Minimizing the Cost of Placing and Sizing Wavelength Division Multiplexing and Optical Cross-Connect Equipment in a Telecommunications Network

BELÉN MELIÁN
University of La Laguna
mbmelian@ull.es

MANUEL LAGUNA
University of Colorado
laguna@colorado.edu

JOSÉ A. MORENO-PÉREZ
University of La Laguna
jamoreno@ull.es

Latest revision: July 17, 2004

ABSTRACT

Cost reduction is a major concern when designing optical fiber networks. Multiwavelength optical devices are new technology for increasing the capacity of fiber networks while reducing costs, when compared to installing traditional (e.g., SONET) equipment and new fiber. In this paper we discuss the development of a metaheuristic method that seeks to optimize the location of Wavelength Division Multiplexing (WDM) and Optical Cross-Connect (OXC) equipment in fiber networks. The procedure combines ideas from the scatter search, tabu search and multi-start methodologies. Computational experiments with both real-world and artificial data show the effectiveness of the proposed procedure. The experiments include a comparison with a permutation-based approach and with lower bounds generated with CPLEX.

1. INTRODUCTION

In the last four decades, problems related to planning and routing in telecommunications networks have become a fertile ground for developing and applying optimization techniques. Two main events have driven these efforts: (1) the large investment in telecommunications, which offers significant opportunities for reducing costs and improving network designs, and (2) the rapid changes in technology.

Optical fiber technology has become the preferred choice for building telecommunications networks, due to its cost effectiveness, reliability and its almost unlimited capacity. Technology advances have motivated the development of appropriate planning tools that respond to the characteristics of a new structural and operational environment.

Within this context, we deal with the *network synthesis* or *provisioning* problem, which consists of minimizing the total cost of installing capacity on links of a given network so that demand requirements are satisfied. In these problems, both the physical network topology and the demand requirements are given and the decision variables relate only to adding capacity to links at minimum cost. When the problem includes also the design of the network topology, that is, determining which links to install, then a complete graph and installation costs are considered.

A given network topology, a cost structure and a set of demand requirements characterize a typical instance of a network synthesis problem. The cost structure depends on each situation as well as on the available technology. It is customary to assume that the system does not add routing costs once the equipment has been installed. A requirement between two nodes is a single commodity flow requirement. Multi-commodity flow requirements are also considered as long as the commodities involve different origins and destinations while sharing the capacity of the network.

Another typical assumption is that the optical traffic is expressed in OC-48 units, i.e., Optical Carrier level 48 SONET channels. Each such channel carries 2.488×10^9 bits per second, equivalent to 48×672 voice-grade digital channels digitized at 64,000 bits per second each, after subtracting out overhead bits used for routing and control.

When the set of requirements consists of a single demand, the synthesis problem reduces to solving a shortest path problem on a graph with incremental costs as arc weights. The incremental costs are associated with the equipment required to route the smallest allowed demand increment (i.e., 1 OC-48 if demand splitting is allowed, and the entire demand if no splitting is allowed). More likely, however, the set of demand requirements consists of several origin-destination pairs. In this case, the demand requirements are routed taking into account the spare capacity in the current network. The spare capacity problem must be studied carefully because the equipment installed on links and nodes to route a given demand under consideration can also be used to route another demand to be considered later, making the design more cost effective. Synthesis problems also consider non-simultaneous demand requirements, where the capacity installed to route one demand at a given time can be used to route another demand at a different time without additional cost.

1.1 Technology

Wavelength Division Multiplexing is the transmission of multiple laser signals at different wavelengths (colors) in the same direction, at the same time and over the same strand of fiber. WDM with more than eight frequencies enables a low cost per bit [Ram99]. These high-capacity WDMs are called Dense Wavelength Division Multiplexing (DWDM) and create multiple bi-directional 'virtual fibers' per physical fiber. DWDM solves the bandwidth bottleneck resulting from growth in data traffic, because it is an emerging technology that increases transportation capacity while preserving optical fiber systems previously installed. Hence, DWDM provides carriers the flexibility and scalability they need to deploy capacity when and where it is needed.

The DWDM system is equipped with amplifiers that allow transmission in one channel. However, most of the capacity cost is related to the channel cards. Channel cards are added as needed and their cost is charged to the design accordingly. This means that a system capable of handling up to 96 channels can be installed where only eight channels are active, and the design would only consider the cost of equipping the eight active channels [Lin99]. To use WDM technology, an equipment "unit" must be placed at both endpoints of each fiber link. For each wavelength, or channel in use, channel equipment must also be placed at both endpoints of the link. WDM channels are bidirectional and have the same capacity as a pair of fibers. *Amplification* is the process of restoring the optical signal to its original optical power and without distortion after the signal has lost power when passing through a strand of fiber. This process is particularly important in DWDM environments. The typical amplifiers are completely optical in that sense that they do not have electronic elements. Consequently, they do not require the classical electrical-optical and optical-electrical conversions, thereby avoiding the associated need for additional bandwidth.

Optical cross-connects (OXCs) are small space-division switches that can switch an optical signal from one wavelength to another on multi-fiber WDM systems or on a single fiber [Ram99]. Providers offer several OXC sizes, such as 16×16 (i.e., 16 incoming wavelengths that can be fully switched among 16 outgoing wavelengths) and up to 64×64 or more, with 512×512, 1024×1024 and even larger sizes contemplated for future product offerings. In this paper, we consider OXCs of two sizes only (namely, 32×32 and 512×512) but the formulation is general and can be expanded to incorporate larger OXC sizes.

