End-To-End Availability-Dependent Pricing of Network Services

Héctor Cancela[2], Pablo Rodríguez-Bocca[2], Bruno Tuffin[1] *
[1] IRISA-INRIA, Rennes, France
[2] Departamento de Investigación Operativa, Instituto de Computación, Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay
cancela@fing.edu.uy , prbocca@fing.edu.uy, btuffin@irisa.fr

Abstract:
We discuss how a new pricing scheme can be integrated within a communication network. The pricing scheme is based on the availability of end-to-end communications, and is an alternative to congestion pricing, which is not applicable when communication capacity is higher than demand (as happens in most communication backbone networks). We also investigate how, based on this scheme, an optimization algorithm for updating the network topology can be applied. The network update problem is modeled as a combinatorial optimization problem, which is approximately solved using a Genetic Algorithm. The good results obtained in a case study show that the method is robust and can be applied even when end-to-end availability measures can only be computed approximately (in this case, using a Monte Carlo method).

Keywords: *network design, pricing schemes, network availability, metaheuristics, genetic algorithms.*

---

* This research is part of the PAIR associated research project, supported by the INRIA, France, and has also received the support of ECOS-Sud, under Action U03E02. The participation of Pablo Rodríguez was supported by the French Embassy in Uruguay as part of the French Ministère des Affaires Étrangères scientific cooperation program; and by the "Programa de Jóvenes Investigadores" of CSIC, UDELAR, Uruguay.

Introduction

Devising new charging schemes for telecommunication networks has become a hot topic in the scientific community, as it is often said that the current flat-rate charges used in the Internet are an incentive for overusing the resources and that traffic continues to grow exponentially. Also, the network has to deal with applications having different quality of service (QoS) requirements. For instance, voice and video over IP require small delays and jitter, but can support some losses, whereas e-mail or file transfers do not support losses but are not delay sensitive. To ensure QoS for the different applications *in case of congestion*, a service differentiation scheme has to be devised, like in IntServ or Diffserv architectures, and a pricing scheme has to be attached to it, otherwise all customers will choose the best available service class, and the service differentiation becomes meaningless. This can lead to a substantial gain in revenue, see for instance Fishburn and Odlyzko (1998). The exhaustive surveys by Da Silva (2000), Falkner et al (2000), Henderson (2001), and Tuffin (2003) present and discuss the different approaches proposed in the recent literature.

We consider here an alternative viewpoint, based on the increasingly admitted observation that with the widespread use of optic fiber and improving technologies, the backbone networks are and probably will be over-dimensioned, so that in general congestion will not occur (see Fraleigh et al. (2003)). Our paper then focuses on backbone networks (with a possibility to include access links to represent access networks, wireless or wirelined, see the examples in Section 5). In this context, it could be more suitable to charge the network access, based on connection availability, representing the steady-state probability (i.e. proportion of time) a connection between two points is available.

We consider a network topology, where each link is assumed to have an infinite capacity (corresponding to over-dimensioning) but may not be available (due for example, to equipment or software failures) with a given probability. Each pair of nodes has then a probability to be connected. The price for each connection between a source s and a destination t depends on the availability of the connection between s and t. Of course the demand is also varying with this price, so that a first goal, discussed in Section 1, is to set up a price that maximizes the network revenue; for this we need to estimate availability measures, discussed in Section 2. Since those availabilities are high, rare event simulation is required. Note that all availabilities are relatively high (as can be checked in practice), but critical applications, such as medical ones, may require still better ones. In a second stage, the problem is to extend or in general modify the topology of an existing network in order to increase the service provider's net profit (the revenue minus the cost of modifying the network); this problem is formulated in Section 3, and can be approximately solved using a genetic algorithm described in Section 4. Finally, in Section 5 we present a test case inspired by the VTHD (Very High Broadband IP/WDM test platform) French network topology.

Our goal here is thus not to compare the proposed pricing scheme with congestion pricing, due to their different nature, but to propose an alternative as yet unstudied. Note also that the purpose of pricing is here to upgrade the network and propose a computationally feasible algorithm (that we do not claim to be the best possible). Pricing for network upgrade has been also proposed in congestion pricing by a fixed charge (two-parts tariff), see for instance Wang and Schulzrinne (2000). Our scheme proposes to finance this upgrade by imposing a larger cost to users having a better QoS (that is a higher availability in our paper; note that the same methodology can be employed with other QoS related parameters, such as jitter, delay, etc.).

