# Revenue Maximization for Broadband Service Providers Using Revenue Capacity

Haleema Mehmood, Madeleine Udell and John Cioffi
Department of Electrical Engineering
Stanford University, Stanford, CA 94305
{hmehmood, udell, cioffi}@stanford.edu

***Abstract*—Broadband internet service providers (ISPs) are for-profit companies: they provide high speed internet service to customers in order to maximize revenue. Their revenue depends on prices they are able to charge the customers, which in turn depends on the data rates provided. This paper defines the *revenue capacity* of a communication channel as the maximum revenue an ISP can achieve by allocating data rates to consumers, given an exogenous price function and an exogenous underlying physical layer model. This paper proposes an algorithm to compute the revenue capacity for multiuser gaussian channels with staircase price functions, and to produce feasible data rates achieving the revenue capacity. The paper concludes with an analysis of incremental migration of VDSL to vectored VDSL, quantifying the increase in revenue capacity with the increase in the percentage of vectored users.**

## I. Introduction

The demand for high speed data services is growing rapidly and is expected to increase exponentially in coming years. With the explosion of smart devices, every member of the household now individually requires a high data rate connection. This translates to a sharp increase in data and bandwidth demand for residential broadband services. Internet service providers face the challenge of providing such services at affordable prices. ISPs today are confronted with the obstacles of sustainability and economic viability [1]. The economic viability challenge necessitates the development of meaningful metrics and assessment tools for joint performance and revenue optimization. This paper introduces the *revenue capacity* of a communication link as a useful and meaningful metric for revenue optimization.

The information theoretic capacity of a communication channel gives an upper bound on the rate at which information can be reliably transmitted over the channel. In a similar manner, revenue capacity gives an upper bound on the monetary worth of a communication link given a price function. A price function is a mapping from the data rate provided to a customer to the price charged per month by the service provider. Depending upon the physical layer technology, different users can be provided different rates with tradeoffs between users' rate. Revenue capacity gives the maximum revenue generating rate tuple and therefore the maximum possible revenue of the shared communication link given the underlying physical layer constraints.

This paper presents a general revenue capacity formulation for multiuser channels. A last-mile/residential broadband service provider maintains a number of links that can be twisted pair copper telephone wire, coaxial cable or optical fiber. Each shared link to the customer premises represents a multiuser communication channel. In coaxial cable systems, the users share a single cable up to a drop point. For digital subscriber line (DSL) transmission over twisted pair copper cables, although a single line is not physically shared between different users, crosstalk between lines sharing the same binder or cable makes it a shared communication medium. The ISP can find the revenue capacity for each of its links using the appropriate physical layer model.

Broadband residential services are generally sold in fixed rate tiers, with a fixed price per month for each tier. In a competitive market, the prices are exogenous and are not determined by each ISP alone. In many countries however, the broadband market is considered a monopoly or oligopoly where the service providers have a bigger control over pricing of services. The service providers also have some control over the rate tiers they offer to their customers. In any case, each ISP has a price function that maps the data rate provided to the price charged per month to the customer. These price functions do not change over long periods of time. The price functions are also the same over certain geographical regions and do not change from one link to another. Revenue capacity is the maximum revenue for a link given such a price function.

The capacity of a communication channel gives an upper bound on the rate of information transmission. Similarly, revenue capacity gives an upper bound on the monthly revenue that a link can provide. It returns the rates for each user that correspond to the revenue capacity point. It does not incorporate user preference and assumes that each user will accept whatever rate is offered and pay for it. The formulation in the paper however can easily be extended to incorporate user preference. Each user can specify a minimum desired rate and a maximum monthly budget. The solution space for the revenue capacity problem can simply be restricted to lie in this region.

The revenue capacity characterizes the maximum revenue that can be derived from a multiuser communication channel. This characterization can be useful in many ways.

- Revenue capacity gives a single rate tuple which is optimal in a meaningful way instead of a multi-dimensional rate region [2]. The rate region for a multiuser communication channel is not only computationally complex to

build, it is also hard to visualize and difficult to use as a basis for making operating decisions. Revenue capacity however returns a single point in the multi-dimensional rate region that is optimal in terms of its monetary return.
- Revenue capacity can be used to compare competing physical layer technologies. It can be also used to make technology or infrastructure upgrade decisions. This can be done by comparing the cost of upgrade with the increase in the revenue capacity resulting from the upgrade.
- Revenue capacity can be used to evaluate different price functions. Experiments can be done with different rate tiers and monthly prices, evaluating the benefits of offering different packages, without rolling them out.
- Revenue capacity gives the rate tuple that is optimal from the ISPs point of view. Equipped with this knowledge, the ISP can do targeted advertisement for plan upgrades to existing customers and for new plans to potential customers.
- Revenue capacity can enable the ISPs to make better economic decisions for themselves. This in turn ensures provision of low cost, high speed internet services to end users, resulting in overall benefit for both the ISPs and the customers.
- The revenue capacity solution presented in this paper can be used for general utility maximization where the utility functions are the similar but mean something other than the monthly price of a broadband connection. For example, a staircase price function can be used as a utility function for multimedia services like the scalable video coding extension of H.264 standard, where the video quality only improves at certain step changes in data rate.

