

Optimizing shared data plans for mobile data access

Georgios Cheirmpos, Merkouris Karaliopoulos, Iordanis Koutsopoulos

Department of Informatics

Athens University of Economics and Business, Athens, Greece

Email: {mkaralio, jordan}@aueb.gr

Abstract—The demand for mobile data has been steadily increasing over the last decade, forming an ever increasing portion of the overall Internet traffic. A considerable portion of this demand is served through capped cellular data plans that charge a fixed fee for data consumption respecting the cap and a typically higher penalty rate for additional consumption. Although shared data plans have been identified as a way to better utilize capacity that is paid for but not used, they are largely restricted to closed groups (*e.g.*, family members) or the devices of a single user.

In this paper we advocate the extension of shared data plans towards more *open* groups of users. We take the viewpoint of a platform that seeks to recommend optimal data plans to users subscribing to it and address the two main algorithmic tasks it faces: the partitioning of users to subscription groups and the selection of data plans that maximize their cost savings. We devise three algorithms that leverage clustering techniques. One of them addresses simultaneously the two tasks, whereas the other two decompose the problem and solve the two tasks sequentially. Our evaluation results suggest that the savings in subscription charges with shared plans are significant, ranging from 20% up to 80% of what users would pay with the cost-optimal individual data plans. They also highlight properties of the three algorithms and trade-offs they present involving the achieved cost savings, the intensity of under utilization and their sensitivity to deviations from the predicted users' data consumption.

Index Terms—Shared cellular data plans, mobile data, clustering, pricing.

I. INTRODUCTION

Reports by Internet commercial actors [1] and independent regulators [2] highlight similar trends about mobile data consumption. The total volume of data traffic grows fast, at more spectacular pace in non-saturated markets, video data being its main component. Mobile data has become the main vehicle for voice and text services substituting traditional voice and SMS services that steadily drop by millions (minutes/unit) per year. As a result, mobile data steadily increases its shares in the operators' revenue breakdown.

A great part of mobile data traffic is still realized through capped data plans that charge a fixed fee for consuming up to a predefined volume of data (*cap*) and a typically higher penalty (*overage charges*) for data volumes that exceed this cap. The stochasticity in the user data consumption patterns together with the relatively scarce offer of distinct cellular data plans generate inefficiencies in the actual usage of these plans. Many of those plans end up under utilized at the moment that others may systematically charge overage charges. The relevant literature has highlighted these inefficiencies and has identified the sharing of data plans as a promising way to

mitigate them. Nevertheless, so far shared data plans have been mainly considered in the context of closed groups (*e.g.*, family members) or for individual subscribers owning multiple mobile devices. In this paper we advocate the extension of shared data plans towards more open groups of users. We take the viewpoint of an online platform that issues recommendations to mobile subscribers for data plan sharing opportunities that maximize savings in their subscription charges. We particularly focus on the two algorithmic tasks that are implicit in such a platform: the formation of subscription sharing groups out of individual users and the identification of the most cost-effective data plans for them.

A. Related work

Researchers have looked into various aspects of data plan design. Hence, in [3] the authors study different pricing metrics and identify conditions under which volume-based pricing is preferred to access speed based pricing. In [4] a quite elaborate model about the strategic way a user adapts her daily data consumption to the residual data quota drives the design of data plans on behalf of an ISP. In [5] the authors take a contract-theoretic approach to the derivation of optimal data plan caps and subscription fees, having in mind data plan structures (rollover and credit data plans) that provide end users with time flexibility.

The sharing of data plans is studied in [6] [7] and [8]. In the first two cases, the emphasis is more on sharing across a user's devices. The authors of [6] compute optimal caps for data plans shared by two devices under simple assumptions about the users' consumption. On the contrary, Jin and Pang [7] work with unlimited data plans and follow the bundling model in [9] to estimate conditions about the unit cost of service under which sharing turns out to be profitable. Finally, [8] lies closer to the work in this paper since it is more focused on the sharing of data plans between multiple users. Besides unfolding motivation for shared data plans, they are the single study we are aware of that addresses the grouping of users into subscription groups. They postulate that good groupings comprise users who feature similar average values of data consumption without elaborating this further into an algorithm.

B. Our contributions

The main contributions of our work are the following:

- We introduce a new cost-sharing scheme, called double proportional cost sharing (DPCS), for splitting the subscription charges of the shared data plan “fairly” between

the subscription group members. Contrary to existing popular cost-sharing schemes, DPCS satisfies all four axiomatic requirements we have introduced for data plan cost sharing schemes.

- We devise three algorithms that leverage clustering techniques to solve the joint problem of partitioning users into subscription sharing groups and assigning cost-optimal data plans to them. One of them addresses simultaneously the two tasks, whereas the other two decompose the problem and solve the two tasks sequentially. All algorithms run in polynomial time and exhibit moderate complexity.
- We show that important savings in subscription charges are achievable when sharing is facilitated for open groups of users, even under the worst of the three algorithms. These savings range from 20% up to 80% of what users would pay with cost-optimal individual data plans.

Furthermore, we provide insights to properties of the algorithms, including their resilience to deviations of the actual consumption of users from what their demand profiles predict. The latter are used for deriving the subscription groups and choosing the optimal shared data plans.

II. USER DATA CONSUMPTION AND (SHARED) CAPPED DATA PLANS

The focus of our work is on capped data plans that stand in offer in the mobile data market. We first review the relationship between data consumption and data plan cost and then motivate shared data plans.

Capturing how users' data consumption is affected by the caps each data plan introduces is not a trivial problem. Both intuition and experience suggest that data consumption is elastic, *i.e.*, users adapt their data consumption patterns to the provisions of the plan (cap, overage charges) they subscribe to. Namely, what they consume when subscribing to a given plan is only part of their actual *a priori* demand for data, their censored demand, much as in the airlines' industry the realized bookings in a fully-booked flight are a censored version of the *actual* demand for a flight (see *e.g.*, [10]).