1 Pricing model

We consider the network of an Internet Service Provider (ISP), represented by an undirected communication network $G = (N,E)$ consisting of a set of nodes N and a set of connecting links E (it is equally acceptable to assume a directed or mixed network). Let m be the number of links and n the number of nodes of G. We assume that each link will be over-provisioned, so that it is considered with infinite capacity. We consider that for each link $l \in E$, we can choose between different technology types, which have different costs and probabilities of failure. Design by T the set of types and, for each link $l \in E$, let $T(l) \subseteq T$ be the set of all the possible technology types applicable to link l. So, for each link $l \in E$ and each technology type of this link $t \in T(l)$, we assign an (independent from others) probability of failure $q_l(t)$ and a cost $c_l(t)$ (which can depend on link length, geography, technology amortization, operation and management associated costs, etc.). For simplicity, we assume that nodes do not have costs and that they do not fail.

Only a subset of nodes $K \subseteq N$ has connection demands, we call these nodes *terminals*. To each pair of terminals (s,t) with $s,t \in K$, we associate a total connection demand rate $\tilde{\lambda}_{s,t}$, a connection duration (assumed to be exponential with rate $\mu_{s,t}$), and an availability measure $r_{s,t}$, which corresponds to the probability that there is a path in the network connecting nodes s and t. Note that the assumption of exponential durations is introduced for tractability reasons, but that Poisson arrivals is a common assumption at the session level (see Ben Fredj et al. (2001)). Arrivals and connection durations are assumed to be independent for tractability reasons. To each pair of nodes (s,t) is also associated the utility function of getting a connection with availability r, modeled by a random variable $U_{s,t}(r)$, expressed in monetary units. The overall level of satisfaction is then $U_{s,t}(r)-p$ where p is the connection price. A customer will enter the network if and only if $U_{s,t}(r) \geq p$. The random variable $U_{s,t}(r)$ is characterized by its distribution: we denote by $F_{s,t}$ its cumulative distribution function and we define $\overline{F}_{s,t} = 1 - F_{s,t}$. The actual arrival rate of connections between s and t, $\lambda_{s,t}(p_{s,t})$, is given by $\lambda_{s,t}(p_{s,t}) = \tilde{\lambda}_{s,t} P(U_{s,t}(r_{s,t}) \geq p_{s,t}) = \tilde{\lambda}_{s,t} \overline{F}_{s,t}(p_{s,t})$.

Our goal is then to find out the optimal prices, for each pair (s,t), in terms of the availability measure $r_{s,t}$, maximizing the network revenue $H(G) = \sum_{(s,t)\in K} n_{s,t} p_{s,t}$ over the set of prices $p_{s,t} \geq 0$, for all s,t, with $n_{s,t}$ mean number of online (s,t)-connections. According to classical queuing theory for the M/M/∞ queue, we have $n_{s,t} = \frac{\lambda_{s,t}}{\mu_{s,t}}$ so that $H(G) = \sum_{(s,t)\in K} \frac{\tilde{\lambda}_{s,t}}{\mu_{s,t}} p_{s,t} \overline{F}_{s,t}(p_{s,t})$.

If we use first order conditions while maximizing this revenue (as the price $p_{s,t}$ is necessarily positive otherwise the revenue between s and t would be zero, meaning that the Lagrange multiplier is zero), i.e., $\partial G/\partial p_{s,t}=0$. s,t, (assuming that it gives the solution) we get $\frac{\partial}{\partial p_{s,t}}\left(p_{s,t}\overline{F}_{s,t}(p_{s,t})\right) = 0$, that $\overline{F}_{s,t}(p_{s,t}) + p_{s,t}\frac{\partial \overline{F}_{s,t}(p_{s,t})}{\partial p_{s,t}} = 0$.

