An Unconventional Clustering Problem: User Service Profile Optimization

Fabio D'Andreagiovanni
Dept. of Mathematical Optimization, Zuse Institute Berlin
Institute for System Analysis and Computer Science, CNR, Italy

Giuseppe Caire
Chair of Communications and Information Theory
Technische Universität Berlin, Germany

***Abstract*—We consider the problem of clustering N users into K groups such that users in the same group are assigned a common service profile over M commodities. The profile of each group k sets for each commodity m the maximum of the service quality that users in the k-th group are willing to pay. The objective is to find the clustering that maximizes the total service user quality, which corresponds to the revenue of the service provider. This Service Profile Optimization Problem (SPOP) emerges in various applications, as for example the bit-loading in Hybrid Fiber Coax data distribution systems. We propose a Mixed Integer Linear Programming (MILP) model for the problem, that allows to use state-of-the-art MILP solvers as the core tool in an original powerful heuristic. We show complexity and performance gains with respect to previously proposed methods and a direct application of a state of the art MILP solver.**

***Index Terms*—User Service Profile, Clustering, Mixed Integer Linear Programming.**

I. Problem definition and motivation

Consider M commodities and N users. Each user $n \in \{1, \dots, N\}$ is characterized by a vector $\mathbf{b}_n = (b_{1,n}, \dots, b_{M,n})^\mathsf{T}$, such that $b_{m,n}$ is the maximum price that he/she is ready to pay for buying the commodity $m \in \{1, \dots, M\}$ (we call such price *willingness to pay* and we will refer to it by the acronym WTP). [1] However, users cannot buy each commodity at the price that they want. Instead, the commodities are offered in packages (*user service profiles*), specifying the possible pricing options. In particular, we define K user service profiles $\mathbf{c}_1, \dots, \mathbf{c}_K$, where $\mathbf{c}_k = (c_{1,k}, \dots, c_{M,k})^\mathsf{T}$, and where $c_{m,k}$ is the price of commodity m in user profile k. A profile k is *admissible* for user n if $c_{m,k} \leq b_{m,n}$ for all $m \in \{1, \dots, M\}$, i.e., if the user can afford all the commodities in the price specified by the profile. We assume that each user n will choose the admissible profile $\mathbf{c}_k$ that maximizes the total (affordable) price, since this corresponds to the best service that he/she can afford. Under this "greedy user" assumption, the goal of the service provider is to design K user service profiles such that the total price over all profiles and user is maximized.

Letting $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_K] \in \mathbb{R}_+^{M \times K}$ denote the matrix of user service profiles, and $\{\mathcal{S}_1, \dots, \mathcal{S}_K\}$ denote a partition of the user set such that $\mathcal{S}_k$ is the set of users that choose profile k, the user service profile design aims at maximizing the function:

$$f(\mathbf{C}, \{\mathcal{S}_1, \dots, \mathcal{S}_K\}) = \sum_{k=1}^{K} |\mathcal{S}_k| \mathbf{1}^\mathsf{T} \mathbf{c}_k, \tag{1}$$

subject to $\mathbf{c}_k \leq \mathbf{b}_n$ (the inequality is intended component-wise), for all $n \in \mathcal{S}_k$, where the maximization is over $\mathbf{C}$ (the prices) and the partition $\{\mathcal{S}_1, \dots, \mathcal{S}_K\}$ (the clusters of users allocated to the profiles). We refer to this optimization problem as the **Service Profile Optimization Problem** (SPOP).

We are not aware of when and in what context the above clustering problem was proposed for the first time. To our knowledge, it arises in the optimization of a recently defined standard for cable digital transmission called DOCSIS 3.1 for *Hybrid Fiber Coax* (HFC) systems (see [1], [2]). In particular, the downstream transmission from the fiber termination box (e.g., in the basement of an apartment building) to the cable modems in the user homes is characterized by different cable transfer functions. DOCSIS 3.1 make use of OFDM modulation with constant power allocation and bit-loading, such that the size of the QAM constellations used on the different subcarriers is matched to the subcarrier- and user-specific SNR.

