Optimizing shared data plans for mobile data access

Georgios Cheirmpos, Merkouris Karaliopoulos, Iordanis Koutsopoulos

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

ITC 33, Avignon, France, 1/9/2021

H.F.R.I. Hellenic Foundation for Research & Innovation

GSRI GENERAL SECRETARIAT FOR RESEARCH AND INNOVATION

Mobile data trends

- Average usage per device >10GB,
  - to reach 35GB by end of 2026.
- Easily accessible video content on mobile devices accounts for 2/3 of total traffic,
  - predicted to rise to 77% by 2026.
- Proliferation of 5G and beyond networks will only lead to further data traffic growth
  - at the expense of traditional voice and message services

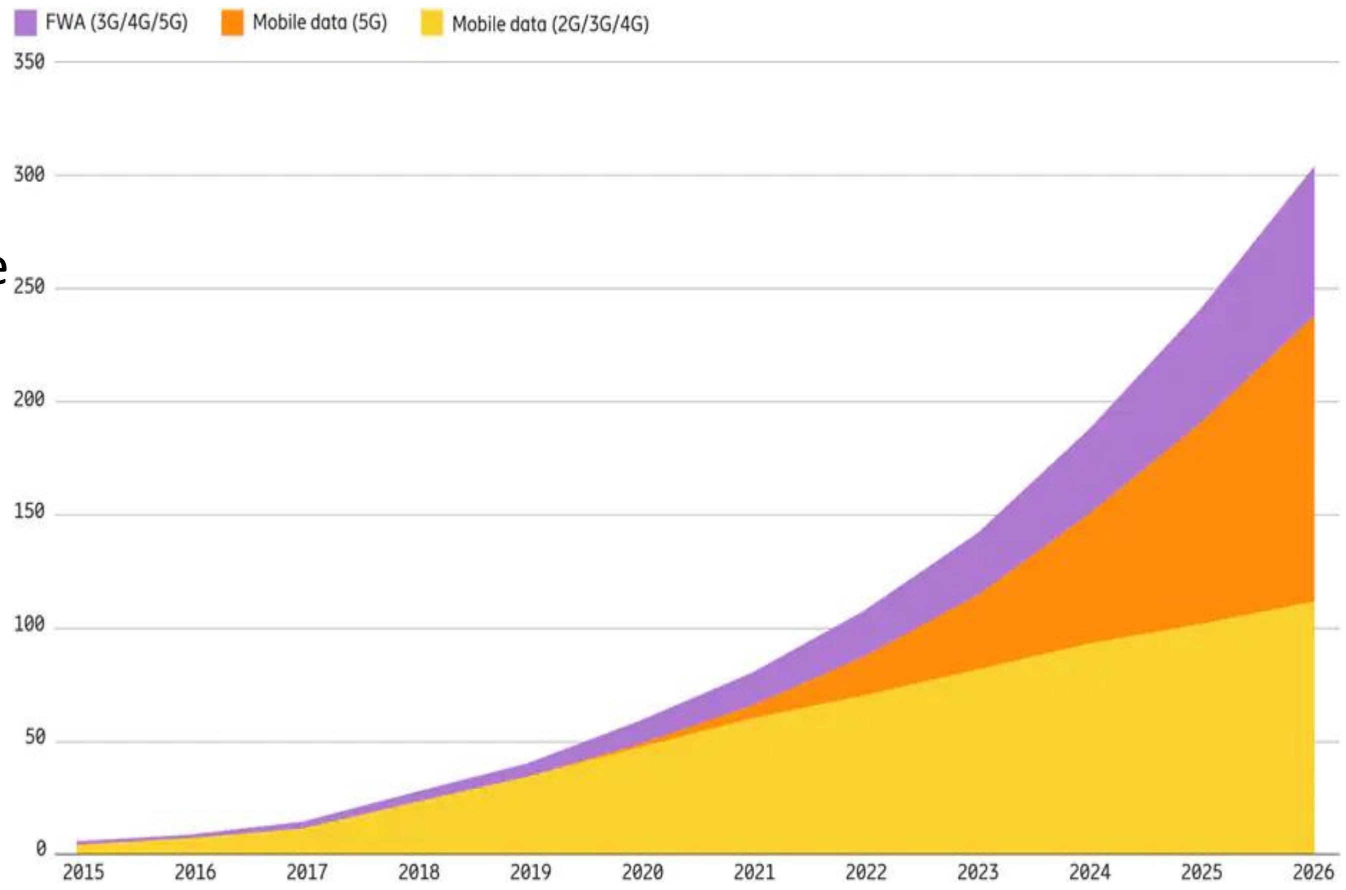

https://www.ericsson.com/en/mobility-report/dataforecasts/mobile-traffic-forecast

