, we potentially can miss the global minimum solution. This differs from the LP solution which does find the global minimum solution (subject to the equally powered wavelengths constraint). On

the other hand, the new NLP method is able to come up with a better solution for the smaller network (Fig. 12). In fact, as we predicted, our new method was able to take advantage of the smaller network environment. The unequally powered wavelengths solution was able to use zero amplifiers compared to four for the equally powered wavelengths solution, which was a savings of four amplifiers. In the reference network (Fig. 10), the unequally powered wavelengths solution was able to use two amplifiers, compared to six for the equally powered wavelengths solution, which was also a savings of four amplifiers.

The network in Fig. 1 is also examined here because both of the previous studies [14], [20] examined this particular network.[8] This network has many nodes, and we predicted that our new method might not perform better than the equally powered wavelengths solution. We predicted this because the more nodes a network has, the more variables the solver is manipulating and the more local minima the solver can get stuck at. As Table II shows, the solver was unable to come up with a better solution than the LP solution, even when given

[8] The number of nodes for group 3 was reduced from 35 to 28 nodes because the original network, as proposed in [14], was infeasible because signals exited the star of degree 35 with power below $p_{sen} = -30$ dB (1 mW).
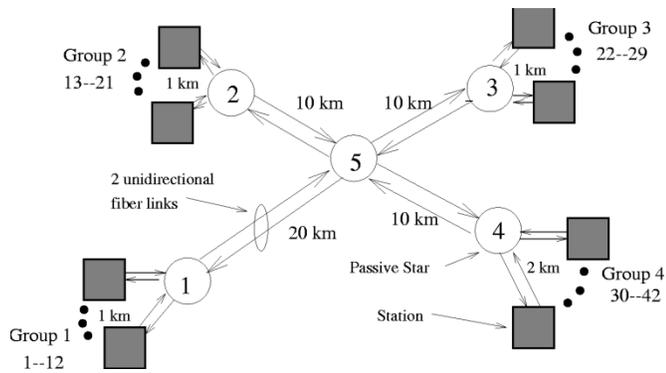
Fig. 13. A denser version of the MAN network in Fig. 10 with 12 additional stations.

Fig. 14. A sample switched network.

multiple feasible starting points including the solutions found in [14] and [20].

The "denser MAN" network in Fig. 13 differs from the MAN network in Fig. 10 in that there are 12 additional stations in it, three in each of the four groups of stations. This example shows the effect of scaling up the network by adding more stations on the number of amplifiers needed. All three amplifier-placement schemes require additional amplifiers to operate. Note, however, that the NLP method performs better than the other two schemes and remains closest to the absolute lower bound on the number of amplifiers.

For each of the previous example networks, column 6 in Table II shows the total CPU time taken by the nonlinear solver running on an otherwise-unloaded DEC 5000/240. In general, the running time is found to increase with: 1) increasing number of network components (which leads to more constraints) and 2) increasing link spans (which leads to a greater choice in feasible solutions). The running time of the solver can be potentially reduced by modifying the stopping criteria (see Section II-B); however, this can also affect the quality of the solution. Note that, due to the characteristics of our NLP solution process, no polynomial-form time complexity can be specified.

IV. FUTURE WORK

A. Switched Networks

The algorithms described in this paper were designed to operate on "loopless" networks where there is only one path from a source to a destination. In a switched network, there can potentially be multiple paths between a source and a destination. Since the above algorithms operate knowing how many wavelengths are on a given link, they assume that all wavelengths that can possibly reach a link could all be present on that link simultaneously. This approach has the potential to place more amplifiers in the network than is absolutely necessary. The example switched network given in Fig. 14 includes multiple paths between any source–destination pair. When examining the permutation of connections that use the *WR1-PS3* link, notice that all eight of the stations could each have a connection set up that use this "half" of L1. Actually,
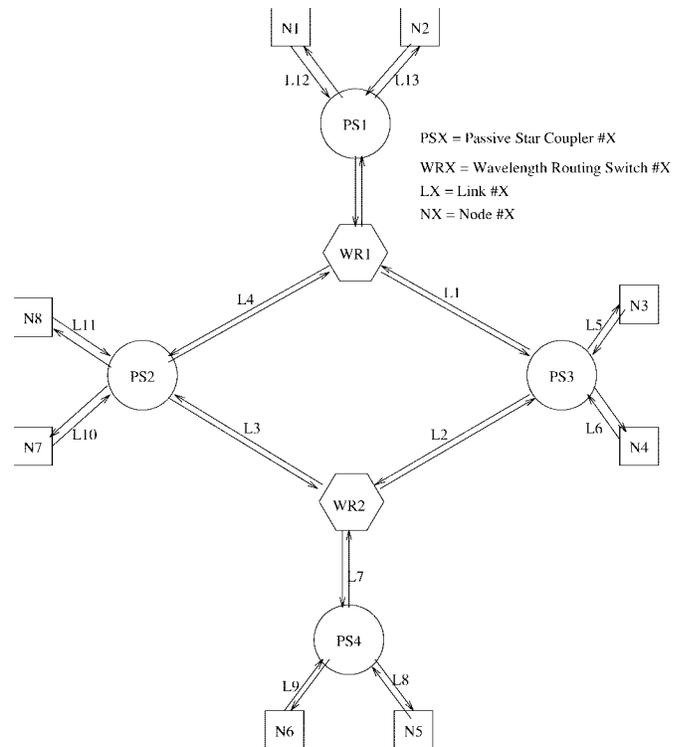
there is always at least one permutation of connections that would cause any of the "halves" of links *L1, L2, L3,* and *L4* to carry eight connections. Now, if amplifiers were placed in this network to allow any possible configuration of connections, all "halves" of links *L1–L4* would have to be designed with enough amplifiers to carry eight connections in the worst case. Now, it is fairly easy to see that, if the connections are setup in a "smart" fashion, a link never has to carry eight connections. In fact, a link should never have to carry more than two connections in this network. Designing links to carry two connections instead of eight, since the network will then potentially need only one-fourth the power on these links, can result in a significant savings in the number of amplifiers. We believe that it will be possible to modify our current algorithms to allow them to exploit this phenomenon that occurs in switched networks. This is a topic of our future work.

