Models and Algorithms for the Design of Service Overlay Networks

Antonio Capone, Jocelyne Elias, Fabio Martignon

***Abstract*— Service Overlay Networks (SONs) can provide end-to-end Quality of Service guarantees in the Internet without requiring significant changes to the underlying network infrastructure. A SON is an application-layer network operated by a third-party Internet Service Provider (ISP) that owns a set of overlay nodes, residing in the underlying ISP domains, interconnected by overlay links.**

The deployment of a SON can be a capital-intensive investment, and hence its planning requires careful decisions, including the overlay nodes' placement, the capacity provisioning of overlay links as well as of access links that connect the end-users to the SON infrastructure.

In this paper, we propose two novel optimization models for the planning of SONs. The first model minimizes the SON installation cost while providing full coverage to all network's users. The second model maximizes the SON operator's profit by further choosing which users to serve, based on the expected gain, and taking into consideration budget constraints. We also introduce two efficient heuristics to get near-optimal solutions for large-scale instances in a reasonable computation time.

We provide numerical results of the proposed models and heuristics on a set of realistic-size instances, and discuss the effect of different parameters on the characteristics of the planned networks. We show that in the considered network scenarios the proposed heuristics perform close to the optimum with a short computing time.

***Index Terms*: - Service Deployment, Network Planning, Overlay Networks, Service-Level Agreements, Optimization, Heuristics.**

I. Introduction

THE Internet has experienced a tremendous growth in its size and complexity in the last few years; it connects today thousands of Autonomous Systems operated by different Internet Service Providers (ISPs), companies and universities.

The Internet was originally designed to provide a best-effort delivery service, but nowadays Internet users often require services that need end-to-end Quality of Service (QoS) guarantees over multiple domains. Although several approaches have been proposed in the literature to support QoS in the Internet, like Integrated Services [1] and Differentiated Services [2], such approaches are far from being widely implemented. Indeed, achieving a large scale deployment is challenging, as cooperation among multiple network operators is difficult to arrange in practice since it involves business and legal issues in addition to technical problems.

Service Overlay Networks (SONs) have recently emerged as alternative and very promising architectures able to provide end-to-end Quality of Service guarantees in the Internet, while leaving the underlying Internet infrastructure unchanged [3], [4], [5], [6], [7].

A SON is an application-layer network built on top of the traditional IP-layer networks. In general, the SON is operated by a third-party ISP that owns a set of overlay nodes residing in the underlying ISP domains. These overlay nodes perform service-specific data forwarding and control functions, and are interconnected by virtual overlay links which correspond to one or more IP-layer links [3].

The service overlay architecture is based on business relationships between the SON, the underlying ISPs, and the users. The SON establishes bilateral service level agreements with the individual underlying ISPs to install overlay nodes and purchase the bandwidth needed for serving its users. On the other hand, the users subscribe to SON services, which will be guaranteed regardless of how many IP domains are crossed by the users' connection. The SON gains from users' subscriptions. Although the quality requirements that a SON must satisfy may be different (e.g. bandwidth, delay, delay jitter, packet loss), we assume they are mapped to an equivalent bandwidth [3], [7]. To assure the bandwidth for the SON, the underlying ISPs have several technical options: they can lease a transmission line to the SON, use bandwidth reservation mechanisms or create a separate Label Switched Path if MPLS [8] is available in their networks.

Obviously, the deployment of Service Overlay Networks can be a capital-intensive investment. It is therefore imperative to develop network design tools that consider the cost recovery issue for a SON. The main costs of SON deployment include the overlay nodes installation cost and the cost of the bandwidth that the SON must purchase from the underlying network domains to support its services.

Very few works consider the problem of topology design for Service Overlay Networks [7], [9], [10]. However, all these works assume that the number and location of overlay nodes are pre-determined, while the overlay nodes placement is a critical issue in the deployment of the SON architecture. These works further assume that a full coverage of all traffic demands must be provided, while the main goal of a SON operator would be to maximize its profit by choosing which users to serve based on the expected revenue. Finally, previous works often do not impose bounds on overlay links capacities, assuming that the underlying ISPs will always be able to

Manuscript received January 16, 2008; revised May 12, 2008; accepted June 21, 2008; approved by IEEE Transactions on Network and Service Management Editor Dr. Nikos Anerousis.

Antonio Capone and Jocelyne Elias are with the Department of Electronics and Information of Politecnico di Milano. E-mail: {capone, elias}@elet.polimi.it

Fabio Martignon is with the Department of Information Technology and Mathematical Methods, University of Bergamo, Dalmine (BG) 24044, Italy. E-mail: fabio.martignon@unibg.it

provide bandwidth to the SON.

In this paper we overcome these limitations by proposing two novel overlay network design models that select the optimal number and position of overlay nodes, as well as the capacity reserved on each overlay link, while taking into account in an accurate way traffic routing.

The first model minimizes the network installation cost while providing full coverage to all network's users. The second model maximizes the SON operator's profit by further choosing which users to serve, eventually considering a budget constraint that the SON operator could specify to limit its economic risks in the deployment of the overlay network.

Although these problems are NP-hard, we show that they can be solved to the optimum for small and medium-size instances.

To tackle large-size instances, we propose two simple but effective heuristic approaches able to provide near-optimal solutions in a reasonable computation time. The proposed algorithms are based on the decomposition of the model into sub-problems and on the solution of the continuous relaxation. 0-1 feasible solutions are then obtained using a randomized rounding technique.

We provide numerical results for a set of realistic-size instances and investigate the impact of different parameters on the SON design problem, such as number and installation cost of overlay nodes, bandwidth costs, traffic demands and SON operator's budget. Finally, we determine bounds to the performance achievable by any optimization algorithm solving continuous relaxations of the proposed models. The numerical results show that in the considered network scenarios the proposed heuristics perform very close to the optimal solution provided by the integer models with a short computing time.

The paper is structured as follows: Section II discusses related work. Section III introduces two novel overlay network design models, while Section IV illustrates the proposed heuristics to plan large-scale overlay networks. Section V presents numerical results that show the effect of different parameters on the characteristics of the planned network. Finally, Section VI concludes this paper.

II. Related Work

Several works have appeared in the literature with the purpose of providing optimal topology design in different contexts, such as wired backbone networks [11], [12], [13], [14] and recently Service Overlay Networks [7], [9], [10], [15], [16], [17], [18], [19].

An adaptive topology design framework for SONs is presented in [7] to assure inter-domain QoS, and a set of heuristics is proposed to solve the least-cost topology design problem. A similar problem is investigated in [9], where end-systems and overlay nodes are connected through ISPs that support bandwidth reservations; simulated annealing is used as heuristic to provide solutions for large-sized networks. Another set of heuristics for SON design is proposed in [10]; these heuristics aim to construct an overlay topology maintaining the connectivity between overlay nodes under various IP-layer path failure scenarios. However, all these works formulate the design problem considering full coverage of all traffic demands and assuming that locations of overlay nodes are given and the underlying ISPs are always able to provide resources to the SON.

Reference [15] deals with dynamic topology construction to adapt to the underlying network topology changes. An architecture for topology-aware overlay networks is proposed to enhance the availability and performance of end-to-end applications by exploring the dependency between overlay paths. Several clustering-based heuristics for overlay node placement and a routing mechanism are also introduced.

The dynamic overlay network reconfiguration issue is addressed in [16], where the main goal is to find the optimal reconfiguration policies that can both accommodate time-varying communication requirements and minimize the total overlay network cost.

The problem of overlay node placement is addressed in [17], [18], [19]. In [17] the authors consider how to place service nodes optimally in a network, balancing the need to minimize the number of nodes and to limit the distance between users and service nodes. This work, however, only proposes optimization algorithms for the uncapacitated version of the coverage problem. The work in [18] focuses on designing an overlay network that maximizes the number of unicast and multicast connections with deterministic delay requirements, without considering link costs. Finally, the overlay node placement problem is investigated in [19] to improve routing reliability and TCP performance. This paper, however, assumes that overlay nodes and links have infinite capacities, and does not take into account the costs involved in the deployment of the overlay network.

The construction of efficient multicast trees in overlay networks is the primary focus of several studies [20], [21], [22], [23], and is not considered in this paper since we focus our analysis on unicast traffic.

In summary, the above cited techniques are less general than our current work since they deal with the design problem considering at least one of the following special cases: 1) the number and location of overlay nodes are pre-determined, 2) there are no capacity constraints on overlay links, and 3) full coverage of all network users is provided without considering the SON profit maximization issue. In our work, on the contrary, we take into consideration all these issues in the formulation of the overlay network design problem. In addition, we introduce a budget constraint in one of our models to limit the economic risk that the SON operator can face when deploying its network.