The rest of the paper is organized as follows: Section II presents a review of the relevant literature. Section III sets up the revenue capacity problem followed by solution methods in Section IV. Section V presents an application of revenue capacity to mixed vectored-unvectored VDSL deployment.

## II. Literature Review

The revenue capacity problem is closely related to the network utility maximization (NUM) problems found in literature. NUM problems generally assume a network of connected links with fixed link capacities. In [3], the authors present a generalized framework for cross layer optimization which can be applied more broadly. They present generalized NUM as cross-layer optimization problems where utility is a higher layer metric and the optimization variables belong to a lower layer in the internet protocol stack. If utility is a function of data rate and the optimization variables are link transmission powers, NUM is equivalent to utility-based power control. The author in [4] presents a comprehensive discussion of utility maximization for wireless systems. In case of concave utility functions, dual decomposition is used to separate the problem into two subproblems, a weighted sum rate maximization subproblem and a concave utility maximization subproblem [5]. This approach is intended for physical layer models that provide efficient ways of finding maximum weighted sum rates

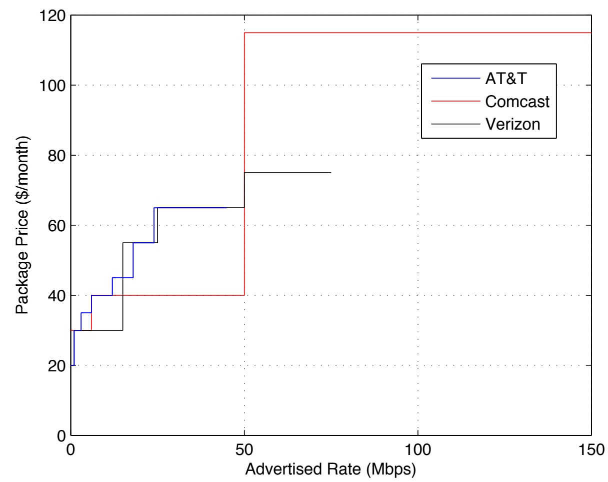


Fig. 1. Offered Packages

and have convex capacity regions. For non-concave utility functions, monotonic optimization from [6] is employed. The authors in [6] use the polyblock algorithm for power control over an interference channel for non-concave utilities that are increasing functions of data rates.

The revenue capacity problem in this paper extends the dual decomposition approach from [4] to non-concave, staircase utility functions. It differs from existing polyblock algorithms for non-concave utility maximization by employing branch and bound methods for global maximization.

Revenue capacity also builds on the literature on network pricing and congestion control. The works [7]–[9] define utility as a somewhat abstract end-user notion. They assume that the user (or an intelligent agent in the user's modem) knows the utility it derives from a certain data rate. The ISP sets a price per unit of data rate and shares it with the users. The end users optimize their utility minus cost, while the service provider maximizes its revenue. The prices remain in flux until equilibrium prices are found. This formulation does not necessarily give real prices that can be charged for bandwidth supplied. Although *dynamic pricing* for internet services has been proposed, there are numerous reasons why such dynamic pricing models cannot realistically be applied to internet services. A survey of literature on pricing for residential broadband is presented in [1], along with the challenges to dynamic pricing. Because of these challenges, the most widely used pricing model today is fixed priced rate-tiers. The rate-tiers are priced by the complex interaction of demand vs. supply in the market, along with service differentiation and smart marketing practices by the ISP, among other factors. For this reason, this paper considers a price function that is a staircase function of advertised rates. The function is assumed to be decided in advance and does not change over time or over different links, as would be the case in a real broadband market. Figure 1 shows the price functions for three major service providers in the US as an example.

