Maximum Coverage at Minimum Cost for Multi-Domain IP/MPLS Networks

M. Yannuzzi, X. Masip-Bruin
R. Serral-Gracià, E. Marin-Tordera
Department of Computer Architecture
Technical University of Catalonia
Vilanova i la Geltrú, Catalonia, Spain

A. Sprintson
Department of Electrical and
Computer Engineering
Texas A&M University
College Station, Texas, USA

A. Orda
Department of Electrical
Engineering, Technion
Israel Institute of Technology
Haifa, Israel

Abstract—At present, service providers have several incentives to extend the reach of long-lived MPLS paths across domains. Providers, however, will face a number of trade-offs while choosing the optimal set of MPLS paths to be established. In this paper, we focus on the multi-objective decision problem of maximizing the traffic demands to be covered by long-lived MPLS paths from a source domain S to its major destination domains, while minimizing the monetary costs incurred. The problem is formulated subject to a budget constraint, which assures the minimum expected revenue for the provider in S. A major advantage of the analysis and solution proposed in this paper is that it can be easily generalized, and applied in other settings where constrained problems considering maximum coverage vs. cost are critical.

I. INTRODUCTION

One of the major motivations for extending IP/MPLS paths beyond domain boundaries is the expansion of the Virtual Private Network (VPN) offer of service providers. Another clear incentive is for content providers, since they can exploit aggregate MPLS paths to reach geographically spread groups of consumers, without maintaining a large number of distributed replicas of their delivery platforms; it is also easier for them to adapt when the distribution of major consumers changes. An additional incentive is that service providers are trying to offer value added services to their customers, by guaranteeing the quality and reliability of their communications, even when the other end-point of the communication is outside their administrative domain.

A limitation, however, is that end-to-end MPLS connectivity is only feasible if intermediate transit providers support the establishment of MPLS paths through them. To this end, transit providers will negotiate special peering agreements with their MPLS customers, for which the latter will be charged. Independently of the incentives that providers might have for extending MPLS beyond their boundaries, providers will naturally incur additional costs for the establishment and maintenance of long-lived MPLS paths through transit providers. Therefore, a key problem to be faced by providers is how to optimally solve the trade-off of exploiting as much as possible the advantages of having long-lived MPLS coverage, against the extra cost that this coverage will represent. In

practice, providers will have an estimation of the additional income that can be obtained due to the MPLS coverage, as well as an expectation in terms of revenue. The difference between these two determines the budget that can be spent.

A provider will often have multiple candidate paths to reach the destination domains that it is willing to cover, and these paths can have associated different monetary costs. In practical settings, providers might not always have the chance to choose the cheapest alternative for their coverage strategy. This could be either because some of the targeted domains are not reachable through the cheapest alternatives (e.g. due to peering or routing policies), or because the available capacity is not enough to allocate the traffic demands.

The subject of this paper is to formulate and efficiently solve the multi-objective decision problem of how a domain can maximize the MPLS coverage of traffic demands with minimum cost, subject to a budget and network capacity constraints. The problem is studied in the context of a Path Computation Element (PCE)-based multi-domain MPLS network. In our model, the PCEs are used for both the computation of the paths and the advertisement of routing information, taking advantage of the aggregated representation for a multi-domain network that we developed in a previous work [1].

We propose an offline solution, based on the knowledge of the aggregated representation of the network provided by the PCE in the domain and the monetary costs of the candidate paths. The details about how a domain is aware of the monetary costs of the paths offered by its providers are out of the scope of this paper (they could be either attached to the advertisements exchanged between the PCEs, or they could be handled offline). This study makes the following contributions.

- Based on realistic assumptions and with only minor knowledge about the pricing schemes of transit providers, we propose a new *Evolutionary Multi-objective Algorithm (EMA)* that is able to find candidate solutions in the Pareto sense in a highly effective way.
- It can be shown that the search and update strategy of our EMA guarantees the elitism of the candidates monotonically, and that our algorithm converges to an ε-Pareto Set [2].
- We show that our EMA is capable of finding much

better solutions than another powerful and highly utilized EMA, namely, the *Strength Pareto Evolutionary Algorithm* version 2 (SPEA2) [3]–[5], while complying with the network capacity and budget constraints.

