

Network Neutrality in Content Cache Sharing: A Bankruptcy Problem Formulation

Angelos Antonopoulos
NearbyComputing S.L.

Barcelona, Spain
aantonopoulos@nearbycomputing.com

Prodromos-Vasileios Mekikis
Aristotle University of Thessaloniki

Thessaloniki, Greece
vmekikis@auth.gr

Elli Kartsakli

Barcelona Supercomputing Center (BSC)
Barcelona, Spain
elli.kartsakli@bsc.com

Abstract—The deployment of fifth generation (5G) mobile networks along with the massive penetration of multimedia applications render the edge cache capacity a valuable but limited network resource. Edge caching has been lately widely adopted, as it can bring to the network important benefits, such as better network utilization and higher Quality of Experience (QoE). However the cache sharing among different Content Service Providers (CSPs) is a non-trivial problem due to: i) the limited nature of cache resources, and ii) the network neutrality concept that requires equal treatment of the CSPs. In this paper, we study the cache sharing problem in the presence of multiple CSPs, focusing both on the fairness and the network performance. Taking into account the importance of network neutrality and the scarcity of the cache resources, we formulate the cache sharing as a bankruptcy problem and we propose a set of different approaches for its solution. The performance of the proposed approaches is assessed through a series of simulation experiments in different scenarios (i.e., assuming different popularity distributions and market shares), identifying the potential trade-offs.

Index Terms—5G, content caching, network sharing, bankruptcy problem, CDN

I. INTRODUCTION

The introduction and the early deployment of fifth generation (5G) mobile networks brings advanced capabilities and introduces new trends in the telecommunications world. This has been driving the transformation of several vertical sectors, such as the multimedia and entertainment industries, where novel applications and services emerge, with demanding requirements in terms of performance and bandwidth that, in turn, pose new challenges to the network. Recent studies have indicated that, globally, video traffic will constitute 82% of all Internet traffic by 2022, i.e., a four-fold growth from 2017 to 2022 that corresponds to a Compound Annual Growth Rate (CAGR) of 33% [1].

To face this unprecedented increase in network load, the Internet Service Providers (ISPs) have been increasingly adopting edge caching. By placing the most popular content close to the end users, clear benefits can be obtained for all the interested stakeholders, i.e., ISPs, end users and Content Service Providers (CSPs), as the users receive a higher Quality of Experience (QoE), while the backhaul network traffic is significantly reduced [2], [3].

Despite the aforementioned advantages, caching comes with several challenges. In particular, the increase of network traffic

is disproportionate with the availability of caching resources, which is quite limited. In addition, the number of CSPs is constantly increasing, thus creating a more complex and dynamic ecosystem with several variables to consider. These challenges have motivated the research community to study the cache sharing problem in mobile networks. Several works applied game and auction theory tools to formulate the problem of assigning parts of the cache capacity to different CSPs [4]–[8]. One common characteristic of these works is that they focus on network performance and mainly on the profit and the incentives of the different stakeholders, without taking into account the network neutrality principle [9].

According to the network neutrality, ISPs should treat all Internet traffic equally, without any discrimination or different charging [10]. Nevertheless, with edge caching some contents are given priority over the rest, thus resulting to potentially preferential treatment towards specific CSPs (especially the major ones). In this context, there have been some recent studies that have proposed cache sharing solutions [11]–[13]. In [11], different sharing methods have been considered, however the claims are defined by the CSPs, something that could affect the neutrality of the solution. In [12], the authors discuss the network neutral caching, however they suggest that the caching should be aligned with the popularity, something that can be considered as prioritization and, therefore, discrimination. Finally, the idea of formulating the cache sharing as a bankruptcy problem [14] has recently appeared in [13], where the authors proposed a solution based on cooperative game theory. Their approach has stressed the suitability of the bankruptcy problem in cache sharing scenarios, while leaving open the possibility of applying different methods for its solution.

In this context, considering the limited nature of caching resources along with the importance of ensuring network neutrality, our contribution can be summarized in the following points:

- 1) Taking into account the limited cache resources, we formulate the content cache sharing as a bankruptcy problem, where the claims are determined by the ISP (following the CSPs' data volumes) to respect the network neutrality.
- 2) Based on the existing solutions in the literature for the bankruptcy problem, we propose three different cache

sharing schemes, i.e., Equal Sharing (ES), Proportional Sharing (PS) and Shapley Value Sharing (SVS).