Nevertheless, there is no consensus as to how this demand elasticity should be captured. Existing literature approaches this differently. At one extreme, in [4], the authors come up with a detailed model of how a fully rational and strategically acting user optimizes her consumption daily depending on the residual data cap and the instantaneous utility that data consumption bears. At the other extreme, in [6], the user suppresses a fixed portion of her *a priori* demand for data, if this exceeds the data plan cap.

In this work, we make the assumption that each user u is described by a demand profile $\{d_{um}, m \in [1..T]\}$, where T is the number of charging periods (typically the 12 months of the year). The user's demand profile is an estimate of her expected monthly consumption that may rely on records of users' past data consumption and other factors such as trends of demand growth for mobile data. In section V-B, we let the actual users' data consumption differentiate from their demand profile and explore the impact this has on the efficiency of data plans.

A. Capped data plans

A capped data plan $p = (c_p, f_p, e_p)$ typically comes with a consumption cap c_p , monthly fee f_p and a penalty fee rate e_p in €/MB for excess consumption beyond the monthly cap. Alternatively, instead of charging a fixed penalty rate per MB of excess consumption, some operators sell cap extensions¹, that is supplement capped data plans of c_e MB at fixed price f_e . We denote with \mathcal{P} the set of all individual plans that are available as subscription options to users, with $P = |\mathcal{P}|$.

Formally data plans are cost functions $C(q)$ of consumed data q . The two types of functions corresponding to the two main data plan options are:

$$C_1(q) = f_p + \max(0, q - c_p) \cdot e_p \quad \text{and} \quad (1)$$

$$C_2(q) = f_p + f_e \cdot \lceil \max(0, (q - c_p)/c_e) \rceil \quad (2)$$

These two functions are shown in Fig. 1. Both are weakly increasing; $C_1(q)$, hereafter called type-1 data plan, is continuous, whereas $C_2(q)$, hereafter called type-2 data plan, is piecewise constant. Neither of the two data plan types is differentiable.




Fig. 1. Example data plans of type-1 (p_1 and p_2) and type-2 (p_3): $f_{p1} < f_{p2} < f_{p3}$, and $e_{p1} > e_{p2}$.

Hence, if q_{um} is the amount of data that user u consumes at month m , she is charged with:

$$C_p(q_{um}) = f_p + \max(0, (q_{um} - c_p)) \cdot e_p \quad (3)$$

in case of a type-1 data plan, and

$$C_p(q_{um}) = f_p + f_e \cdot \lceil \max(0, (q_{um} - c_p)/c_e) \rceil \quad (4)$$

under a type-2 data plan.

In either case, she pays a total amount of

$$C_p(q_u) = \sum_{m=1}^T C_p(q_{um}) \quad (5)$$

over T charging periods.

¹Cap extension plans are data packages that are not sold separately but only in conjunction with a "main" data plan.

B. Why sharing data plans?

Consider two users, x and y , and let q_{xt} and q_{yt} , $1 \leq t \leq T$ be the time series of their monthly data consumption (e.g., in GB). To simplify the argument, the two time series q_x and q_y are assumed to be sampling continuous uniform distributions in $[v_1, v_2]$ and $[w_1, w_2]$, respectively: $q_x \sim \mathcal{U}[v_1, v_2]$ and $q_y \sim \mathcal{U}[w_1, w_2]$.

To find out what this consumption implies for charging, we need to distinguish between two cases, *i.e.*,

$$\begin{aligned} v_2 - v_1 &\geq w_2 - w_1 & \text{and} \\ v_2 - v_1 &\leq w_2 - w_1 \end{aligned} \quad (6)$$

The expected amount that the first user will pay over one period under a dataplan $p = (c_p, f_p, e_p)$, for $c_p > v_2$, is $C_p(q_x) = f_p$, while for $c_p < v_2$, her consumption exceeds the plan consumption cap c_p with probability $p_e = \frac{v_2 - c_p}{v_2 - v_1}$. Overall, if $v_l \doteq \max(c_p, v_1)$, the expected amount that x pays over T months when subscribing to a specific data plan p is

$$C_p(q_x) = \begin{cases} T \cdot f_p, & c_p > v_2 \\ T \cdot [f_p + \int_{v_l}^{v_2} e_p \cdot \frac{v - c_p}{v_2 - v_1} dv] & \text{otherwise} \end{cases} \quad (7)$$

Likewise, the expected amount that y pays is given by:

$$C_p(q_y) = \begin{cases} T \cdot f_p, & c_p > w_2 \\ T \cdot [f_p + \int_{w_l}^{w_2} e_p \cdot \frac{w - c_p}{w_2 - w_1} dw] & \text{otherwise} \end{cases} \quad (8)$$

where $w_l \doteq \max(c_p, w_1)$.

Under a shared plan, the aggregate consumption q_{zt} of the two users over a single month is the sum of the two random variables q_{xt} and q_{yt}

$$q_{zt} = q_{xt} + q_{yt} \quad (9)$$

and assuming independence in the consumption patterns of the two users, its distribution is given by

$$f_Z(q_{zt}) = f_X(q_{xt}) * f_Y(q_{yt}) \quad (10)$$

Carrying out the convolution in (10) for the two cases in (6), we get

$$f_Z(z) = \begin{cases} 0 & z \leq v_1 + w_1 \\ \frac{z - v_1 - w_1}{(v_2 - v_1)(w_2 - w_1)} & v_1 + w_1 < z \leq v_1 + w_2 \\ \frac{1}{v_2 - v_1} & v_1 + w_2 < z \leq v_2 + w_1 \\ \frac{v_2 + w_2 - z}{(v_2 - v_1)(w_2 - w_1)} & v_2 + w_1 < z \leq v_2 + w_2 \\ 0 & z > v_2 + w_2 \end{cases} \quad (11)$$

when $v_2 - v_1 \geq w_2 - w_1$, and

$$f_Z(z) = \begin{cases} 0 & z \leq u_1 + w_1 \\ \frac{z - v_1 - w_1}{(v_2 - v_1)(w_2 - w_1)} & v_1 + w_1 < z \leq v_2 + w_1 \\ \frac{1}{w_2 - w_1} & v_2 + w_1 < z \leq v_1 + w_2 \\ \frac{v_2 + w_2 - z}{(v_2 - v_1)(w_2 - w_1)} & v_1 + w_2 < z \leq v_2 + w_2 \\ 0 & z > v_2 + w_2 \end{cases} \quad (12)$$