II. THE NETWORK MODEL

The multi-domain MPLS network model proposed here is supported by a PCE-based architecture. Our focus is on a rather small set of MPLS domains Ω, in which pairs of domains have already negotiated peering agreements supporting the transit or termination of MPLS paths. In this scenario, the multi-domain MPLS network is captured by the representation $G(V, E)$ that includes an aggregated representation of each domain in Ω, as well as inter-domain links. In this study, we use the aggregated representation scheme for a multi-domain network that we proposed in [1]. This aggregated representation captures the transitional characteristics of the network, while guaranteeing that the confidentiality and administrative limits between domains are respected. A key advantage of this approach is that it allows the PCE located in a source domain to optimally compute entire paths towards any reachable destination domain in Ω, avoiding coarse-grained solutions, like those relying on crankback (notice that only destination domains in Ω for which the source domain has already negotiated a peering agreement will be reachable by the source).

In this general setting, a source domain S would like to establish long-lived MPLS paths to K destination domains $D_1, \ldots, D_K$. The total traffic volume to be covered between S and each D_i, $1 \leq i \leq K$, is known and it is denoted as d_i. Each link e in E has a bound $\hat{C}_e$ on the total amount of traffic that can be forwarded through e.

Based on the aggregated representation of the network, the PCE in the source domain S will typically have multiple alternative paths for the establishment of an MPLS path towards a destination domain D_i (see Fig. 1), each of which may have a different cost. For each destination domain D_i, the PCE in S has a set of feasible paths $P_i^1, \ldots, P_i^{N_i}$ that connect S and D_i, where N_i denotes the number of such paths. Here, the term "feasible" indicates that each of the N_i candidate paths will support the establishment of an MPLS path for the expected traffic demand d_i between S and D_i. Each path P_i^j, $1 \leq j \leq N_i$, is associated with a cost $c(P_i^j)$ that captures the total amount and reliability of network resources that need to be allocated in order to support forwarding traffic of volume d_i along P_i^j. We assume that domain S is multi-homed to M ISPs (see Fig. 1), each of them with its own pricing scheme. Typically, the difference in the cost of the feasible paths will depend on the reliability offered by the different ISPs of S, as well as on their charging schemes (total volume, percentile-based, etc.) [6]. The source S has an overall budget B that can be spent on all paths.

In this framework, the goal of the source domain S is to maximize the coverage of traffic demands using MPLS paths with minimum cost, subject to the network capacity and the budget restriction B.
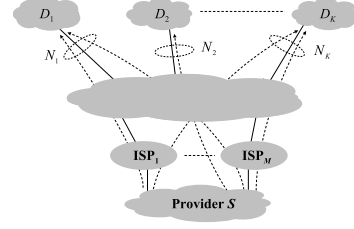

Fig. 1. The Network Model.

III. PROBLEM FORMULATION

This section formulates the problem, and introduces a basic – and realistic – set of assumptions that will let us build the key contributions in this work.

Assumption 1: The pricing functions of S's ISPs are increasing and concave in the traffic demand (with decreasing unit cost as the traffic volume increases). This reflects the fact that MPLS transit providers will expect decreasing marginal returns as the traffic demand of customers grow. This is precisely the case of the typical pricing schemes offered by ISPs at the moment [6].

Assumption 2: The traffic demands that S would like to cover are all of the same nature. In other words, the traffic demands are only distinguishable by their volume. In a more complex scenario, S might have incentives to cover some of the traffic demands explicitly using more expensive paths. For the sake of simplicity, we consider that the traffic demands are equally treated by S.

Assumption 3: As in [6], we assume that traffic volume and reliability are the variables that MPLS providers will consider in their pricing functions. Each of S's ISPs has its own pricing function. In [6], the assumption was that reliability is a feature tightly linked to the ISP selection. In a multi-domain MPLS framework, reliability might be a different feature, since it depends on the peering agreements between S's ISPs and the transit domains that need to be crossed to reach each destination D_i. The model proposed here is more general, in the sense that reliability is not linked to the ISP subscription, and the same ISP might offer more reliable or less reliable paths depending on its agreements. Therefore, for a given traffic volume d_i, the variation in the cost $c(P_i^j)$ of the feasible paths P_i^j is due to the difference in their reliability as offered by S's ISPs. We note that the costs $c(P_i^j)$ are denoted without an ISP sub-index, given that they are functions of the paths per-se, rather than the ISP that is offering the candidate path.