- 3) We perform extensive simulation experiments to evaluate the network performance and the fairness of the proposed solutions, while we compare them against a widely adopted caching policy, where the most popular contents are cached. In addition, we study the impact of various parameters (e.g., content virality, number of CSP files, etc.) on the performance of the proposed schemes.

The remainder of the paper is organized as follows. Section II discusses the related work in cache sharing. Section III presents the system model of our work, while Section IV introduces the bankruptcy problem formulation along with the proposed strategies. In Section V, we provide the performance evaluation of our approach. Finally, Section VI concludes this paper.

II. RELATED WORK

In this section, we present a literature review on cache sharing focusing on i) works that neglect network neutrality and ii) works that take into account network neutrality.

A. Cache Sharing without Network Neutrality

In [4], the authors first study a single CSP optimization problem, where, under a given spatial distribution of the base stations, the CSP decides how to share the cache among the different classes of contents to minimize the cache miss rate. In addition, they formulate the competition among CSPs as a Stackelberg game, proving the existence of a Nash Equilibrium.

In [5], Xiong *et al.* introduce the concept of sponsored content (i.e., content that is paid by the CSPs and does not count in the data consumption of the mobile users) in edge caching scenarios. The authors formulate a novel hierarchical three-stage Stackelberg game to i) model the interactions among the ISPs, the CSPs and the mobile users, and ii) to jointly maximize their individual payoffs.

In [6], Fang *et al.* theoretically analyze the collaboration framework between ISPs and CSPs in content delivery services, taking into account the impact of content popularity and edge caching. The proposed formulation aims at maximizing the network profit considering both offline and online caching strategies.

In [7], Mitra and Sridhar study the coalition formation among ISPs and CSPs in content caching scenarios. In their attempt to evaluate the cache and cost sharing, they consider fixed costs (e.g., infrastructure, equipment, building, power supply, etc.) and marginal costs (e.g., number of subscribers, cache and enhanced network capacity), while they assume a non-cooperative model, where the CSPs compete for profit.

In [8], Dehghan *et al.* study the tradeoff between cache partitioning and cache sharing in two different scenarios where: i) CSPs serve disjoint files and ii) some content is served by multiple CSPs. In particular, they formulate cache partitioning as an optimization problem with constraints on the

cache storage and develop decentralized algorithms to partition the cache in an online manner.

B. Network Neutral Cache Sharing

In one of the seminal works that focus on the fair cache sharing among CSPs [11], Hoteit *et al.* propose different allocation strategies to maximize the fairness in the network. However, as pricing is involved, their work is not fully compatible with the concept of network neutrality, while the authors also assume a specific fixed content popularity.

More recently, Andreoletti *et al.* [12] have proposed different strategies for network neutral content caching. However, their proposed solutions are either quite simplistic (i.e., sharing the cache in equal parts) or not fully compliant with the neutrality notion (i.e., sharing the cache according to the content popularity). More specifically, regarding their second approach, even though the particular CSPs are not considered during the cache sharing, such policy may benefit the most popular CSPs, thus violating the network neutrality rules.

Lately, Chai *et al.* [13] have considered the limited nature of the cache resources in heterogeneous networks, formulating a bankruptcy problem for the cache sharing. Despite the novel insights that this formulation provides, the authors consider only one solution (i.e., Shapley Value) for this problem, neglecting the potential trade-offs among the possible solutions.

III. SYSTEM MODEL

We consider a scenario where an ISP offers caching services to a set of N CSPs (Fig. 1). We focus on a single cache server of limited storage capacity S (expressed in bytes), which must be shared among the CSPs. Each CSP has a content catalogue defined as $F_i = \{f_{i,1}, f_{i,2}, \dots, f_{i,K_i}\}$, where $f_{i,j}$ with $j \in [1, K_i]$ are the K_i content files of the i th CSP. The content files of each CSP have a size $s_{i,j}$ (in bytes) that follows a normal distribution with average \bar{s}_i and standard deviation σ_i .