## III. Problem Statement

Consider an ISP providing service to $K$ users over a single, shared communication channel. Let $r_k$ be the rate provisioned for user $K$, with $\mathbf{r} = [r_1 \; r_2 \; \cdots r_K]^T$ the vector of user data rates. The ISP has a price function $u : \mathbf{R} \rightarrow \mathbf{R}$, that maps a rate $r$ to a price $u(r)$ that is charged to a customer per month. The function $u$ is exogenous to the problem, and is a monotonically increasing function of rate. The ISP seeks to maximize its revenue for the link under consideration by maximizing the sum of the prices charged to all users. This maximum revenue depends on the set of achievable rate vectors over the underlying communication channel and is characterized by the capacity region $\mathcal{C}$ of the channel. A general revenue capacity problem can be formulated as

$$\begin{array}{ll} \text{maximize} & \sum_{k=1}^{K} u_k(r_k) \\ \text{subject to} & \mathbf{r} \in \mathcal{C}. \end{array}$$

This paper considers capacity regions of Gaussian vector channels with $n$ independent subcarriers/tones. There are three basic configurations of interest: the multiple access channel (MAC), the broadcast channel (BC) and the interference channel (IC). The MAC has a single receiver that accepts signals from $K$ separate transmitters. The BC has a single transmitter that communicates with $K$ uncoordinated receivers. The IC has $K$ transmitters and $K$ receivers and no coordination is possible on either transmit or receive ends. User $k$ receives/sends $b_{k,n}$ bits on tone $n$ based on the power loading and the channel and noise characteristics on that tone. A configuration agnostic capacity relation for the $n$th tone can be written as,

$$\mathbf{b}_n \in c_n(\mathbf{P}_{k,n(k=1,\cdots,K)}, \mathbf{H}_n),$$

where $\mathbf{H}_n$ is the noise-normalized channel matrix, $\mathbf{b}_n = [b_{1,n} \; b_{2,n} \; \cdots b_{K,n}]^T$ is the vector of user bit capacities on tone $n$ and $\mathbf{P}_{k,n}$ is the transmit covariance matrix for user $k$[1], on tone $n$ (in case of multiple transmit and/or receive antennas). The total data rate for user $k$ is the sum of $b_{k,n}$ over all tones, i.e. $\hat{r}_k = f_s \sum_{n=1}^{N} b_k^n$ where $f_s$ is the transmit symbol rate. Each transmitter has a total power budget of $P_t$, i.e.,

$$\sum_{n=1}^{N} \text{trace}(\mathbf{P}_{k,n}) \leq P_t.$$

Let $\mathcal{R}$ denote the set of all achievable rate vectors $\hat{\mathbf{r}}$ corresponding to transmit covariance matrices that satisfy the power constraints. Then the capacity region $\mathcal{C}$ is the convex hull of $\mathcal{R}$. The revenue capacity problem finds the maximum revenue rate vector over this convex capacity region. For rate vectors that are in $\mathcal{C} \backslash \mathcal{R}$ the physical layer can achieve those rates by time sharing between at most $K$ rate vectors in $\mathcal{R}$ [4]

The capacity region $\mathcal{C}$ is characterized by its boundary. The rate vectors on the boundary of the capacity region can be found by solving a weighted sum rate maximization (WSRMax) problem. Assume $f_s = 1$ for further discussion. For a weight vector $\mathbf{w} \in \mathbf{R}^K$, WSRMax solves the following problem:

$$\begin{array}{ll} \text{maximize} & \sum_{k=1}^{K} w_k \sum_{n=1}^{N} b_k^n \\ \text{subject to} & \mathbf{b}_n \in c_n(\mathbf{P}_{k,n(k=1,\cdots,K)}, \mathbf{H}_n) \quad \forall n \\ & \sum_{n=1}^{N} \text{trace}(\mathbf{P}_{k,n}) \leq P_t \quad \forall k. \end{array}$$

The procedure returns the optimal transmit covariance matrices and corresponding tonal bit capacities. These can be used to find a weighted sum rate maximizing rate vector $\hat{\mathbf{r}}$ corresponding to $\mathbf{w}$, that lies on the boundary of $\mathcal{R}$. Different weights $\mathbf{w}$ correspond to different rate tuples that lie at the capacity region boundary. The boundary of $\mathcal{C}$ is a convex combination of these rate tuples.

### A. MAC Capacity

A successive decoding implementation is assumed for the MAC. The tonal capacity relation for the MAC is,

$$\begin{array}{l} \mathbf{b}_n \in c_n(\mathbf{P}_{k,n(k=1,\cdots,K)}, \mathbf{H}_n) = \\ \{\mathbf{b}_n \,|\, 0 < \sum_{k \subseteq K} b_{k,n} \leq \log_2 |\sum_{k \subseteq K} \mathbf{H}_{k,n}\mathbf{P}_{k,n}\mathbf{H}_{k,n}^* + I|\}. \end{array}$$

Let $\pi(k)$ be a decoding order for the MAC. Then,

$$\begin{aligned} b_{k,n} = & \log_2(|\sum_{i=1}^{k} \mathbf{H}_{\pi(i),n}\mathbf{P}_{\pi(i),n}\mathbf{H}_{\pi(i),n}^*| + I|) \\ & - \log_2(|\sum_{i=1}^{k-1} \mathbf{H}_{\pi(i),n}\mathbf{P}_{\pi(i),n}\mathbf{H}_{\pi(i),n}^*| + I|) \end{aligned}$$

In the WSRMax problem for the MAC, the weights directly translate to a decoding order. This results in a concave objective function. The WSRMax problem can be solved using dual decomposition [2].