This basic set of assumptions is all that is needed in order to establish the rest of our analysis. For each path P_i^j we associate a decision variable δ_i^j that is equal to 1 if the path P_i^j is selected, and 0 otherwise. Then, the problem can be formulated as the following multi-objective integer program (Table I introduces the notation used):

$$max\left[\sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j d_i \ , \ -\sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j c(P_i^j)\right] \quad (1)$$

$$s.t. \quad \forall i \in \{1,\ldots,K\} : \ \sum_{j=1}^{N_i}\delta_i^j \leq 1 \quad (2)$$

$$\forall e \in E : \ \sum_{P_i^j : e \in E}\delta_i^j d_i \leq \hat{C}_e \quad (3)$$

$$\sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j c(P_i^j) \leq B \quad (4)$$

$$\forall i \in \{1,\ldots,K\} \quad and \quad \forall j \in \{1,\ldots,N_i\} : \ \delta_i^j \in \{0,1\} \quad (5)$$

Expression (1) represents the objective vector, where minimizing the total cost is equivalent to maximizing its negative. Expression (2) ensures that at most one path is selected per domain. Expression (3) ensures that the total traffic on all selected paths that use link e does not exceed its capacity $\hat{C}_e$. Note that the summation here is over all paths that use that link. Expression (4) ensures that the total cost of all selected paths does not exceed the total budget B. Finally, (5) ensures that each δ_i^j is either 0 or 1.

Independently of the technique chosen for solving a multi-objective optimization problem like (1), the outcome is a set of candidate solutions, which, typically, will not contain one that is better than the rest in each of the objectives. In our problem, while some candidate solutions will increase the coverage (being more costly), others will be cheaper but at the price of providing poorer coverage. Therefore, an additional step is typically needed in a multi-objective optimization problem, namely, a *Decision Maker* (DM), which, based on the set of candidate solutions, chooses the one that better fits the problem that is being solved. A general and widely used approach for the decision maker is to use a linear combination of the normalized objectives. Our specific approach is to use:

$$DM = sup\left[\frac{w_1}{D}\sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j d_i + w_2 - \frac{w_2}{B}\sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j c(P_i^j)\right] \quad (6)$$

where, $w_k \in [0,1]$, $k = 1,2$ are the weights, and $w_1 + w_2 = 1$. It is important to note that, in (6), each of the normalized objectives varies between 0 and 1, since D is the total traffic volume that S would like to cover (see Table I), and the normalized cost objective is given under the budget constraint in (4).

The most general trade-off to solve is when $w_1 = w_2 = \frac{1}{2}$, i.e., when the objectives are equally important for the decision maker. Accordingly, we consider:

$$DM = DM \mid_{w_1 = w_2 = \frac{1}{2}} \quad (7)$$

It is worth highlighting some properties of (7). On the one hand, candidate solutions with $DM < \frac{1}{2}$ are of no interest to our purpose. From (7) it is easy to see that zero coverage (i.e.

Symbol	Description
Known Data	
$G(V,E)$	The aggregated representation of the multi-domain MPLS network
S	The source Provider
K	Number of destination domains for which S is willing to provide MPLS coverage
D_i	One of the destination domains to be covered: $i \in \{1,\ldots,K\}$
d_i	Aggregate traffic demand to be covered between S and D_i
D	Total traffic demands that S is willing to cover: $\sum_{i=1}^{i=K} d_i$
N_i	Number of feasible paths between S and D_i
$c(P_i^j)$	Cost of establishing and maintaining an MPLS path between S and D_i using path P_i^j
B	Budget of the source provider S
$\hat{C}_e$	The available capacity of link e before the coverage. Capacities are known by the PCE in S by means of the distributed routing process running between PCEs [1]. This routing process is what allows the PCE in S to assemble $G(V,E)$
Unknown Data	
δ_i^j	Decision variable $\{0,1\}$ depending if path P_i^j is chosen or not
$\hat{D}$	Total demand finally covered: $\hat{D} = \sum_{i=1}^{K}\sum_{j=1}^{N_i}\delta_i^j d_i$

TABLE I

$\delta_i^j = 0 \ \forall \ i,j$) offers the decision maker a value of $DM = \frac{1}{2}$. Thus, all candidate solutions with $DM < \frac{1}{2}$ are of no practical interest, given that the decision maker will assess them worse than having no coverage at all. It is clear that the problem we are trying to solve is to find the best possible trade-off between coverage and cost in a setting where blocking might occur. In other words, rather than spending the entire budget B, the provider in domain S is seeking the most beneficial combination of coverage vs. cost.