TABLE I

YEARLY SUBSCRIPTION FEES PAID BY 4 USERS UNDER THE BEST INDIVIDUAL AND SHARED DATA PLANS (IN €), WHEN PAIRED IN ALL THREE POSSIBLE WAYS. THE USERS' MONTHLY CONSUMPTION IS UNIFORMLY DISTRIBUTED IN $[v_1, v_2]$ AND $[w_1, w_2]$, RESPECTIVELY.

Monthly consumption $[v_1, v_2], [w_1, w_2]$ in MBs	Best plan for user 1	Best plan for user 2	Best shared plan total cost
[800, 1300], [1600, 2100]	172.56	232.56	250.8
[3000, 4000], [600, 700]	250.8	162	250.8
[800, 1300], [600, 700]	172.56	162	222
[3000, 4000], [1600, 2100]	250.8	232.56	250.8
[600, 700], [1600, 2100]	162	232.56	250.8
[3000, 4000], [800, 1300]	250.8	172.56	250.8

when $w_2 - w_1 \geq v_2 - v_1$.

If we set $s_1 = v_1 + w_1$ and $s_2 = v_2 + w_2$, the probability that the aggregate consumption of the two users will exceed the data cap of the shared dataplan $p_s = (c_{ps}, f_{ps}, e_{ps})$ is

$$p_{es} = \int_{c_{ps}}^{s_2} f_Z(z) dz \quad (13)$$

Then, the expected amount that has to be paid by the two users, *i.e.*, the “subscriber” z of the shared dataplan, equals

$$C_{ps}(z) = \begin{cases} T \cdot f_{ps}, & c_{ps} > v_2 + w_2 \\ T \cdot [f_{ps} + \int_s^{s_2} e_{ps}(z - c_s) f_Z(z) dz] & \text{otherwise} \end{cases} \quad (14)$$

where $s \doteq \max(c_{ps}, s_1)$.

Table I compares the subscription fees paid by four users with indicative uniform distributions of monthly data consumption when they are paired with each other in all three possible ways. For each possible pair, the comparisons are made between the best individual data plans they can individually subscribe to, *i.e.*, the minimum cost data plans, and the plan that minimizes the cost for their aggregate data consumption. In either case, the data plan alternatives presented to them correspond to data-only plans in offer in a middle-sized European country.

It is worth making the following remarks.

- First, savings are achievable with data plan sharing in all three ways that users can be paired with each other into subscription groups. These savings may exceed 50% of the original cost.
- Secondly, the savings with data plan sharing depend on how users are paired.
- Thirdly, for users with high data consumption, the same data plan remains the preferred one even after one more user is added (rows 2,4,6). These users essentially offload significant parts of their “unnecessary” charges by letting users with lower consumption utilize unused residuals of their data plan capacity.

This simple example highlights the cost-saving potential of shared plans and its dependence on how users are combined in subscription-sharing groups. In what follows, we formalize these dependencies and seek optimal responses to them.

III. THE SHARING DATA PLAN PROBLEM

A. Modeling shared data plans

A shareable data plan p is described by the same three parameters, c_p, f_p, e_p (c_e, f_e for type-2 data plans), but on top of them there may be an additional charge, o_p for each additional user who joins the plan. If $g \subseteq \mathcal{U}$ is the group of users that share the data plan, the full charge that has to be paid by its members at month m is

$$C_p(q_{gm}) = f_p + \max(0, q_{gm} - c_p) \cdot e_p + (|g| - 1) \cdot o_p \quad (15)$$

for type-1 data plans and

$$C_p(q_{gm}) = f_p + f_e \cdot \lceil \max(0, (q_{gm} - c_p)/c_e) \rceil + (|g| - 1) \cdot o_p \quad (16)$$

for type-2 data plans, where $q_{gm} = \sum_{u \in g} q_{um}$ is the monthly data consumption of the subscription (sharing) group g .

Then each group member u pays a share of this charge depending on her own consumption, the data consumption of the other group users, and the specific *cost-sharing scheme* that is used to split the overall plan cost into the members of the subscription group g .

1) *Splitting the plan cost among the subscribers:* The cost-sharing scheme is a function ξ_p that determines how the overall plan cost $C_p(q_{gm})$ is split among the users in g . ξ_p maps a vector of users' data consumption values $\{q_{um}\}_{u \in g}$ to cost shares $\{y_{um}\}$ for each user u in the sharing group g .

The cost-sharing scheme ξ should satisfy the following axiomatic requirements:

- (R1) It should always compute cost shares that sum exactly to the data plan cost, i.e., $\sum_{u \in g} y_{um} = C_p(q_{gm})$. This includes both the fixed cost f_p and the penalty cost due to consumption beyond the data plan cap².
- (R2) ξ has to be symmetric in all data consumption variables. Namely, if $\{y_u\}_{u \in g}$ are the cost shares that ξ computes for an initial set $\{q_{um}\}_{u \in g}$ of data consumption values per group user, when we generate arbitrary permutations of the latter across the members of the sharing group, the resulting cost shares that ξ computes should be the respective permutations of the cost shares $\{y_u\}_{u \in g}$. This ensures that ξ does not discriminate against any group user.
- (R3) The cost share that ξ computes for given user should be a non-decreasing function of her own consumption. It should not be possible for any user to utilize the data plan more heavily (hence, either leave intact or increase the overall data plan cost that is charged to the group) and, at the same time, reduce her own cost share.
- (R4) Finally, and less trivially, ξ should split the fixed fee f_p of the plan in proportion to users' contributions to the cost and it should not penalize a user with excess fees if she is not responsible for excess data consumption.