B. Modeling Device Characteristics

In the near future, we plan to try and further improve on the optical amplifier gain model. We expect to be able to create a reasonably accurate gain model of the popular EDFA. Analytical methods for modeling the amplifier gain, gain saturation, and noise described in [6] will be incorporated in the model. We also plan to expand our amplifier gain model to handle per-wavelength gain. This would allow us to model an amplifier that has a nonflat gain spectrum. It would also allow us to model the small gain for wavelengths that are normally considered to lie outside of the "amplifier bandwidth." The formulation of the problem would have to be changed to handle per-wavelength gain, too.

Further, we plan to consider the effect of crosstalk and the received BER of the signals on the amplifier placement in such networks. A study describing the computation of BER in the presence of crosstalk and amplifier-generated ASE noise in switched networks can be found in [4].

V. CONCLUSION

We considered the problem of minimizing the number of optical amplifiers in an optical LAN/MAN. This study departed from previous studies by allowing the signal powers of different wavelengths on the same fiber to be at different levels. Although this increases the complexity of the amplifier-placement algorithm, numerical results show that certain networks do benefit significantly from this method by requiring fewer amplifiers.

ACKNOWLEDGMENT

REFERENCES

[1] G. P. Agrawal, *Fiber-Optic Communication Systems.* New York: Wiley, 1992.

[2] S. B. Alexander *et al.*, "A precompetitive consortium on wide-band all-optical networks," *J. Lightwave Technol.*, vol. 11, pp. 714–735, May/June 1993.

[3] A. R. Chraplyvy, J. A. Nagel, and R. W. Tkach, "Equalization in amplified {WDM} lightwave transmission systems," *IEEE Photon. Technol. Lett.*, vol. 4, pp. 920–922, Aug. 1992.

[4] D. Datta, B. Ramamurthy, H. Feng, J. P. Heritage, and B. Mukherjee, "BER-based call admission in wavelength-routed optical networks," in *Tech. Dig. OFC'98*, San Jose, CA, Feb. 1998, vol. 2, pp. 92–93.

[5] A. F. Elrefaie, E. L. Goldstein, S. Zaidi, and N. Jackman, "Fiber-amplifier cascades with gain equalization in multiwavelength unidirectional inter-office ring network," *IEEE Photon. Technol. Lett.*, vol. 5, pp. 1026–1031, Sept. 1993.

[6] C. R. Giles and E. Desurvire, "Modeling erbium-doped fiber amplifiers," *J. Lightwave Technol.*, vol. 9, pp. 271–283, Feb. 1991.

[7] C. R. Giles and D. J. Giovanni, "Dynamic gain equalization in a two-stage fiber amplifier," *IEEE Photon. Technol. Lett.*, vol. 2, pp. 866–868, Dec. 1990.

[8] L. Grippo, F. Lampariello, and S. Lucidi, "A nonmonotone line search technique for newton's method," *SIAM J. Numer. Anal.*, vol. 23, pp. 707–716, 1986.

[9] O. K. Gupta and A. Ravindran, "Nonlinear integer programming algorithms: A survey," *OPSEARCH*, vol. 20, pp. 189–206, 1983.

[10] P. Hansen, "Methods of nonlinear 0-1 programming," in *Discrete Optimization II*, P. L. Hammer *et al.*, Eds. New York: North Holland, 1979.

[11] J. Iness, "Efficient use of optical components in WDM-based optical networks," Ph.D. dissertation, Dep. Comput. Sci., Univ. California, Davis, Nov. 1997.

[12] K. Inoue, T. Kominato, and H. Toba, "Tunable gain equalization using a Mach-Zehnder optical filter in multistage fiber amplifiers," *IEEE Photon. Technol. Lett.*, vol. 3, pp. 718–720, 1991.