In general, solving these equations can easily be carried out numerically, using Newton algorithm for instance. Nevertheless, we will make some additional hypothesis, leading to analytical results. In particular, we will suppose that, as in many economic applications, the utility is linear in its argument (here the availability) so that $U_{s,t}(r) = -D_{s,t}+\gamma_{s,t}r$ with $\gamma_{s,t}$ translating the availability in financial terms, as the monetary value of an availability unit (so that the utility increases with r) and $D_{s,t}$ being a random variable not depending on r and representing a dis-utility. Let $F^{*}_{s,t}$ be the distribution function of random variable $D_{s,t}$ not depending on r. Then the previous first order conditions become $F^{*}_{s,t}\left(\gamma_{s,t} r - p_{s,t}\right) + p_{s,t} \dfrac{\partial F^{*}_{s,t}\left(\gamma_{s,t} r - p_{s,t}\right)}{\partial p_{s,t}} = 0$.

In particular, if for all s, t $F^{*}_{s,t}(p) = \left(\dfrac{p}{M_{s,t}}\right)^{\alpha_{s,t}+1}$ with $0 \leq p \leq M_{s,t}$ and $\alpha_{s,t} > 0$ (so that the dis-utility increases with the price), from the first order conditions we obtain that $p_{s,t} = \gamma_{s,t} r_{s,t} / (\alpha_{s,t} + 2)$ provides the optimal prices and the maximum revenue. In the rest of this paper, we work with this demand function.

2 Computation of Network Availability Measure

The scheme described above takes as input the network availability measures for each pair of terminal nodes. Computing the availability measure $r_{s,t}$, is an NP-hard problem (see Provan and Ball, 1983, and Valiant, 1979), but there are efficient estimation methods which can be used to estimate its value. In this paper, we will employ Monte Carlo simulation with the Generalized Antithetic Variable method proposed by El Khadiri and Rubino (1992), as this variant performs better in computational time and in precision than the standard Monte Carlo technique.

As before, we consider the network as a graph $G = (N,E')$. We define a binary random variable X_l for each link l, called the state of the link: $X_l = 0$ means that link l is unavailable and $X_l = 1$ means that the link is available. The state of the network is completely characterized by the vector X whose components are the values X_l. If we fix two nodes s and t, the connection availability can be formalized employing a binary function $\Phi_{s,t}$, called the structure function, such that $\Phi_{s,t}(X) = 1$ if and only if s and t are connected in the graph defined by X. Finally, we denote by $r_{s,t}$ the availability measure between s and t, such that $r_{s,t} = \Pr(\Phi_{s,t}(X) = 1) = \mathrm{E}(\Phi_{s,t}(X))$ (this measure is also called in graph theory literature the source-terminal reliability of a graph, see Barlow and Proschan (1981)).

The Generalized Antithetic Variable method generates B independent blocks of L samples each one (notation: $X^{(b,1)},\ldots,X^{(b,L)}$ samples of block b). The L samples of a block are chosen in a dependent way that decreases the global variance (respect to the standard Monte Carlo). In order to estimate the measure $r_{s,t}$ we use the unbiased estimator

$$\hat{r}_{s,t} = \frac{1}{B}\sum_{b=1}^{B}\frac{1}{L}\sum_{i=1}^{L}\Phi_{s,t}\left(X^{(b,l)}\right).$$

To estimate the variance $\mathrm{Var}(r_{s,t})$ we use the unbiased estimator

$$\hat{v}_{s,t} = \frac{1}{B(B-1)}\sum_{b=1}^{B}\left(\frac{1}{L}\sum_{i=1}^{L}\Phi_{s,t}\left(X^{(b,l)}\right)\right)^{2} - \frac{1}{B-1}\hat{r}_{s,t}^{\,2}.$$

3 Extending the Network, based on Requests

The next step is to plan the capacity of the network. The idea is as follows: consider a family F of graphs such that for every graph G such that $G=(N,E') \in F$, the set N of nodes is the same, but the set of links E' is a subset of the set of feasible links E ($E' \subseteq E$).

In order to completely define a network $G=(N,E')$ in our model, we have to choose a technology type for each link $l \in E'$, this is modeled with the assignment function $a : E' \to T$ (where $a(l)$ means the technology type chosen for the l link, $a(l) \in \mathrm{T}(l)$).

From the network point of view, the goal is to determine the topology G=(N, E') (and the assignment function a) maximizing the benefits $H(G)-\sum_{l\in E'} c_l(a(l))$.