III. Service Overlay Network Design Models

A common approach to the network design problem is to consider feasible positions of traffic concentration points in the service area (Test Points, TPs), which generate traffic towards one or more Destination Nodes (DNs), and feasible positions where overlay nodes can be installed (Candidate Sites, CSs) [11]. The placement of TPs and DNs depends on the expected traffic distribution, while that of CSs on the underlying network topology and the agreements of the SON operator with ISPs. Although the concept of *test point* is distinguished from *end-user* (formally, the end-user is the

TABLE I
BASIC NOTATION

c_j^I	Cost for installing an overlay node in CS j
c_{jl}^B	Cost for buying one bandwidth unit between CSs j and l
c_{ij}^A	Access cost per bandwidth unit between TP i and CS j
c_{jk}^E	Egress cost per bandwidth unit between CS j and DN k
d_{ik}	Traffic generated by TP i towards DN k
u_{jl}	Maximum capacity that can be reserved on overlay link (j, l)
v_j	Maximum capacity of the access link of CS j
h_{jk}	Maximum capacity that can be reserved on egress link (j, k)
a_{ij}	0-1 parameter that indicates if TP i can access the SON through CS j
e_{jk}	0-1 parameter that indicates if CS j can be connected to DN k
b_{jl}	0-1 parameter that indicates if CSs j and l can be connected with an overlay link
g_i	Revenue per bandwidth unit obtained for serving TP i
B	SON operator's budget
x_{ij}	0-1 variable that indicates if TP i is assigned to CS j
z_j	0-1 variable that indicates if an overlay node is installed in CS j
w_{jk}	0-1 variable that indicates if CS j is connected to DN k
f_{jl}^k	Flow variable which denotes the traffic flow routed on link (j, l) destined to DN k
f_{jk}	Flow variable which denotes the traffic flow routed on egress link (j, k)

traffic generation agent that is placed in a TP), we will use the two terms as synonyms throughout the paper. *Destination nodes* can represent either terminal nodes or access points to other networks.

Let $S = 1, \ldots, m$ denote the set of CSs, $I = 1, \ldots, n$ the set of TPs, and $D = 1, \ldots, p$ the set of destinations. The basic notation used in this paper is presented in Table I.

The cost associated with installing an overlay node in CS j is denoted by c_j^I; c_{jl}^B denotes the cost for the SON operator to buy one bandwidth unit between CSs j and l from the underlying ISPs, and c_{ij}^A is the access cost per bandwidth unit required between TP i and CS j; finally, c_{jk}^E represents the cost per bandwidth unit for the traffic transmitted on the egress link between CS j and destination node $k \in D$.

The traffic generated by TP i towards destination node k is given by the parameter d_{ik}, $i \in I, k \in D$. The maximum capacity that can be reserved by the SON operator between CSs j and l on the overlay link (j, l) is denoted by u_{jl}, $j, l \in S$, while the maximum capacity of the access link of CS j is denoted by v_j, $j \in S$ and that of the egress link between the installed overlay node j and destination node k is denoted by h_{jk}.

According to TPs, DNs and CSs geographic location and the underlying physical topology, the following connectivity parameters can be calculated.

Let a_{ij}, $i \in I, j \in S$ denote the test point coverage parameters:

$$a_{ij} = \begin{cases} 1 & \text{if TP } i \text{ can access the SON through an} \\ & \text{overlay node installed in CS } j \\ 0 & \text{otherwise} \end{cases}$$

and e_{jk}, $j \in S, k \in D$ the destination nodes coverage parameters:

$$e_{jk} = \begin{cases} 1 & \text{if CS } j \text{ can be connected with destination node } k \\ 0 & \text{otherwise} \end{cases}$$

Obviously, a_{ij} depends on the proximity of TP i to CS j, that is on the access coverage provided by the SON operator with CS j through agreements with local network operators. Similarly, e_{jk} is related to the distance between DN k and CS j.

Finally, let b_{jl}, $j, l \in S$ denote the connectivity parameters between two different CSs, which may depend on the proximity of the overlay nodes j and l in the underlay network, as well as on the agreements between the SON and the different ISPs:

$$b_{jl} = \begin{cases} 1 & \text{if CS } j \text{ and } l \text{ can be connected with an overlay link} \\ 0 & \text{otherwise} \end{cases}$$

Decision variables of the problem include TP assignment variables x_{ij}, $i \in I, j \in S$:

$$x_{ij} = \begin{cases} 1 & \text{if TP } i \text{ is assigned to CS } j \\ 0 & \text{otherwise} \end{cases}$$

overlay nodes' installation variables z_j, $j \in S$:

$$z_j = \begin{cases} 1 & \text{if an overlay node is installed in CS } j \\ 0 & \text{otherwise} \end{cases}$$

destination assignment variables w_{jk}, $j \in S, k \in D$ (if $z_j = 1$, w_{jk} denotes if j is connected to destination node k):

$$w_{jk} = \begin{cases} 1 & \text{if CS } j \text{ is connected to destination node } k \\ 0 & \text{otherwise} \end{cases}$$

and finally flow variables f_{jl}^k which denote the traffic flow routed on link (j, l) destined to destination node $k \in D$. The special variables f_{jk} denote the traffic flow on the egress link between CS j and destination node k.

Given the above parameters and variables, we propose two different Service Overlay Network design formulations. The first, called Full Coverage SON Design model (FCSD), minimizes the total network cost while assuring full coverage of all end-users. The second formulation, called Profit Maximization SON Design model (PMSD), maximizes the total profit, selecting which users to serve based on the revenue generated by their subscription to the SON services and the cost necessary to satisfy their traffic demand.

A. Full Coverage SON Design Model

The Full Coverage SON Design model (FCSD) minimizes the total network cost while assuring full coverage of all network users.

$$Minimize \quad \Big\{ \sum_{j \in S} c^I_j z_j + \sum_{j,l \in S} \sum_{k \in D} c^B_{jl} f^k_{jl} + \\ + \sum_{i \in I, j \in S, k \in D} c^A_{ij} d_{ik} x_{ij} + \sum_{j \in S, k \in D} c^E_{jk} f_{jk} \Big\} \quad (1)$$

$$s.t.$$

$$\sum_{j \in S} x_{ij} = 1, \quad \forall i \in I \quad (2)$$

$$x_{ij} \leq z_j a_{ij}, \quad \forall i \in I, j \in S \quad (3)$$

$$\sum_{i \in I} d_{ik} x_{ij} + \sum_{l \in S} (f^k_{lj} - f^k_{jl}) - f_{jk} = 0, \quad \forall j \in S, k \in D \quad (4)$$

$$\sum_{k \in D} f^k_{jl} \leq u_{jl} b_{jl} z_j, \sum_{k \in D} f^k_{jl} \leq u_{jl} b_{jl} z_l, \quad \forall j, l \in S \quad (5)$$

$$\sum_{i \in I, k \in D} d_{ik} x_{ij} \leq v_j, \quad \forall j \in S \quad (6)$$

$$f_{jk} \leq h_{jk} w_{jk}, \quad \forall j \in S, k \in D \quad (7)$$

$$w_{jk} \leq e_{jk} z_j, \quad \forall j \in S, k \in D \quad (8)$$

$$x_{ij}, z_j, w_{jk} \in \{0, 1\}, \quad \forall i \in I, j \in S, k \in D \quad (9)$$

The objective function (1) accounts for the total Service Overlay Network cost, including installation costs and the costs related to the connection of overlay nodes, users' access and egress costs.

Constraints (2) provide full coverage of all TPs, while constraints (3) are coherence constraints assuring respectively that a TP i can be assigned to CS j only if an overlay node is installed in j and if i can be connected to j.

Constraints (4) define the flow balance in node j for all the traffic destined towards node k. These constraints are the same as those adopted for classical multicommodity flow problems. The term $\sum_{i \in I} d_{ik} x_{ij}$ is the total traffic generated by the assigned TPs destined towards destination node k, $\sum_{l \in S} f^k_{lj}$ is the total traffic received by j from neighboring nodes, $\sum_{l \in S} f^k_{jl}$ is the total traffic transmitted by j to neighboring nodes, and f_{jk} is the traffic transmitted towards the destination node k.

Constraints (5) impose that the total flow on the link between overlay nodes j and l does not exceed the capacity of the link itself (u_{jl}); at the same time, they define the existence of an overlay link between CS j and CS l, depending on the installation of nodes in j and l and the connectivity parameters b_{jl}.

Constraints (6) impose for each overlay node that the ingress traffic serviced by such network device does not exceed the capacity of the link used for the access, whilst constraints (7) force the flow between node j and the destination node k to zero if node j is not connected to k, and impose that such flow does not exceed the maximum capacity (h_{jk}) of the egress link between the installed overlay node j and k.

Constraints (8) are coherence constraints assuring that a CS j can be connected to a destination node k only if an overlay node is installed in j and if k can be connected to j. Finally, constraints (9) are the integrality constraints for the binary decision variables.

Note that we can consider alternative formulations to the FCSD model. For example, we might want end-users to be connected to more than one overlay node, for redundancy. This is easily accomplished by modifying constraints (2) as:

$$\sum_{j \in S} x_{ij} = \eta, \quad \forall i \in I \quad (10)$$

where η is the number of overlay nodes per end-user.

It is easy to observe that the above model is NP-hard since it includes the set covering and the multicommodity flow problems as special cases. It is worth noting that the structure of this model is quite different from classical flow models usually adopted for the design of physical networks [12] since it includes the traffic routing and the placement of nodes.

B. Profit Maximization SON Design Model

The Profit Maximization SON Design model (PMSD) maximizes the total profit, choosing which users to serve based on the revenue generated by their subscription to the SON services and the cost necessary to the SON operator to cover them.

The objective function (1) is therefore modified as follows:

$$Maximize \sum_{i \in I, j \in S, k \in D} g_i d_{ik} x_{ij} - \Big\{ \sum_{j \in S} c^I_j z_j + \\ + \sum_{j,l \in S} \sum_{k \in D} c^B_{jl} f^k_{jl} + \sum_{i \in I, j \in S, k \in D} c^A_{ij} d_{ik} x_{ij} + \sum_{j \in S, k \in D} c^E_{jk} f_{jk} \Big\} \quad (11)$$

where $g_i, \forall i \in I$, represents the revenue per bandwidth unit that the SON operator gets for serving Test Point i. Here we assume for simplicity that the price paid by the ith user is proportional to the amount of traffic the user introduces in the SON, $\sum_{k \in D} d_{ik}$, with g_i being the proportionality coefficient, but some general pricing models can be easily accounted for.