A WDM system must originate and terminate at an OXC or DCS (Digital Cross-Connect Signal) port. Also, an OXC or DCS port is needed to add or drop traffic at the origin and destination of each demand carried by the network. The OXC and DCS ports are bidirectional. The installed base of DCS machines generally lacks OC-48 ports, making each DCS OC-48 port a set of 48 DS3 ports. Since we will be modeling at the OC-48 level, we can consider OXC and DCS equipment to be functionally equivalent. Optical signals originate and terminate at network nodes, which typically are SONET (Synchronous Optical Network) ring nodes carrying traffic expressed in OC-48 units.

1.2 Problem Description

In order to increase the capacity of a network at a minimum cost, it is necessary to decide:

- Where to place WDM and OXC systems;
- How to route the traffic within the resulting network; and
- How to restore the network in the event of a single link failure.

DWDM telecommunications network planning is often divided into four main phases: design of the network, routing of the demands, multiplexing and survivability. We assume that there is a current network design such that our problem consists of dealing with the remaining steps of the planning process. We do not tackle the problem of restoring the network after a link failure. However, since the restoration problem can also be treated as a provisioning problem, the proposed formulation for the service network is essentially the same for the restoration network, which normally is obtained after the service network has been configured.

Cox et al. [Cox00] proposed the planning problem that simultaneously addresses the provisioning, routing and survivability problems. The problem was approached using a genetic algorithm (GA). The procedure is based on incrementally adding equipment to minimize the cost of routing each demand. The GA uses permutations to represent solutions. A permutation represents the ordering in which the demands are considered, one by one, for routing purposes. Therefore, a permutation is mapped into an actual solution by a procedure that uses the given order to route the demands in the most cost-effective way. Since the equipment is added to satisfy the current demand without considering the demands that are yet to be routed, each permutation is expected to result in a different network design. (It is possible, but unlikely, for two different permutations to be mapped to the same network design.) The approach cannot guarantee the existence of an ordering of the demands that would result in an optimal design. In other words, even an exhaustive search of all permutations may result in a sub-optimal network design.

The optimization problem deals with a set of demands to be routed through the existing optical network. Associated with each demand is an origin node, a destination node, and a size, expressed in OC-48 units. Optical fiber joining pairs of nodes is used to route demands through the network. Each demand can be routed either entirely through one or more bare fibers, over one or more channels of a WDM system or it can be switched from a WDM to another through OXCs. The goal of the network planner is to minimize the total cost, which consists of the cost of additional fiber, WDM systems and OXC equipment.

The existing physical network design (i.e., the set of exiting links) constrains where new optical fibers and WDM systems can be placed. A *segment* is defined as a sequence of individual links that do not pass through any OXC system. In this case, any intermediate node will be called *glass-through node*, meaning that fiber or a WDM system passes through the node without adding or dropping traffic and without requiring additional equipment. Each OC-48 unit uses two bare fibers or a channel of a WDM system. For convenience, we will refer to the capacity required for an OC-48 unit as a *channel,* regardless of whether a pair of fibers or a channel of a WDM system actually is used. All links within a segment must carry the same amount of traffic running from its origin to its destination. In an optimal network design, each segment should follow a least-cost path (with respect to fiber cost) from its origin to its destination. Since the shortest path from any node to any other is treated as a potential segment, the network of segments results in a complete graph, which is intractable in most cases. Therefore, it is useful to generate a subset of promising segments as one of the search strategies.

Once an OXC is reached, wavelengths and fibers can be rearranged. Therefore, the capacity constraints on each segment simply are that enough fiber and WDM equipment must be available on the segment to handle the number of OC-48 units assigned to it. Each individual link must have enough channel capacity to cover all demands routed over segments that uses it.

Kennington et al. [Ken01] formulate and solve a similar problem. They model the wavelength division multiplexing routing and provisioning problem with uncertain demands and a fixed budget as a multicriteria optimization problem. The primary objective is to minimize a quadratic regret function that models the total amount of over and/or under provisioning of the network resulting from uncertainty in the demand forecast. The secondary objective is to minimize the equipment cost that achieves the optimal value for regret. They propose a two-phase robust optimization strategy based on mixed integer linear programs. In the basic provisioning model for each scenario, the objective is to minimize the total cost for provisioning the network with terminal equipment located at each node and optical amplifiers and regenerators associated with the needed links. Their resulting mixed integer linear programs have a large number of continuous variables and two integer variables per link, representing the number of fibers and the number of channels. Their models are implemented using the AMPL modeling language and the problems are solved with CPLEX. Kennington et al. do not use the concept of segments because neither OXC equipment nor glass-through nodes are considered.

2. PROBLEM FORMULATION

In this section we provide a segment-based formulation for the synthesis problem with multicommodity flow requirements. Our formulation is based on Cox et al. [Cox99] with the difference that we do not tackle simultaneously the provisioning, routing and survivability problems but instead deal with survivability after provisioning and routing of the service network. We will show that our formulation of the problem using a path-assignment approach to represent solutions results in improved outcomes when compared to tackling the whole problem with the permutation based approach proposed by Cox et al.

Consider an undirected graph $G = (N, E)$, where N is the set of nodes and E is the set of segments. A non-simultaneous multicommodity flow requirement, consists of a set of demands $D = \{(s_1, t_1, d_1), (s_2, t_2, d_2), \ldots, (s_k, t_k, d_k)\}$ to be routed through the graph. Each single demand consists of an origin node, s_i, a destination node, t_i, and a size, d_i. Our formulation uses the following definitions.