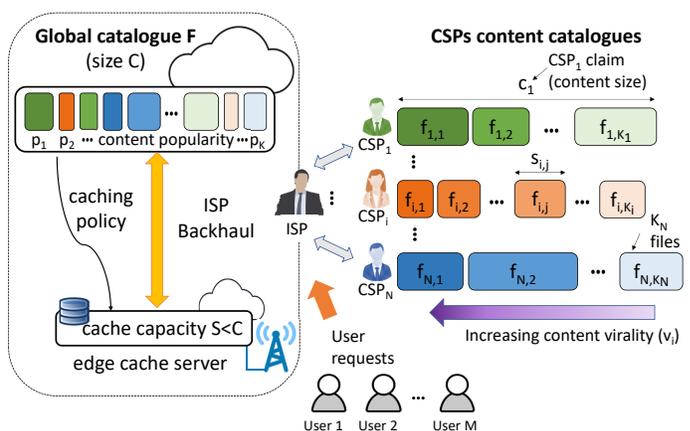


Fig. 1: System Model

From the user perspective, we consider a pool of end users M that request content from the complete content catalogue $F = \bigcup_{i=1}^N F_i$ containing a total of $K = \sum_{i=1}^N K_i$ content files. However, as in real scenarios, not all content types

have the same appeal to the end users. For instance, content from a given provider (e.g., YouTube videos) may generally be more viral (i.e., attracting higher user downloads) than other types of contents (e.g., Spotify songs). To capture this diversity in content attractiveness, we introduce the concept of *virality*, which is employed to rank the contents of the different CSPs in the global content catalogue F . We define virality as an exponential distribution with a mean value v_i which characterizes the attractiveness of the content of the i th CSP. A higher value v_i corresponds to CSPs with more viral content on average (e.g., YouTube videos in the previous example).

In order to generate the global content catalogue, the individual CSP catalogues are first sorted by descending order of their content virality. Then, the catalogues are merged into the global content catalogue F by order of the virality metric. As a result, highly viral content will be more likely to occupy the first positions of the global catalogue. Once the global catalogue is formed, we calculate the global popularity of the content which follows a Zipf distribution as

$$p_k = \frac{1}{k^\rho} / \left(\sum_{j=1}^{j=K} \frac{1}{j^\rho} \right), \quad (1)$$

where p_k is the probability with which the content occupying the k th position in the global catalogue is requested by the end users. The parameter $\rho > 0$ is the skew parameter of the Zipf distribution, with lower ρ values corresponding to a more uniform popularity distribution.

IV. PROBLEM FORMULATION AND CACHE SHARING STRATEGIES

In this section we formulate the cache sharing problem and we provide a set of different strategies for its solution.

A. Cache Sharing: Problem Formulation

As explained in Sec. III, we consider the sharing of a cache server of size S among N different content providers. To be inline with network neutrality, we assume that the ISP considers as CSPs' claims their whole individual catalogues F_i , hence the cache claim of the i th CSP is $c_i = \sum_{j=1}^{K_i} s_{ij}$. In addition, we focus on realistic scenarios, where the total demand of storage resources C exceeds the storage capacity of the server, i.e., $C = \sum_{i=1}^N c_i > S$, and therefore, only a subset of the available content can be stored at the cache server. Hence, the problem can be formulated as a bankruptcy problem, with c_i representing the individual claims of each CSP.

Once the ISP defines all the claims for storage capacity, a policy must be defined in order to determine the percentage of storage resources to be allocated to each CSP. In the following subsection, we describe three different solutions for this problem.

B. Equal Sharing

Equal sharing is probably the simplest form of sharing a commodity among a set of interested parties. Equal sharing

does not take into account the individual claims of the involved parties and divides the commodity into equal shares that are allocated to each agent, respectively. In our problem, this can be translated to

$$\hat{c}_i^{ES} = \frac{S}{N}, \quad (2)$$

where \hat{c}_i^{ES} denotes the allocated cache portion to the i th CSP following the equal sharing approach.

One important advantage of this solution is the simplicity thanks to its low complexity, while it has to be noted that, following this approach, it is possible that one agent may receive more resources compared to their actual needs.