Constraints (2) are modified as follows:

$$\sum_{j \in S} x_{ij} \leq 1, \quad \forall i \in I \quad (12)$$

while all the other constraints are the same as in the FCSD model. Such formulation maximizes the SON operator profit, which is obtained by subtracting the total cost necessary to deploy an overlay network that satisfies the users' requirements from the total revenue achieved by serving a subset of the Test Points. Note that, differently from constraints (2) in the FCSD

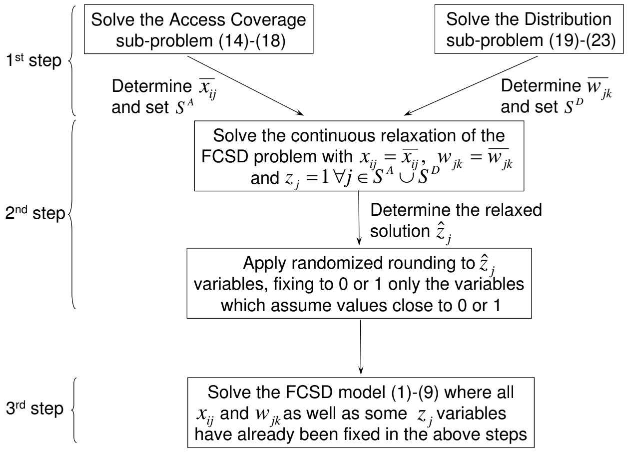

Fig. 1. Flow diagram of the H-FCSD heuristic.

model, in this formulation constraints (12) do not impose full coverage of all TPs.

The SON planner may be required to specify a certain cost budget to limit the economic risks in the deployment of its network due to traffic predictions. To this end, the PMSD formulation can be easily modified to account for cost limitations. With B the budget, this can be done simply by the addition of the following constraint:

$$\sum_{j \in S} c_j^I z_j + \sum_{j,l \in S} \sum_{k \in D} c_{jl}^B f_{jl}^k + \\ + \sum_{i \in I, j \in S, k \in D} c_{ij}^A d_{ik} x_{ij} + \sum_{j \in S, k \in D} c_{jk}^E f_{jk} \leq B \quad (13)$$

An alternative formulation of the PMSD model can be considered to provide different service levels for end-users. The model could select the optimal amount of traffic to admit in the network for each user in order to maximize the SON profit. This can be easily obtained relaxing x_{ij} variables, i.e. modifying the integrality constraints (9) as follows:

$$x_{ij} \in [0,1], z_j, w_{jk} \in \{0,1\}, \quad \forall i \in I, j \in S, k \in D$$

The sum $\sum_{j \in S} x_{ij}$ now represents the fraction of traffic generated by TP i which is admitted in the network; such quantity can assume any value in the $[0,1]$ interval. Note that solving this alternative problem is easier than solving the original PMSD problem, since the assignment variables x_{ij} are no longer binary.

IV. Heuristics to Solve the SON Design Problems

The computing time required to obtain an optimal solution for both the FCSD and PMSD problems might be very long for medium-to-large-size instances. As an example, the average computing time required to solve at optimum random network instances with the FCSD model grows exponentially from 10 s for CS=30 up to 4600 s for CS=50, when run on an Intel Pentium 4 (TM) processor with CPUs operating at 3 GHz and with 1024 Mbyte of RAM, which is the workstation used to obtain the numerical results reported in this paper. For this reason, we propose hereafter two simple but effective heuristics, named H-FCSD and H-PMSD, that provide near-optimal solutions in a reasonable amount of time for the Full Coverage and the Profit Maximization SON Design problems, respectively.

A. H-FCSD: a Heuristic to solve the Full Coverage SON Design Problem

In the case of the full coverage problem, we apply a problem reduction technique which sets some variables solving first the access coverage and distribution sub-problems, then considering a continuous relaxation of the FCSD formulation, and a randomized rounding procedure to obtain an integer solution.

The H-FCSD heuristic is composed of three steps, which are described in detail in the following and are also illustrated in the flow diagram of Figure 1.

First Step

In this step we solve separately the access coverage and the distribution sub-problems, whose formulations are reported below. All the variables and parameters have the same definition as described in the previous Section for the FCSD model.

The minimum-cost access coverage sub-problem consists in computing the optimal location of the overlay nodes through which end-users can access the SON with minimum cost, and it is formulated as follows:

$$Minimize \sum_{i \in I, j \in S, k \in D} c_{ij}^A d_{ik} x_{ij} + \sum_{j \in S} c_j^I z_j \quad (14)$$

$$s.t.$$

$$\sum_{j \in S} x_{ij} = 1, \quad \forall i \in I \quad (15)$$

$$x_{ij} \leq z_j a_{ij}, \quad \forall i \in I, j \in S \quad (16)$$

$$\sum_{i \in I, k \in D} d_{ik} x_{ij} \leq v_j, \quad \forall j \in S \quad (17)$$

$$x_{ij}, z_j, \in \{0,1\}, \quad \forall i \in I, j \in S \quad (18)$$

The objective function (14) accounts for the access cost, composed of the costs related to the users' access and the installation cost of access nodes. Constraints (15), (16) and (17) are the same as constraints (2), (3) and (6) in the FCSD problem formulation, while (18) are the integrality constraints for the decision variables.

The minimum-cost distribution sub-problem consists in determining the optimal location of the overlay nodes that can deliver with minimum cost the SON traffic to destination nodes, and it is formulated as follows:

$$Minimize \sum_{j \in S, k \in D} c_{jk}^E f_{jk} + \sum_{j \in S} c_j^I z_j \quad (19)$$

$$s.t.$$

$$f_{jk} \leq h_{jk} w_{jk}, \quad \forall j \in S, k \in D \tag{20}$$

$$w_{jk} \leq e_{jk} z_j, \quad \forall j \in S, k \in D \tag{21}$$

$$\sum_{j \in S} f_{jk} = \sum_{i \in I} d_{ik}, \quad \forall k \in D \tag{22}$$

$$z_j, w_{jk} \in \{0, 1\}, \quad \forall j \in S, k \in D \tag{23}$$

The objective function (19) calculates the egress cost, including the costs related to distribute egress traffic and the installation cost of egress overlay nodes. Constraints (20) and (21) are the same as constraints (7) and (8). Constraints (22) are coherence constraints which impose that all the traffic destined to destination node k ($\sum_{i \in I} d_{ik}$) is effectively routed in the egress assignment sub-problem. Finally, (23) are the integrality constraints for the decision variables.

Solving the access coverage problem (14)-(18) we determine the optimal assignment of each TP to a corresponding CS, $\overline{x_{ij}}$; at the same time, we select the subset $S^A \subseteq S$ of CSs where overlay nodes must be installed to cover all TPs (i.e. $S^A = \{j \in S | z_j = 1\}$).

In the same way, solving the distribution problem (19)-(23) we determine both the optimal egress connections between CSs and DNs, $\overline{w_{jk}}$, and the subset $S^D \subseteq S$ of CSs where overlay nodes must be installed to guarantee a connection with the egress nodes (i.e. $S^D = \{j \in S | z_j = 1\}$).

Let $S^F = S^A \cup S^D$ be the subset of CSs where an overlay node must be installed as determined either in the access coverage or distribution sub-problems.

Second Step

In this step we solve a continuous relaxation of the FCSD problem (1)-(9), i.e. the FCSD problem (1)-(8) with the integrality constraints (9) replaced as follows:

$$x_{ij}, z_j, w_{jk} \in [0, 1], \quad \forall i \in I, j \in S, k \in D \tag{24}$$

and where, in addition, the x_{ij}, w_{jk} and $z_j | j \in S^F$ variables are constrained to assume the values determined in the first step. To this aim, the following constraints are added:

$$x_{ij} = \overline{x_{ij}}, \quad \forall i \in I, j \in S \tag{25}$$

$$w_{jk} = \overline{w_{jk}}, \quad \forall j \in S, k \in D \tag{26}$$

$$z_j = 1, \quad \forall j \in S^F \tag{27}$$

Let $\hat{z}_j$ be the optimal solution of the relaxed problem (1)-(8), (24)-(27). We then apply randomized rounding on $\hat{z}_j$:

- if $\widehat{z_j} = 0$, then the corresponding z_j variable is set to 0;
- otherwise, the variable z_j is set to 1 with probability $P[z_j = 1] = \widehat{z_j}$; to this end, we extract a random value v uniformly distributed in $[0, 1]$: if $v \leq \widehat{z_j}$ the z_j variable is set to 1 (that is, an overlay node is installed in CS j). If $v > \widehat{z_j}$, no decision is taken in this step on the installation of an overlay node in CS j, and the z_j value will be determined in the third step of the heuristic.

The rational behind such procedure is that we strive for a balance between reducing the problem complexity (i.e. the number of binary variables in the integer problem solved in the third step) and the possibility of obtaining a feasible and close to the optimum solution.

Third Step

Finally, in this step we solve the original integer FCSD problem (1)-(9), with all the x_{ij}, w_{jk} and z_j variables set as described in the first two steps, obtaining both integer solutions for the remaining z_j variables and optimal values for the routing variables f_{ij}^k.

Comments

Randomized rounding is a general technique first proposed in [24] to solve 0-1 optimization problems, which consists in solving the continuous relaxation of the integer problem and then transforming the optimal solution of the relaxed problem into a feasible solution for the integer problem. It has been demonstrated in [24] that such technique provides provably good solutions, in the sense that with high probability this algorithm provides an integer solution in which the objective function assumes a value close to the optimum of the continuous relaxation.