In this application, M represents the number of subcarriers in the OFDM modulation, $\mathbf{b}_n$ is the bit-loading that user n can support, and $\mathbf{c}_k$ is a bit-loading profile that the system can transmit. Following [1], [2], we have $b_{m,n} = \log(1 + \alpha |H_{m,n}|^2 / \sigma_{m,n}^2)$, where $\alpha = P_0/\Gamma$, P_0 denotes the frequency-flat transmit power spectral density, Γ is an empirical "gap to capacity" factor, and $H_{m,n}$ is the channel transfer function for user n on subcarrier m. Since for each bit-loading profile the system must perform buffering and packetization of the downstream data in order to match the data packets to the physical layer coded modulation format. Since the number of user N is potentially very large, in order to simplify these data formatting operations the number of allowed bit-loading profiles K is limited.

Other application examples.
1) Consider the LTE Media Broadcast/Multicast Service (MBMS), where multiple base stations send the same multicast data stream to a large number of users using the same frequency channel (single-frequency network) [3]. Due to multipath, distance-dependent pathloss and shadowing, users

[1] Notice that $b_{m,n}$ indicates the *absolute* WTP, not the price per unit quantity. For example, one may assume that all commodities have the same unit cost, and therefore $b_{m,n}$ reflects the quantity of commodity m that user n wishes to buy.

978-1-5090-1806-2/16/$31.00 ©2016 IEEE

may have very different frequency-selective channels and received average SNR. In this case, it may be convenient to define K user classes, characterized by their bit-loading profiles, such that the overall multicast aggregate received rate is maximized. Users in class $\mathcal{S}_k$ should be able to decode the (multicast) data packets transmitted with bit-loading profile $\mathbf{c}_k$ with aggregate multicast rate $|\mathcal{S}_k|\mathbf{1}^\mathsf{T}\mathbf{c}_k$ bits per OFDM symbol, such that the system spectral efficiency is given by $f(\mathbf{C}, \{\mathcal{S}_1, \ldots, \mathcal{S}_K\})/(KM)$ bit/s/Hz, where again $f(\cdot)$ is given by (1).

2) Consider a telecommunication service provider offering service "packages" including a certain amount of different services, such as wireless and wired voice, text messages, LTE high-rate data, on-demand video streaming, wired internet, TV, and so on. After a customer survey, the provider knows how much each customer is ready to pay for each given service. However, offering a personalized contract tailored individually to each customer would be impractical for a number of reasons (e.g., advertizing, accounting). Hence, the provider wishes to create $K \ll N$ packages in order to maximize its revenues. This is obviously exactly the problem described before.

Our contribution.

1) We introduce a new alternative mathematical optimization model for the SPOP considered in [1], [2]. Our model has the desirable property of being *linear*, in contrast to that introduced in [1], [2]. More specifically, we formulate the SPOP as a *Mixed Integer Linear Programming* (MILP) problem.

2) Since SPOP is challenging and is hard to solve even when applying state-of-the-art solvers like IBM ILOG CPLEX [5] to the MILP formulation, we propose a new fast solution algorithm. Our approach is centered on a fast heuristic that first builds an initial solution through a procedure inspired by a *k-means* algorithm [9] and then improves the initial feasible point through a sequence of Very-Large Neighborhood Search (VLNS) (e.g., [4]). A distinctive feature of our algorithm is that the neighborhood search is performed optimally and exactly by formulating itself as a lower-dimensional MILP problem.

3) We present computational experiments over a set of realistic instances of the SPOP that refer to a wireless network with MBMS and that show the superior performance of the new algorithm with respect to the one proposed in [2].