Capped cellular data plans

https://www.itsworthmore.com/blog/post/how-much-data-do-you-need

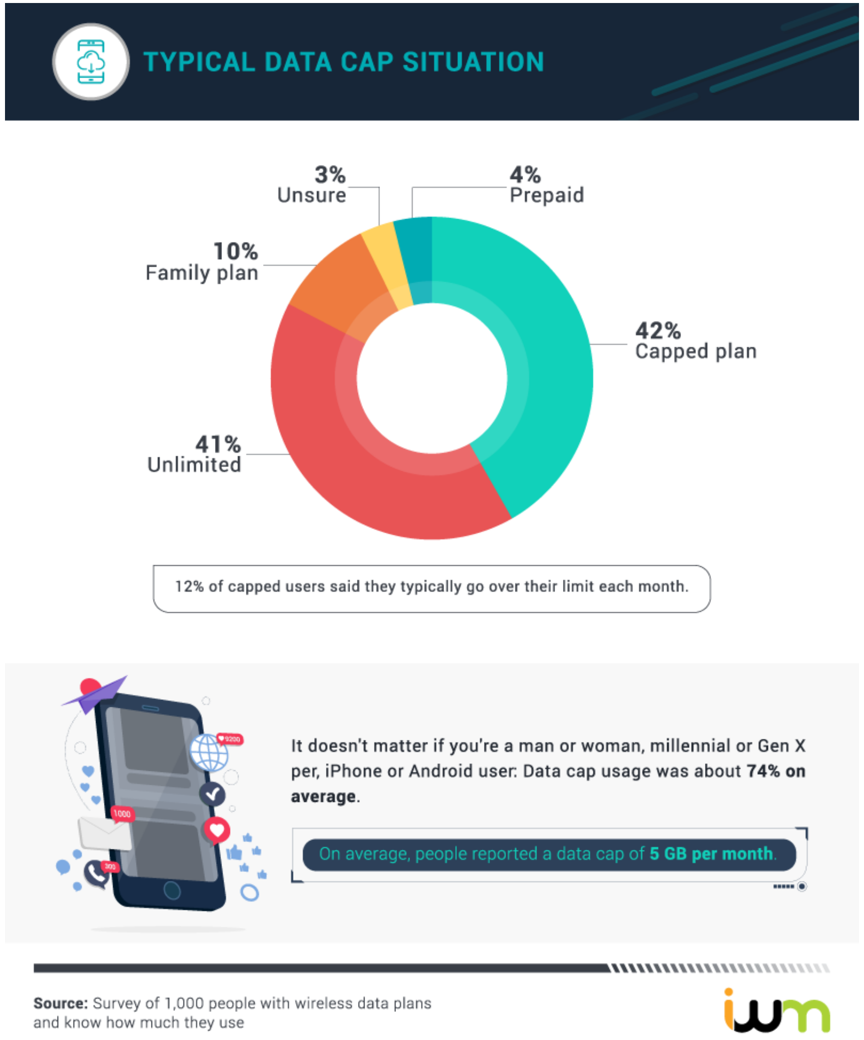

- Capped data plans remain standard practice for mobile data subscriptions
  - *unlimited* data plans often come with hidden restrictions (e.g., slower speed after an X amount of data is consumed or daily caps)
- Limited choice with inadequate granularity results in
  - either excess consumption and overage fees;

    or
  - plan underutilization and unnecessary subscription costs

Shared cellular data plans

- Subscription sharing is around for almost a decade
  - users share the data cap and possible excess consumption fees
  - data plans better utilized, less unused capacity

- “Closed” groups of users (e.g., family) or more than two devices of the same subscriber
  - data tethering to devices, ‘’hidden sharing’’,

  sometimes blocked by sim operator

Figure 4. Barriers to acquiring a new shared data plan

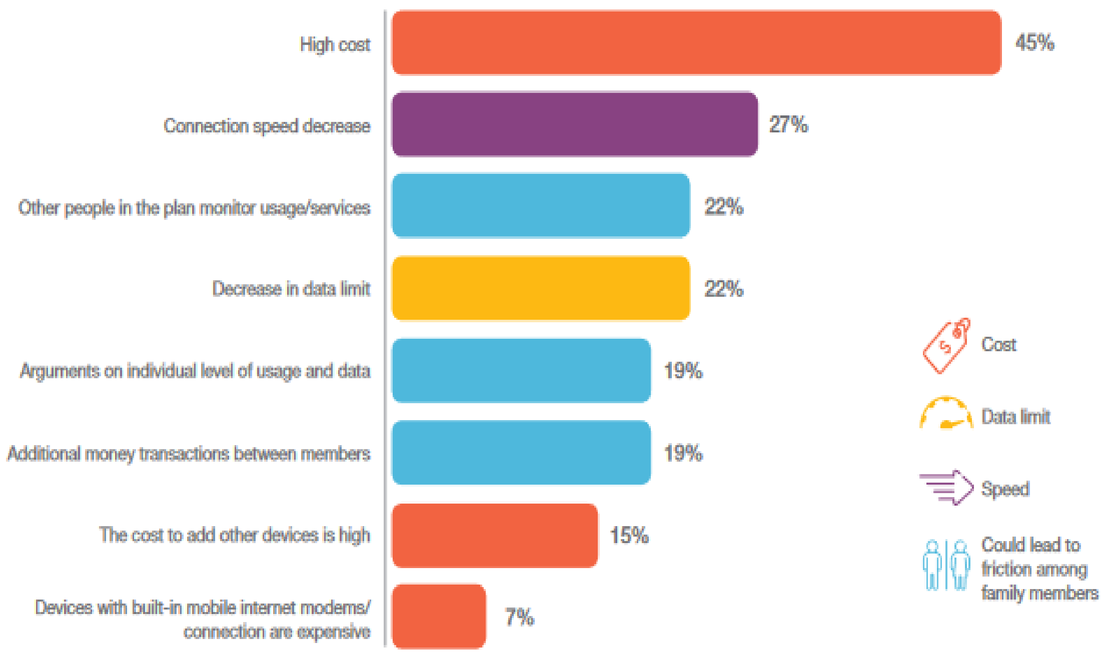

Source: Ericsson ConsumerLab, Embracing data sharing, 2015
Base: Smartphone users in Brazil, India, Japan, South Korea, the UK and the US, who are aware of shared data plans but do not have one