The first requirement is a prerequisite for the efficiency and practical implementation of the subscription groups. The other

²We assume that any membership cost $(|g| - 1) \cdot o_p$ is shared equally among the subscribers of the data plan p .

three embody the notion of "fairness" against all members of the sharing group.

Cost-sharing schemes have been proposed in the economics and computer science area; see, for example, [11] and [12]. Three of the most exhaustively studied schemes are the *Average Cost Pricing* (ACP), where

$$y_{um} = \frac{q_{um}}{q_{gm}} C(q_{gm}) \quad (17)$$

the *Incremental Cost Sharing* (ICS), where

$$y_{um} = C(q_{gm}) - C\left(\sum_{v \in g \setminus u} q_{vm}\right) \quad (18)$$

and the *Serial Cost Sharing* (SCS) scheme demanding that

$$y_{jm} = \frac{1}{|g| - j + 1} C(q^j) - \sum_{k=1}^{j-1} \frac{1}{|g| - k + 1} C(q^k) \quad (19)$$

where the group's user consumption values q_{um} are arranged in increasing order, $q_{1m} \leq q_{2m} \leq \dots \leq q_{|g|m}$ and $q^j = (|g| - j + 1)q_{jm} + \sum_{k=1}^{j-1} q_{km}$. We can show that

Proposition 1. *None of three cost-sharing schemes, ACP, ICS, and SCS, satisfy all four requirements (R1)-(R4) for the cost functions (15) and (16).*

Proof. All three schemes trivially satisfy (R2) and (R3) and two of them, ACP and SCS also satisfy (R1). The ICS scheme fails, at least, (R1). For example, for a group of two users with monthly data consumption values that sum below the cap of the data plan, ICS computes zero cost shares. The ACP and SCS schemes fail (R4) in different ways. ACP shares the penalty fee between all users, even when the data plan cap is exceeded because only one user consumes aggressively; whereas, it takes a few more algebraic computations to show that SCS will split the fixed fee f_p of the plan equally between users, irrespective of their individual data consumption. \square

Hence, we devise a custom cost-sharing scheme, the *double proportional cost sharing* (DPCS) scheme, which satisfies all four requirements (R1)-(R4) and is shown in Algorithm 1

Algorithm 1: Implementation of the DPCS scheme for type-1 data plans.

Input: User data consumption vector $\{q_{um}\}$ and profile demand vector $\{d_{um}\}$, $u \in g$, data plan $p = (f_p, c_p, e_p, o_p)$
Output: Individual cost shares $\{y_{um}\}$, $u \in g$

for user u in g **do**

$$y_{um} \leftarrow f_p \frac{d_{um}}{\sum_{u \in g} d_{um}}$$

if $\sum_{u \in g} q_{um} > c_p$ **then**

$$excData(u) \leftarrow \max(0, q_{um} - c_p \frac{d_{um}}{\sum_{u \in g} d_{um}})$$

$$y_{um} \leftarrow y_{um} + \frac{excData(u)}{\sum_{u \in g} excData(u)} \cdot (C_p(\sum_{u \in g} q_{um}) - f_p)$$

return $\{y_{um}\}$

The reference for judging each user's contribution to excess consumption are the profile demands $\{d_u\}$ that are used for constructing subscription groups. As far as no excess consumption exists, users share the fixed fee in proportion to their demand profiles. If the data cap is exceeded and penalty costs occur, each user is charged for the amount of data that exceed her share of the cap (*excData*), which is also proportional to her demand profile.

Note that a user who underspends at some charging period is still charged according to her profile demand. This is the standard practice with individual capped plans as well, where the user is charged a fixed fee even if she does not consume any data. In shared plans, this practice prevents the unfair penalization of active group users due to one or more non-active ones, without whom they could subscribe to a data plan with smaller cap.

B. Problem formulation

The mission of the platform is then to partition users into a finite number of subscription groups, each of size $|g|, 1 \leq |g| \leq g_{max}$, and assign them to data plans such that their achieved savings in data plan subscription fees are maximized.

Let G be a user partition and g_1, g_2, \dots, g_n the subscription groups that make it up. The monetary savings of a user u , when she is member of subscription group g_i are:

$$s_u = \sum_{m \in T} C_{p_i}(q_{um}) - \xi_{p_g}(q_{um}, \mathbf{q}_{-um}) \quad (20)$$

where $\mathbf{q}_{-um} := \sum_{u \in g \setminus u} q_{um}$,

$$p_i = \arg \min_{p \in \mathcal{P}} \sum_{m \in T} C_p(q_{um}) \quad (21)$$

is the plan minimizing what the user is paying under the best standalone data plan, and

$$p_g = \arg \min_{p \in \mathcal{P}_s} \sum_{m \in T} C_p(\sum_{u \in g} q_{um}) \quad (22)$$

is the plan that minimizes what the group collectively pays under a shared data plan³.

Taking into account that savings of a fixed amount of money is of higher value to someone paying lower subscription fees as opposed to someone paying higher fees we normalize these savings against what users pay under the most cost-effective individual plan

$$s_{u,n} = \frac{s_u}{\sum_{m \in T} C_{p_i}(q_{um})} \quad (23)$$

The platform then seeks to partition users into subscription groups and assign data plans to them such that their normalized cost savings are maximized. Formally, the platform is after

³An alternative is to define as optimal plan for a given subscription group the one that maximizes the minimum savings over all users in the group. The choice of MAXSUM (MINSUM) vs. MAXMIN (MINMAX) optimization criterion, or efficiency vs. fairness objective, is a recurring theme in all assignment/allocation problems that involve multiple players [13].