Now, we can summarize the formal problem used throughout the rest of the paper. Inserting in the revenue equation the optimal prices and the demand distribution functions we have

$$H(G)=\sum_{(s,t)\in K}\frac{\tilde{\lambda}_{s,t}}{\mu_{s,t}}\frac{\gamma_{s,t}r_{s,t}}{\alpha_{s,t}+2}\left[\frac{\gamma_{s,t}r_{s,t}-\gamma_{s,t}r_{s,t}/(\alpha_{s,t}+2)}{M_{s,t}}\right]^{(\alpha_{s,t}+1)}=\sum_{(s,t)\in K}\frac{\tilde{\lambda}_{s,t}\gamma_{s,t}r_{s,t}}{\mu_{s,t}(\alpha_{s,t}+2)}\left[\frac{(\alpha_{s,t}+1)\gamma_{s,t}r_{s,t}}{(\alpha_{s,t}+2)M_{s,t}}\right]^{(\alpha_{s,t}+1)}.$$

We then arrive at the following combinatorial optimization problem:

$$\textit{Maximize}\quad \sum_{(s,t)\in K}\frac{\tilde{\lambda}_{s,t}\gamma_{s,t}r_{s,t}}{\mu_{s,t}(\alpha_{s,t}+2)}\left[\frac{(\alpha_{s,t}+1)\gamma_{s,t}r_{s,t}}{(\alpha_{s,t}+2)M_{s,t}}\right]^{(\alpha_{s,t}+1)}-\sum_{l\in E'} c_l(a(l)).$$

where the decision variables are E', the edge set of G=(N,E'); and $a : E' \to T$, the technology assignment function; Note that although we have not made it explicit in the formulation, $r_{s,t}$ depends on both E' and a.

This formulation does not seem easy to exploit by an exact or numerical algorithm, as we have two sources of difficulty. On one hand, we have the combinatorial nature of the problem itself; on the other, as we already discussed, the fact that computing the availability measure $r_{s,t}$, is an NP-hard problem implies that to exactly compute even a single value of $H(G)$ is also NP-hard. An alternative is to employ metaheuristics, which have been used with success to solve many difficult combinatorial optimization problems. Among these techniques, we have chosen to employ Genetic Algorithms (GA), which, although computationally expensive, have given good results in other network design problems, as reported for example by Dengiz, Altiparmak and Smith (1997), Deeter (1998), Burgos et al (2003), Duarte, Barán and Benítez (2003). An important property of GA is that they are robust with respect to the computed values of the objective function, that is, they obtain good solutions even in the presence of (small) errors in the evaluation of the function to optimize.

4. Genetic algorithm design

We have followed a rather standard GA algorithm design. We describe (briefly, for space reasons) the components of the proposed algorithm:

i. *Encoding*: the genotype (solution encoded) is an array of size given by the number of edges of E, where we have an allele for each possible link in the network. The alphabet of each allele is an integer between zero and the maximum number of technology types of this link, where zero means that this link does not appear in the solution, and any other value that the link is present and that we use the corresponding technology type in this link.

ii. *Fitness function*: the objective function defined in Section 3 (the benefits of the network) plus the sum of the maximum costs of all the links (so that fitness is always positive). The availability measures are computed as discussed in Section 2.

iii. *Initial population*: generated randomly. Existing links are always included; each non-existing link is selected (or not) according to a Bernoulli variable of parameter 0.8. The type of included links is determined uniformly choosing between possible technology types.

iv. *Stopping criterion*: we tested two criterions, either to fix the number of generations; or to iterate while there is an "improvement" in the solution. The first one was selected, on the basis on tests over calibration problems.

v. *Selection*: "roulette wheel" selection with elitism; the best individual is always included in the next generation, and for the other individuals, the probability of including them in the next generation is proportional to their fitness over the population total fitness.

vi. *Crossover:* single point *Crossover* selecting the crossover point uniformly, and swapping all alleles of the parents between the sampled position and the end of the string. Crossover is applied to two randomly selected strings with a probability p_c (if this does not happen, the parents are copied exactly to the next generation).

vii. *Mutation*: the new value for the current allele is chosen uniformly between zero and the maximum number of technology types of this link. The value zero corresponds to removing the link; the other values include the link with a given technology type.

A remark is that the operations preserve the feasibility of the solutions; this is useful, because feasibility can be hard to maintain in a genetic algorithm when the problem has many constraints.