*https://www.ericsson.com/491d9c/assets/local/news/2015/1/ericsson-consumerlab-embracing-data-sharing.pdf

Prior work on (shared) data plans

- **Work on individual mobile data plans has addressed**
  - comparing different pricing strategies (e.g., volume-based *vs.* access speed-based pricing) [1]
  - modeling censored data demand and how users adapt to the remaining plan quota and tariffs [2][4]
  - deriving optimal caps and subscription fees (with rollover) [3]

- **Work on shared mobile data plans is related to**
  - comparing shared vs. individual plans, with unlimited data plans and bundling models [5][7]
  - computing optimal cap values for sharing between 2 devices under a simple model for user demand elasticity [4]
  - assessing cost savings with shared plans and "open" subscription groups [6]

[1] Y.-J. Chen and K.-W. Huang, "Pricing data services: Pricing by minutes, by gigs, or by megabytes per second?" Info. Sys. Research, 2016.
[2] L. Zheng, C. Joe-Wong, M. Andrews, and M. Chiang, "Optimizing data plans: Usage dynamics in mobile data networks," IEEE INFOCOM 2018
[3] Z. Wang, L. Gao, and J. Huang, "Multi-dimensional contract design for mobile data plan with time flexibility," ACM MOBIHOC 2018
[4] S. Sen, C. Joe-Wong, and S. Ha, "The economics of shared data plans," WITS 2012
[5] Y. Jin and Z. Pang, "Smart data pricing: To share or not to share?" INFOCOM 2014 WKSHPS
[6] J. C. Cardona, R. Stanojevic, and N. Laoutaris, "Collaborative consumption for mobile broadband: A quantitative study," CoNEXT 2014
[7] R. Venkatesh. W. Kamakura, "Optimal bundling and pricing under a monopoly: Contrasting complements and substitutes from independently valued products," Journal of Business, 2003

Our contributions

- Take the viewpoint of a data plan recommendation platform and look closer into the two main “algorithmic” tasks related to shared data plans for “*open*” groups
    - how to partition users into subscription groups
    - how to share the cost of a shared data plan between the subscription group members

- New cost sharing scheme for “fairer” sharing of subscription charges between users
    - **Double Proportional Cost Sharing**

- Three clustering-based algorithms for the joint user grouping and data plan assignment problem
    - **Agglomerative cost-minimization clustering -ACMC**, simultaneously identifies groups and computes the optimal plan for them
    - **Agglomerative uniform-consumption clustering (AUCC)** and **Double greedy maximal uniform-consumption clustering (DGMC)** decompose the joint problem

- Achievable savings in subscription charges in the 20% - 80% range under perfect prediction of users’ data consumption
    - algorithms turn out to be robust when user demand profiles deviate from actual data consumption

Modeling : user demand profile

- **User demand profile :** Each user *u* is described by a demand profile $\{d_{um}\}$, m=1..*T*, *T* the number of charging periods (typically months)
  - may come from historical data on past data consumption