an optimal partition \mathbf{G}^* and assignment \mathbf{x}^* of data plans to partitions that solve the following optimization problem

$$\max_{\mathbf{x}, \mathbf{G}} \sum_{u \in \mathcal{U}} s_{u,n} \quad (OPT)$$

$$s.t. \quad 1 \leq |g_i| \leq g_{max}, \quad g_i \in G \quad (24)$$

$$g_i \cap g_j = \emptyset \quad \forall g_i, g_j \in G \quad (25)$$

$$s_{u,n} \geq 0 \quad u \in \mathcal{U} \quad (26)$$

IV. SOLVING THE SHARING DATA PLAN PROBLEM

The problem (OPT) is not trivial. It entails two interrelated tasks, the partitioning of users into subscription groups and the assignment of data plans to them. Hereafter, we present three algorithms for it. All three algorithms leverage clustering starting with singleton clusters corresponding to individual users and working with their demand profiles, $\{d_{um}\}_{u \in \mathcal{U}}$. However, the first one attacks the user partitioning and the data plan assignment tasks simultaneously, whereas the other two decompose the problem: first, they partition users into subscription groups and then, in a second simpler step, they identify optimal shared data plans for them.

A. Agglomerative cost-minimization clustering

The algorithm first identifies the optimal data plan for each user, *i.e.*, the plan that minimizes her expected charge over T charging periods under her demand profile. It then initiates the agglomerative clustering process. At each step in this process, the algorithm merges those existing clusters g_k, g_l , with $|g_k| + |g_l| \leq g_{max}$ that maximize the normalized subscription cost savings for the members of the two clusters

$$score(g_k, g_l) = \frac{\sum_{m \in T} \left(C_{p_k}(d_{g_k m}) + C_{p_l}(d_{g_l m}) - C_{p_{kl}}(d_{g_{kl} m}) \right)}{\sum_{m \in T} \left(C_{p_k}(\sum_{u \in g_k} d_{um}) + C_{p_l}(\sum_{u \in g_l} d_{um}) \right)} \quad (27)$$

under the assumption that cost-optimal data plans

$$p_k = \arg \min_{p \in \mathcal{P}} \sum_{m \in T} C_p(\sum_{u \in g_k} d_{um}) \quad \text{and} \quad (28)$$

$$p_{kl} = \arg \min_{p \in \mathcal{P}} \sum_{m \in T} C_p(\sum_{u \in g_k \cup g_l} d_{um}) \quad (29)$$

are chosen in each case.

The clustering process ends when either the subscription group limit g_{max} is reached for each cluster or no further cost savings are possible for any subscription group. The pseudocode of the algorithm is shown in Algorithm 2.

Algorithm complexity: From a complexity point of view the algorithm carries out $O(U)$ steps and in each one of those, it searches for the best one out of $O(U^2)$ pairs of clusters to merge. For each one of the candidate cluster-pairs, the algorithm computes the minimum-cost plan, which requires $O(P)$ time. Hence, the overall time-complexity of the algorithm is $O(U^3P)$.

Algorithm 2: Agglomerative cost-minimization clustering

Input: User demand profiles $\{d_u\}, u \in \mathcal{U}$; group size limit, g_{max} ; data plan cost functions $C_p(q), p \in \mathcal{P}$

Output: Sharing groups, $\{g\}$, and data plan assignments, $p_{opt}(g), \cup g = \mathcal{U}$

Start with one cluster for each user: $g_u \leftarrow u \in \mathcal{U}$

Compute $p_{opt}(u) \leftarrow \min_{p \in \mathcal{P}} C_p(D_u), \forall u \in \mathcal{U}$

while there are clusters with size $< g_{max}$ **do**

For each cluster pair (g_k, g_l) compute $\text{score}(g_k, g_l)$ from (27), (28)

Merge the two clusters $(g_{k'}, g_{l'})$ with the highest positive score

$$p_{opt}(g_{k'} \cup g_{l'}) \leftarrow \arg \min_{p \in \mathcal{P}} \sum_{m \in T} C_p(\sum_{u \in g_{k'} \cup g_{l'}} d_{um})$$

B. Agglomerative uniform-consumption clustering

The cost-minimization clustering algorithm simultaneously constructs subscription groups and assigns optimal data plans to them. On the contrary, this algorithm decomposes the problem into its two subproblems: first, it clusters users into subscription groups, and, in a second step, it identifies optimal plans for them.

The metric that scores clusters throughout the process is the *normalized fluctuation of the group demand* over the period for which the users' demand profiles are available. Therefore, the demand fluctuation for a cluster g is measured by

$$d_F(g) = \frac{\max_{m \in T} \sum_{u \in g} d_{um} - \min_{m \in T} \sum_{u \in g} d_{um}}{\min_{m \in T} \sum_{u \in g} d_{um}} \quad (30)$$

and in each step of the algorithm's execution, we merge existing clusters (g'_k, g'_l) such that

$$(g'_k, g'_l) = \arg \min_{g_k, g_l} d_F(g_k \cup g_l) \quad (31)$$

The intuition behind the cluster score is that cost savings are achieved when we group together users who can absorb the temporal fluctuation in each others' demands and together present a flatter profile that can be more easily matched to a data plan. This is reminiscent of the statistical multiplexing gains achieved when aggregating smaller traffic streams from end user access links to traffic aggregates in higher capacity links.

The algorithm terminates when either g_{max} is reached for each cluster or no further improvement is feasible in the normalized fluctuation of demand for any cluster pair during a merging step. Finally, the data plan p_k assigned to a group g_k is given by (28).