5 Numerical Illustration

The VTHD (Very High Broadband IP/WDM test platform; see http://www.vthd.org) network is a French project, whose main goal is to investigate the applications of a new generation of Internet and Intranet networks. We use this network as an illustrative application of our method. The VTHD network uses two main technologies types for its links: the backbone part of the network uses a IP/WDM architecture, with STM1/4 and STM16 links (in this work we suppose a 0.01 probability of non-availability for these links); the access part of the VTHD network uses Giga-Ethernet links (with a probability of failure of 0.1). Figure 1 shows the network for our illustrative example. The same network is shown schematically in Figure 2, also representing some feasible additional links (shown as dotted lines).

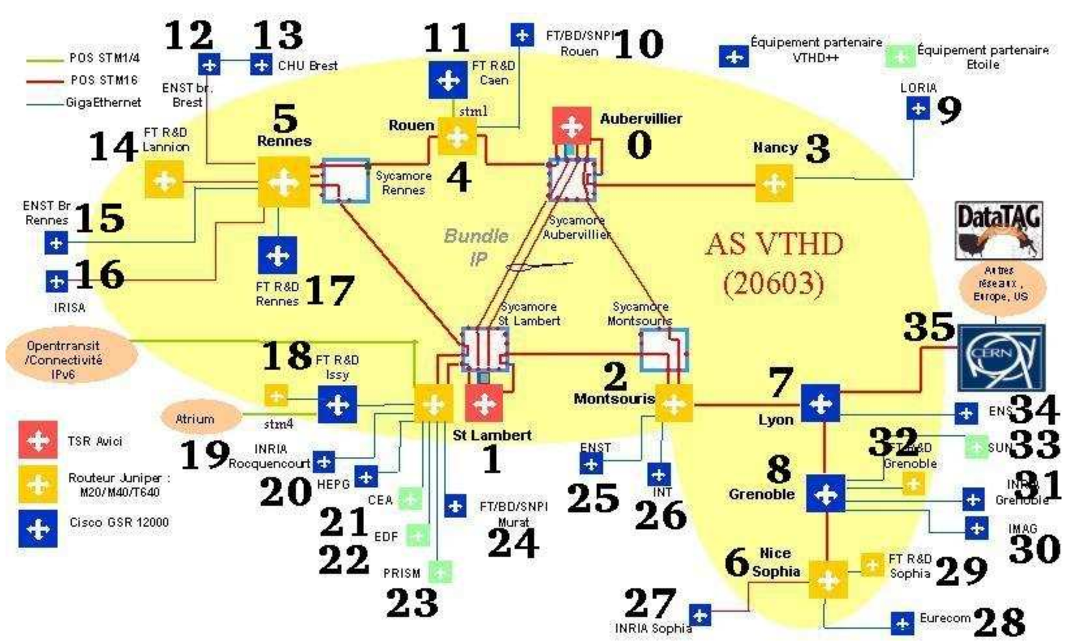

Figure 1: Validation Problem: Very High Broadband IP/WDM test platform.

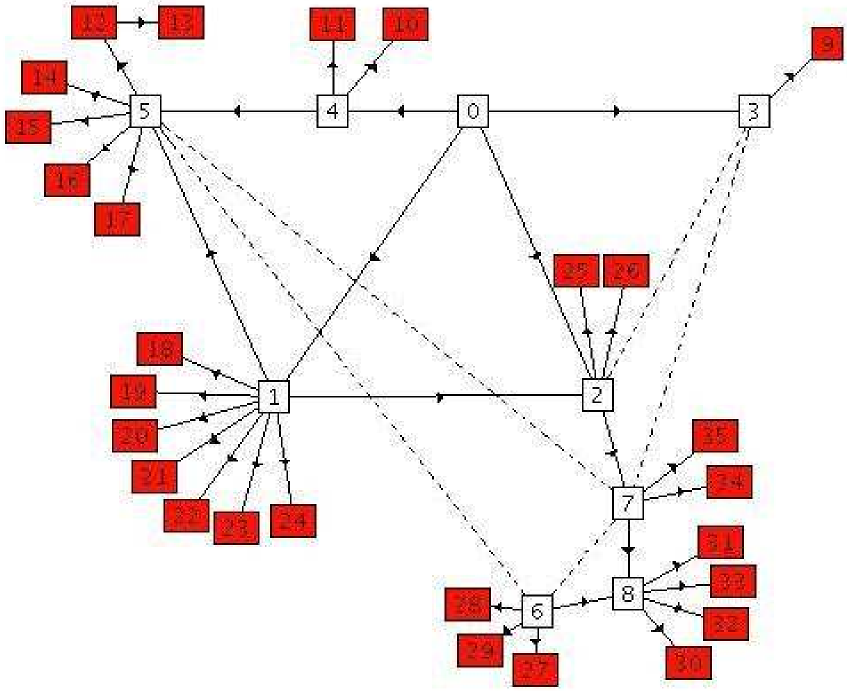

Figure 2: Graph corresponding to the Very High Broadband IP/WDM test platform.