II. MILP FORMULATION FOR THE SPOP

In [1], the SPOP is formulated by introducing the assignment matrix $\mathbf{Y}$ with elements:

$$y_{k,n} = \begin{cases} 1 & \text{if user } n \text{ is assigned to profile } k \\ 0 & \text{otherwise} \end{cases} \qquad \forall \ k, n$$

representing the assignment of users to profiles. Then, the SPOP is written as the *Non-linear* Binary Program [1]:

$$\begin{aligned} \max \quad & \mathbf{1}^\mathsf{T}\mathbf{C}\mathbf{Y}\mathbf{1} && \text{(2a)} \\ \text{s.t.} \quad & \mathbf{0} \leq \mathbf{C}\mathbf{Y} \leq \mathbf{B} \quad \text{(componentwise)} && \text{(2b)} \\ & \mathbf{Y}^\mathsf{T}\mathbf{1} = \mathbf{1} && \text{(2c)} \\ & \mathbf{Y} \in \{0,1\}^{K \times N} && \text{(2d)} \end{aligned}$$

This formulation is a *Non-Linear* Mixed Integer Programming (Non-LinMIP) problem with non-linearities both in the objective function and in the constraints. Modeling the SPOP as a MILP yields a crucial advantage: it allows to exploit state-of-the-art MILP solvers, which are dramatically more advanced that Non-LinMIP solvers and can be used as tools (or "building blocks") for developing faster and more computationally efficient solutions algorithms. We first eliminate the non-linear constraints by the so-called "big-M" method (e.g., [6]). This yields the equivalent problem:

$$\begin{aligned} \max \quad & \sum_{k=1}^{K} \left(\sum_{m=1}^{M} c_{m,k} \right) \left(\sum_{n=1}^{N} y_{k,n} \right) && \text{(3a)} \\ \text{s.t.} \quad & c_{m,k} \leq b_{m,n} + B_{m,n}\,(1 - y_{k,n}) \qquad \forall \ k, m, n && \text{(3b)} \\ & \sum_{k=1}^{K} y_{k,n} = 1 \qquad \forall \ n && \text{(3c)} \\ & c_{m,k} \geq 0 \qquad \forall \ m, k && \text{(3d)} \\ & y_{k,n} \in \{0,1\} \qquad \forall \ k, n && \text{(3e)} \end{aligned}$$

where the objective function (3a) is simply the "exploded" version of the matrix form (2a), and we have replaced the non-linear constraints (2b) with the *linear* constraints (3b). The new constraints are obtained by recalling the following fact: if a user n is assigned to profile k, then the price $c_{m,k}$ that profile k assigns to commodity m must not exceed the WTP of user n (i.e., $c_{m,k} \leq b_{m,n}$); if instead n is not assigned to k, then this inequality needs not be enforced. Hence, equivalence of (2) and (3) follows by imposing $B_{m,n} \geq b_m^{\max} - b_{m,n}$ with $b_m^{\max} = \max_n b_{m,n}$.

Next, we tackle the non-linearity in the objective function by introducing additional variables and constraints [8]. Specifically, we replace each product $c_{m,k}\, y_{k,n}$ with one *continuous* variable $v_{n,k,m}$ such that:

$$\begin{aligned} & v_{n,k,m} \geq 0 && \text{(4a)} \\ & v_{n,k,m} \leq b_m^{\max}\, y_{k,n} && \text{(4b)} \\ & v_{n,k,m} \leq c_{m,k} && \text{(4c)} \\ & v_{n,k,m} \geq c_{m,k} - b_m^{\max}\,(1 - y_{k,n}) && \text{(4d)} \end{aligned}$$

where we have used the fact that $0 \leq c_{m,k} \leq b_m^{\max}$ for each m, k, by the definition of $b_m^{\max}$. These four *linear* inequalities can replace the *non-linear* product $c_{m,k}\, y_{k,n}$ since:

- if $y_{k,n} = 0$, then (4a) and (4b) implies $v_{n,k,m} = 0$; additionally, (4c) becomes $0 \leq c_{m,k}$ and (4d) becomes $0 \geq c_{m,k} - b_m^{\max}$, which are both satisfied recalling that $0 \leq c_{m,k} \leq b_m^{\max}$ for each m, k;
- if $y_{k,n} = 1$, (4a) and (4b) imply $0 \leq v_{n,k,m} \leq b_m^{\max}$, which is correct since by (4c) and (4d) we have $v_{n,k,m} = c_{m,k}$ and $0 \leq c_{m,k} \leq b_m^{\max}$.