### B. BC Capacity

For the scalar BC, dirty paper coding is the optimal transmission scheme. Rate tuples achieving the capacity can be found using duality between the BC and the MAC [2]. This solution method is extended to parallel vector broadcast channels in [10]. The WSRMax problem can be solved as a dual MAC problem.

### C. IC Capacity

For the IC, single transmit and recieve antennas are considered. The bits for user $k$ on tone $n$ are given by

$$b_{k,n} = \log_2\left(1 + \frac{|H_n^{k,k}|^2 p_{k,n}}{1 + \sum_{j \neq k} |H_n^{k,j}|^2 p_{j,n}}\right),$$

where $H_n^{k,k}$ is the noise normalized channel gain of user $k$ on subcarrier $n$ and $H_k^{k,j}$ is the noise normalized crosstalk channel gain from user $j$ to user $k$ on subcarrier $n$. The bits are a non-convex function of the powers.

[1] in case of BC, the dual MAC has $K$ transmit covariance matrices

Optimum spectrum balancing (OSB) [11] is a method to solve this problem for a discrete number of bits. If a system supports only discrete bits, $b_{k,n} \in \{0, \cdots, b_{max}\}, \forall n$ on each subcarrier (this is the case in real systems), then the PSD combinations are also limited to discrete values. OSB uses dual decomposition to break the problem into per subcarrier dual subproblems. The tonal subproblems are solved by exhaustive search over all possible user bit combinations on each subcarrier. Separating the problem into tonal subproblems significantly reduces the complexity of exhaustive search. The solution is optimum for a multicarrier system with infinite number of tones. The authors in [12] show that under time sharing, the duality gap of the optimization problem is always zero, regardless of the convexity of the objective function.

## IV. Solution Methods

This section first discusses the revenue capacity problem for concave price functions. The algorithm for non-concave, staircase, price functions builds on the concave function approach.

### A. Dual Decomposition

For concave utilities, a dual approach for solving the revenue capacity problem over convex capacity regions has been discussed in [4]. Dual decomposition is used to divide the problem into two subproblems that can each be solved to optimality for fixed dual variables. An outer loop iterates over the dual variables to find the optimal solution. The problem is first modified by introducing additional variables,

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{k=1}^{K} u(s_k) \\ \text{subject to} & 0 \le s_k \le r_k \quad \forall k \\ & \mathbf{r} \in \mathcal{C}. \end{array}$$

The Lagrangian for the above problem $S$ is given by,

$$L(\mathbf{r}, \mathbf{b}) = \sum_{k=1}^{K} (u(s_k) + \lambda_k (r_k - s_k))$$

For fixed dual variables $\lambda_k$, the problem can be decomposed into two subproblems $S_1$,

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{k=1}^{K} (u(s_k) - \lambda_k s_k) \\ \text{subject to} & s_k \ge 0 \quad \forall k, \end{array}$$

and $S_2$,

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{k=1}^{K} \lambda_k r_k \\ \text{subject to} & \mathbf{r} \in \mathcal{C}, \end{array}$$

$S_1$ is convex by assumption and can be maximized to obtain $s^*(\lambda)$. $S_2$ is a weighted sum rate maximization (WSRMax) problem which can be solved for the physical layer models under consideration. However, the WSRMax procedure returns a point in $\mathcal{R} \cap \mathcal{C}$. The following discussion shows why this is not a problem for dual decomposed problem.

The dual problem is written as

$$\begin{array}{ll} \text{minimize} & g(\lambda) = g_{S_1}(\lambda) + g_{S_2}(\lambda) \\ \text{subject to} & \lambda \ge \mathbf{0} \end{array}$$

Because of time sharing over the rate region, $g$ is not differentiable. The dual problem therefore can be solved using a cutting plane method, a subgradient method or an ellipsoid method. The procedure has an outer loop that iterates over the dual variables and an inner loop that solves the two subproblems for the current value of $\lambda$. At each iteration, a subgradient of $g$ is needed to find the next iterate of $\lambda$. The subgradient is given $\mathbf{r} - \mathbf{s}$. Since each point in $\mathcal{C}$ is a convex combination of points in $\mathcal{R}$,

$$\max \lambda^T \mathbf{r}, \mathbf{r} \in \mathcal{C} = \max \lambda^T \hat{\mathbf{r}}, \hat{\mathbf{r}} \in \mathcal{R}$$

A point in $\mathcal{R}$ therefore returns a valid subgradient at each iteration. Problem $S_2$ can simply be written as the original WSRMax problem from Section III.