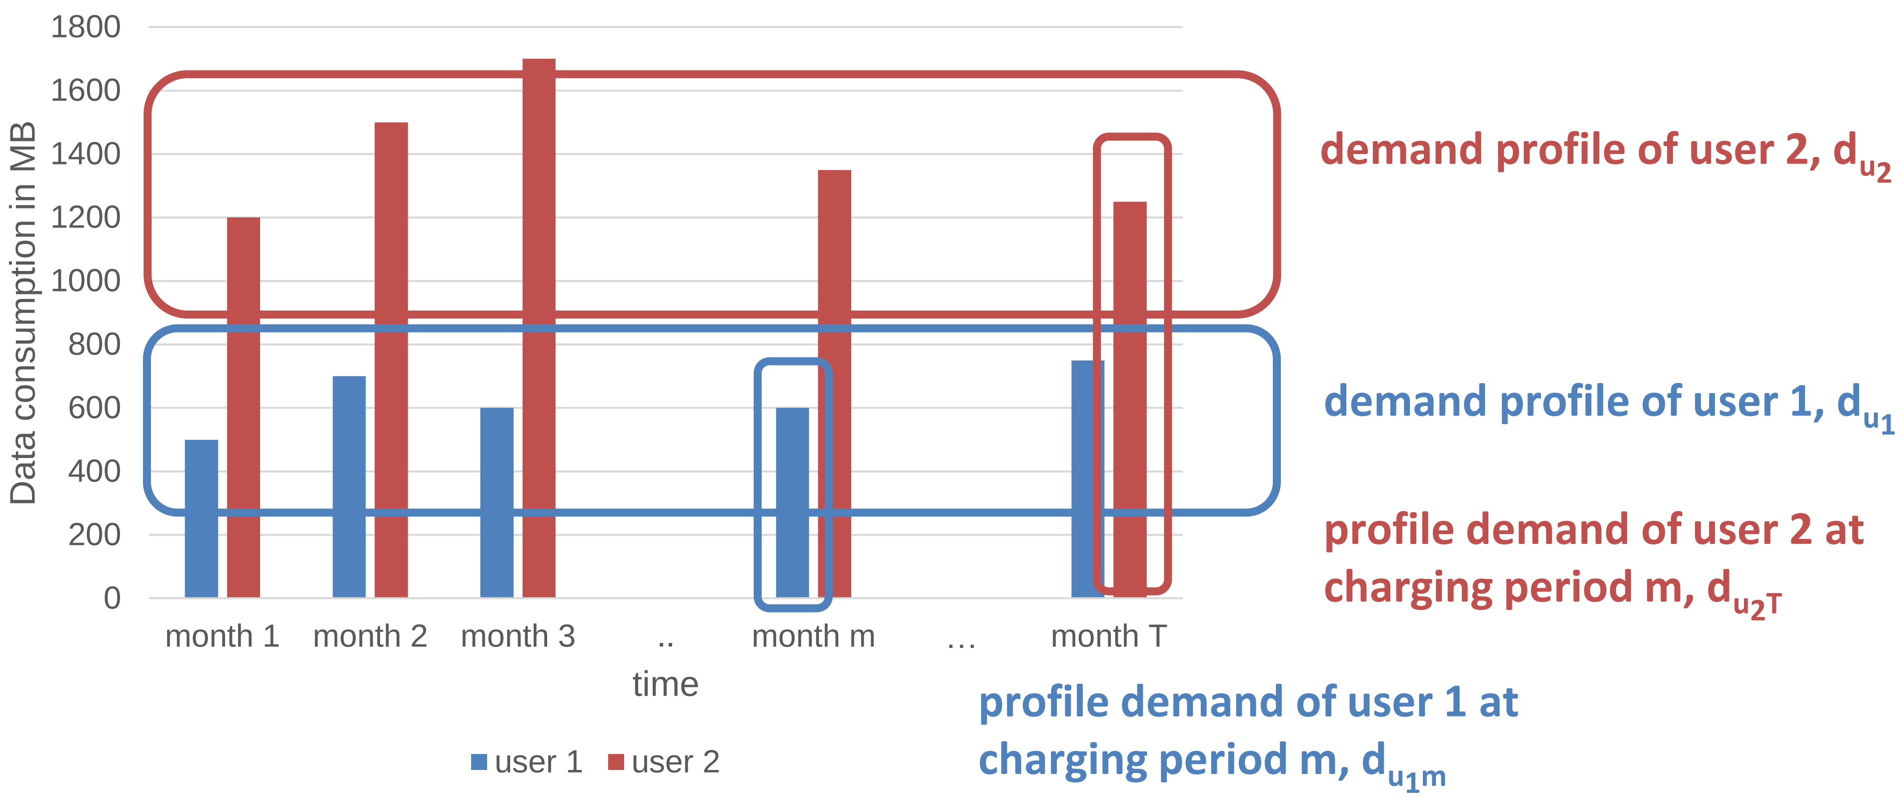

Modeling : capped plans

- Each capped plan p is described by a tuple (c_p, f_p, e_p), where:
    - c_p : plan cap, *i.e.*, amount of data credit at the expense of a fixed monthly fee
    - f_p : fixed monthly fee
    - e_p : overage fee rate (usually per MB of excess consumption beyond c_p)

- Cost function $C_p(q)$ of plan p under consumption of q : $C_p(q) = f_p + \max(q - c_p) \cdot e_p$