Finally, the sought MILP formulation of the original SPOP is given by:

$$
\begin{array}{lll}
\max & \displaystyle\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{m=1}^{M} v_{n,k,m} & \text{(SPOP-MILP)} \\
\text{s.t.} & c_{m,k} \le b_{m,n} + B_{m,n}\,(1-y_{k,n}) & \forall\ k,m,n \\
& \displaystyle\sum_{k=1}^{K} y_{k,n} = 1 & \forall\ n \\
& c_{m,k} \ge 0 & \forall\ m,k \\
& y_{k,n} \in \{0,1\} & \forall\ k,n \\
& v_{n,k,m} \ge 0 & \forall\ k,m,n \\
& v_{n,k,m} \le b_m^{\max}\, y_{k,n} & \forall\ k,m,n \\
& v_{n,k,m} \le c_{m,k} & \forall\ k,m,n \\
& v_{n,k,m} \ge c_{m,k} - b_m^{\max}\,(1-y_{k,n}) & \forall\ k,m,n
\end{array}
$$

III. A Low Complexity MILP-Heuristic Method

We assess the performance of our algorithm and the quality of its results by comparing with a heuristic approach proposed in [1]. Specifically, we refer to the algorithm called CoPA (Clustering for Profile Assignment), which present two main phases: 1) an "initialization" phase, where an initial feasible solution is built by randomly assigning one single user to each profile and then completing the solution by assigning the remaining users to the profiles according to a greedy procedure; 2) an improvement phase called "group assignment update", which tries to find better solutions by executing a sequence of "swap-based" local searches, where it is evaluated the effect of swapping one user from one profile to another (for a detailed description of CoPA, we refer the reader to [1]).

The proposed alternative heuristic is also based on an initialization phase and an improvement phase, but it differs from that of [1] in the following features: 1) our initialization is a non-random assignment which take into account the users' WTP $\{b_{m,n}\}$; 2) our improvement phase performs a *very-large neighborhood search*, which considers a much larger space of feasible solutions, thus increasing the probability of finding a better solution. Additionally, this neighborhood search is done exactly, using a state of the art MILP solver as a building block.

A. Initial feasible solution

To build an initial feasible solution, we propose an algorithm that is partially inspired by *k-means* algorithms [9], commonly used for clustering problems. As first step, we introduce the sum WTPs $\Sigma b_n = \sum_{m=1}^{M} b_{m,n}$ for evaluating the overall level of service that each user n can afford: the idea is that users of "better quality" have a higher sum WTP and that users of "similar quality" have close sum WTP and should be put in the same profile. In what follows, we assume without loss of generality that users $n = 1,\dots,N$ are sorted in non-decreasing order of Σb_n, i.e. $\Sigma b_n \le \Sigma b_{n+1}$ for $n = 1,\dots,N-1$ (to have this, it is sufficient to sort the users according to their values Σb_n). Hence, the sum WTP range is $[\Sigma b_1, \Sigma b_N]$ and we identify K (= no. of profiles) equidistant points β_k over this range. To this end, we define the step $\Delta\Sigma b = (\Sigma b_N - \Sigma b_1)/(K-1)$ and the generic point $\beta_k \in [\Sigma b_1, \Sigma b_N]$ is equal to $\beta_k = \Sigma b_1 + (k-1)\Delta\Sigma b$, for $k = 1,\dots,K$. We use the points β_k as reference for building a complete assignment of users to profiles, according to the procedure presented in Algorithm 1. This provides: 1) a reference user for each profile k, defined as the closest user to the reference point β_k; 2) assigning each non-reference user to the closest reference user with respect to the usual squared Euclidean distance in $\mathbb{R}^M$. The initial feasible assignment corresponds to a specific choice of the decision matrix $\mathbf{Y}$ of SPOP-MILP. It should be noted that, as evident from (1), for given profile assignment sets $\mathcal{S}_1,\dots,\mathcal{S}_K$, the SPOP objective function is trivially maximized with respect to $\mathbf{C}$ by letting $c_{m,k} = \min_{n\in\mathcal{S}_k} b_{m,n}$. Therefore, in the following we describe the solution only in terms of the assignment sets $\mathcal{S}_1,\dots,\mathcal{S}_K$ or, equivalently, of the decision matrix $\mathbf{Y}$. Also, with a slight abuse of notation, we denote the SPOP objective function in (1) simply as $f(\mathbf{Y})$.