The authors in [4] show that strong duality holds for concave utility function maximization over convex and proper rate regions $\mathcal{C}$. Even when the capacity is a non-convex function of powers, the convex hull operation ensures that the overall rate region is convex. The dual decomposition procedure therefore returns the optimal objective function value along with the optimal dual variables. These can be used to recover the optimal primal variables using primal recovery methods [13], [14].

It should be mentioned here that most widely used WSRMax algorithms for parallel vector channels use dual decomposition as well. The WSRMax problem is decoupled across tones and can be separated into $N$ tonal Lagrangian terms. The utility maximization problem can therefore be split into $N + 1$ subproblems that can be solved in parallel.

Revenue capacity is applicable to downstream transmission over wired communication systems. Downstream data rates usually determine the price of a broadband package. Also, the assumption that channel characteristics don't change over time is better suited to wired communication systems. Such systems today do not use multiple antennas or non-linear precoding. However, the revenue capacity analysis is still valid as an upper bound on the revenue.

### B. Staircase Price Functions

Price functions for broadband services used by the industry today are mostly staircase functions. Broadband services are offered in rate tiers with a fixed monthly price per tier. This section shows how find revenue capacity of a multiuser channel using a staircase price function.

A staircase function can be represented as a sum of sigmoidal functions. The authors in [15] present a sigmoidal programming framework that solves the problem of maximizing a sum of sigmoidal functions over a convex constraint set. They use a branch and bound technique to globally optimize sigmoidal programming problems by solving a sequence of convex subproblems. This paper draws inspiration from that work and uses a branch and bound method for utility maximization over a staircase utility function.

*1) Branch and Bound:* The branch and bound method starts with an initial rectangle $Q_{init}$ and recursively divides the rectangle into smaller rectangles. It constructs a lower bound as well as an upper bound on the value of the objective function over each rectangle. Let $p^*(Q)$ be the optimal value of the problem

$$\begin{array}{ll}\text{maximize} & \sum_{k=1}^{K} u_k(r_k) \\ \text{subject to} & \mathbf{r} \in \mathcal{C} \cap Q\end{array}$$

for any rectangle $Q$. If $U(Q)$ and $L(Q)$ are the upper and lower bounds on $p^*(Q)$, they satisfy

$$L(Q) \leq p^*(Q) \leq U(Q).$$

The bounds become tight as the rectangles shrink. By recursively branching into smaller rectangles and computing upper and lower bounds, the algorithm obtains global bounds on the value of the solution.

In order to find an upper bound on the objective function value over rectangle $Q$, the function $u$ is replaced by its concave envelope $\hat{u}_k : [l_k, u_k] \to R$ along each dimension. An upper bound on the objective function value is found by solving the revenue capacity problem over rectangle $Q$ with $\hat{u}_k$ as the price functions.

$$\begin{array}{ll}\text{maximize} & \sum_{k=1}^{K} \hat{u}_k(r_k) \\ \text{subject to} & \mathbf{r} \in \mathcal{C} \cap Q,\end{array}$$

where $\mathbf{r} \in \mathcal{C}$ is the capacity region constraint from the multiuser channel. This problem has a concave utility function and can be solved using the dual decomposition approach of problem $S$. Subproblem $S_1(Q)$ is now given by,

$$\begin{array}{ll}\text{maximize} & \sum_{k=1}^{K} (\hat{u}_k(s_k) - \lambda_k s_k)) \\ \text{subject to} & \mathbf{s} \in Q,\end{array}$$

while problem $S_2$ remains the same.

As mentioned earlier, problem $S_Q$ is solved by an outer subgradient method that updates the dual variables $\lambda$ at each iteration and passes them to the WSRmax problem as weights. Each inner iteration solves a WSRMax problem and a small linear program. The procedure converges to give the optimal objective function value and the optimal dual variables. The solution to $S_2$ in each iteration lies on the boundary of $\mathcal{C}$ while the solution to $S_1(Q)$ lies in Q. After solving the dual problem, the primal solution $\mathbf{r}^*(Q) \in \mathcal{C} \cap Q$ can be obtained using primal recovery techniques. This solves the maximization problem for the concave envelope of the price function over rectangle $Q$ and provides the upper bound $U(Q)$.

The lower bound is constructed by evaluating the original objective function at $\mathbf{r}^*(Q)$. A bound on $U(Q) - L(Q)$ can be given in terms of the *nonconvexity* $\rho(u)$ of the functions $\rho_k(u)$. In [16], the authors define $\rho_k(u) = \sup(\hat{u}_k(s) - u_k(s)),\ s \in Q$. For the $U(Q)$ and $L(Q)$ chosen above, in each rectangle $Q$,

$$U(Q) - L(Q) \leq \sum_{k=1}^{K} \rho(u_k),$$

As the algorithm proceeds, the size of the rectangles decreases and the bound becomes increasingly tight.