Cost for user u in one charging period m for consumption q_{um}

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

over T charging periods

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

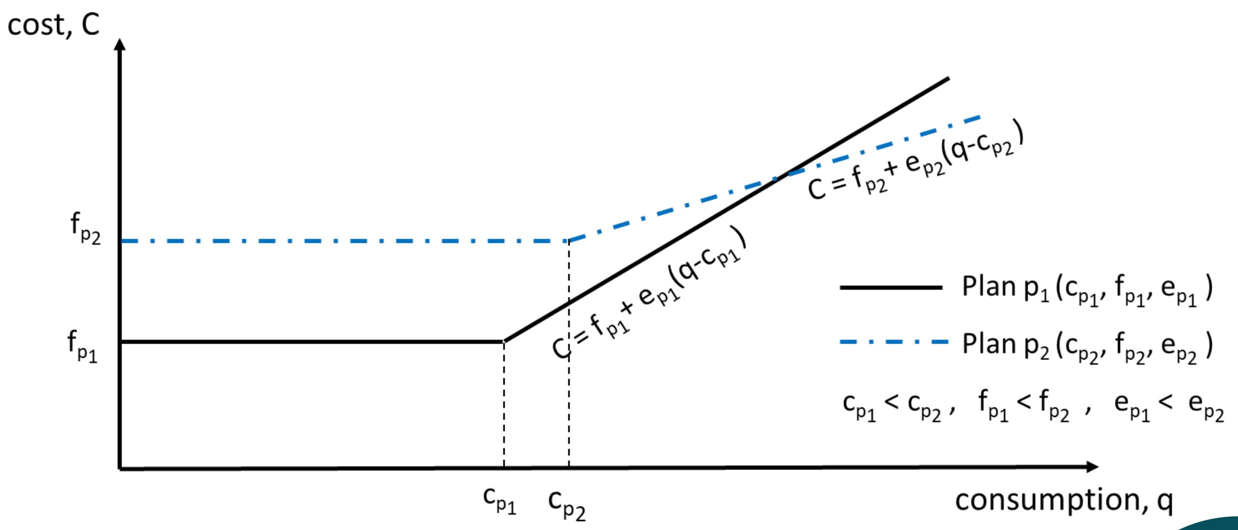

Shared capped data plans and cost sharing rule

Shared plans are described by the same parameters (c_p, f_p, e_p) with the individual plans:

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

o_p : extra fixed fee per user/device joining the shared data plan.

How is the subscription cost shared across subscribers? - four axiomatic requirements for the *cost-sharing rule* :

- (R1) Cost shares should sum exactly to the data plan cost
- (R2) Unanimity : the sharing should not discriminate against any user
- (R3) Cost shares should be non-decreasing functions of actual user consumption
- (R4) Any overage fees should be "fairly" undertaken by the users who are responsible for the excess consumption

Known cost-sharing rules fail in one or more of (R1)-(R4), hence we introduce a custom group sharing rule, called **Double Proportional Cost Sharing (DPCS)**

Double proportional cost sharing (DPCS)

- Rough idea:
  - as far as the overall group consumption lies below the cap, split the cost proportionally to the demand profiles
  - if the overall group consumption *q* exceeds the cap, split the overage fees in proportion to the extent to which each user's consumption exceeds its profile demand
- The computation of cost shares is repeated at the end of each charging period *m*

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**

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

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

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

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

return $\{y_{um}\}$

Joint subscriber partitioning and plan assignment problem

- **Objective :** Partition *U* users in a number of subscription groups and assign them with one out of *P* shared data plans so that their relative savings over their best possible individual data plans are maximized

If p_i is the regular data plan that minimizes the subscription cost for user *u* and ξ_g its cost share under the shared plan that minimizes her monetary contribution, the relative subscription savings for user *u* are

$$s_{u,n} = \frac{s_u}{\sum_{m \in T} C_{p_i}(q_{um})} \qquad s_u = \sum_{m \in T} C_{p_i}(q_{um}) - \xi_{p_g}(q_{um}, \mathbf{q}_{-um})$$

- **Constraints**
  - finite group size, g_{max}
  - assignment constraint (each user to one subscription group)
  - savings should be non-negative for all users

$$\begin{aligned} \max_{\mathbf{x},\mathbf{G}} \quad & \sum_{u \in \mathcal{U}} s_{u,n} \qquad (OPT) \\ s.t. \quad & 1 \leq |g_i| \leq g_{max}, \quad g_i \in G \\ & g_i \cap g_j = \emptyset \quad \forall g_i, g_j \in G \\ & s_{u,n} \geq 0 \quad u \in \mathcal{U} \end{aligned}$$

- **Complexity:** could be viewed as an instance of the Generalized Assignment Problem (items to bins) with costs that are interdependent across items
  - we know for these problems that they are hard even to approximate
  - we resort to heuristic approaches

G **:** set of possible user partitions

x(G), G ∈ **G**: |G|x|*P*| size matrix, $x_{gp} = 1$ when plan *p* is assigned to group *g* ∈ G

Solving (OPT)

We introduce 3 heuristic algorithms to approximate (OPT), all based on clustering techniques

- **Agglomerative cost-minimization clustering (*ACMC*)**
  - simultaneously computes groups and optimal shared plans
  - agglomerative clustering process: at each step merge the two clusters that result in maximum relative savings of the subscription cost
  - stopping criterion : no possible improvement or group limit reached for all clusters
  - O(U^3*P) complexity

- **Agglomerative uniform-consumption clustering (*AUCC*)**
  - decomposes the problem: first derive groups, then assign them the optimal shared subscription plan
  - agglomerative clustering process: at each step merge the two clusters that minimize the variation (over time) of the demand profile
  - stopping criterion : no possible improvement or group limit reached for all clusters
  - assign the optimal plan to each subscription group
  - O(U^3) complexity

Solving (OPT)

We introduce 3 heuristic algorithms to approximate OPT, all based on clustering techniques

- **Double greedy maximal uniform-consumption clustering (*DGMC*)**
  - iterative clustering process
  - generate maximal clusters of users so that the variation of intra-cluster demand is minimized
  - greedily identify disjoint maximal clusters that exhibit the minimum variation of demand and promote them to subscription groups
  - repeat the process with the remaining users till no further clustering is possible
  - assign the optimal plan to each subscription group
  - Complexity: $O(U^{1+g_{max}}) + O(U \cdot P)$