On the other hand, a $DM \geq \frac{1}{2}$ is where the potential solutions can be found. In particular, $DM^{max} = 1$ represents the optimal case (i.e., maximum coverage for free, see (7)).

It can be shown that the size of the problem is:

$$\sum_{n=0}^{K-1}\binom{K}{K-n}\prod_{i=1}^{K-n}(N_i+1) \quad (8)$$

which, with an average of $N_i = 5$ candidates paths per-destination and only 20 destination domains, gives more than 7.6e+16 alternatives for the coverage.

We note that a single objective version of the problem in (1) is a special case of the *Multi-dimensional Knapsack Problem* (*MKP*). It has been recognized that the MKP problem is a particularly difficult problem that requires sophisticated methods to be solved. We also note that in our problem the constraints matrix is relatively sparse, while the MKP problem is typically characterized by a dense constraint matrix. In this paper, we exploit this property to design an algorithm that finds a set of near-optimal solutions in an efficient manner. Our focus here is on the more general multi-objective problem.

Evolutionary multi-objective optimization techniques offer a suitable and efficient way of solving large problems like the MKP problem, and like our problem. The key of their efficiency lies in the search and update strategy adopted. In [2], we show that it is possible to take advantage of the minor knowledge about the concave shape of the pricing functions of S's ISPs (see *Assumptions* 1-3), and exploit it to develop a highly effective search and update strategy for the evolutionary algorithm that we propose in Section IV.

IV. EVOLUTIONARY MULTI-OBJECTIVE ALGORITHM

Our EMA, named *Maximum Coverage at minimum Cost* (MC^2), is shown in Algorithm 1. As many EMA do, we keep an *archive A* composed by the best potential solutions found throughout the evolutionary process. In our model, a *vertex V_L* in the objective space Y, represents the maximum coverage at minimum cost of the traffic demands for L out of the K domains D_i. Moreover, a vertex V_L is said to be *feasible* if the cost associated with its coverage does not exceed the total budget B, and there exists a combination of feasible paths with enough capacity to cover the traffic demands with $DM \geq \frac{1}{2}$. Clearly, the feasible vertices are always potential candidates for the optimal solution to our problem, given that they encode a feasible coverage at minimum cost. Therefore, for the initial population we choose the set of *feasible vertices* in the objective space – if any – or we randomly generate a set of candidates close to the vertices in case they were unfeasible.

Our algorithm performs the usual set of operations carried out by evolutionary techniques, namely, *Mating Selection*, *Variation* and *Environmental Selection*. The reader is referred to [7] for a comprehensive description about these operations.

In MC^2, the Mating Selection is performed between the maximum feasible vertex V_L and the $(L + n)$ subsequent non feasible vertices, where $0 \leq n \leq (K - L)$. The outcome of this mating produces individuals with similar characteristics to V_L, and also to $V_{L+n} \forall n$.

The Variation corresponds to a mutation towards more expensive paths, which are likely to support the traffic coverage.

On the other hand, the Environmental Selection compares the candidates found and only the bests are kept in the archive A. With each new generation we obtain the *Best_Candidate*, which is determined as the candidate with the highest DM found up to that moment. The iso-objective lines[1] through the successive *Best_Candidates* found along the evolutionary process are the basis for steering the search towards better solutions in the objective space (see Algorithm 1).

A detailed description of the overall search and update strategy used in MC^2, together with its elitism and convergence properties can be found in [2].

V. PERFORMANCE EVALUATION

To validate the performance of MC^2, we performed a series of simulations using the PAN European network [8], which

---

Algorithm 1   EMA $MC^2(G(S, D_i), c(P_i^j), d_i)$

Input: $G(S, D_i)$ - a graph with all the feasible paths between S and $D_i \forall i = 1, \ldots, K$
$c(P_i^j)$ - cost of the paths between S and $D_i \forall i, j$
d_i - Traffic volume to be covered for domain $D_i \forall i$

Output: The set of paths chosen, the covered demand, and its cost: $\{\delta_i^j, \sum \sum \delta_i^j d_i, \sum \sum \delta_i^j c(P_i^j)\}$