For branching, the method in [15] is followed. The algorithm branches by splitting the rectangle with the largest upper bound along the coordinate with maximum gap between $u_k(r_k^*(Q))$ and $\hat{u}_k(r_k^*(Q))$ into two subrectangles that meet at $r_k^*(Q)$. Such a branch maximally reduces the error at the previous solution.

*2) Concave Envelope:* The staircase function with $M$ stairs is represented as a sum of $M$ sigmoidal functions,

$$u(s) = \sum_{i=1}^{M} f_m(s)$$

where $f_m(s) : [l, u] \to R$ is a sigmoidal function parameterized by $\alpha$, $a_m$ and $b_m$. $\alpha$ controls the steepness of the transition between the minimum and maximum value of the sigmoidal function, and $b_m$ determines the inflection point of for each stair.

$$f_m(s) = \left( \frac{a_m}{1 + \exp(-\alpha(s - b_m))} \right)$$

The concave envelope is generated in the following way: For each function $f_m$, there is a point $w_m >= b_m$ such that $f(\hat{s}) = f(s)$, for $w_m \leq s \leq u$. The point $w_m$ can easily be found by bisection. Given $l$, the slope $\delta_m$ between each point $w_m$ and $l$ can be determined. Let $\tilde{m}_0$th step have the maximum slope $\delta_{\tilde{m}_0}$. Then for $l \leq s \leq w_{\tilde{m}_0}$

$$\hat{u}(s) = u(l) + \delta_{\tilde{m}_0}(s - l).$$

Now let $l = w_{\tilde{m}_0}$. Find the new $w_m$s for the new $l$. The stairs with $b_m \leq w_{\tilde{m}_0}$ can be ignored. For the remaining stairs, find the new maximum slope $\delta_{\tilde{m}_1}$. Repeat the process to find complete concave envelope,

$$\hat{u}(s) = \begin{cases} u(l) + \delta_{\tilde{m}_0}(s - l) & l \leq s \leq w_{\tilde{m}_0} \\ u(w_{\tilde{m}_0}) + \delta_{\tilde{m}_1}(s - w_{\tilde{m}_0}) & w_{\tilde{m}_0} \leq s \leq w_{\tilde{m}_1} \\ \vdots & \vdots \\ u(s) & w_M \leq s \leq u \end{cases}$$

Since the transitions of the staircase function are sharp, it is more convenient to form a piecewise linear upper bound on the last stair as well. Figure 2 shows the piecewise linear concave envelope of $u$ formed by the above procedure.

*3) Convergence:* The convergence of this algorithm follows from the proof of convergence for general sigmoidal programs in [15], since each function $u_k$ can be expressed as a sum of $M$ sigmoidal functions $f_{m,k}$.

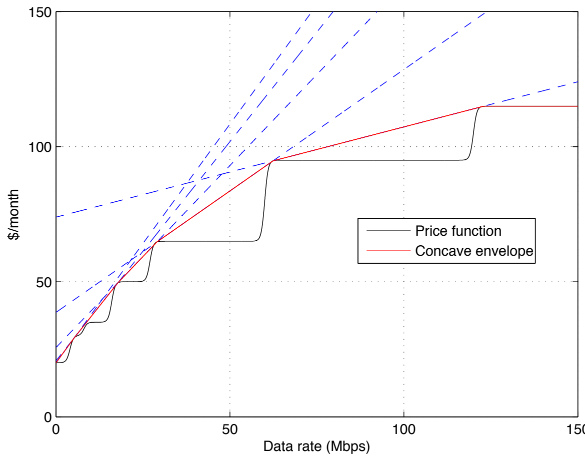


Fig. 2. Concave envelope

## V. Application Example

This section presents the merits of revenue capacity using a problem very relevant to DSL broadband service providers. As ISPs make the decision to migrate to vectored VDSL, the problem of coexistence with legacy or unvecoted VDSL arises. Dynamic spectrum management (DSM) can enable coexistence of vectored and unvectored VDSL but the rollout of vectored VDSL has been slow. A revenue capacity analysis can present a picture of the economic benefits of migration to vectored VDSL, assuming efficient DSM solutions are used to ensure acceptable performance for all customers.

Figure 1 shows the current broadband packages offered by three of the biggest service providers in the US. Based on the technology used, the offered packages differ in their advertised speeds. AT&T provides services mainly over twisted pair copper cables using DSL technologies, Comcast offers data services over coaxial cables, whereas Verizon offers lower rate tiers over DSL and higher rate tiers over optical fiber. According to a 2014 FCC report on *Consumer Fixed Broadband Performance in the U.S.* [17], one-third of the ISPs delivered only 60 percent or better of advertised speeds 80 percent of the time to 80 percent of the consumers. This performance was exceeded by most of the ISPs providing data rates very close to advertised rates even in peak hours. Out of the popular broadband packages considered in the 2014 report, the only package with speeds above 50 Mbps was provided by Verizon over optical fiber. Vectored VDSL can easily attain rates higher than 50 Mbps and even higher than 100 Mbps over short lines. DSL service providers can offer these high rate packages to their customers by migrating to vectored VDSL.