The construction of the first feasible solution is followed by the execution of a *random Swap-Improve* heuristic with time limit $\tau > 0$. This is a simple yet effective improvement heuristic based on randomly choosing a user n, assigned to a profile $k(n)$, and a profile $k \ne k(n)$ and then evaluating the effects of moving n from profile $k(n)$ to k. If this improves the objective function, the change of profile for n (swap) is adopted.

Algorithm 1 Defining an initial feasible solution

```
Define 𝒩 = {1,...,N},   S_k = ∅, ∀k = {1,...,K}
for k = 1 to K do
    n̄ = arg min_{n∈𝒩} {|β_k − Σb_n|}
    Let Ref(k) := n̄ be the reference user of profile k
    S_k := S_k ∪ {n̄}
    𝒩 := 𝒩\{n̄}
end for
while 𝒩 ≠ ∅ do
    Select n̄ ∈ 𝒩
    k̄ = arg min_{k∈{1,...,K}} DIST(n̄, Ref(k))
    S_k̄ := S_k̄ ∪ {n̄}
    𝒩 := 𝒩\{n̄}
end while
execute Swap-Improve with time limit τ > 0
return S_k, ∀k = {1,...,K}
```

B. Improving solutions by an exact large variable neighborhood search

In order to improve a feasible solution returned by Algorithm 1, we rely on an MILP heuristic based on a the execution of a *very-large variable neighborhood search*, which combines the mechanisms of *large* [4] and *variable* neighborhood search [10]. A distinctive feature of our heuristic is to formulate the neighborhood search as an *exact MILP search*: the exploration of the neighborhood is formulated as a MILP problem that is *exactly* solved by a state-of-the-art MILP solver.

Specifically, given a feasible solution to the problem identified by the corresponding user-profile assignment matrix $\mathbf{Y}$, we consider the neighborhood constituted by all the feasible solutions that can be obtained by modifying at most δ components of $\mathbf{Y}$. This can be expressed by a version of SPOP-MILP where we impose the additional *Hamming distance constraint*:

$$\sum_{(k,n):\, \bar{y}_{k,n}=0} y_{k,n} + \sum_{(k,n):\, \bar{y}_{k,n}=1} (1 - y_{k,n}) \; \leq \; \delta \qquad (5)$$

which counts the number of elements that can change value in $\mathbf{Y}$ with respect to a reference assignment matrix $\bar{\mathbf{Y}}$.

Additionally, we exploit the valuable information that comes from the *linear relaxation* of SPOP-MILP, obtained by relaxing the integrality constraint on $\mathbf{Y}$ (i.e., by replacing $y_{k,n} \in \{0,1\}$ with $y_{k,n} \in [0,1]$. This yields a *Linear Program* (LP) denoted in the following by SPOP-LP. The optimal solution of SPOP-LP can be obtained with very low complexity and can be exploited to refine the neighborhood of solutions. As well-known from the theory of Linear Programming, the optimal solution of SPOP-LP may provide useful information about how to get a good quality feasible solution of the original problem [7]. In particular, let $\mathbf{Y}^{LR}$ be an optimal solution of SPOP-LP and let $\epsilon > 0$ be a small value (in our experiments, we adopt $\epsilon = 0.1$). If $y_{k,n}^{LP} \leq \epsilon$ (resp., $y_{k,n}^{LP} \geq 1-\epsilon$) and $\bar{y}_{k,n} = 0$ (resp., $\bar{y}_{k,n} = 1$), we have a reasonably good indication that $y_{k,n} = 0$ (resp., $y_{k,n} = 1$) is a good choice of the new assignment variable $y_{k,n}$ (see [7] for a detailed discussion about this argument). In our specific case, we adopt such procedure just for variables that are "sufficiently" close to 1, so if $y_{k,n}^{LP} \geq 1-\epsilon$ and $\bar{y}_{k,n} = 1$, we impose $y_{k,n} = 1$. We remark that this choice of the value of $\mathbf{Y}$ constitutes a variant of the heuristic presented in [7].