1:   Generate an empty archive A and the initial population
2:   Search for the max feasible vertex V_L and $\{V_{L+1}, \ldots, V_K\}$
3:   *Best_Candidate* $\leftarrow V_L$
4:   Compute iso-objective line through *Best_Candidate*
5:   **for** each new generation g until stop condition **do**
6:     Perform Mating Selection
7:     Perform Variation
8:     /* Perform Environmental Selection */
9:     **for** each individual o in offspring O **do**
10:       **if** $DM(o) <$ iso-objective through *Best_Candidate* **then**
11:        discard o
12:       **else**
13:        **if** $\exists$ individuals $a \in$ archive A / $o \succ a$ **then**
14:         replace *all* dominated individuals a by o
15:        **else**
16:         add o to A /* since $o \nsucc a \forall a$ in archive A */
17:        **end if**
18:        *Best_Candidate* $\leftarrow o$
19:        Compute iso-objective line through *Best_Candidate*
20:       **end if**
21:     **end for**
22:     /* End of Environmental Selection */
23:     $g \leftarrow O$ /* Update generation */
24:     Update stop condition
25:   **end for**

---

is composed by 28 domains and 41 inter-domain links. We considered 1 domain as the source S, 7 transit domains, and 20 destination domains. The simulations are divided into three different sets: i) variable traffic demands; ii) variable network capacity; and iii) variable budget B. The results that we show here are the outcome of 100 simulation rounds for each of the three different sets of evaluations. More specifically, we run: i) 100 simulations by randomly changing the traffic demands d_i to be covered for each of the 20 destination domains; ii) 100 simulations by randomly changing the network capacity; and iii) another set of 100 simulations by progressively increasing the budget constraint.

The trials are aimed at providing supporting evidence of the quality of the solutions that MC^2 is capable of finding. To this end, we contrast the performance of MC^2 against the powerful and widely used EMA SPEA2 [3]. SPEA2 is an improved version of SPEA (Zitzler and Thiele (1999) [9]). The strengths of SPEA2 lie in its elitism preservation, its fitness assignment, its density preservation technique, and as MC^2, it uses an external archive where non-dominated solutions are stored. SPEA – and particularly SPEA2 – has become a reference EMA as it is being actively used to solve large multi-objective problems in many different fields, including medicine, engineering, etc [4], [5].

---

[1]We define an iso-objective curve as the set $\{P \in Y \ / \ DM(P) = constant\}$. It can be easily shown that in our case, the iso-objective curves are lines with slope $(\frac{D}{B})$ (see (7)).

	Mean Coverage		STD Coverage	
	Variable Traffic Demands	Variable Capacity	Variable Traffic Demands	Variable Capacity
SPEA2	43%	55%	0.05	0.05
MC^2	52%	80%	0.04	0.15

	Mean Cost		STD Cost	
	Variable Traffic Demands	Variable Capacity	Variable Traffic Demands	Variable Capacity
SPEA2	51%	53%	~ 0	~ 0
MC^2	53%	55%	0.01	~ 0

TABLE II

MEAN AND STD FOR THE COVERED DEMAND AND ITS COST.

The simulation results shown here have a preconfigured limit of 3000 generations for both MC^2 and SPEA2, and SPEA2 is set up with the usual archive size of 100 individuals. All the results provided here can be reproduced, as the source code of MC^2 used during the tests is available from [10].

Table II shows the mean and standard deviation values for the variable traffic demands and variable capacity tests. Our results confirm that MC^2 is able to find much better solutions than SPEA2, since it considerably increases the coverage with comparable costs. For example, for the variable capacity case, while SPEA2 is only able to cover 55% of the traffic demands on average, spending 53% of the budget, MC^2 is able to cover 80% of the traffic demands, with just a 2% increase in the overall cost. Likewise, in the variable traffic demand case, MC^2 is able to pass from covering 43% to 52% of the demand, again with just a 2% increase in the cost. Table II also shows that, for the set of tested traffic demands and network capacity constraints, an investment of around half of the budget is where the best trade-off – hence the maximum revenue – is found.