Specifically, we apply the proposed GA method to three different scenarios. The three problems have the same specification (the same parameters of the demand, utility, etc.), but they differ in the possibilities of network extension. In the first problem (called VTHD1), we evaluate the benefits of extending the backbone of the network with the links shown as dotted lines; in this case the best solution of this problem can be found by enumeration because the network extension has only 32 possibilities. The second and third problems (called VTHD2 and VTHD3 respectively) add the possibility to upgrade the access network with IP/WDM links, that is to increase the availability of links connecting red nodes to the backbone in Figure 2, since they are less reliable. The difference between these two problems is the cost of updating the technology of the access links (in VTHD2 problem we use reasonably moderated costs, and in VTHD3 problem we consider very high costs). It is very hard to obtain the optimal solution for these two problems, as the solution space is very large (exactly 2^{25} possibilities). The VTHD1 optimal solution is also feasible for these problems, its value providing then a lower bound for their optima.
The experiments were run on a SunFire 280R, with two 1.2 GHz UltraSPARC III Cu processors, 2 GB of main memory, and Solaris™ 8 operating system. The parameters of the algorithm were chosen as follows: mutation rate p_m=0.01, crossover rate p_c=0.95, population size P=100, generation number G=100, generalized antithetic Monte Carlo block size B=100 and number of blocks L =50. These values were chosen on the basis of calibration experiments over a set of ten smaller problems.
The execution times for the three problems are similar. Each mutation takes in average 75.26 milliseconds, and each crossover takes 1.16 seconds. The mutation and the crossover are often executed in the execution of a genetic algorithm (exactly 10000 times the mutation and 5000

times the crossover, because we have a population of size 100, and 100 generations). The selection operator needs the fitness of the population to be computed, therefore, before each selection, we have to calculate the fitness of the new individuals, that implies an availability estimation. This estimation takes in average 1660.99 milliseconds, the consequence is that the algorithm execution time is approximately 5 hours.

Problem	Benefit of known feasible solution	Best Fitness of GA	Maximum Cost	Cost	Benefit
VTHD1	1630	1628	160	0	1468
VTHD2	1630	2137	370	80	1767
VTHD3	1630	3892	2260	120	1632

Table 1: VTHD solutions. The maximum cost of the network is the sum, for all links, of the most expensive technology type costs. The cost corresponds to $\sum_{l \in E'} c_l(a(l))$, i.e. the newly included and the upgraded links. The benefit is the objective function, i.e. $H(G) - \sum_{l \in E'} c_l(a(l))$.

Table 1 shows the results of the genetic algorithm. We found that for VTHD1 the GA obtains a solution very close to the optimal solution (with a difference well within the statistical error induced by the computation by Monte Carlo). In the VTHD2 problem, the solution found is quite better than the lower bound, and the selected links are in general quite different. The results of the GA for the VTHD3 problem are also quite better than the ones obtained for VTHD1, but a bit below the VTHD2 ones, this is expected as in the VTHD3 case, the technology upgrade costs are much higher.

In Figures 3 and 4 we show the evolution of the average fitness and best fitness of the population respectively. For problems VTHD1 and VTHD2, the initial population has already quite good fitness values; that is not the case for VTHD3. All the same, the GA attains quickly good values in the three cases; for VTHD1, the best value seems to be found in less than forty generations. In the case of VTHD2 and VTHD3, it is difficult to know if the optimum has been attained, but the evolution seems to have stopped after 80 generations.

An important point is the influence on the optimization procedure of the availability estimation error. We have estimated that the deviation coefficient of a single fitness evaluation in our experiments is about 3%. In order to evaluate its impact, we did five experiments for the VTHD2 case, using different random number seeds in the Monte Carlo procedure that estimates the availability measures. The deviation coefficient of the fitness values of the solutions obtained by the GA was about 1%, smaller than the deviation of a single fitness evaluation, and showing the robustness of the method. All the same, this value, even if small, can correspond to significant monetary amounts in the context of network design.