Data

Facilities Input Data

N = set of nodes at which demands originate and terminate
L = set of links between pairs of nodes
E = set of potential segments

J = types of DWDM systems
I = types of OXC systems

Demand Input Data

D = set of demands to be carried, expressed in OC-48 units, between origin-destination node pairs
s_i = origin node of demand i
t_i = destination node of demand i
d_i = the amount of demand i

Cost Input Data

a_e = cost of a fiber on segment e (sum of costs per link along that segment), $e \in E$
b_j = cost of a type j WDM unit, $j \in J$
c_l = cost of a type l OXC unit, $l \in I$
q_j = channel cost of a type j WDM unit, $j \in J$
r_l = port cost of a type l OXC unit, $l \in I$

Capacity Data

o_j = capacity of a type j WDM unit
p_l = capacity of a type l OXC unit

Existing Infrastructure

g_e^j = spare WDM channels on WDM systems of type j on segment e, $j \in J$, $e \in E$
h_n^l = spare OXC ports on OXC systems of type l at node n, $l \in I$, $n \in N$

Decision Variables

x_{ie} = amount of demand i routed on segment e
x_{ie}^F = amount of demand i routed on segment e in the forward direction
x_{ie}^R = amount of demand i routed on segment e in the reverse direction
f_e = number of stand-alone (no WDM) fiber pairs on segment e
w_e^j = number of type j WDM units on segment e
v_e^j = number of channels on type j WDM units on segment e
y_n^l = number of type l OXC units installed at node n
u_n^l = number of ports on type l OXC units installed at node n

Objective Function

The following objective function minimizes the sum of fiber costs (first term), WDM costs (second term) and the OXC costs (third term).

$$Min \sum_{e \in E} a_e f_e + \sum_{e \in E} \sum_{j \in J} \left((a_e + b_j) w_e^j + q_j v_e^j \right) + \sum_{n \in N} \sum_{l \in I} (c_l y_n^l + r_l u_n^l)$$

Constraints

The following constraints require that all demand must be carried, that no link should be assigned more demand than its capacity allows it to carry and that no switching element should be assigned more traffic than its capacity allows.

Conservation of Service Flow

$$\sum_{\substack{e \in E \\ s_i = start(e)}} x_{ie}^F + \sum_{\substack{e \in E \\ s_i = end(e)}} x_{ie}^R - \sum_{\substack{e \in E \\ s_i = end(e)}} x_{ie}^F - \sum_{\substack{e \in E \\ s_i = start(e)}} x_{ie}^R = d_i \ , \ \forall i \in D$$

$$\sum_{\substack{e \in E \\ t_i = end(e)}} x_{ie}^F + \sum_{\substack{e \in E \\ t_i = start(e)}} x_{ie}^R - \sum_{\substack{e \in E \\ t_i = start(e)}} x_{ie}^F - \sum_{\substack{e \in E \\ t_i = end(e)}} x_{ie}^R = d_i \ , \forall i \in D$$

$$\sum_{\substack{e \in E \\ j = end(e)}} x_{ie}^F + \sum_{\substack{e \in E \\ j = start(e)}} x_{ie}^R - \sum_{\substack{e \in E \\ j = start(e)}} x_{ie}^F - \sum_{\substack{e \in E \\ j = end(e)}} x_{ie}^R = 0, \ \ \forall i \in D, \forall j \in N, j \neq s_i, t_i$$

$$x_{ie}^F + x_{ie}^R = x_{ie} \qquad \forall i \in D, \forall e \in E$$

Segment Capacity

$$\sum_{i \in D} x_{ie} \leq f_e + \sum_{j \in J} v_e^j, \forall e \in E$$

$$v_e^j \leq w_e^j o_j + g_e^j, \forall e \in E \ \ , \forall j \in J$$

Switch Requirements

$$\sum_{n = end(e)} (f_e + \sum_{j \in J} v_e^j) \leq \sum_{l \in I} u_n^l, \forall n \in N$$

$$u_n^l \leq p_l y_n^l + h_n^l, \forall n \in N, \forall l \in I$$

<u>Integrality Constraints</u>. All variables are nonnegative integer.

This formulation assumes an undirected graph. For directed graphs, the variables x_{ie}^{F} and x_{ie}^{R} can be eliminated. That is, the *R*-variables are entirely eliminated and *x* is used in place of the *F*-variables. With this definition, each segment needs to be listed only once, but can be used in either direction. A WDM system on that segment can also be used in either direction.

The preceding objective function, decision variables, and constraints specify a formal version of the synthesis problem, where the provisioning and routing problems are tackled simultaneously to solve either the service or restoration problem. In practice, however, demands and costs often are uncertain, while available technology options, such as OXC and WDM system capacities, change frequently as new products are introduced. Therefore, this formal version of the problem only approximates a more complicated synthesis problem with uncertainty in key data and changing constraints. Rather than pursuing a true multi-period optimization approach with formulations that explicitly model uncertainty, many designers prefer to work with a simpler formulation and re-run the associated optimization procedure frequently as conditions change. Previous values of decision variables may then become initial conditions for a new optimization run, which may involve changes in costs and an expanded set of technically feasible options. Repeatedly running a static (one-period) optimization procedure with changing inputs is, in principle, a sub-optimal approach to adaptive planning. However, in practice, such "rolling optimization" is often preferred to theoretically more realistic dynamic formulations for which required input data cannot be estimated with an appropriate accuracy level.