Figure 3 presents a price function that will be used for the example. This function builds on the packages presented in Figure 1. The ISPs advertise packages as being 'upto' certain Mbps, as shown in the 'Original' curve in figure 1. However, following the results in the FCC report, a 60 to 120 percent restriction is imposed on the data rate provided to the customer, i.e. the ISP can charge the amount for a certain package only if the rate provided is within 60 to 120 percent of the advertised rate for that price. This is shown in the 'Shifted staircase' function in figure 2. This staircase function will be the price function for the results below. A concave approximation (of the form $a - be^{-x/c}$) of the price function is also used for comparison of results, shown in the figure as the '1-exp approximation'.

The simulations here calculate VDSL2 downstream bit rates with both vectored VDSL2 and un-Vectored VDSL2. Downstream VDSL2 Profile 17a is simulated. The transmit PSD is at most 3.5 dB below the VDSL2 profile 998ADE17-M2x-A PSD limit mask defined in Annex B of G.993.2 [18]. The margin is 6 dB and the total coding gain is 3 dB. Bit rates are calculated by summing the capacity calculation of each 4.3125 kHz tone with a 9.75 dB SNR gap, with bits per Hz per sub-carrier lower limited to at least one bit and upper limited to 16 bits per Hz per subcarrier. Simulations use 1 percent worst-case same-binder FEXT plus -140 dBm/Hz noise. There are 25 active lines in the binder and the transmit PSDs of all lines are shaped by OSB. There is no crosstalk between the vectored lines, with the unvectored lines causing crosstalk into the vectored lines, and with vectored lines as well as non-vectored lines causing crosstalk to the unvectored lines. Note that since 1 percent worst-case crosstalk is used, the results here have worse crosstalk than a typical case with a filled 25-pair binder. For the simulations, it is assumed that the un-vectored lines and the vectored lines all originate at the same cabinet and are of the same length.

Figure 4 shows the revenue capacity for the sigmoidal utility function as well as the approximate concave utility function as a function of the loop length. The decrease in revenue capacity as a function of loop length is expected since the data rates go down as loop length increases. The figure gives several curves for the revenue capacity of the binder as the percentage of vectored lines in the binder increases from 20 percent to 80 percent. As more customers migrate to vectored VDSL, the expected revenue in dollars/month/line increases. For instance, for the case of 1000 feet loop length, the revenue capacity increases from about 78$ per month to 106$ per month as the percentage of vectored users increases from 20 to 80 percent. This is a 36% increase in revenue per line per month.

Figure 5 shows the rates assigned to the vectored and unvectored lines at revenue capacity. These rates can be used as guidelines for offering service packages as vectored VDSL is rolled out.

## VI. Conclusion

This paper presents a general framework for revenue maximization for physical layer models that have two major properties. First, they provide a way to compute a maximum weighted sum rate given any weights. Second, the physical layer can use time sharing to achieve rate tuples that it cannot achieve otherwise, making the capacity region convex. As future work, particular solutions for different types of channel models can be studied in more detail. The revenue capacity