To summarize, given a feasible solution $\bar{\mathbf{Y}}$, the MILP solver explores a neighborhood of feasible assignment matrices $\mathbf{Y}$ such that:

1) they are within a distance δ from $\bar{\mathbf{Y}}$, expressed by the constraint (5);
2) they have fixed components according to the choice based on $\mathbf{Y}^{LP}, \bar{\mathbf{Y}}$ as described above.

The exploration of such neighborhood is performed efficiently and *exactly* by the state-of-the-art MILP solver CPLEX, used as a building block. This problem is denoted as SPOP-MILP($\mathbf{Y}^{LP}, \bar{\mathbf{Y}}, \delta, \epsilon$) in the following. We stress that we apply a time limit for solving SPOP-MILP($\mathbf{Y}^{LP}, \bar{\mathbf{Y}}, \delta, \epsilon$), since CPLEX can typically find good quality solutions for problems of reduced size in small amount of time, but it may require a large time to certify that a feasible solution is optimal (closing the so-called optimality gap). The proposed search heuristic, referred to as *MILP-VNS*, is summarized in Algorithm 2.

Given an incumbent solution $\bar{\mathbf{Y}}$, we repeatedly solve SPOP-MILP($\mathbf{Y}^{LP}, \bar{\mathbf{Y}}, \delta, \epsilon$) until a global time limit is reached. In each run of solution of SPOP-MILP($\mathbf{Y}^{LP}, \bar{\mathbf{Y}}, \delta, \epsilon$), we impose a "local" time limit for solving the single problem. If a better solution is found, the incumbent solution $\bar{\mathbf{Y}}$ is updated and the distance δ is reset to its initial value δ^0. If instead the "local" time limit is reached without finding a better solution, we enlarge the neighborhood by increasing the distance limit by a step $\Delta\delta > 0$ (i.e., we impose $\delta := \delta + \Delta\delta$) and start a new exact search until the global time limit is reached. The steps of MILP-VNS are detailed in Algorithm 2.

Algorithm 2 MILP-VNS

```
INPUT: an incumbent feasible solution Ȳ, an optimal
  solution of the linear relaxation Y^LP, values δ^0, Δδ, ε
set δ := δ^0
while a global time limit is not reached do
    solve SPOP-MILP(Y^LP, Ȳ, δ, ε) with time limit τ >
    0 and let Y* be the best feasible solution found
    if f(Y*) > f(Ȳ) then
        update the incumbent solution Ȳ := Y*
        δ := δ^0
    else
        δ := δ + Δδ
    end if
end while
return Ȳ
```

IV. Computational results

We tested the performance of our algorithm on 20 instances based on realistic data of an LTE Media Broadcast/Multicast Service wireless system, including 32 subcarriers and 100 users. The channel instances are generated according to a standard wireless channel model, typical of a medium-size cell in a urban environment, with the following parameters: bandwidth 20 MHz, delay spread 0.5 μs, base station transmit power (normalized to the power spectral density at the receivers) $\mathsf{SNR}_0 = 30$ dBm, users uniformly distributed over a disk of radius 500 m, with pathloss law $\chi/(1+(d/d_0)^\alpha)$, with pathloss exponent $\alpha = 3.5$, 3 dB loss distance $d_0 = 50$m, and long-normal shadowing χ with variance 6 dB.

According to the terminology and notation of this paper, the subcarriers correspond to commodities and thus $M = 32$, and for the users we have $N = 100$. We ran the experiments on a 2.70 GHz Windows machine with 8 GB of RAM and using the solver IBM ILOG CPLEX 12.5. The code was written in C/C++ and interfaced with CPLEX through Concert Technology.