Additional, and perhaps more realistic, formulations are obtained by constraining the optimization as follows:

- *Maximum length allowed for WDM systems*. The signal reach stimulated by WDM equipment spans a maximum of about 400 miles without electronic regenerators. If the length of the segment, with an installed WDM system, exceeds that distance, additional WDM systems must be placed back to back along the segment. Our proposed metaheuristic procedure handles this constraint, and the additional WDM systems and channel cards are subsumed in the formulation by adding the appropriate cost.

- *Allowed technologies.* The environment in which the network design problem arises may limit the fiber, WDM systems, and OXC equipment that can be assumed that are available for use.

3. A METAHEURISTIC SOLUTION APPROACH

This section summarizes the development of a metaheuristic procedure that searches for optimal solutions to the provisioning and routing problems. As indicated above, the restoration problem is tackled as a synthesis problem, which is solved after the service problem, so that a similar procedure can be used to design the service and restoration networks. The MIP model presented in the previous section has a very large number of variables and constraints, making it impractical for the exact solution of real instances of moderate or large size. For small planning problems, the MIP formulation can be solved in reasonable amount of computer time, as shown in our computational experiments. However, the exact solution of the model is only a lower bound on the optimal solution to the real problem, as stated by Birkan et al. (2002).

"A node-arc model has all paths implicitly available and hence produces a true lower bound for this problem. The disadvantage of the node-arc model is that routing may be circuitous and require an unacceptably large hop count."

Our solution procedure employs the notion of a *base network*, which initially consists of the current network design. A *base* is an incomplete network design that does not satisfy the set of demand requirements that a complete design should be capable of handling. As the process iterates, the base network evolves and the estimated cost of routing a demand becomes more accurate. An evolved base network includes additional equipment, which has been tentatively added to the original base. When a demand is considered for routing on an evolved base network, this demand can share the additional capacity with other demand requirements, making the cost estimates more accurate, due to a decreasing fraction of the capacity that is not shared for costing purposes. The evolution of the base network is linked to an adaptive memory mechanism that keeps track of where new equipment is added in the best solutions recorded during the search. The solution approach that we propose builds a list of paths for each demand by making use of an efficient implementation of the *k*-shortest path algorithm. This procedure identifies a controlled set of feasible paths for each demand [Glo93] and is a variant of the *k*-shortest path algorithm reported in [Law76]. The paths for a given demand are found calculating the incremental cost of routing the entire demand in the base network. For example, one of the possible paths would be to add the necessary fiber and WDM systems to create a segment from the origin to the destination of a given demand. Other paths are created using alternative ways of carrying the demand from origin to destination, which would most likely imply adding WDMs and OXCs.

Four basic elements are common to heuristic search, regardless of the specific methodology or strategic design choices: (1) a solution representation, (2) an objective, (3) an evaluation function, and (4) a move mechanism. The specifications for our proposed search procedure are:

Solution representation. The construction of a solution starts with the selection of a path for each demand requirement. Once each demand is assigned to a path, the cost of the resulting design is calculated. The cost is associated with the equipment that is required to satisfy the demands using the chosen paths. A solution is fully determined by a data structure that stores the path assignments and the equipment required in each element of the original network.

Objective. The goal of the DWDM planning problem is to minimize the sum of additional fiber cost, WDM equipment cost and its terminal equipment (OXC units) cost, subject to the appropriate technology constraints.

Evaluation. Once each demand has been assigned to a path in its list of potential paths, the evaluation of the solution consists of calculating the increase of capacity required in the elements of the network that route the demands through the assigned paths. The increased capacity is then translated into cost of installing additional fiber and adding WDMs and OXCs.

Move mechanism. Every solution has a neighborhood, which consists of all the feasible solutions that are reached by changing a demand from one path to another.

Our overall solution strategy consists of an adaptive metaheuristic method that combines ideas from scatter search [Lag02], multi-start [Glo00], and tabu search [Glo97]. The hybrid metaheuristic takes advantage of strategies that can explore a large solution space effectively. Specifically, tabu search contributes with a short term memory component that is designed to avoid cycling. Scatter search adds a mechanism to generate new solutions from the combination of solutions in an updated reference set of solutions. Finally, the multi-start component uses a long term memory that forces construction of new solutions in a wider range of the solution space.

Figure 1 shows the main steps of our proposed procedure. The procedure starts with the generation of a set of promising segments using the shortest path algorithm (with distances as weights). Specifically, we first decide the density with respect to number of segments that we wish to obtain in the network. For instance, suppose that the original network consists of 12 nodes and 17 links. The density of this network is 17/66 = 25.75%. If we wish to augment this network to a density of 50% then we need to add 33-17 = 16 segments, because the original links are considered segments corresponding to existing WDM systems and are also included in the promising set. The 16 new segments are the shortest paths between any two nodes using distances as weights.

The procedure uses these segments to execute the *k*-shortest path algorithm for each demand (with incremental costs from a base network *B* as weights). After the execution of this step (line 5) each demand has a set of paths that are used as the basis for building solutions. Given the network of segments, the spare capacity on the segments and nodes is determined. Obtaining spare capacities allows the procedure to assess incremental costs of routing demands in each segment.