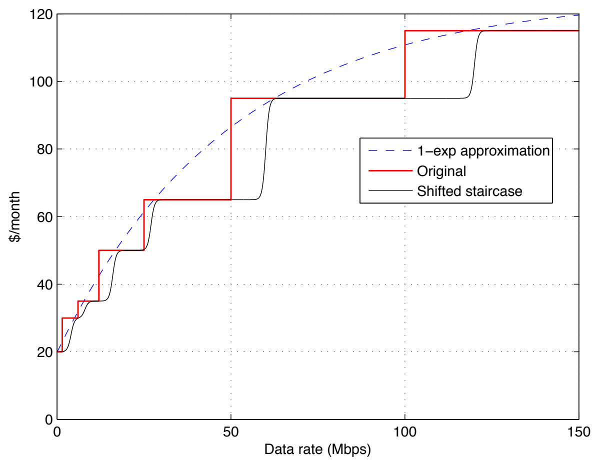


Fig. 3. Price function

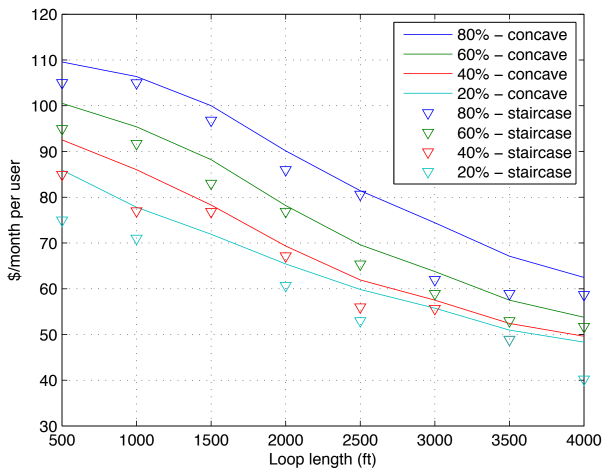


Fig. 4. Revenue capacity vs. loop length

algorithm presented here finds the maximum revenue rate vector. It does not return the corresponding input autocorrelation matrices in case of time sharing solutions. The current formulation also does not take user preference into account. The authors intend to address this in future.

Revenue capacity can have many real world applications. This paper presents a single example. The authors intend to apply the concept to different technologies and channel models. Revenue capacity is also useful for comparison of pricing strategies. Furthermore, it can be useful for provision of on-demand bandwidth.

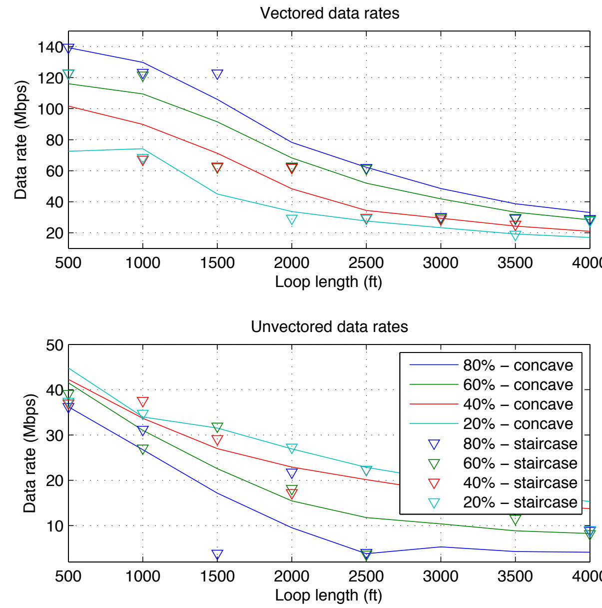


Fig. 5. Vectored and unvectored rates at revenue capacity

## References


[1] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "A survey of smart data pricing: Past proposals, current plans, and future trends," *ACM Computing Surveys (CSUR)*, vol. 46, no. 2, p. 15, 2013.

[2] J. M. Cioffi, *Multi-channel Modulation*, [Online].

[3] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, 2007.

[4] J. Brehmer, "Utility maximization in nonconvex wireless systems," 2010.

[5] J. Liu and Y. T. Hou, "Cross-layer optimization of mimo-based mesh networks with gaussian vector broadcast channels," *arXiv preprint arXiv:0704.0967*, 2007.

[6] L. P. Qian and Y. Jun, "Monotonic optimization for non-concave power control in multiuser multicarrier network systems," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 172–180.

[7] H. Yaiche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 8, no. 5, pp. 667–678, 2000.

[8] P. Hande, M. Chiang, R. Calderbank, and S. Rangan, "Network pricing and rate allocation with content provider participation," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 990–998.

[9] F. Kelly, "Charging and rate control for elastic traffic," *European transactions on Telecommunications*, vol. 8, no. 1, pp. 33–37, 1997.

[10] M. Mohseni, "Capacity of gaussian vector broadcast channels," 2006.

[11] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, and T. Bostoen, "Optimal multiuser spectrum balancing for digital subscriber lines," *Communications, IEEE Transactions on*, vol. 54, no. 5, pp. 922–933, 2006.

[12] W. Yu and R. Lui, "Dual methods for nonconvex spectrum optimization of multicarrier systems," *Communications, IEEE Transactions on*, vol. 54, no. 7, pp. 1310–1322, 2006.

[13] M. S. Bazaraa, H. D. Sherali, C. Shetty *et al.*, "Nonlinear programming: theory and applications," 1993.

[14] J. Brehmer and W. Utschick, "Utility maximization strategies in the multi-user mimo downlink," in *Cross Layer Design, 2007. IWCLD'07. International Workshop on*. IEEE, 2007, pp. 86–90.

[15] M. Udell and S. Boyd, "Maximizing a sum of sigmoids," 2013.

[16] ——, "Bounding duality gap for problems with separable objective," *arXiv preprint arXiv 1410.4158*, 2014.

[17] F. O. of Engineering, Technology, Consumer, and G. A. Bureau, "2014 measuring broadband america fixed broadband report," 2014.

[18] ITU-T, "Very high speed digital subscriber line transceivers (vdsl2), itu-t recommendation g.993.2," no. G.993.2, 2011.