However, it may be difficult to obtain a good integer solution from the fractional one. For this reason, to design an efficient heuristic, we introduced the first two steps described above since we observed that applying directly randomized rounding to the fractional variables of the relaxed FCSD problem leads very often to unfeasible solutions where several constraints are violated; then, the computation of a feasible and near to the optimum solution is not very efficient even applying scaling techniques as proposed in [24].

Note that in the first step several optimal solutions may exist to the access coverage and distribution sub-problems, and it is therefore interesting to extend the H-FCSD algorithm considering different initial solutions. This can also increase the probability that a feasible solution is obtained in the third step. Obviously, there exists a trade-off between the execution time of the heuristic (which grows with the number of different initial solutions considered) and the improvement in the total network cost. To evaluate this issue we performed several tests, considering up to 15 initial solutions, and we found that in all the network scenarios considered in this paper the maximum improvement obtained in the cost of the planned SON was less than 2%. For this reason, in Section V we report numerical results obtained considering only one optimal solution in the first step of the H-FCSD heuristic.

B. H-PMSD: a Heuristic to solve the Profit Maximization SON Design Problem

In general, determining the optimal subset of end-users to cover in order to maximize the SON operator's profit is a more difficult problem than designing the minimum cost overlay network that provides full coverage to a given set of users.

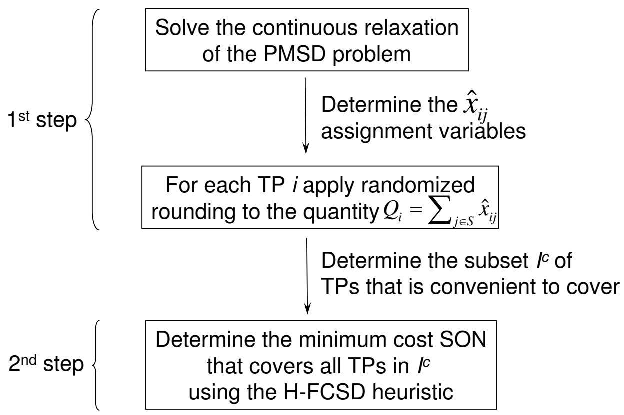

Fig. 2. Flow diagram of the H-PMSD heuristic.

The H-PMSD heuristic tackles such difficulty determining preliminarily which users to serve, and then designing the minimum cost SON that covers such users. H-PMSD is therefore composed of two steps, as illustrated in the flow diagram of Figure 2 and detailed in the following.

First Step: determining which users to cover

In this step, a continuous relaxation of the PMSD problem is solved. Let $\hat{x}_{ij}$ be the optimal solution for the TPs assignment variables. Then, for each TP i we consider the quantity $Q_i = \sum_{j \in S} \hat{x}_{ij}$, which assumes values in the $[0, 1]$ interval and can be interpreted intuitively as the probability with which the ith Test Point should be covered by the SON. We then perform randomized rounding on Q_i: a random value v uniformly distributed in $[0, 1]$ is extracted; if $v \leq Q_i$, then TP i is selected to be covered by the SON; otherwise, TP i is not selected. Let $I^c \subseteq I$ be the subset of TPs chosen in this step to be covered by the SON.

Second Step: designing the minimum cost SON

In this step we use the H-FCSD heuristic to design the minimum cost SON that covers all the TPs $\in I^c$ chosen in the previous step.

Note that the heuristic may be unable to find any feasible solution of the problem since the cost for covering the selected TPs is higher than the revenue.

H-PMSD with Budget constraints

A cost budget constraint can be taken into account in the first step of the H-PMSD heuristic, introducing constraint (13) in the continuous relaxation of the PMSD problem. With B the budget, if the cost of the SON planned in the second step is not greater than B, then the computed solution is acceptable. Otherwise, to obtain a feasible solution we apply a scaling technique [24] in the randomized rounding procedure of the first step. Such technique consists in multiplying the solution of the relaxed problem by a factor $\gamma < 1$, which corresponds to using γQ_i in the first step of H-PMSD. If γ decreases, the probability that user i is covered by the SON, and as a consequence the total network cost, is reduced until the budget constraint is not violated.

In our work we perform a simple iterative procedure that proceeds as follows:

1) Initialize $\gamma = 1$
2) Solve the first step of H-PMSD using γQ_i
3) Solve the second step of H-PMSD
4) If the cost of the SON planned in the second step of H-PMSD is $\leq B$ then STOP (a feasible solution has been obtained). Otherwise reduce the γ value by 0.1 and go to 2)

However, a finer tuning of the γ parameter could be performed, using for example a binary search technique, to find the γ value which guarantees at the same time feasibility and a good quality of the solution.

V. NUMERICAL RESULTS

In this section we test the sensitivity of the proposed models and heuristics to different parameters like the number of candidate sites and test points, the traffic demands, the installation costs as well as the revenue obtained by covering end-users and the SON operator's budget. We compare the performance of the exact and heuristic approaches in terms of the obtained results and computing time. We also provide bounds to the performance achievable by any optimization algorithm solving continuous relaxations of the integer models.

To this end we consider both randomly generated network instances and real ISP topologies mapped by the Rocketfuel tool [25], [26]. Random network topologies are obtained using a custom generator as well as hierarchical (Transit-Stub) models generated by the GT-ITM topology generator [27], [28], and finally using a degree-based generator (BRITE [29], [30]) to obtain topologies with node degree power laws.

To generate random network instances, we have implemented a topology generator which considers a square area with edge equal to 1000, and randomly extracts the position of m Candidate Sites (CSs), n Test Points (TPs) and p Destination Nodes (DNs). The area is divided into N Internet Service Providers (ISPs); for sake of simplicity in this paper we consider $N = 25$ ISPs obtained dividing the whole area into $L \times L$ squares, with $L = 200$.

We assume that each TP and DN can be connected to a CS only if the CS is at a distance not greater than 100 from the TP or DN. As for the connectivity parameters between different CSs, we assume that each CS can be directly connected with an overlay link to any other CS (i.e., $b_{jl} = 1, \forall j, l \in S$); this allows our models to investigate all possible link configurations to find the optimal overlay topology.

The cost matrix for bandwidth (c^B_{jl}) is then generated. If CSs j and l belong to the same ISP, we assume that c^B_{jl} is fixed and equal to 1 monetary unit per Mb/s. On the other hand, if CSs j and l belong to different ISPs, c^B_{jl} depends on the peering agreements between such ISPs. For the sake of simplicity, we assume that in this case c^B_{jl} is a random variable uniformly distributed between $C/2$ and $3C/2$, with C being equal to $\frac{L_{jl}}{L}$, that is the distance between j and l (L_{jl}) divided by the width of an ISP domain (L), i.e. 200 with the above settings.

If not specified differently, the installation cost of an overlay node is equal to 10 monetary units. As for the access and

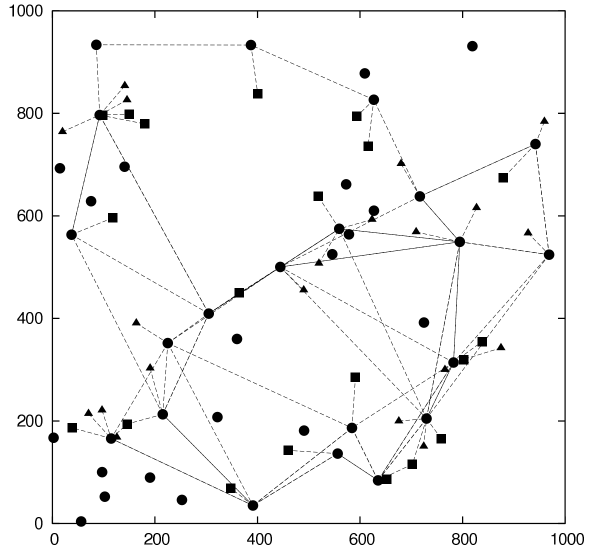

(a) 500 kb/s

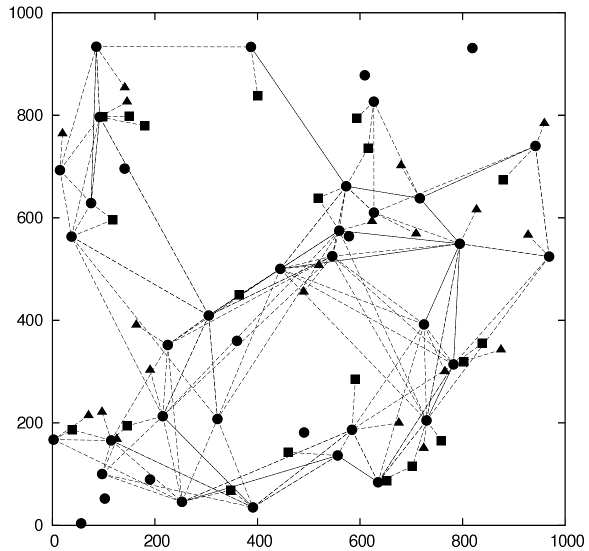

(b) 1000 kb/s

Fig. 3. Sample SONs planned by the FCSD model with increasing traffic demands (500 and 1000 kb/s). The number of TPs and DNs is 20, while the number of CSs is 40. CSs, TPs and DNs are represented with circles, triangles and squares, respectively.

egress cost, we assume they are fixed and equal to 1 monetary unit per Mb/s.

The maximum capacity that can be reserved between CSs j and l on the overlay link (j, l) u_{jl}, $j, l \in S$ is set equal to 50 Mb/s, as well as the maximum capacity of the access link of CS j, v_j, $j \in S$. The capacity of the egress links connecting overlay nodes to destination nodes is $h_{jk} = 100$ Mb/s, for all $j \in S$ and $k \in D$.