The initial reference set is constructed in lines 6 and 7. The set is populated using a constructive procedure (line 7) that attempts to assign demands to paths in order to utilize efficiently the spare capacity in the original base network. The rationale behind this initialization is that spare capacity for channels in the final network design should be zero except for channels on WDM systems covering a segment without slack. The strategy acknowledges that spare capacity in the original network simply accounts for existing network infrastructure. The solutions in the reference set are ordered according to their total cost, where the first solution, labeled $RefSet_1$, is the one with lowest cost. The reference set is updated as the process iterates (lines 16 and 20). The notion of a reference set is the same as the one used in the scatter search methodology [Lag03], where it is used as a repository of solutions that are submitted to a combination method. In some basic designs, the reference set contains the best solutions (according to the objective function value) found during the search. However, in more advanced designs, the reference set strategically mixes high quality solutions and diverse solutions, as explained below.

```
1.    procedure DWDM_planning
2.    {
3.      generate (segments);
4.      B = initial_base;
5.      Paths = find_k-shortest_ paths ( B );
6.      for ( i = 1, …, |RefSet| )
7.          RefSet_i = constructive ( B, Paths );
8.      S = RefSet_1;
9.      do {
10.         find_demand_order ( S, B );
11.         while ( improving move ) {
12.             move = find_next_improving_move;
13.             if ( move ) {
14.                 S = execute ( move );
15.                 if ( S better that RefSet_last )
16.                     RefSet = update ( S );
17.             }
18.         }
19.         if ( equal (RefSet ) )
20.             rebuild ( RefSet );
21.         B = evolve ( RefSet );
22.         Paths = find_k-shortest_paths ( B );
23.         S = constructive ( B, Paths );
24.     } until (stopping criterion)
25.   }
```

Figure 1. Proposed Metaheuristic Procedure.

We use the current solution *S,* which at the beginning is the first solution in the reference set (line 8), to obtain an ordering of the demands according to their unit cost (where a unit is an OC-48). In this ordering, the demand that contributes most (per OC-48) to the total cost of the design is first and the one that contributes least (per OC-48) is last. The demand ordering is important, because the local search, which is based on changing one demand from its current path to another, starts with the demand that has the largest unit cost. To calculate the unit cost, the demands are examined one by one. The examination consists of deleting the demand from the current solution and calculating the cost reduction. The cost reduction is then divided by the bandwidth requirement of the demand under consideration. Once all demands have been examined, the unit cost associated with each demand is known.

The neighborhood search (line 12) within the local search in lines 11 to 18 examines moves employing the ordering of the demands determined in line 10. That is, the first candidate move is to reassign the demand that is at the top of the unit cost list. If reassigning this demand leads to an improving move, the move is executed to change the current solution (lines 13 and 14). If the new solution is better than the worst in the current reference set, then the reference set is updated (line 16). If an improving move that involves reassigning the first demand in the list cannot be found, then the second demand is considered. The process continues until a demand is found for which a reassignment of paths leads to an improving move. If all the demands are examined and no improving move is found, the local search is abandoned. Once the local search is abandoned, the procedure compares the current reference set with the reference set before the last time the local search was executed (line 19). If the reference set did not change after the last execution of the local search, the set is rebuilt (line 20). The process of evolving the base network from a reference set is mainly deterministic and therefore if the reference set does not change, then the base does not evolve properly. Rebuilding of the reference set entails keeping the top |*RefSet*|/2 solutions and generating new solutions to substitute the worst |*RefSet*|/2 in the set, as generally done in implementations of scatter search.

Numerous studies show that effective metaheuristic procedures keep a balance between search intensification and diversification, that is, between reinforcing attributes associated with good solutions and driving the search into regions not visited yet. To achieve this balance, the original base is evolved (line 21) employing the information embedded in the reference set. One of the main criteria used to evolve the base network relates to the number of times a segment has appeared in the paths assigned to the demands in the *RefSet* solutions. The procedure also uses global (referred to the whole search process) and local (referred to the current reference set) information in the form of counters that keep track of the number of channels used in each segment in order to decide where to add equipment to the current base. The difference between the maximum global and the maximum local number of channels used in each segment shows its importance. The smaller the difference the more important the segment is in the final network design.

For each demand, the *k*-shortest paths are again calculated (line 22) by using the incremental costs of routing the demand through the new base network. This step updates the list of best candidate paths according to the current base network. The new paths take advantage of the additional channels included in the evolved base network that can be used without increasing the total cost of the design. The local search now starts from an initial solution constructed to best utilize the spare resources in the new base (line 23). The procedure includes intensification and diversification strategies in the evolution of the base network and in the utilization of the spare capacity during the construction of a starting solution for the local search. The procedure terminates after a pre-specified number of iterations.

4. COMPUTATIONAL RESULTS

In this section, we present and discuss our computational experiments. We first describe the problem instances that were used to carry out the experimentation. Then we report the results of our experiments. All programs were implemented in C and compiled with Microsoft Visual C++ 6.0. All experiments were performed on a PC with one Pentium 4 processor at 2.53 GHz.

The problem instances used for testing are both real (shared by Dr. Leonard Lu of AT&T Labs) and randomly generated. The random instances are based on the networks corresponding to the real instances, with the demands and existing equipment randomly generated. The motivation for generating random instances is to study the performance of our methods on instances with various characteristics. We consider five different network sizes (with number of nodes varying from 11 to 113) and we generate links to create several densities. Finally, several sets of uniform and clustered sets of demands are randomly created. Uniform demands are generated by randomly selecting an origin and a destination, where each pair has the same probability of being selected. Clustered demands are generated selecting a subset of nodes as "high traffic" locations and then generating a demand pattern that has a higher density around those nodes. The magnitude of the demands ranges between 1 and 105 OC-48s. The problem instances used for testing are summarized in Table 1. For each set, Table 1 shows the name, the number of nodes $|N|$, links $|L|$, and the different number of demands $|D|$ corresponding to the clustered, uniform and real instances.