[13] C. Lawrence, J. L. Zhou, and A. L. Tits, "User's guide for cfsqp version 2.4: A C code for solving (large scale) constrained nonlinear (minimax) optimization problems, generating iterates satisfying all inequality constraints," Institute for Systems Research, Univ. Maryland, College Park, Tech. Rep. No. TR-94-16rl, Feb. 1996.

[14] C.-S. Li, F. F.-K. Tong, C. J. Georgiou, and M. Chen, "Gain equalization in metropolitan and wide area optical networks using optical amplifiers," in *Proc. IEEE INFOCOM'94*, Toronto, Ont., Canada, June 1994, pp. 130–137.

[15] C.-S. Li, F. F.-K. Tong, C. J. Georgiou, and M. Chen, "A near-far compensation scheme for all-optical WDMA/WDM networks with arbitrary topology," IBM T. J. Watson Res. Center, Yorktown Heights, NY, Tech. Rep., 1994.

[16] H.-L. Li, "An approximate method for local optima for nonlinear mixed integer programming problems," *Comput. Operations Res.* , vol. 19, no. 5, pp. 435–444, July 1992.

[17] H.-D. Lin, "Gain splitting and placement of distributed amplifiers," IBM T. J. Watson Res. Center, Yorktown Heights, NY, Tech. Rep. RC 16216 (#72010), Oct. 1990.

[18] K. Magari, M. Okamoto, and Y. Noguchi, "1.55 μm polarization insensitive high gain tensile strained barrier MQW optical amplifier," *IEEE Photon. Technol. Lett.*, vol. 3, pp. 998–1000, Nov. 1991.

[19] E. R. Panier and A. L. Tits, "On combining feasibility, descent and superlinear convergence in inequality constrained optimization," *Math. Programming*, vol. 59, pp. 261–276, 1993.

[20] B. Ramamurthy, J. Iness, and B. Mukherjee, "Optimizing amplifier placements in a multi-wavelength optical LAN/MAN: The equally powered-wavelengths case," *J. Lightwave Technol.*, vol. 16, pp. 1560–1569, Sept. 1998.

[21] A. E. Siegman, *Lasers.* Mill Valley, CA: Univ. Science, 1986, pp. 298–301.

[22] V. Tandon, M. Wilby, and F. Burton, "A novel upgrade path for transparent optical networks based on wavelength reuse," in *Proc. IEEE INFOCOM'95*, Boston, MA, Apr. 1995, pp. 1308–1315.

[23] A. E. Wilner and S. M. Hwang, "Passive equalization of nonuniform EDFA gain by optical filtering for megameter transmission of 20 WDMA channels through a cascade of EDFA's," *IEEE Photon. Technol. Lett.*, vol. 5, pp. 1023–1026, Sept. 1993.

Byrav Ramamurthy (S'97) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology, Madras, India, and the M.S. and Ph.D. degrees in computer science from the University of California, Davis, in 1993, 1995, and 1998, respectively.

Since August 1998, he has been an Assistant Professor in the Department of Computer Science and Engineering, University of Nebraska-Lincoln. His research interests include high-speed computer networks, distributed systems, and telecommunications.

Prof. Ramamurthy was a recipient of the Indian National Talent Search Scholarship and was a Fellow of the Professors for the Future Program at the University of California.

Jason Iness received the B.S. degree (highest honors) in computer science and engineering and the M.S. and Ph.D. degrees in computer science from the University of California, Davis, in 1992, 1995, and 1997, respectively.

While at the University of California, he conducted research in WDM-based lightwave networks, focusing primarily on efficient design of such networks (including sparse amplifier placement and sparse wavelength conversion). Other research interests included ATM networks and wireless communications. In 1997, he joined Intel Corporation, Hillsboro, OR.

Dr. Iness was the recipient of the University of California, Davis, Fellowship in the 1992–1993 and 1995–1996 academic years.

Biswanath Mukherjee (S'82–M'83) received the B.Tech. (Hons.) degree from the Indian Institute of Technology, Kharagpur, India, in 1980 and the Ph.D. degree from the University of Washington, Seattle, in 1987.

In July 1987, he joined the University of California, Davis, where he has been a Professor of Computer Science since July 1995 and Chairman of Computer Science since September 1997. His research interests include lightwave networks, network security, and wireless networks. He is the author of *Optical Communication Networks* (New York: McGraw-Hill, 1997), a book which received the Association of American Publishers, Inc. 1997 Honorable Mention in Computer Science. He is a member of the editorial boards of *ACM/Baltzer Wireless Information Networks*, *Journal of High-Speed Networks*, and *Photonic Network Communications*.

Prof. Mukherjee is a co-winner of paper awards presented at the 1991 and the 1994 National Computer Security Conferences. While at the University of Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship. He serves on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING and *IEEE Network*. He served as the Technical Program Chair of the IEEE INFOCOM '96.