Obviously, none of the above assumptions affects the proposed models and heuristics which are general and can be applied to any problem instance and network topology.

Identifying Candidate Sites can be a difficult task in real ISPs networks. To solve this problem, a topology-aware node placement heuristic could be used, as proposed in [15], to decide the potential locations for overlay nodes inside an ISP. Such techniques can be used together with our heuristics and models, thus representing a further research topic worth pursuing.

All the results reported hereafter are the optimal and approximate solutions of the considered instances obtained, respectively, by formalizing the proposed models in AMPL [31] and solving them with CPLEX 9.1 [32], or using the proposed heuristics. For each network scenario, the results are obtained averaging each point on 10 network instances.

a) Effect of the Traffic Demands: Random network instances

We first consider the Full Coverage SON Design problem in a random network scenario with $n = 20$ TPs and $p = 20$ DNs. Each test point offers the same amount of traffic d_{ik} to all destination nodes.

Figure 3 reports an example of the planned networks when applying the FCSD model to the same instance with $m = 40$ candidate sites and with two different requirements on the end-user traffic, $d_{ik} = 500$ kb/s and $d_{ik} = 1$ Mb/s for all TPs and DNs. CSs, TPs and DNs are represented with circles, triangles and squares, respectively. As expected, increasing the traffic demands forces the model to install a higher number of overlay nodes and links to convey the traffic towards the destination nodes.

Table II analyzes the characteristics of the solutions of the FCSD model, its continuous relaxation and the H-FCSD heuristic in the same scenario when varying the number of candidate sites.

For each couple (m, d_{ik}) and optimization algorithm, the Table reports the number of installed overlay nodes (N_R) and links (N_L), the total network cost and the computing time (measured in seconds) to obtain the solution.

For small instances we could obtain the exact solutions with the FCSD model, which enabled a comparison between the H-FCSD heuristic and optimal integer solutions; column gap_I reports the percentage gap between the cost provided by H-FCSD and the optimal cost obtained with the FCSD model. This gap shows how close is H-FCSD to the optimum obtained by FCSD, in terms of the objective function value.

We also reported the results obtained solving the continuous relaxation of the FCSD model, which provides a lower bound on the network cost that can be obtained with any optimization algorithm; column gap_B shows the percentage gap between the cost provided by the H-FCSD heuristic and such lower bound. Note that such relaxation is used exclusively to provide a reference point, especially for large network instances where solutions using FCSD cannot be obtained.

Finally, column gap_L shows the percentage gap between the cost provided by the FCSD model and that obtained solving the continuous relaxation of FCSD. Thus, gap_L measures how far is the optimal integer solution from the relaxed one.

Three main results come from the observation of the Table: first, the very same effect of traffic increase observed in Figure 3 is evident also on averaged results; in fact, the number of installed nodes and links increases when increasing the traffic demands.

TABLE II

SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC, WITH 20 TPS AND 20 DNS.

d_{ik}=500 kb/s

	FCSD					FCSD Continuous Relaxation				H-FCSD					
m	N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
30	19.6	150.1	1001.3	12.4	7.72	13.8	205.3	929.5	0.3	19.8	146.3	1032.0	0.5	3.07	11.03
40	19.2	146.6	993.2	244.8	9.08	13.9	269.5	910.5	0.7	21.0	154.2	1025.9	1.3	3.29	12.67
50	19.5	148.3	981.9	4665.8	9.71	13.9	321.2	895.0	1.8	21.4	153.7	1022.6	1.6	4.15	14.26
60						13.8	380.6	879.8	3.2	22.2	160.0	1018.2	3.2		15.73
80						13.7	491.5	858.3	10.9	23.3	163.5	1017.9	8.2		18.59
100						13.7	603.4	845.5	28.0	23.6	163.1	1013.4	15.9		19.86

d_{ik}=1000 kb/s

	FCSD					FCSD Continuous Relaxation				H-FCSD					
m	N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
30	21.7	167.8	1813.4	5.6	4.54	15.5	208.2	1734.6	0.3	22.2	165.8	1830.8	0.6	0.96	5.55
40	22.7	174.8	1795.1	65.3	5.72	15.9	266.3	1697.9	0.7	23.2	173.2	1829.9	1.3	1.94	7.77
50	22.9	175.1	1776.1	3655.5	6.60	15.8	324.1	1666.2	1.9	24.6	178.3	1829.5	3.5	3.01	9.80
60						15.8	380.6	1639.6	3.2	26.1	187.0	1823.2	6.3		11.20
80						15.4	492.7	1593.7	11.0	27.2	187.8	1821.3	13.2		14.28
100						15.5	595.5	1568.0	27.7	28.0	190.5	1817.5	30.0		15.91

TABLE III

LARGE-SIZE INSTANCES: SOLUTIONS PROVIDED BY THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC, WITH 100 TPS AND 10 DNS.

d_{ik}=20 kb/s

	FCSD Continuous Relaxation				H-FCSD				
m	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_B %
100	23.3	570.1	303.7	7.3	27.3	241.9	351.1	4.8	15.61
200	21.7	1036.0	284.8	128.9	27.2	236.3	350.0	58.8	22.89
300	20.8	1564.9	274.7	792.1	27.1	236.3	348.2	241.8	26.76

d_{ik}=40 kb/s

	FCSD Continuous Relaxation				H-FCSD				
m	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_B %
100	23.4	570.8	374.8	7.6	27.4	242.3	429.8	4.8	14.67
200	21.8	1035.1	353.5	131.3	27.4	237.6	429.7	57.9	21.56
300	20.9	1567.6	342.1	979.9	27.1	236.3	425.8	248.4	24.47

Second, for a given traffic value, increasing the number of CSs (m) increases the solution space; as a consequence, the model favors the solutions providing connectivity that have a lower impact on the network cost, which in turn decreases with m.

Finally, the proposed heuristic performs very close to the optimal solutions (less than 4.2% for $d_{ik} = 500$ kb/s and less than 3.1% for $d_{ik} = 1$ Mb/s), and in all cases the computation time is below 30 seconds for all the tested instances. Even when compared to the bound provided by the continuous relaxation, the performance gap is less than 20% in all network scenarios.

Since for small network sizes ($m = 30, 40, 50$) it can be observed that gap_I is significantly smaller than gap_B, we can guess that even for larger topologies (where FCSD cannot be solved and consequently gap_I cannot be computed) H-FCSD performs considerably closer to the optimal integer solution than what is indicated by the gap_B value.

We further observe that the gap_L value is quite small (less than 10%) for all network instances, which motivates the choice to use the continuous relaxation in our heuristics as a basis to obtain good integer solutions applying the randomized rounding technique.

We then considered a variation of this network scenario with large-scale instances containing $n = 100$ TPs and $p = 10$ DNs, which can be seen as acting like concentrator nodes or access points towards other networks.

The results obtained with the H-FCSD heuristic and the continuous relaxation of the FCSD model are shown in Table III with m ranging from 100 to 300 and for different d_{ik} values, and they are in line with the observations reported above.

A variation of the above scenarios is further considered, where the offered traffic is not constant, but uniformly distributed at random between $\frac{1}{2}d_{ik}$ and $\frac{3}{2}d_{ik}$. We observed that the results obtained are very close to those reported in Tables II and III, with a maximum gap inferior to 1%.

b) Effect of the Traffic Demands: Transit-Stub topologies

We then generated Transit-Stub topologies using GT-ITM [27]. In such scenarios, the Internet is modeled as a collection of interconnected routing domains, which can be classified as either Transit domains (that contain backbone nodes) or Stub domains (which have one or more gateway nodes that are connected to transit domains).

We considered 10 random Transit-Stub topologies with 100 nodes and an average of 550 links, including access and egress links. All nodes can be selected to install overlay nodes (i.e. $m = 100$ CSs) and each link can be selected as an overlay link. For each topology we generated 10 random distributions of $n = 100$ TPs and $p = 100$ DNs, where each TP offers the same amount of traffic $d_{ik} = 10$ kb/s to all destination nodes.

TABLE IV

TRANSIT-STUB TOPOLOGIES: SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC, WITH 100 TPS, 100 DNS, 100 CSS AND d_{ik}=10 KB/S.

FCSD					FCSD Continuous Relaxation				H-FCSD					
N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
38.0	338.1	1272.4	484.1	7.61	27.5	508.8	1182.4	5.3	48.5	347.2	1434.1	1.5	12.71	21.29

TABLE V

POWER LAW TOPOLOGIES: SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC IN THE TOPOLOGIES WITH NODE DEGREE POWER LAWS GENERATED USING BRITE, WITH 100 CSS, 20 TPS AND 20 DNS.

	FCSD					FCSD Continuous Relaxation				H-FCSD					
d_{ik}	N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
500	17.1	87.8	987.9	70.1	11.58	11.3	146.0	885.4	0.6	22.5	96.6	1272.7	0.3	28.83	43.74
1000	24.4	110.7	2022.8	48.8	12.18	15.3	155.9	1803.2	0.7	29.6	119.6	2583.2	0.4	27.70	43.26

All other parameters are the same as in the previous network scenarios. The numerical results obtained with the FCSD model, its continuous relaxation and the H-FCSD heuristic, averaged over all network topologies and random TPs/DNs distributions, are shown in Table IV. It can be seen that due to the hierarchical structure of Transit-Stub topologies, a large number of overlay links is selected in the planned SON. Moreover, even in this scenario H-FCSD performs close to the optimal solution provided by the FCSD model.

c) Effect of the Traffic Demands: Power Law topologies

To better capture the power law distributions for node degrees that characterize Internet topologies, we further considered BRITE, a degree-based topology generator [29], [30]. BRITE was used to generate power law topologies with 100 CSs, 20 TPs and 20 DNs on a square area with edge equal to 1000; the Barabási $-$ Albert model [33] was adopted with default parameters provided by BRITE.