The last study performed to validate the behavior of MC^2 involves changing the maximum budget with a constant increase of 1 unit per test. The budget units are left generic on purpose, in order to draw attention to the fact that its usage may be generalized to any currency or budget schema of a service provider. Figure 2 shows the comparison between both algorithms (recall that the higher the value of DM, the better is the quality of the solution obtained). MC^2 presents a smooth increase of the DM value, which is due to the fact that it
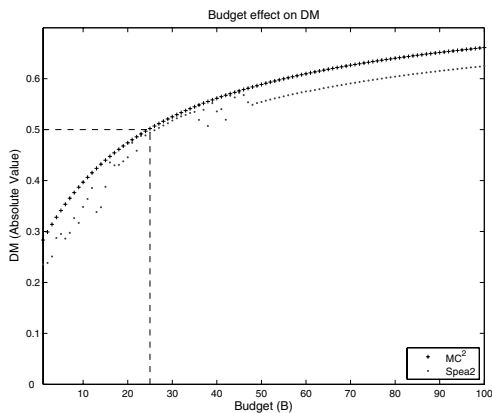
is able to locate the best candidates in the objective space. Clearly, the DM is limited by the budget at the beginning (recall that values less than $\frac{1}{2}$ are of no practical interest), but when enough budget is available, the network capacity is what actually limits reaching higher values of DM. Figure 2 also shows that a minimum budget is always required to have a feasible solution. Our simulations show that investing less that 25 units of budget will not provide the minimum revenue worth the effort, given that $DM < \frac{1}{2}$.

VI. CONCLUSIONS

In this paper we have formulated and solved the problem of how a domain can maximize the MPLS coverage of traffic demands with minimum cost, subject to a budget and network capacity constraints. The problem was formulated as a multi-objective integer program, and solved by means of a new evolutionary algorithm named *Maximum Coverage at minimum Cost* (MC^2) that we proposed and tested in this work. We have shown that MC^2 is able to find much better solutions than the powerful and widely used EMA SPEA2 [3]–[5]. The key is the way in which we exploit the knowledge about the concavity of common pricing functions of ISPs. The most promising outcome of this work is that the contributions are general in scope and can be applied to other problems. In particular, our proposals can be applied in settings where constrained problems considering maximum coverage vs. cost are critical, given that costs associated with concave pricing functions are widely used in practice.

Fig. 2.   Comparison of SPEA2 and MC^2 for different budgets.

REFERENCES

[1] A. Sprintson, M. Yannuzzi, A. Orda, and X. Masip-Bruin. "Reliable Routing with QoS Guarantees for Multi-Domain IP/MPLS Networks". In *Proc. of IEEE Infocom 2007*, Anchorage, Alaska, USA, May 2007.

[2] M. Yannuzzi, X. Masip-Bruin, R. Serral-Gracià, E. Marin-Tordera, A. Sprintson, and A. Orda. "Coverage vs. Cost for PCE-based MPLS Multi-Domain Networks". In *UPC-DAC-RR-CBA-2008-1, Broadband Communications Systems, Technical University of Catalonia*, Spain, January 2008. http://elio.ac.upc.edu/~dacsecre/reports/.

[3] E. Zitzler, M. Laumanns, and L. Thiele. "SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multi-objective Optimization". In *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)*. International Center for Numerical Methods in Engineering (CIMNE), 2002.

[4] T. Hiroyasu, M. Miki, S. Nakayama, and Y. Hanada. "Multi-Objective Optimization of Diesel Engine Emissions and Fuel Economy using SPEA2+". In *Proc. of 2005 Genetic and Evolutionary Computation Conference (GECCO'2005), ACM Press*, New York, USA, June 2005.

[5] J. Li et al. "Automatic Colonic Polyp Detection using Multiobjective Evolutionary Techniques". In *Proceedings of the SPIE, Vol. 6144, pp. 1742-1750, Medical Imaging 2006: Image Processing. Ed. by J. M. Reinhardt, J. P. W. Pluim,*, California, USA, January 2006.

[6] H. Wang, H. Xie, L. Qiu, A. Silberschatz, and Y. R. Yang. "Optimal ISP Subscription for Internet Multihoming: Algorithm Design and Implication Analysis". In *Proc. of IEEE INFOCOM 2005*, Miami, FL, USA, March 2005.

[7] E. Zitzler, M. Laumanns, and S. Bleuler. *"A Tutorial on Evolutionary Multi-objective Optimization"*. in X. Gandibleux et al., editors, Meta-heuristics for Multi-objective Optimisation, Lecture Notes in Economics and Mathematical Systems. Springer, 2004.

[8] R. Hulsermann et al. "A Set of Typical Transport Network Scenarios for Network Modelling". *ITG Fachbericht*, issue 182, 2004.

[9] E. Zitzler and L. Thiele. "Multiobjective Evolutionary Algorithms: A comparative case study and the strength Pareto approach". *IEEE Transactions on Evolutionary Computation*, 3(4), November 1999.

[10] MC2:. http://www.ccaba.upc.edu/mc2-results/index.html.