Algorithm complexity: Each step of the clustering algorithm takes $O(U^2 + O(U)) = O(U^2)$ time since we need to compute the score d_F for all possible cluster pairs and merge the two that minimize it. Overall, the time complexity of the

Algorithm 3: Agglomerative uniform-consumption clustering

Input: User demand profiles $\{d_u\}, u \in \mathcal{U}$; group size limit, g_{max} ; data plan cost functions $C_p(q), p \in \mathcal{P}$

Output: Sharing groups, $\{g\}$, and data plan assignments, $p_{opt}(g), \cup g = \mathcal{U}$

Start with one cluster for each user: $g_u \leftarrow u \in \mathcal{U}$

while there are clusters with size $< g_{max}$ **do**

For each pair of clusters (g_k, g_l) compute $\text{score}(g_k, g_l)$ after (30)

Merge the two clusters $(g_{k'}, g_{l'})$ after (31)

for every group g in the resulting cluster structure **do**

$$p_{opt}(g) \leftarrow \min_{p \in \mathcal{P}} \sum_{m \in T} C_p(\sum_{u \in g} d_{um})$$

Algorithm 4: Double greedy maximal uniform-consumption clustering

Input: User demand profiles $\{d_u\}, u \in \mathcal{U}$; group size limit, g_{max} ; data plan cost functions $C_p(q), p \in \mathcal{P}$

Output: Sharing groups, $\{g\}$, and data plan assignments, $p_{opt}(g), \cup g = \mathcal{U}$

Start with one cluster for each ungrouped user,

$g \leftarrow u \in \mathcal{U}$

while $\mathcal{U} \neq \emptyset$ **do**

build d_F -maximal clusters $cl(u), u \in \mathcal{U}$

Rank the U maximal clusters $cl(u)$ in order of non-decreasing d_F , see (30).

while $|\{cl(u)\}| > 1$ **do**

Add the top disjoint clusters $\{m\}$ to the clustering structure

$\mathcal{U} \leftarrow \mathcal{U} \setminus \{u \in \{m\}\}$

for every group g in the resulting cluster structure **do**

$$p_{opt}(g) \leftarrow \min_{p \in \mathcal{P}} C_p(\sum_{u \in g} d_{um})$$

algorithm is $O(U^3)$, for the clustering part, plus $O(U \cdot P)$ for the data plan assignment part, *i.e.*, $O(U(U^2 + P)) = O(U^3)$ since typically $P \ll U^2$.

C. Double greedy maximal uniform-consumption clustering

Similar to the agglomerative uniform-consumption clustering, this algorithm decomposes the original problem into the grouping and the data plan subproblems and uses the normalized fluctuation of demand measure, d_F , to score clusters. However, the algorithm is no longer agglomerative. It rather searches iteratively and more exhaustively for possible subscription groups.

The algorithm starts from each individual user u and greedily builds U different u -maximal clusters. These clusters contain u and they are maximal in the sense that they cannot increase any more either because their size is g_{max} or because no addition of another user can further decrease the

TABLE II
SET OF PLANS AND CORRESPONDING CAPS (c_p), FIXED FEE (f_p) AND OVERAGE CHARGES (e_p).

ID	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	P16	P17
c_p (GB)	0.5	2	5	10	20	100	0.4	1	2	4	5	10	1	3	7	15	30
f_p (€)	4.85	9.8	14.76	19.71	24.66	44.47	3.84	7.72	11.6	15.48	17.42	23.24	4.99	9.99	19.99	29.99	49.99
e_p (€/MB)	0.019	0.019	0.019	0.019	0.019	0.019	0.039	0.039	0.039	0.039	0.039	0.039	0.19	0.19	0.19	0.19	0.19

cluster's d_F value. Since these U clusters typically overlap, the algorithm, greedily again, retains disjoint clusters with the minimum d_F value. Namely, the algorithm ranks the clusters in order of non-decreasing d_F and picks as many as possible disjoint ones.

Users who are included in those clusters are removed from consideration in the second iteration of the algorithm, which builds maximal clusters from scratch for the remaining users. A new set of disjoint clusters with minimum d_F scores is chosen and the corresponding users are removed from consideration. This doubly greedy process of maximal cluster formation and selection of disjoint clusters continues until all users are clustered. Note that some users may end up standalone if no pairing with another user can decrease the fluctuation in their aggregate consumption.

The resulting groups are then matched, in a separate step, with the shared data plan that minimizes the subscription fees they need to pay as a group. The overall algorithm is shown in Algorithm 4.

Algorithm complexity The first iteration of the algorithm includes all $U_1 \equiv U$ users and requires $O(U_1^{g_{max}})$ steps for building maximal clusters plus $O(U_1 \ln U_1)$ time for sorting the clusters and picking the maximum possible number of disjoint ones. Subsequent iterations involve reduced sets of users $U_k < U$ and require time $O(U_k^{g_{max}} + O(U_k \ln U_k))$. The overall time for the clustering step is $\sum_k \left(O(U_k^{g_{max}}) + O(U_k \ln U_k) \right) \subseteq O(U^{1+g_{max}})$. An additional $O(U \cdot P)$ time is needed for the data plan assignment step.

V. EVALUATING THE THREE ALGORITHMS

A. Methodology

We carry out a comparative study of the three algorithms presented in section IV: the agglomerative cost-minimization clustering (ACMC), the agglomerative uniform-consumption clustering (AUCC) and the double greedy maximal uniform-consumption clustering (DGMC). Our input datasets include:

1) *User Data*: Each user is represented by a synthetically generated T -dimensional profile demand vector $d_u = [d_{u1}, d_{u2}, \dots, d_{uT}]$, where T is the number of charging periods covered by the profile of u . We fix the average user monthly demand values \bar{d}_u according to the mobile data plan distributions reported in [14] and generate the T values by sampling normal distributions $\mathcal{N}(\bar{d}_u, \sigma_u)$. Unless otherwise stated, the demand values are in MB, $T = 12$ and $\sigma_u = 0.2 \cdot \bar{d}_u$, $u \in \mathcal{U}$.

2) *Data plans*: We have identified and collected information about 17 different cellular data plans that stand in offer by operators in various European countries. These data plans, as




Fig. 2. Data plan cap vs. cost per MB for each of the 17 cellular data plans.

listed in Table II, feature various data caps (400MB-100GB), fixed fees and overage charges and they form the set \mathcal{P} .