The area was divided into $N = 25$ ISPs, each occupying an $L \times L$ square, with $L = 200$. Moreover, the bandwidth cost matrix was generated as for the custom topology generator.

Table V reports the numerical results obtained with such topologies, with an offered traffic d_{ik} equal to 500 and 1000 kb/s. The results confirm the behavior already observed for the random instances obtained with the custom generator, where the number of installed overlay nodes and links increases with increasing traffic demands. H-FCSD plans SON topologies with a reasonable cost in a very short computing time, which is almost independent of the traffic offered to the network.

d) Effect of the Traffic Demands: Rocketfuel topologies

We considered four diverse ISP backbone topologies, listed in Table VI, collected by Rocketfuel [25], [26], an Internet mapping tool.

In such topologies, all routers can be selected to install overlay nodes and all links can be selected as overlay links. The cost of each overlay link was set proportional to its length. For each topology we generated 10 random distributions of $n = 20$ TPs and $p = 20$ DNs, where each TP offers the same amount of traffic d_{ik} to all destination nodes. All other parameters are the same as in the previous network scenarios.

TABLE VI

ROCKETFUEL-INFERRED ISP TOPOLOGIES: NUMBER OF BACKBONE ROUTERS AND LINKS (INCLUDING ACCESS AND EGRESS LINKS)

Network	Location	Routers	Links
Ebone	EU	88	362
Tiscali	EU	164	696
Exodus	US	80	334
Abovenet	US	145	788

The numerical results, averaged over all network topologies and random TPs/DNs distributions, are shown in Table VII for $d_{ik} = 500$ and 1000 kb/s.

For these topologies we further measured the average length of the paths chosen by the FCSD model to route traffic flows; the path length is expressed in terms of the number of hops traversed by the flow (Hop Count, H_C). Such performance figure is reported in the last column of Table VII.

Note that in all such real scenarios the proposed heuristic performs very well and close to the optimum, thus representing a practical solution for the planning of SONs. Furthermore, the average path length increases with increasing traffic demands, since an increasing fraction of the offered traffic must be routed on longer (and generally costlier) paths due to capacity constraints.

e) Effect of the Cost

We evaluate the effect of the bandwidth cost as well as the nodes' installation cost on our models' performance, considering a random network scenario with $n = p = 20$ TPs and DNs and $m = 50$ CSs. The offered traffic d_{ik} is equal to 500 kb/s. The solution, and in particular the number of installed nodes and links, intuitively depends on two factors: the ratio β between the overlay nodes' installation cost and the bandwidth reservation cost, and the ratio δ between the cost of a link connecting two CSs which belong to different ISPs and that of a link connecting two CSs belonging to the same ISP.

Tables VIII and IX report the results obtained varying, respectively, the parameters β and δ. Table VIII shows that if the cost of installing an overlay node decreases with respect to the bandwidth reservation cost, the proposed models and heuristics tend to install more overlay nodes.

TABLE VII

ROCKETFUEL TOPOLOGIES: SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC, WITH 20 TPS AND 20 DNS.

d_{ik}=500 kb/s

	FCSD					FCSD Continuous Relaxation				H-FCSD						
Network	N_R	N_L	Cost	Time	gap_L	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I	gap_B	H_C
Ebone	26.8	115.6	2150.3	105.4	7.43	15.8	171.2	2001.6	0.4	34.9	135.0	2418.0	0.3	12.45	20.80	4.96
Tiscali	26.7	111.4	2021.3	79.2	7.72	16.3	184.5	1876.5	0.6	34.5	128.5	2394.4	0.4	18.46	27.60	5.14
Exodus	25.8	116.3	3432.9	36.3	4.32	14.2	159.1	3290.8	0.4	32.4	131.3	3729.6	0.3	8.64	13.33	4.33
Abovenet	24.0	114.0	5071.2	193.9	2.49	13.9	215.2	4948.0	1.0	33.5	135.6	5605.9	0.7	10.54	13.30	4.74

d_{ik}=1000 kb/s

	FCSD					FCSD Continuous Relaxation				H-FCSD						
Network	N_R	N_L	Cost	Time	gap_L	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I	gap_B	H_C
Ebone	32.1	135.5	4069.3	8.3	4.83	20.4	175.0	3907.4	0.6	41.5	155.9	4438.6	0.3	8.36	13.59	5.17
Tiscali	32.1	135.4	3813.7	20.7	4.80	20.7	183.8	3639.0	0.6	41.5	156.1	4408.7	0.5	15.60	21.15	5.30
Exodus	30.3	134.0	6659.0	5.4	2.55	19.1	160.1	6493.4	0.5	37.5	152.1	7170.5	0.2	7.68	10.43	4.55
Abovenet	29.8	138.5	9996.8	27.4	1.79	17.1	223.3	9821.0	1.0	40.4	166.4	11057.5	0.6	10.61	12.59	4.92

TABLE VIII

VARIABLE COST RATIO β: SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC WITH 20 TPS, 20 DNS, 50 CSS AND d_{ik}=500 KB/S.

	FCSD					FCSD Continuous Relaxation				H-FCSD					
β	N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
10	19.5	148.3	981.9	4665.8	9.71	13.9	321.2	895.0	1.5	21.4	153.8	1018.2	1.6	3.70	13.77
1	30.8	209.9	783.5	108.9	2.45	16.0	266.5	764.8	1.3	31.4	192.7	800.6	1.3	2.18	4.68
$1/10$	50.0	239.6	748.5	1.7	0.00	50.0	239.4	748.5	1.1	50.0	237.8	762.2	1.6	1.83	1.83

TABLE IX

VARIABLE COST RATIO δ: SOLUTIONS PROVIDED BY THE FCSD MODEL, THE FCSD CONTINUOUS RELAXATION AND THE H-FCSD HEURISTIC WITH 20 TPS, 20 DNS, 50 CSS AND d_{ik}=500 KB/S.

	FCSD					FCSD Continuous Relaxation				H-FCSD					
δ	N_R	N_L	Cost	Time	gap_L %	N_R	N_L	Cost	Time	N_R	N_L	Cost	Time	gap_I %	gap_B %
1	19.5	148.3	981.9	4665.8	9.71	13.9	321.2	895.0	1.5	21.4	153.8	1018.2	1.6	3.70	13.77
2	21.8	154.3	1355.9	13853.6	9.50	14.5	329.7	1238.3	1.4	24.2	162.4	1404.7	2.8	3.60	13.44
5	26.6	168.8	2356.0	7060.3	7.48	15.1	281.3	2192.0	1.5	30.3	174.3	2441.7	4.7	3.64	11.39
10	31.1	182.9	3925.7	2780.5	5.46	15.6	256.5	3722.4	1.3	35.2	184.5	4058.6	3.4	3.39	9.03

Note that the computing time of the H-FCSD heuristic is very short and the approximate solution is very close to the optimal solution provided by the FCSD model. Moreover, when β diminishes, the performance gap between H-FCSD and both the FCSD model and the lower bound greatly decreases.

Similarly, when δ increases, we observed that our model and heuristic tend to deploy more overlay nodes and links; the majority of these links connect overlay nodes belonging to the same ISP, thus diminishing the number of longer (and costlier) links that traverse different ISPs in the planned SON. This behavior is reflected in the numerical results reported in Table IX.

f) Effect of the Gain parameter

We then evaluate the Profit Maximization SON Design problem, considering a random network scenario with 20 TPs, 20 DNs, $m = 40$ CSs and offered traffic d_{ik}=500 kb/s. We assume that the gain per bandwidth unit that the SON operator obtains for serving an end-user (the parameter g_i in the objective function (11)) is a random variable with average equal to G and a uniform distribution between $G/2$ and $3G/2$, where G ranges between 0 and 0.01 monetary units per Mb/s.

Figure 4 reports an example of the planned networks when applying the PMSD model to the same instance considered in Figure 3 with $m = 40$ CSs and with two different requirements on the end-user traffic, $d_{ik} = 500$ kb/s and $d_{ik} = 1$ Mb/s for all TPs and DNs; G is equal to 0.005 monetary units per Mb/s. CSs, TPs and DNs are represented with circles, triangles and squares, respectively. In this case the planned SONs cover only a subset of the TPs to maximize the operator's profit.

Figure 5 shows the number of end-users covered by the SON as a function of G for both the PMSD model and the H-PMSD heuristic. Note that the two curves almost overlap, as the H-PMSD heuristic is quite accurate in determining the optimal number of users to serve to maximize the SON operator's profit. Obviously, for small G values, the SON is not profitable enough to cover any of the end-users; as G increases, the SON covers more end-users, eventually all of them. Similar results have been observed in several network topologies with a varying number of CSs.