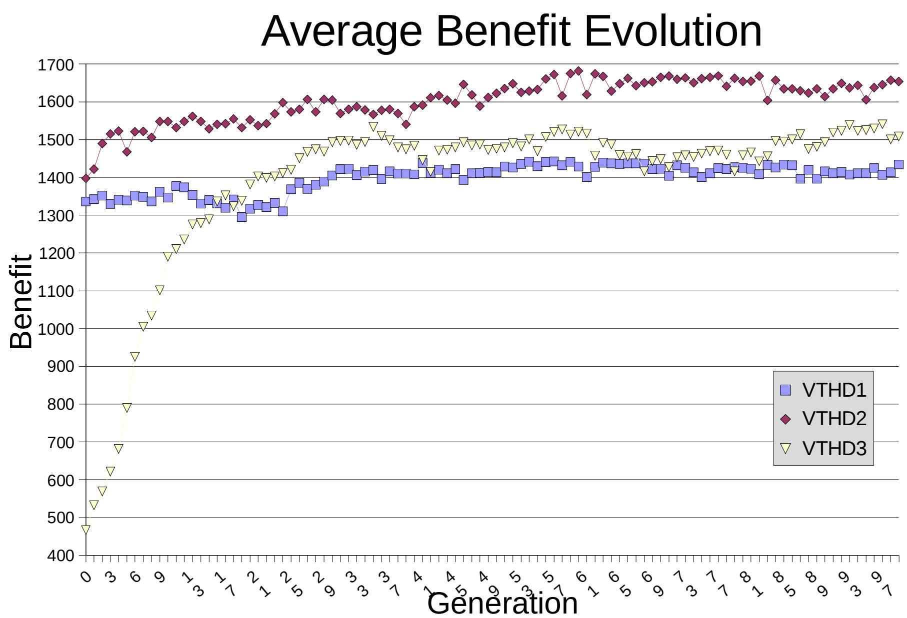

Figure 3: Population average for each GA generation.

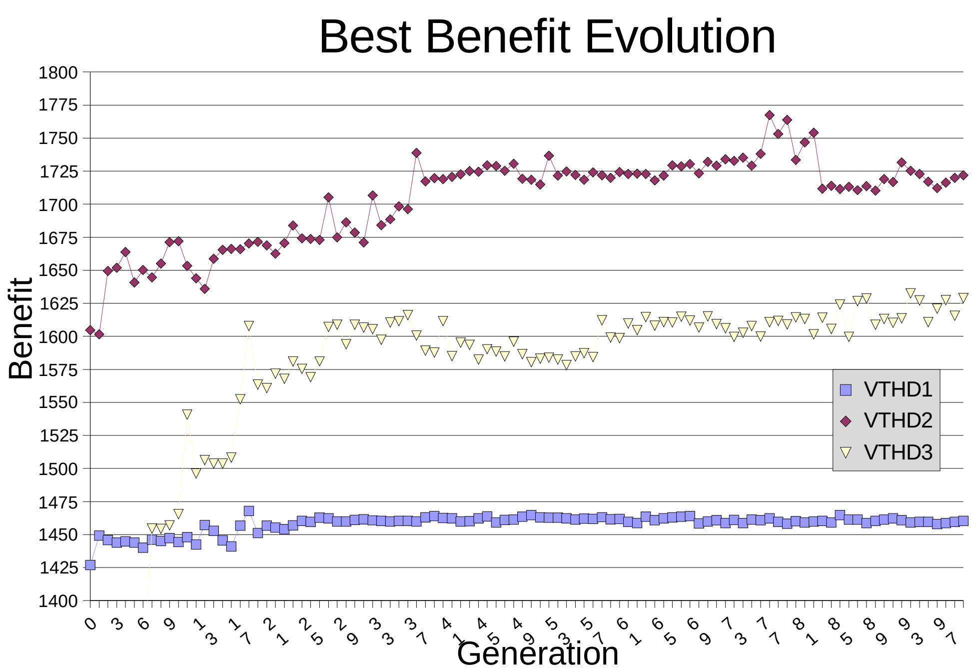

Figure 4: Population best benefit for each GA generation.