3) *Performance measures*: We compare the three algorithms along various dimensions. The ultimate performance measure for all algorithms are the cost savings they achieve for the mobile users. We measure such savings for each user in absolute terms (in €)

$$sav(u) = \min_{p \in \mathcal{P}} \sum_{m \in \mathcal{T}} C_p(d_{um}) - \min_{p \in \mathcal{P}} \sum_{m \in \mathcal{T}} \xi_p(d_{um}, d_{-um}) \quad (32)$$

and relative terms, *i.e.*, as ratios of the fee savings over the charges under the optimal individual plan

$$nsav(u) = \frac{sav(u)}{\min_{p \in \mathcal{P}} C_p(\sum_{m \in \mathcal{T}} d_{um})} \quad (33)$$

We report histograms and empirical cumulative distribution functions of these savings over the user population. We also compute the portions of users who experience normalized savings beyond $\alpha \in [0, 1]$ as

$$perc(\alpha) = \frac{\sum_{u \in \mathcal{U}} 1_{nsav(u) > \alpha}}{U} \quad (34)$$

where where 1_x is the indicator function that equals one when condition x is true.

Moreover, a number of statistics yield further insights into the way the three algorithms assign users to subscription groups. The first one is the *distribution of subscription group sizes*, each algorithm generates. A second one, relates to how well each subscription group utilizes the data plan it is assigned to, *i.e.*, how much data remain unused and how much excess consumption takes place.

B. Results

1) *Subscription cost savings*: Fig. 3 reports the predicted cost savings per user, according to (32), (33), and (34), when




Fig. 3. Distribution of per user subscription cost savings under the three algorithms, computed based on their demand profiles.

TABLE III
DISTRIBUTION OF (NORMALIZED) SUBSCRIPTION COST SAVINGS UNDER THE THREE ALGORITHMS ($U = 1400$)

	sav(u) in €					nsav(u) in %					$\sum sav(u)$
	[0,75]	[75,151]	[151,226]	[226,302]	[302,377]	[0,0.19]	[0.19,0.37]	[0.37,0.56]	[0.56,0.75]	[0.75,1]	
DGMC	26.57%	63.35%	6.35%	1.85%	1.64%	1.42%	9.5%	35.92%	44.35%	8.57%	137061.27€
AUCC	26.28%	61.92%	6.92%	1.85%	3%	0.5%	9.85%	30.64%	44.64%	14.35%	142971.3€
ACMC	24.64%	64.42%	4.64%	0.14%	6.07%	0%	0.92%	31.5%	48.85%	18.71%	157502.1€

the three algorithms derive subscription groups and assign data plans to them according to the user demand profiles.

In absolute terms, the savings with the three algorithms appear to be comparable. The ACMC algorithm distinguishes from the other two in securing higher annual subscription savings, beyond 300€, for distinctly more subscribers (more than 6%) than the other two algorithms. This results in aggregate savings that are 13% (10%) higher than the DGMC (AUCC) algorithms, as shown in the last column of Table III.

The performance advantage of the ACMC algorithm is more evident in terms of normalized savings, the measure that this algorithm actually tries to optimize (see Algorithm 2). The algorithm consistently tends to produce subscription groups that save more with respect to what their members paid under individual plans. This trend is clearer in the third plot of Fig. 3, where we see that practically $\forall \alpha \in [0, 1]$

$$perc^{ACMC}(\alpha) > \max(perc^{AUCC}(\alpha), perc^{DGMC}(\alpha))$$

implying a stochastic dominance relationship of ACMC over the other two algorithms in terms of achievable normalized subscription savings.

As a last interesting note, all three algorithms result in subscription charge savings, $sav(u) > 0$, for all mobile subscribers⁴. This is important since all algorithms avoid checking subscriber-level constraints (*i.e.*, constraint (26)), that would greatly burden the run times of the algorithms. They rather cluster users by computing scores and checking constraints at the group-level (see Algorithms 2, 3, and 4).

2) *Subscription group size*: Table. V yields more insights to the way the three algorithms work. The two algorithms that use the monthly fluctuation of user demand (consumption) as a proxy measure to make user-grouping decisions are strongly biased towards large subscription groups. The bias is slightly

⁴In all runs and subscription groups produced by the three algorithms for 1500 users, we could only count two users who ended up being charged higher with the shared plan derived with the DGMC algorithm than under the optimal individual plan; none for the other two algorithms.

stronger for the AUCC, which gathers almost all (99.13%) of the users into maximum size subscription groups. The respective number is around 10% smaller for the DGMC, whereas the ACMC spreads the users in more balanced way among subscription groups of size three to five. Although larger subscription groups reduce the subscription charges leveraging the economy of scale properties of data plans (see Fig. 2), they do not necessarily do it in the optimal manner. Simultaneously solving the subscriber grouping and the data plan assignment tasks, the ACMC algorithm reaches better decisions about the number and size of subscription groups that maximize the benefits for the subscribers.

3) *Sensitivity of cost savings to data consumption prediction accuracy*: The assignment of shared data plans to subscribers is made on the basis of their demand profiles. These profiles rely on data about their data consumption in past charging periods and form a predictor for their future consumption. How would the significant cost savings reported in Fig. 3 and Table III be affected by different amounts of data consumption?

We recompute the cost savings in subscription charges when the amount of data user u consumes each charging period $m \in \mathcal{T}$ is sampled from a normal distribution $\mathcal{N}(1.1d_{um}, 0.05)$; namely, her actual data consumption is systematically underestimated by her demand profile and, on top of this, there is a mild fluctuation around the actual mean consumption.

The new cost savings, realized with the shared data plans assigned on the basis of the user demand profiles, are reported in Table IV. First of all, a non-negligible part of subscribers now end up paying more than they did under individual plans (see the two columns reporting negative savings). For those users, the assigned data plans under their demand profiles are no longer optimal and the process should be repeated under updated data on their actual consumption. This number is higher for the two clustering algorithms (AUCC, DGMC) that work with demand fluctuation over charging periods and much smaller for the ACMC, which is more robust in this respect.