Set Name	$\|N\|$	$\|L\|$	$\|D\|$		
			Clustered	Uniform	Real
MetroD	11	16	10,20,30	54	
		27	10,20,30	54	
		42	10,20,30	54	48
Extant0D	12	17	15,21,44	66	19
		33	15,21,44	66	
		46	15,21,44	66	
Example2D	17	26	27,36,81	135	79
		68	27,36,81	135	
NationalD	50	63	45,65,91		112
108_annealed-3D	113	137	150,200,250,400		130

Table 1 : Data characteristics

Table 2 summarizes the data regarding the equipment cost used in the solution of the problem instances listed in Table 1.

Constant	Cost	Description
a_e	\$1,400 * length(*e*)	Cost of a fiber on a segment *e*
B	\$95,000	cost of a WDM unit
C	\$120,000	cost of an OXC unit
Q	\$18,000	channel cost of a WDM unit
R	\$10,000	port cost of an OXC unit

Table 2 : Description of costs

For comparison purposes, we have implemented a permutation-based algorithm that follows the same structure as the one proposed in [Cox00]. In this approach, a permutation represents the ordering in which the demands are considered for routing. A permutation is mapped into a solution by a procedure that uses the given order to route the demands in the most cost-effective way. When the first demand is considered for routing the current design consists of the original network. The demands are considered one by one as specified by the order in the current permutation. Additional equipment is added as required and the design is updated. The permutation is fully mapped when all the demands have been considered. The approach has the goal of locally minimizing the addition of equipment as each demand is routed through the network. More details of the approach can be found in [Cox00] where the permutation search is conducted using a genetic algorithm. Since the genetic algorithm used in [Cox00] is a proprietary code of Cox and Associates, Inc., we employ OptQuest [Opt01], a commercial scatter search solver that is capable of searching a permutation space.

Our first experiment consists of comparing the solutions obtained by the permutation-based metaheuristic and our hybrid metaheuristic approach applying the Wilcoxon Signed Ranks Test [Dan90]. The objective is to determine if we may conclude from sample evidence that there is a significant difference between these two procedures. We apply both procedures to the problems in Table 1 and record the objective function values obtained by each procedure. Then, we compute the absolute objective function value differences (without regard of the sign) for each problem and all differences of zero are omitted. Let the number of pairs remaining be denoted by n. Ranks from 1 to n are assigned to these n pairs according to the relative size of the absolute difference, as follows. Rank 1 is given to the pair with the smallest absolute difference; rank 2 is given to the pair with the second smallest difference; and so on, until rank n is assigned to the pair with the largest absolute difference. If several pairs have absolute differences that are equal to each other, we assign to each of these several pairs the average of the ranks that they would have been assigned if ties were broken arbitrarily.

Wilcoxon suggested a T statistic, which has the approximate quantiles given by the normal distribution, under the null hypothesis that there are no significant differences between the two compared procedures. The critical region of approximate size $\alpha = 0.001$ corresponds to all values of T less than -3.0902. Since in our case T = -3.756, the null hypothesis is rejected and we may conclude that there are significant differences between the two metaheuristic procedures.

We have now established that our procedure performs significantly better than the permutation-based approach, as indicated by the statistical test. In our second experiment, we assess the quality of the solutions obtained by the application of our hybrid metaheuristic. For this experiment, we give the MIP formulation presented in section 4 to the CPLEX 8.0 MIP solver. The solution of this model provides a lower bound because the number of intermediate nodes for paths between origin and destination pairs is not bounded. The first column in Table 3 identifies the problem set. Columns 2, 3, and 4, contain the number of nodes, number of segments, and number of demands. Under the headings "PERM" and "METAH" we report the total costs in millions of dollars and CPU times in seconds corresponding to the permutation based procedure and the proposed metaheuristic, respectively. Under the heading CPLEX we report the total cost obtained by solving our MIP formulation and the solution time if CPLEX with its default parameters is capable of finding and confirming the optimal solution in 24 CPU hours. The last three columns in this table show the deviation between the PERM solution (C_P) and the METAH solution (C_M), between the METAH solution and CPLEX solution (C_C) if the optimal was found, and between the METAH solution and the best lower bound (C_{BLB}) if the optimal solution to the MIP model was not found.