The global time limit for heuristic execution is of 900 seconds. In the case of our heuristic, defining the first feasible solutions can be made instantaneously and we devote 60 seconds to run the Swap-Improve procedure of Algorithm 1. The best feasible solution found is then passed to Algorithm 2, where the time limit τ for solving SPOP-MILP($\mathbf{Y}^{LP}, \bar{\mathbf{Y}}, \delta, \epsilon$) is set equal to 240 seconds. Additionally, we set $\delta^0 = 0.95\,N$ and $\Delta\delta = 0.05\,N$.

The computational results are presented in Table I, where, for each instance identified by its ID, we report the number K of possible profiles, the value CoPA-Heu $f(\mathbf{Y}^\star)$ of the best solution found by the benchmark heuristic "CoPA" proposed in [1], the value MILP-Heu $f(\mathbf{Y}^\star)$ of the best solution found

TABLE I
COMPUTATIONAL RESULTS

ID	K	CoPA-Heu $f(\mathbf{Y}^\star)$	MILP-Heu $f(\mathbf{Y}^\star)$	$\Delta f(\mathbf{Y}^\star)\%$ MILP-COPA	CPLEX $f(\mathbf{Y}^\star)$	$\Delta f(\mathbf{Y}^\star)\%$ MILP-CPLEX
U1	5	2650.56	2958.17	+ 11.60	1566.55	+ 88.83
U2	5	2342.89	2588.76	+ 10.49	1499.36	+ 72.65
U3	5	2882.86	3001.00	+ 4.09	1528.07	+ 96.39
U4	5	3014.84	3133.52	+ 3.93	2447.30	+ 28.03
U5	5	3058.30	3236.27	+ 5.81	2073.96	+ 56.04
U6	5	2621.84	2758.89	+ 5.22	1663.03	+ 65.89
U7	5	2792.22	2924.99	+ 4.75	2049.53	+ 42.71
U8	5	3327.23	3469.73	+ 4.28	2504.16	+ 38.55
U9	5	2927.71	3158.48	+ 7.88	2281.11	+ 38.46
U10	5	3250.59	3341.51	+ 2.79	2390.72	+ 39.77
U11	10	3783.00	4010.87	+ 6.02	1870.25	+ 114.45
U12	10	3362.92	3527.83	+ 4.90	1613.54	+ 118.63
U13	10	3777.96	3947.21	+ 4.47	1943.35	+ 103.11
U14	10	4265.47	4568.62	+ 7.10	2243.94	+ 103.59
U15	10	4143.96	4248.72	+ 2.52	2032.02	+ 109.08
U16	10	3565.47	3742.74	+ 4.97	2042.60	+ 83.23
U17	10	4005.39	4060.98	+ 1.38	2108.30	+ 92.61
U18	10	4602.62	4761.59	+ 3.45	2305.11	+ 106.56
U19	10	4605.21	4663.63	+ 1.26	1803.91	+ 158.52
U20	10	3272.1	3430.23	+ 4.83	2010.96	+ 70.57

by our new heuristic, the percentage gain $\Delta f(\mathbf{Y}^\star)\%$ MILP-CoPA of our method over COPA. Additionally, CPLEX $f(\mathbf{Y}^\star)$ yields the value of the best solution found by direct use of CPLEX to solve the complete SPOP-MILP, and $\Delta f(\mathbf{Y}^\star)\%$ MILP-CPLEX shows the percentage gain of our method over CPLEX. We notice that the proposed SPOP-MILP method yields very significant gains with respect to the direct application of CPLEX and consistently better solutions than those provided by CoPA. We consider these results indicative of the fact that the construction and exploration phase of our heuristic are promising and deserve further investigations to be refined and strengthened.

We conclude this section by providing a qualitative and visually appealing snapshot of the clustering operated by our algorithm. For a given realization of $N = 100$ LTE frequency selective channels with $M = 32$ subcarriers, Fig. 1 shows the channel transfer functions in the frequency domain grouped by colors. each color corresponds to one of the $K = 5$ clusters, and the thick curves correspond to the user service profile for each group.