6 Conclusions and Future Work

In this paper, we have studied a new pricing scheme based on connection availability, and used it to set prices and to extend an already existing network in order to increase the service provider's benefit. To validate our method, we have run the method on three problems inspired by the VTHD (Very High Broadband IP/WDM test platform) network. One of the problems consists of modifying the backbone by incorporating new links. The other two, with a large solution space, additionally upgrade the access network technology of the access links. The genetic algorithm finds an almost optimal solution for the first problem, and very good solutions (although we can not be sure if optimal) for the remaining problems.

Note that the fitness evaluation of GA is computationally very expensive since it is based on availability estimation, an NP-hard problem in general. In our work, the availability is estimated by means of the Generalized Antithetic Monte Carlo simulation method. Our GA evaluates the fitness many times, therefore an important improvement in run time can be attained by diminishing the computing time of each evaluation or the total number of evaluations. As future work, we could try different approaches to solve this problem:

a) using efficient upper bounds; this approach is interesting because it can represent Service Level Agreements, based in availability, in a natural way;
b) in the execution of our GA, as some estimations might be computed more than once; savings can be attained by storing previous computations and avoiding repeating them;
c) developing heuristic methods to estimate the availability from previous similar estimations (for example, using a random neural network).

An important point is the impact on the optimization procedure of the availability error introduced by the estimation. A preliminary evaluation of this aspect has been discussed in Section 5; a refinement of the trade-off between uncertainty in genetic algorithm and availability estimation is then an important issue for future work, in order to improve the overall precision of the method.

References

R.E. Barlow and F. Proschan. (1981*). Statistical Theory of Reliability and Life Testing*. Holt, Reinhart and Winston, Silver Spring, Md

S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J.W. Roberts (2001). Bandwidth sharing: a study of congestion at flow level *SIGCOMM 2001*, San Diego, CA, USA.

P. Burgos, H. Cancela, A. Godoy and L. Petingi (2003). Network reliability optimization with source-terminal distance constraints. INOC'2003 (International Network Optimization Conference), pp. 132-137. October 27-29, 2003, Evry/Paris, France.

L.A. DaSilva (2000). Pricing of QoS-Enabled Networks: A Survey. *IEEE Communications Surveys & Tutorials*, 3(2).

D. L. Deeter and A. E. Smith (1998). Economic design of reliable networks. IIE Transactions, 30:1161-1174.

B. Dengiz, F. Altiparmak, and A. E. Smith (1997). Local search Genetic Algorithm for optimal design of reliable networks. IEEE Trans. on Evolutionary Computation, 1(3):179-188.

S. Duarte, B. Barán, and D. Benítez (2003). Telecommunication network design with parallel multiobjective evolutionary algorithms. In *Proceedings of IFIP/ACM Latin American Networking Conference*, pages 1-11.

M. El Khadiri and G. Rubino (1992). A Monte-Carlo method based on antithetic variates for network reliability computations. Technical report, Publication Interne nro 626, IRISA-INRIA, 1992.

M. Falkner, M. Devetsikiotis, and I. Lambadaris (2000). An Overview of Pricing Concepts for Broadband IP Networks. *IEEE Communications Surveys & Tutorials*, 3(2).

P. Fishburn and A. Odlyzko (1998). Dynamic behavior of differential pricing and quality of service options in the Internet. *Proceedings of ICE98,* ACM Press.

C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, C. Diot,. (2003) Packet-level Traffic Measurement from the Sprint IP Backbone". *IEEE Network Magazine.*

T. Henderson, J. Crowcroft, and S. Bhatti (2001). Congestion Pricing. Paying Your Way in Communication Networks. *IEEE Internet Computing*, September/October:85-89.

J.S. Provan and M.O. Ball (1983). The complexity of counting cuts and the probability that a graph is connected. *SIAM Journal on Computing*, 19:777-788.

B. Tuffin (2003). Charging the Internet without bandwidth reservation: an overview and bibliography of mathematical approaches. *Journal of Information Science and Engineering*, 19(5):765-786.

L.G. Valiant (1979). The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8: 410-420.

X Wang and H. Schulzrinne (2000). An Integrated Resource Negotiation, Pricing, and QoS Adaptation Framework for Multimedia Applications. *IEEE Journal on Selected Areas in Communications*, 18(12):2514-2529.