Set Name	|N|	|E|	|D|	PERM		METAH		CPLEX		$\frac{C_P - C_M}{C_P}$ (%)	$\frac{C_M - C_C}{C_M}$ (%)	$\frac{C_M - C_{BLE}}{C_M}$ (%)
				Cost millions	Time seconds	Cost millions	Time seconds	Cost millions	Time seconds			
MetroD	11	16	10	4.09	5.12	4.09	0.13	4.09	0.20	0	0	
			20	4.38	9.77	4.38	4.90	4.38	0.75	0	0	
			30	8.42	15.02	8.42	7.37	8.42	0.39	0	0	
			54	14.13	27.02	14.12	16.60	14.03	6.53	0.07	0.63	
		27	10	2.75	5.21	2.75	1.96	2.75	0.29	0	0	
			20	4.31	10.28	4.26	4.30	4.03	1.09	1.16	5.39	
			30	6.46	15.35	6.46	6.35	6.40	1.76	0	0.92	
			54	11.39	27.52	11.24	14.69	10.84	17.34	1.31	3.55	
		42	10	1.96	5.49	1.96	2.03	1.94	0.38	0	1.02	
			20	3.08	11.16	3.08	3.48	3.06	1.21	0	0.64	
			30	5.40	16.24	5.38	4.86	5.38	7.21	0.37	0	
			48	7.31	26.72	7.11	12.24	6.99	10.89	2.73	1.68	
			54	8.80	29.86	8.50	13.18	8.35	39.44	3.40	1.76	
Extant0D	12	17	15	3.69	8.16	3.69	3.25	3.69	0.75	0	0	
			19	6.26	10.41	6.26	9.99	6.26	8.04	0	0	
			21	6.21	11.43	6.21	9.46	6.21	3.45	0	0	
			44	14.48	23.55	14.36	27.19	14.36	96.11	0.82	0	
			66	11.99	35.71	12.14	41.00	11.83	81.81	-1.25	2.55	
		33	15	3.69	7.82	3.69	3.57	3.69	20.09	0	0	
			21	7.32	11.60	7.32	7.21	6.03	269.68	0	17.62	
			44	13.94	23.69	14.23	19.99	13.66	67709.06	-2.08	-	
			66	11.83	35.81	11.83	29.93	11.87	-	0	-	0.25
		46	15	3.69	7.94	3.69	2.30	3.69	38.61	0	0	
			21	7.32	11.61	7.32	3.83	6.03	770.15	0	17.62	
			44	13.97	24.43	13.95	18.54	14.29	-	0.14	-	11.46
			66	11.83	36.19	11.77	28.82	13.24	-	0.50	-	11.46
Example2D	17	26	27	23.22	26.11	23.22	21.59	22.47	38.42	0	3.22	
			36	81.84	33.34	81.84	20.69	81.84	492.57	0	0	
			79	180.70	78.27	178.05	93.94	182.94	-	1.46	-	2.60
			81	98.80	75.44	97.37	89.57	96.65	5438.10	1.44	0.73	
			135	177.42	127.73	173.03	191.46	182.04	-	2.47	-	2.41
		68	27	24.43	27.12	24.43	18.11	19.27	4406	0	21.12	
			36	68.15	35.19	68.10	20.20	69.51	-	0.07	-	9.45
			81	84.09	80.05	82.65	80.24	102.12	-	1.71	-	15.29
			135	149.71	135.11	144.14	169.60	-	-	3.72	-	13.32
NationalD	50	63	45	37.63	125.12	37.36	61.77	44.08	-	0.71	-	28.13
			65	51.16	192.28	50.87	141.70	56.25	-	0.56	-	26.34
			91	59.18	267.46	59.13	162.75	62.77	-	0.08	-	24.38
			112	44.44	296.16	42.88	230.70	51.51	-	3.51	-	36.14

Table 3 : Summary of Results

Several observations can be made regarding the results shown in Table 3. First, the instance-by-instance comparison between PERM and METAH shows that only in 2 instances (Extand0D-12-17-66 and Extand0D-12-33-44) out of 40 the PERM solutions are better than the METAH solutions. The deviations of the METAH solutions from the optimal solutions found with CPLEX are generally small, ranging from 0 to 5.39%, except for Extant0D-12-33-21, Extant0D-12-46-21 and Example2D-17-68-27. For the problem for which CPLEX could not terminate within 24 CPU hours, our hybrid metaheuristic was able to always improve upon the best upper bound. However, the deviations against the best lower bounds for the largest problems in Table 3 (NationalD) are such that provide little information regarding the solution quality of the solutions obtained with the proposed metaheuristic.

We now compare the performance of METAH and PERM on the largest problems in our data set (i.e., 108_annealed-3D). The results of this experiment are shown in Table 4. The relative deviations shown in the last column of Table 4 favor the solutions found by METAH. Note that although the relative differences decrease with the size of the clustered problems (|*D*| = 150, …, 400), the absolute differences are always more than $30 million. We do not include a column for CPLEX in Table 4 because the bounds found after 24 hours of CPU time do not provide any useful information.

|*N*|	|*E*|	|*D*|	PERM		METAH		$\frac{C_P - C_M}{C_P}$ (%)
			Cost in millions	Time in seconds	Cost in millions	Time in seconds	
113	137	130	118.30	1124.07	107.50	682.47	9.13
		150	142.25	1348.44	108.75	839.24	23.55
		200	181.43	1801.84	141.45	1116.23	22.04
		250	225.22	2239.37	181.71	1787.53	19.32
		400	390.77	3728.42	355.55	3057.21	9.01

Table 4 : Comparative Results for 108_Annealed-3D

One of the main components of our hybrid metaheuristic is the local search. In our third experiment, we use the permutation-based procedure to independently test the effectiveness of the local search. That is, we use the procedure to isolate the local search from other components of our hybrid metaheuristic in order to assess its effectiveness. As described in the previous section, the local optimizer performs a first-improving local search in the neighborhood of the current solution. When a network design is obtained using a permutation of the demands, it is possible to execute our local search to try to improve upon the given solution. Our experiments show that the designs obtained after executing the local search are generally better than the initial designs obtained using the permutation procedure alone. We applied the permutation based approach augmented with the local search to all the instances in Table 1 and use the results to test for significant differences between this approach (PERM+LS) and the one that does not use the local search (PERM).

We once again use Wilcoxon's test with $\alpha = 0.001$. Since for the procedures PERM and PERM+LS, $T = -3.516 < -3.0902$, we conclude that the differences between the permutation based procedure and the permutation based procedure with the local search are statistically significant. Therefore, we may say that the application of the local search to the designs obtained by the permutation based procedure generally improves upon the final network designs by reducing their total costs.

Our previous experiment and associated statistical test have determined that the performance of the permutation-based procedure is enhanced with the application of our local search. In our final comparison, we test if there is a significance difference between PERM+LS and our hybrid metaheuristic (METAH). The application of Wilcoxon's test results in $T = -0.355$. The critical region of size α = 0.05 corresponds to values of T less than -1.6449. Since T > -1.6449, the test leads to the conclusion that the performance of PERM+LS is not significantly different than the performance of METAH. This conclusion indicates that the local search that we have designed is quite effective and can be used to improve upon solutions yielded by construction procedures.