Evaluation of the algorithms

- **Input data**
  - user demand profiles for all users*
  - sets of available data plans
  - (computed) optimal (min cost) individual data plans
  - actual data consumption for all users

- **Output**
  - subscription group memberships
  - min cost shared data subscription plan
  - resulting absolute and relative subscription cost savings

- **Two sets of experiments**
  - demand profiles accurately predict the actual consumption
    - *perfect prediction*
  - actual consumption deviates from profile demand
    - normal distribution $N(\alpha \cdot d_{um}, \sigma)$, where d_{um} data consumption in charging period m according to profile demand

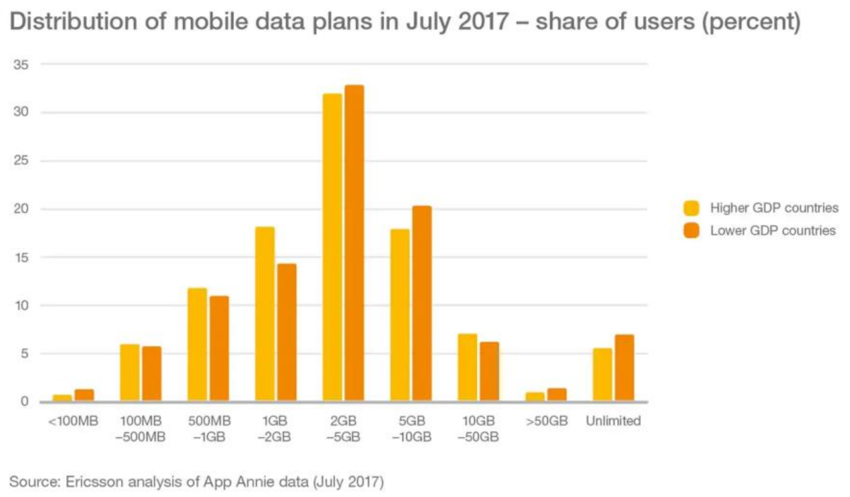

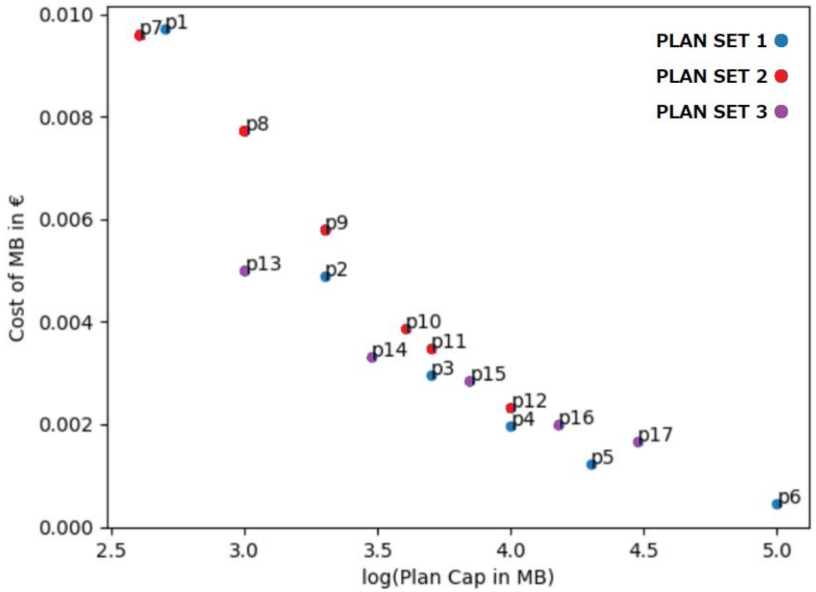

*Ericsson mobility report. (2017) Shifting mobile data consumption and data plans. [Online]. Available: https://www.ericsson.com/en/mobility-report/articles/shifting-mobile-data-consumption-data-plans

Performance measures

- **Per user relative subscription cost savings**

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

- **Percentile of users with normalized savings > α**

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

Additionally, we compute and plot

- **Per user absolute subscription cost savings,** *sav(u)*

- **Group size distribution**

Results : perfect prediction (1/2)

Total annual cost of individual plans for 1400 users under no sharing : 258.616,8 €

- **Total subscription costs with shared data plans**
  - ACMC: 101.114,70 € (▼ ~61%)
  - AUCC: 115.645,51 € (▼ ~56%)
  - DGMC: 121.555,55 € (▼ ~53%)

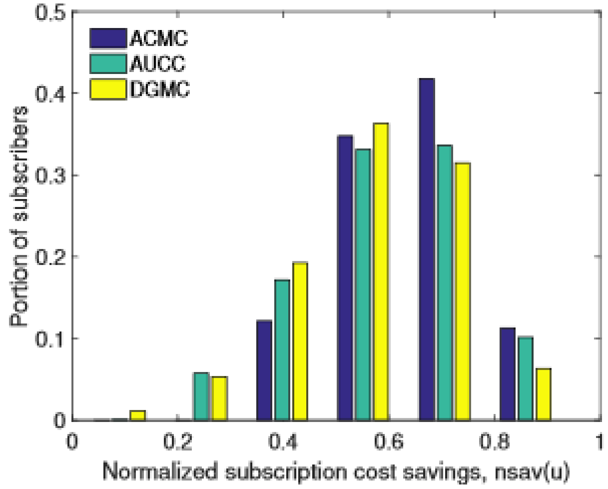

- **Per user savings**
  - Half of the users see above 57% relative savings in DGMC, 60.5% in AUCC and 64% in ACMC
    - 10% of users under ACMC retain relative savings above 78.4%. With DGMC the relative savings are above 74.1% and for AUCC 78.2%
    - ACMC guarantees relative savings of at least 37% for more than 99% of the users