On the other hand, the savings under the ACMC algorithm




Fig. 4. Distribution of per user subscription cost savings under the three algorithms, computed based on their realized data consumption.

TABLE IV
DISTRIBUTION OF (NORMALIZED) SUBSCRIPTION COST SAVINGS UNDER THE THREE ALGORITHMS ($U = 1400$)

	sav(u) in €					nsav(u) in %					$\sum sav(u)$
	< 0	[0,87]	[87,175]	[175,263]	[263,351]	< 0	[0,0.23]	[0.23,0.46]	[0.46,0.69]	[0.69,1]	
DGMC	7%	43.85%	44.28%	3.07%	1.64%	7.07%	4.14%	21.78%	50.64%	16.35%	95959.07€
AUCC	4.57%	47.35%	40.64%	4.28%	3.14%	4.57%	5.07%	23.35%	46.71%	20.28%	108142.48€
ACMC	1.92%	49.14%	42.28%	3.92%	2.71%	1.92%	10.14%	22.78%	48.28%	16.85%	119558.64€

TABLE V
GROUP SIZE DISTRIBUTION

	1	2	3	4	5
DGMC	0%	2.94%	2.52%	2.94%	91.5%
AUCC	0%	0.43%	0.43%	0%	99.13%
ACMC	0%	0.36%	17.71%	39.11%	42.8%

for those users who still benefit from shared plans (Fig. 4), are less profound than in Fig. 3. In fact, the algorithm is no longer the one it achieves the highest savings for more subscribers and is equivalent or even inferior to its two alternatives. In the aggregate, it still exhibits the top cumulative charge savings (last column in Table IV) but it suffers higher loss when compared to the AUCC algorithm that emerges as the most resilient algorithm overall.

TABLE VI
DIFFERENCE IN AGGREGATE SUBSCRIPTION CHARGES PAID BY CELLULAR USERS UNDER PROFILE DEMANDS AND ACTUAL CONSUMPTION

	$\sum savings (D_u)$	$\sum savings (Q_u)$	Difference
DGMC	137061.27€	95959.07€	-41,102.2€
AUCC	142971.3€	108142.48€	-34,828.82€
ACMC	157502.1€	119558.64€	-37,943.46€

VI. CONCLUSIONS

We have looked closely into the fundamental and non-trivial problem of partitioning users into subscription groups that share capped cellular data plans. We have first introduced a cost-sharing scheme that matches the requirements to fairly split the subscription charges between the users sharing the data plan. Then we devised three clustering-type algorithms for efficiently partitioning users into subscription sharing groups. Finally, we have assessed the achievable savings in subscription charges under the three algorithms and extracted insights to the way they operate.

ACKNOWLEDGMENT

The research work of M. Karaliopoulos received funding from the Hellenic Foundation for Research and Innovation

(HFRI) and the Hellenic General Secretariat for Research and Innovation (GSRI), under grant agreement No 892. I. Koutsopoulos acknowledges the support from the CHIST-ERA grant CHIST-ERA-18-SDCDN-004, through GSRI grant number T11EPA4-00056.

REFERENCES

- [1] Cisco Annual Internet Report (2018–2023). White Paper. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.pdf>
- [2] Hellenic telecommunications post commission. (2019) 2019 Annual Report. [Online]. Available: <https://services.eett.gr/areport2019/en/>
- [3] Y.-J. Chen and K.-W. Huang, “Pricing data services: Pricing by minutes, by gigs, or by megabytes per second?” *Info. Sys. Research*, vol. 27, no. 3, p. 596–617, Sep. 2016.
- [4] L. Zheng, C. Joe-Wong, M. Andrews, and M. Chiang, “Optimizing data plans: Usage dynamics in mobile data networks,” in *Proc. IEEE INFOCOM 2018*, Honolulu, HI, USA, April 2018, pp. 2474–2482.
- [5] Z. Wang, L. Gao, and J. Huang, “Multi-dimensional contract design for mobile data plan with time flexibility,” in *Proc. 18th ACM Mobihoc*, Los Angeles, CA, USA, 2018, p. 51–60.
- [6] S. Sen, C. Joe-Wong, and S. Ha, “The economics of shared data plans,” in *22nd Workshop on Information Technologies and Systems, WITS 2012*, Jan. 2012, pp. 157–162.
- [7] Y. Jin and Z. Pang, “Smart data pricing: To share or not to share?” in *Proc. INFOCOM 2014 WKSHPs*, 2014, pp. 583–588.
- [8] J. C. Cardona, R. Stanojevic, and N. Laoutaris, “Collaborative consumption for mobile broadband: A quantitative study,” in *Proc. 10th ACM CoNEXT*, New York, NY, USA, 2014, p. 307–318.
- [9] R. Venkatesh and W. Kamakura, “Optimal bundling and pricing under a monopoly: Contrasting complements and substitutes from independently valued products,” vol. 76, no. 2, 2003.
- [10] L. Weatherford and S. Pöhl, “The price of anarchy of serial, average and incremental cost sharing,” *Journal of Revenue and Pricing Management*, vol. 3, no. 1, pp. 234–254, October 2002.
- [11] H. Moulin and S. Shenker, “Average cost pricing versus serial cost sharing: An axiomatic comparison,” *Journal of Economic Theory*, vol. 64, no. 1, pp. 178–201, 1994.
- [12] H. Moulin, “The price of anarchy of serial, average and incremental cost sharing,” *Economic Theory*, vol. 36, no. 3, pp. 379–405, 2008.
- [13] D. Bertsimas, V. F. Farias, and N. Trichakis, “On the efficiency-fairness trade-off,” *Management Science*, vol. 58, no. 12, p. 2234–2250, December 2012.
- [14] Ericsson mobility report. (2017) Shifting mobile data consumption and data plans. [Online]. Available: <https://www.ericsson.com/en/mobility-report/articles/shifting-mobile-data-consumption-data-plans>