Table X reports, for the PMSD model, its continuous relaxation and the H-PMSD heuristic, the total number of installed nodes and links, the SON operator's profit (i.e., the value of the objective function (11) or the approximate value obtained with H-PMSD) and computing time, as a function of G. Note that when G increases, the planned network covers

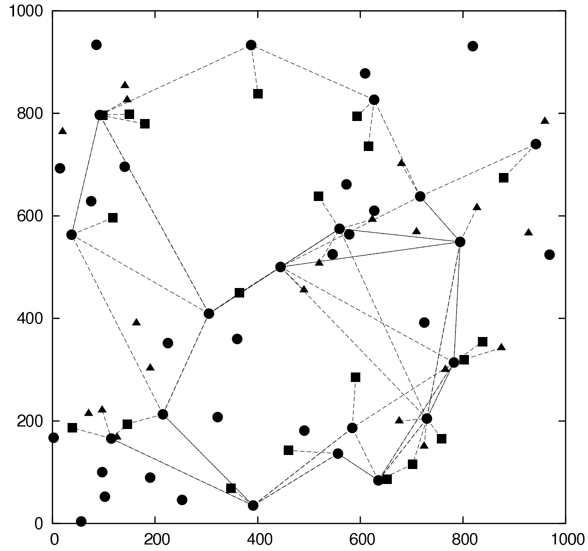

(a) 500 kb/s

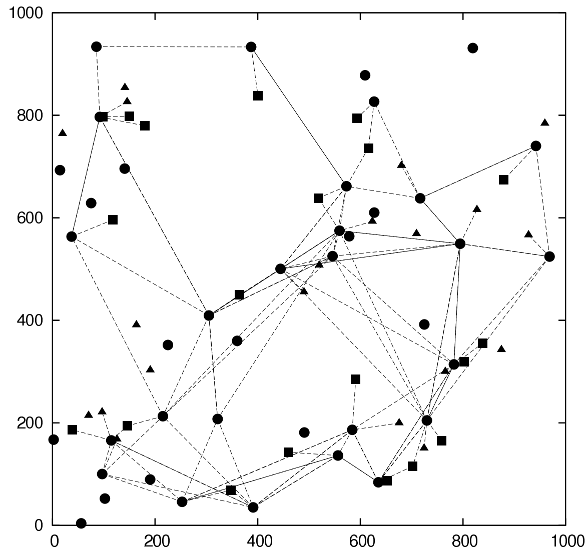

(b) 1000 kb/s

Fig. 4. Sample SONs planned by the PMSD model with increasing traffic demands (500 and 1000 kb/s). The number of TPs and DNs is 20, while the number of CSs is 40. The average gain per bandwidth unit, G, is equal to 0.005 monetary units per Mb/s. CSs, TPs and DNs are represented with circles, triangles and squares, respectively.

more end-users, and as a consequence it comprises more overlay nodes and links.

Column gap_I shows the percentage gap between the profit provided by the H-PMSD heuristic and the optimal profit obtained with the PMSD model, while column gap_B reports the gap between H-PMSD and the upper bound provided by the continuous relaxation of the PMSD model. Column gap_L provides the gap between the PMSD model and its continuous relaxation. The performance gap is slightly larger with respect to the full coverage design problem; this is mainly due to the difficulty in determining the optimal subset of users to

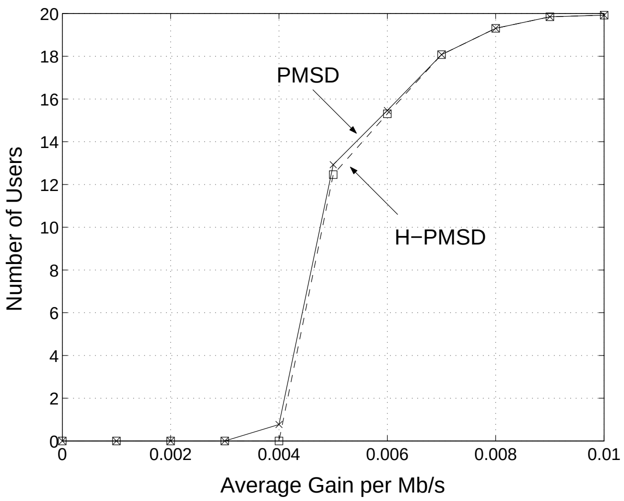

Fig. 5. Number of end-users covered by the SON with the PMSD model and the H-PMSD heuristic as a function of the average gain per bandwidth unit, with 20 TPs, 20 DNs, 40 CSs and d_{ik}=500 kb/s.

cover. However, note that the performance gap is still largely acceptable since in all cases it is less than 30%, and it greatly decreases with increasing values of G.

We further tested the H-PMSD heuristic with large-scale network instances. Table XI reports the numerical results obtained with 50 TPs, 50 DNs, 100 CSs and d_{ik}=20 kb/s for both the H-PMSD and the PMSD continuous relaxation. Even in this network scenario the proposed heuristic provides accurate solutions with a considerably reduced computing time.

g) Effect of the Budget parameter

Finally, to evaluate the effect that a budget constraint has on the planning of a SON, we consider several budget (B) values in the 500 to 1000 range, solving the PMSD model in a random network scenario with 20 TPs and DNs, 40 CSs and d_{ik} = 500 kb/s.

Figure 6 shows the number of end-users covered by the SON as a function of the operator's budget, for different G values. For each value of G, as the budget increases, the number of end-users accepted in the network increases until it reaches its maximum, which can be obtained observing Figure 5.

For small budget values ($B \leq 700$), all curves practically overlap except for that corresponding to $G = 0.005$ monetary units per Mb/s. This is due to the fact that for $G > 0.005$ the number of users covered by the PMSD model is limited by the budget constraint, so that the model chooses the most profitable users that can be accepted in the SON given the budget value. On the other hand, for $G = 0.005$ it is not profitable to cover the same number of users than for $G > 0.005$, even if the budget would allow the SON operator to serve them.

Table XII illustrates in detail the characteristics of the solutions provided in such scenario by the PMSD model, its continuous relaxation and the H-PMSD heuristic, for $G = 0.010$ monetary units per Mb/s and for different budget values.

TABLE X

SOLUTIONS PROVIDED BY THE PMSD MODEL, THE H-PMSD HEURISTIC AND THE CONTINUOUS RELAXATION WITH 20 TPS AND DNS, 40 CSS AND d_{ik}=500 KB/S.

	PMSD					PMSD Continuous Relaxation				H-PMSD					
G	N_R	N_L	Profit	Time	gap_L %	N_R	N_L	Profit	Time	N_R	N_L	Profit	Time	gap_I %	gap_B %
0.005	15.0	103.0	71.9	560.8	56.50	9.7	234.1	165.3	0.6	14.3	115.5	51.0	1.4	29.07	69.15
0.006	17.2	125.7	241.8	359.5	25.81	11.2	248.0	325.9	0.6	17.9	128.2	213.1	1.5	11.87	34.61
0.007	18.2	138.0	422.7	326.1	16.31	12.7	261.8	505.1	0.6	18.6	144.2	390.0	1.6	7.74	22.79
0.008	18.9	142.8	616.3	263.3	11.76	13.4	267.9	698.4	0.6	20.1	149.7	585.0	1.5	5.08	16.24
0.009	19.2	146.5	814.8	190.5	9.17	13.8	270.7	897.1	0.6	20.7	153.4	786.5	1.6	3.47	12.33
0.010	19.2	146.5	1015.0	195.0	7.51	13.8	270.7	1097.4	0.6	20.7	153.5	987.4	1.6	2.72	10.02

TABLE XI

LARGE-SIZE INSTANCES: SOLUTIONS PROVIDED BY THE H-PMSD HEURISTIC AND THE PMSD CONTINUOUS RELAXATION, WITH 50 TPS, 50 DNS, 100 CSS AND d_{ik}=20 KB/S.

	PMSD Continuous Relaxation				H-PMSD				
G	N_R	N_L	Profit	Time	N_R	N_L	Profit	Time	gap_B %
0.005	13.6	127.5	213.0	78.8	27.2	269.6	13.2	123.0	92.38
0.006	16.2	131.3	416.5	71.6	29.0	312.5	215.5	114.6	48.26
0.007	17.9	134.4	642.0	66.3	29.5	331.5	444.3	111.7	30.79
0.008	18.8	134.9	879.5	68.9	30.1	342.4	683.6	114.8	22.27
0.009	19.5	134.3	1120.7	66.3	30.5	351.4	929.3	111.8	17.08
0.010	19.9	134.4	1363.9	67.2	31.0	361.4	1171.8	111.9	14.08

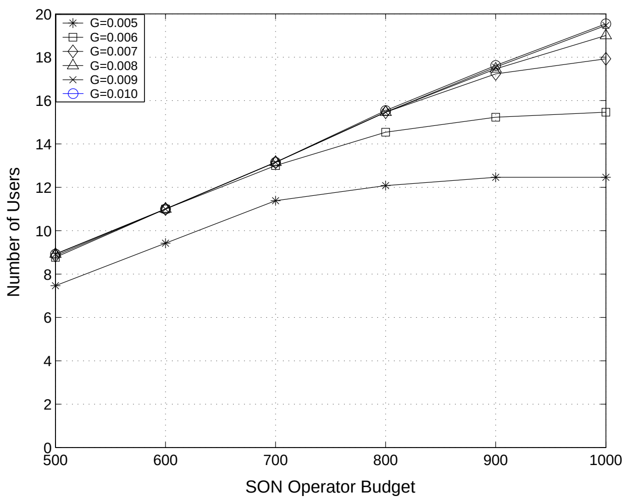

Fig. 6. Number of end-users covered by the SON with the PMSD model as a function of the budget for different values of the average gain per bandwidth unit G, with 20 TPs, 20 DNs, 40 CSs and d_{ik}=500 kb/s.

Note that in this scenario the scaling technique described in Section IV was used for H-PMSD.

The results show that higher profits are achieved by deploying higher-cost SONs. However, this also increases the economic risk the SON operator faces in the deployment of the overlay network.

Finally, note that the performance gap between H-PMSD and PMSD is quite small and decreases with increasing budget values.