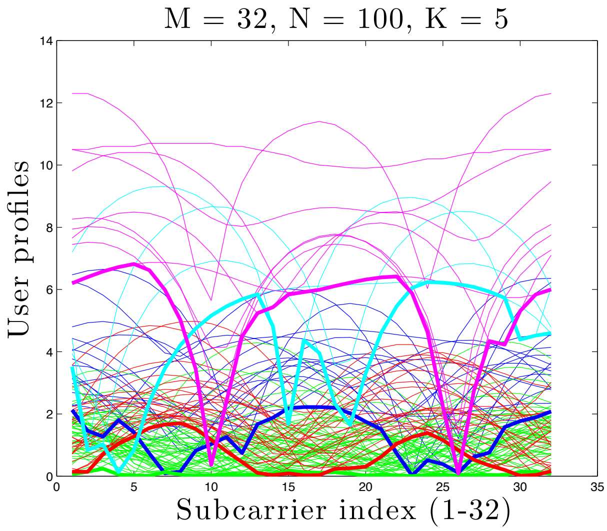

Fig. 1. Example of user grouping and user service profiles.

V. CONCLUSION

We considered the unconventional clustering problem of partitioning a user population into subsets assigned to service profiles, such that the users must be clustered taking into account their willingness to pay for each commodity in the profile. The problem arises in bit-loading in digital transmission systems over coaxial cables, but we believe that it has more applications, some of which have been discussed briefly in this paper. We derived a *linear* optimization model for this problem (precisely a MILP model), which tackles the non-linearities of the immediate direct problem formulation proposed in literature and allows to exploit the power of state-of-the-art MILP solvers. Additionally, we proposed a new heuristic that is based on the exact exploration of sub-regions of the solution space (neighborhood search) through the solution of smaller versions of the underlying MILP, obtained by introducing additional constraints. Computational tests on realistic instances motivated by wireless multicast over a family of frequency selective channels show that our new heuristic can fast provide solutions of better quality w.r.t. a benchmark heuristic proposed in literature, with significantly faster convergence time.

ACKNOWLEDGMENT

The work of Fabio D'Andreagiovanni was partially supported by the *Einstein Center for Mathematics Berlin* (ECMath) through Project MI4 (ROUAN) and by the *German Federal Ministry of Education and Research* (BMBF) through Project VINO (Grant 05M13ZAC).

REFERENCES

[1] H. Mehmood, J. Cioffi, "User clustering for high-speed small cell backhaul over coaxial cable", IEEE International Conference on Communication Workshop (ICCW) 2015, pp. 216-221, 2015
[2] H. Mehmood, S. Rahman, and J. Cioffi, "Bit loading profiles for high-speed data in DOCSIS 3.1", IEEE Communications Magazine, Vol. 53, Number 3, pp. 114-120, 2015.
[3] D. Lecompte and F. Gabin, "Evolved multimedia broadcast/multicast service (eMBMS) in LTE-advanced: overview and Rel-11 enhancements", *IEEE Communications Magazine,* 50.11, pp. 68-74, 2012.
[4] R.K. Ahuja, O. Ergun, J. Orlin, A. Punnen, "A survey of very large-scale neighborhood search techniques", Discrete Appl. Math. 123, 75-102 (2002)
[5] IBM ILOG CPLEX, http://www-01.ibm.com/software
[6] F. D'Andreagiovanni, C. Mannino, A. Sassano, "GUB Covers and Power-Indexed Formulations for Wireless Network Design", Management Science, vol.59 (1), pp- 142-156,2013
[7] E. Danna, E. Rothberg, C. Le Pape, "Exploring relaxation induced neighborhoods to improve MIP solutions", Mathematical Programming, vol. 102 (1), pp, 71-90, 2005
[8] P.L. Hammer, S. Rudeanu, "Boolean methods in operations research and related areas", Springer, Heidelberg, 1968.
[9] J.B. MacQueen, "Some Methods for Classification and Analysis of Multivariate Observations", Proc. Fifth Berkeley Symp. Math. Statistics and Probability, vol. 1, pp. 281-296, 1967
[10] N. Mladenovic, P. Hansen, "Variable neighborhood search", Comput. Oper. Res. 24 (11) 1097-1100 (1997)