Our last experiment consists of comparing the solutions obtained by our hybrid metaheuristic approach and a multistart search that does not include the tabu and scatter search features applying the Wilcoxon Signed Ranks Test [Dan90]. The objective is to determine if we may conclude from sample evidence that the tabu and scatter search elements in our metaheuristic make a significant difference. The critical region of approximate size α = 0.001 corresponds to all values of T less than -3.0902. Since in our case T = -4.516, the null hypothesis is rejected and we may conclude that there is a significant difference between the metaheuristic procedure with and without the tabu and scatter search mechanisms.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have addressed an important and current problem in the telecommunications industry. We have provided the motivation for studying this optimization problem and have discussed the technology behind it. Our segment-based formulation is used as the framework for developing heuristic procedures and as a means for finding lower and upper bounds.

Our experiments with real and randomly generated data show the merit of our proposed solution procedure when compared to a permutation-based approach and to the upper bounds generated by solving an MIP formulation with CPLEX. We used a nonparametric statistical test to compare our hybrid procedure and two variants of a permutation-based approach. The test revealed the effectiveness of our local search, which is capable of improving solutions constructed with the permutation-based approach to a point that the resulting method is statistically comparable to the proposed hybrid metaheuristic. Additional experiments showed that the scatter and tabu search elements significantly contribute to the quality of the solutions found with the proposed metaheuristic.

Although our general approach contemplates solving the protection problem as part of the design process, in the scope of this paper we have not included the implementation and experimentation associated with network survivability. An extension of our work will include solving the restoration problem using the last reference set obtained when the termination criterion (line 24 in Figure 1) is satisfied. Specifically, given a network design (i.e., a solution in the reference set), the restoration problem would consist of finding the most cost-efficient way of routing demands after a link failure. We believe that the lessons learned while tuning the procedure for finding good solutions to the service problem will be valuable in the development of a comprehensive procedure that includes the restoration problem.

REFERENCES

[Bal91] Balakrishnan, A., Magnanti, T. L., Shulman, A. and Wong, R. T., “Models for Planning Capacity Expansion in Local Access Telecommunication Networks”, *Annals of Operations Research*, **33**, pp. 239-284 (1991).

[Bir02] Birkan, G., Kennington, J. L., Olinick, E. V., Ortynski, A., and Spiride, G., "Making a Case for Using Integer Programming to Design DWDM Networks". To appear in *Optical Networks Magazine* (2002).

[Cox00] Cox, A. L., Jr., Sanchez, J., Lu, L. and Sun, X., "Cost Savings from Network Optimization of DWDM Facilities and Optical Switches" Technical Report, Cox Associates (2000).

[Cox01] Cox, L. A., Laguna, M., Melián, B., Moreno-Pérez, J.A., Sánchez, J., "Optimizing Placement and Sizing of WDM and Optical Cross-Connect Equipment", in *Proceedings of the 9th International Conference on Telecommunication Systems, Modeling and Analysis*, pp. 98-107, March (2001).

[Dan90] Daniel, W. W. *Applied Nonparametric statistics*, PWS-Kent Publishing Company, Boston, 1990.

[Gav91] Gavish, B., "Topological Design of Telecommunication Network. Local Access Design Methods", *Annals of Operations Research*, **33**, pp. 17-71 (1991).

[Glo93] Glover, F. and Laguna, M., "Bandwidth Packing: A Tabu Search Approach", Management Science, vol. 39, no. 4, 492-500 (1993).

[Glo97] Glover, F. and Laguna, M., *Tabu Search*, Kluwer Academic Publishers, (1997).

[Glo00] Glover, F., "Multi-Start Strategic Oscillation Methods – Principles to Exploit Adaptive Memory", in *Computing tools for modeling, Optimization and Simulation: Interfaces in Computer Science and Operations Research*, eds. M. Laguna and J. L. González-Velarde, Kluwer Academic Publishers, 1-23 (2000).

[Grö95] Grötschel, M., Monma, C. L. and Stoer, M., "Polyhedral and Computational Investigations for Designing Communication Networks with High Survivability Requirements", *Operations Research*, Vol. 43, No. 6, November-December (1995).

[Ken01] Kennington, J. L., Lewis, K. R., Olinick, E. V., Ortynski, A. and Spiride, G., "Robust Solutions for the WDM Routing and Provisioning Problem: Models and Algorithms", *Optical Networks Magazine*, 4: 74-84 (2003).

[Kra98] Krause, T., "The emerging optical layer", *Telephony,* June (1998).

[Lag03] Laguna, M. and R. Martí, *Scatter Search: Methodology and Implementations in C*, Kluwer Academic Publishers, Boston (2003).

[Law76] Lawler, E.L., *Combinatorial Optimization: Networks and Matroids*,. Holt, Rinehart and Winston, New York (1976).

[Lin99] Lindstrom, A., "It is metro time for DWDM", *America's Network*, (1999).

[Mag84] Magnanti, T. L., Wong, R. T., "Network Design and Transportation Planning: Models and Algorithm", *Transportation Science*, 18, 1-55 (1984).

[Opt01] OptQuest Callable Library, OptTek Systems, Inc. (www.opttek.com), 2001.

[Ram99] Ramani, M. and Keates, R., "The All-Optical Gateway", *Nortel Networks*, (1999).

[Ste99] Stern, T. E. and Bala, K., *Multiwavelength Optical Networks*, Prentice Hall, (1999).