VI. CONCLUSION

In this paper we addressed the topology design problem for Service Overlay Networks in terms of deciding the number and location of the overlay nodes to be deployed, the capacity reserved on each overlay link as well as the optimal subset of end-users to be covered in order to maximize the SON operator's profit.

To this end, we proposed two novel optimization models based on mathematical programming that take into account the individual requirements of the end-users, the connectivity between overlay nodes and the management of the traffic flows. The objective of the first model is the minimization of the total network installation cost while assuring full coverage of all end-users. The second model maximizes the SON profit choosing which users to serve based on the expected gain and budget constraints specified by the SON operator.

We further introduced two efficient heuristics that obtain near-optimal solutions for large-scale instances in a reasonable computation time.

Finally, we provided bounds to the performance achievable by any optimization algorithm solving continuous relaxations of the proposed integer models.

The numerical results we gathered show that our algorithms are able to capture the effect on the overlay topology configuration of all the considered network parameters, providing a promising framework for the design of SONs.

Since the proposed heuristics perform very close to the optimum with a short computing time, they can be also used for periodical SON redesign based on traffic statistics measured online.

REFERENCES

[1] Integrated Services Charter. Available at http://www.ietf.org/html.charters/OLD/intserv-charter.html.

[2] Differentiated Services Charter. Available at http://www.ietf.org/html.charters/OLD/diffserv-charter.html.

[3] Z. Duan, Z.-L. Zhang, and Y.T. Hou. Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning. *IEEE/ACM Trans. Networking*, pp. 870–883, vol. 11, no. 6, Dec. 2003.

[4] Z. Li and P. Mohapatra. QRON: QoS-aware Routing in Overlay Networks. *IEEE J. Sel. Areas Commun.*, pp. 29–40, vol. 22, no. 1, Jan. 2004.

TABLE XII

SOLUTIONS PROVIDED BY THE PMSD MODEL, ITS CONTINUOUS RELAXATION AND THE H-PMSD HEURISTIC WITH 20 TPS AND DNS, 40 CSS, $G = 0.010$ MONETARY UNITS PER MB/S AND d_{ik}=500 KB/S.

	PMSD					PMSD Continuous Relaxation				H-PMSD					
B	N_R	N_L	Profit	Time	gap_L %	N_R	N_L	Profit	Time	N_R	N_L	Profit	Time	gap_I %	gap_B %
500	14.3	87.2	588.1	3956.6	27.76	8.4	230.4	814.1	1.9	16.4	87.0	433.8	2.4	26.24	46.71
600	15.6	101.8	723.1	1394.5	21.07	9.8	246.4	916.1	1.5	17.0	99.0	507.2	2.2	29.85	44.63
700	16.4	111.9	827.4	1420.8	17.05	11.3	251.1	997.5	1.2	18.1	112.0	627.4	2.1	24.17	37.10
800	17.4	126.8	913.5	920.8	13.68	12.5	264.0	1058.3	1.1	18.8	123.7	686.0	1.9	24.91	35.18
900	18.2	136.0	973.2	674.9	11.07	13.5	270.7	1094.4	0.9	20.4	141.2	808.9	1.8	16.88	26.09
1000	19.2	145.6	1008.2	339.9	8.13	13.8	270.7	1097.4	0.7	20.6	147.0	877.1	1.8	13.00	20.08

[5] L. Subramanian, I. Stoica, H. Balakrishnan, and R. H. Katz. OverQoS: Offering Internet QoS Using Overlays. In *Proc. of the 1st Workshop on Hot Topics in Networks HotNets-I*, Princeton, New Jersey, USA, Oct. 2002.

[6] J. Touch and S. Hotz. The X-Bone. In *Proc. of the third Global Internet Mini-Conference*, pp. 75–83, Sydney, Australia, 1998.

[7] H.T. Tran and T. Ziegler. A design framework towards the profitable operation of service overlay networks. *Computer Networks*, pp. 94–113, vol. 51, no. 1, 2007.

[8] E. Rosen, A. Viswanathan, and R. Callon. Multiprotocol Label Switching Architecture. In *IETF RFC 3031*, Jan. 2001.

[9] S. L. Vieira and J. Liebeherr. Topology Design for Service Overlay Networks with Bandwidth Guarantees. In *Proc. of the 12th IEEE International Workshop on Quality of Service, IWQoS*, pp. 211–220, Montreal, Canada, June 2004.

[10] Z. Li and P. Mohapatra. On investigating overlay service topologies. *Computer Networks*, pp. 54–68, vol. 51, no. 1, 2007.

[11] R.R. Boorstyn and H. Frank. Large-Scale Network Topological Optimization. *IEEE Trans. Commun.*, pp. 29–47, vol. 25, no. 1, Jan. 1977.

[12] M. Pioro and D. Medhi. *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann, 2004.

[13] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Topologically-aware overlay construction and server selection. In *Proc. of IEEE Infocom'02*, pp. 1190–1199, vol.3, New York, USA, June 2002.

[14] A. Kershenbaum, P. Kermani, and G.A. Grover. MENTOR: an algorithm for mesh network topological optimization and routing. *IEEE Trans. Commun.*, pp. 503–513, vol. 39, no. 4, Apr. 1991.

[15] J. Han, D. Waston, and F. Jahanian. Topology Aware Overlay Networks. In *Proc. of IEEE Infocom'05*, Miami, FL, 13-17 Mar. 2005.

[16] J. Fan and M.H. Ammar. Dynamic Topology Configuration in Service Overlay Networks: A Study of Reconfiguration Policies. In *Proc. of IEEE Infocom'06*, Barcelona, Spain, Apr. 2006.

[17] S. Shi and J. Turner. Placing servers in overlay networks. In *Proc. of the Int. Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS) 2002*, San Diego, CA, July 2002.

[18] B.D. Vleeschauwer, F.D. Turck, B. Dhoedt, and P. Demeester. On the construction of QoS enabled overlay networks. In *Proc. of the Fifth International Workshop on Quality of future Internet Services (QofIS04)*, pp. 164–173, Barcelona, Spain, Oct. 2004.

[19] S. Roy, H. Pucha, Z. Zhang, Y.C. Hu, and L. Qiu. Overlay Node Placement: Analysis, Algorithms and Impact on Applications. In *Proc. of the 27th International Conference on Distributed Computing Systems*, Toronto, Canada, June 2007.

[20] S.Y. Shi and J.S. Turner. Multicast routing and bandwidth dimensioning in overlay networks. *IEEE J. Sel. Areas Commun.*, pp. 1444–1455, vol. 20, no. 8, Oct. 2002.

[21] D. Pompili, C. Scoglio, and L. Lopez. Multicast algorithms in service overlay networks. *Computer Communications*, pp. 489–505, vol. 31, no. 3, Feb. 2008.

[22] C.K. Yeo, B.S. Lee, and M.H. Er. A survey of application level multicast techniques. *Computer Communications*, pp. 1547–1568, vol. 27, no. 15, Sept. 2004.

[23] S. Fahmy and M. Kwon. Characterizing Overlay Multicast Networks and their Costs. *IEEE/ACM Trans. Networking*, pp. 373–386, vol. 15, no. 2, Apr. 2007.

[24] R. Raghavan and C.D. Thompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, pp. 365–374, vol. 7, no. 4, 1987.

[25] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *Proc. of ACM SIGCOMM 2002*, Pittsburgh, PA, Aug. 2002.

[26] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. Inferring Link Weights using End-to-End Measurements. In *Proc. of the 2nd Internet Measurement Workshop*, Marseille, France, Nov. 2002.

[27] GT-ITM: Modeling Topology of Large Internetworks. Available at http://www.cc.gatech.edu/projects/gtitm/.

[28] E. Zegura, K.L. Calvert, and S. Bhattacharjee. How to model an Internetwork. In *Proc. of IEEE Infocom'96*, pp. 594–602, vol.2, San Francisco, CA, Mar. 1996.

[29] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An Approach to Universal Topology Generation. In *Proc. of MASCOTS 2001*, Cincinnati, OH, Aug. 2001.

[30] A. Medina, I. Matta, and J. Byers. On the Origin of Power-Laws in Internet Topologies. *ACM Communications Review*, pp. 18–28, vol. 30, no. 2, Apr. 2000.

[31] *AMPL: A Modeling Language for Mathematical Programming*. Available at http://www.ampl.com.

[32] ILOG Optimization Products. ILOG CPLEX. http://www.ilog.com/products/cplex/.

[33] A.L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, pp. 509–512, vol. 286, no. 5439, 1999.

Antonio Capone is an Associate Professor at the Information and Communication Technology Department, Politecnico di Milano. His main research activities include protocol design and performance evaluation of wireless access and multi-hop networks, traffic management and quality of service issues in IP networks, network planning and optimization. He received the M.S. and Ph.D. degrees in electrical engineering from the Politecnico di Milano in 1994 and 1998, respectively. He is a Senior Member of the IEEE.

Jocelyne Elias received her Master of Computer Sciences from the Lebanese University of Tripoli in 2002, the Masters Degree from UTT, Troyes, in 2003, and the Ph.D. degree in Computer Science from UPMC, Paris, in July 2006. She is now working with the Department of Electronics and Information at the Politecnico di Milano. Her current research activities include dynamic resource allocation in quality of service networks, network planning and overlay networks optimization.

Fabio Martignon received the Laurea and the Ph.D. degree in telecommunication engineering from the Politecnico di Milano in October 2001 and May 2005, respectively. He is now an assistant professor in the Department of Information Technology and Mathematical Methods at the University of Bergamo. His current research activities include routing and MAC for multi-hop wireless networks, network planning, congestion control and QoS routing over IP networks.