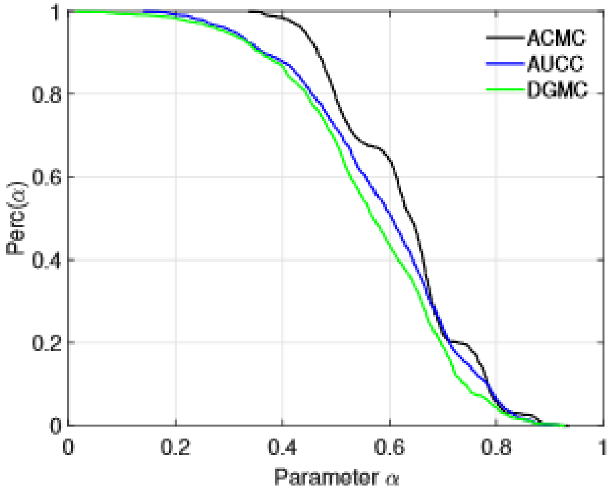

Results : perfect prediction (2/2)

- **Absolute subscription cost savings**
  - The distributions achieved by all three algorithms are more uniform despite the differentiation at higher amounts of cost savings
    - Half of the users see savings above 93€ with DGMC and AUCC, above 110 € with ACMC
    - 10% of users saw savings above 150 € with DGMC, 171€ with AUCC and 160€ with ACMC
  - ACMC distributes savings more proportionally to the reference cost of individual subscriptions

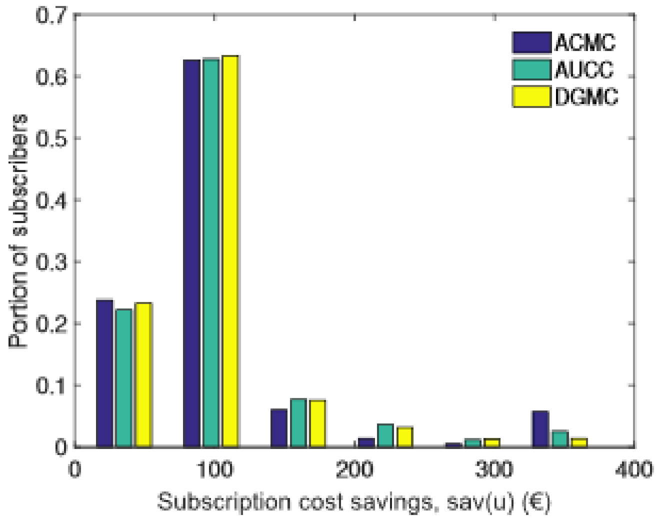

GROUP SIZE DISTRIBUTION

	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%

- **Group size distribution**
  - the two algorithms that decompose the problem are clearly biased towards the highest group size (g_{max} = 5 in all experiments)
  - ACMC demonstrates a more balanced distribution of group sizes with more than half of the generated subscription groups having fewer than g_{max} members

Results : imperfect prediction

$\alpha = 1.1, \sigma = 0.05d_{um}$

- **Total subscription costs with shared data plans after sensitivity**
  - ACMC: 139,058.16 € (▼ ~46%)
  - AUCC: 150,474.32 € (▼ ~42%)
  - DGMC: 162,657.73 € (▼~ 37%)

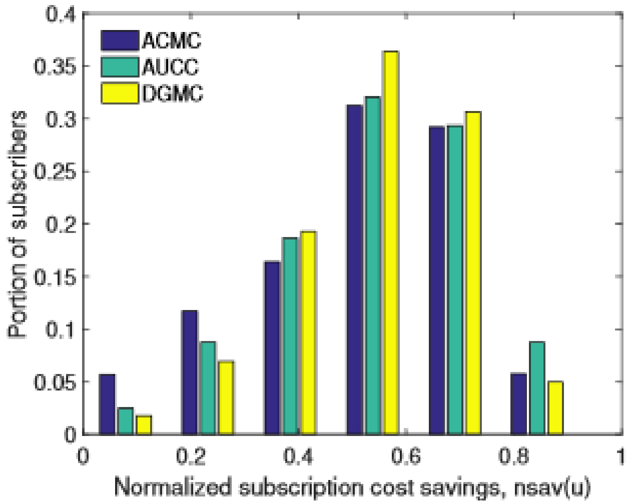

- **Per user savings**
  - Half of the users see relative savings above 54.1% in DGMC, 55.8% in AUCC and 55.3% in ACMC
  - 10% of users retain relative savings above 75.2% with AUCC. With DGMC the relative savings are above 72% and for ACMC 74.3%

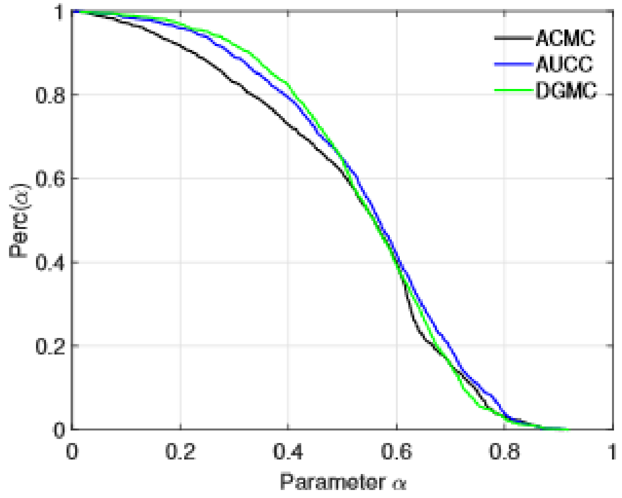

Results : imperfect prediction

$\alpha = 1.1$, $\sigma = 0.1d_{um}$

- **Absolute subscription cost savings**
  - Half of the users retain savings above 85€ across all algorithms
  - Algorithms retain savings above 122€ for 10% of the users, with ACMC above 134€ and AUCC above 147€.

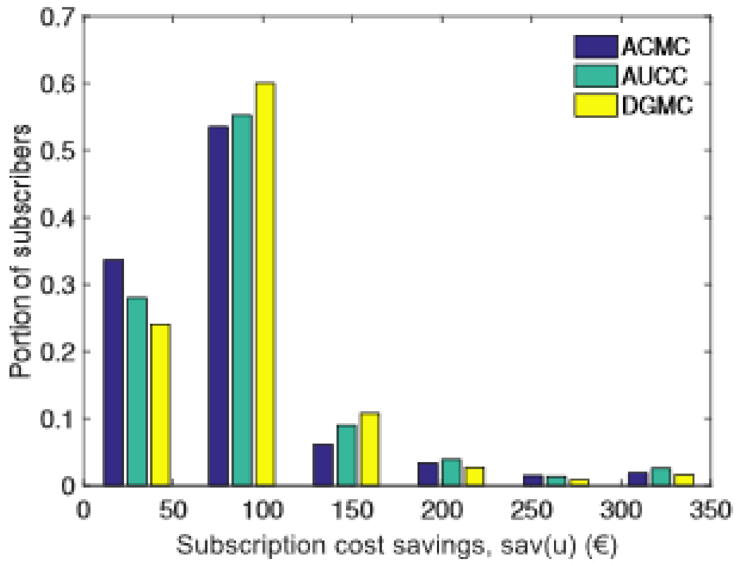

DIFFERENCE IN AGGREGATE SUBSCRIPTION CHARGES PAID BY CELLULAR USERS UNDER PROFILE DEMANDS AND ACTUAL CONSUMPTION

	$\sum$ 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€

- **Savings difference**
  - Highest savings are retained by ACMC
  - …but as the prediction error grows another 10% on average, AUCC and DGMC outperform ACMC

	$\sum$ savings (D_u)	$\sum$ savings ($\alpha = 1.2 \sigma = 0.1$)	Difference
DGMC	137061.27€	33510.13€	-103550.87€
AUCC	142971.3€	52405.46€	-90565.84€
ACMC	157502.1€	28602.98€	-128899.12€

Conclusions

- We have revisited shared cellular data plans
  - around for almost a decade now, still restricted to “closed user groups” (e.g., family plans) and device sharing
  - viewpoint : online recommendation platform issuing recommendations to end users in “open user groups” setting

- We have looked more closely into algorithmic aspects that have not attracted much attention in the past
  - How the plan cost could be shared between the participants subscription group → proposed a custom cost sharing rule that satisfies intuitive requirements for fair and efficient sharing
  - How could users be partitioned into cost-optimal subscription groups → proposed three heuristic clustering-based algorithms to this end

- The potential subscription savings are too important to ignore
  - With all three algorithms and perfect user demand profiling, the subscription costs drop to one half on average
    - Specifically, above 57% even with the ‘worst’ performance algorithm.
  - The savings are lower but persist when the actual consumption deviates “modestly” from the predicted one of the demand profile

Future work

Modeling challenges

- user demand elasticity : how users adapt their consumption patterns over time in response to content, residual cap and idiosyncrasy
  - under different underlying assumptions (e.g., *full* vs. *bounded* rationality)

- strategic behaviors may emerge on the side of users
  - at the demand profiling phase (if demand profiles are built on the basis of users' statements); or
  - during the actual use of the plan (if the aggregate and individual consumption can be monitored)

- the MNO's viewpoint and the viability of shared data plans in the open market

Algorithmic challenges

- problem characterization and approximability properties

Optimizing shared data plans for mobile data access

Georgios Cheirmpos, Merkouris Karaliopoulos, Iordanis Koutsopoulos

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

ITC 33, Avignon, France, 1/9/2021

H.F.R.I. Hellenic Foundation for Research & Innovation

GSRI GENERAL SECRETARIAT FOR RESEARCH AND INNOVATION