C. Proportional Sharing

Proportional sharing explicitly takes into account the individual claims, as the sharing takes place in a proportional to the claims way. Proportional division is a quite appealing approach and it is also adopted in several real life use cases (e.g., in case of natural disasters, insurance companies pay off the losses with a fixed amount per dollar). It can be also easily proven that the proportional sharing of the losses provides the same solution, so we avoid dissatisfaction due to unbalanced losses. According to the proportional sharing scheme, each CSP will be allocated a portion \hat{c}_i of the storage, calculated as

$$\hat{c}_i^{PS} = b_i \cdot S, \quad (3)$$

where $b_i = c_i/C$ denotes the weight of i th CSP's claim with respect to the total claimed amount C .

D. Shapley Value

The Shapley value offers an interesting, but also more complex, alternative method for sharing the caching resources. The concept of the Shapley value is derived by the coalitional game theory and is used to calculate the average expected marginal contribution of each player, after taking into account all possible combinations of arrival. Hence, the allocation at each player i can be expressed as

$$\hat{c}_i^{SV} = \frac{1}{|N|!} \sum_{L \subseteq N \setminus \{i\}} |L|!(|N| - |L| - 1)!(v(L \cup \{i\}) - v(L)), \quad (4)$$

where N is the total number of players and the sum extends over all subsets L of N not containing player i . In addition, the characteristic function $v(L)$ indicates the value of coalition L . It is also worth noting that, due to the nature of the bankruptcy problem, the value of each coalition is limited by the cache size S , as the claim of one or more players cannot exceed S .

V. PERFORMANCE EVALUATION

We have conducted extensive simulation experiments in a custom-based Matlab simulator to evaluate our proposed cache sharing methods, i.e., Equal Sharing (ES), Proportional Sharing (PS) and Shapley Value Sharing (SVS). In addition, as a baseline method we employ the widely adopted Most

Popular Content (MPC) caching method, where the contents are stored according to their popularity without taking into account the content provider and the respective claims.

Regarding the metrics of interest, we focus on three fundamental metrics on caching scenarios, i.e., hit rate, backhaul (BH) traffic and fairness. In particular, the hit rate corresponds to the ratio of requests that are satisfied through the cache server over the total number of requests. The BH traffic indicates how much traffic is transferred through the BH from file requests that do not appear in the cache. With respect to fairness, we adopt the Jain's fairness index defined as

$$J = \frac{\left[\sum_{i=1}^N \frac{\hat{c}_i}{c_i} \right]^2}{\left[N \sum_{i=1}^N \left(\frac{\hat{c}_i}{c_i} \right)^2 \right]}, \quad (5)$$

with c_i being the claim of the i th CSP and \hat{c}_i the respective allocated cache storage.

A. Simulation Scenario

We consider the presence of three $N = 3$ CSPs that produce a global catalogue of $K = 1000$ content files. The occupied positions by the i th CSP in the global catalogue are defined by two metrics: i) the total number of files K_i for CSP_i and ii) the CSP's virality v_i , which is the tendency of the CSP's content to receive a higher number of user requests. A higher mean virality value v_i means that the files of CSP_i are more likely to be placed at the top of the global catalogue. We also assume $M = 100$ end users with an average of 100 content requests per user, i.e., generating a total of 10000 requests that follow the content popularity distribution. The key simulation parameters are summarized in Table I.

For the performance evaluation, we consider two different scenarios, depending on the type of CSPs. In the **first scenario**, we consider three similar CSPs, i.e., of similar service with files that follow the same distribution, having an average size $\bar{s}_i = 500$ MB. Two different cases are examined with respect to the CSP's catalogue sizes and virality: i) all CSPs have the same virality (i.e., $v_1 = v_2 = v_3 = 0.1$) with different number of files in their catalogue (i.e., $K_1 = 800, K_2 = K_3 = 100$), and ii) all CSPs have the same catalogue size (i.e., $K_1 = K_2 = K_3 = 333$) but different virality (i.e., $v_1 = 0.7, v_2 = v_3 = 0.1$).

In the **second scenario**, we consider three CSPs with different content types, namely short YouTube-like (YT) video content (with average size $\bar{s}_1 = 200$ MB), larger files such as Netflix (NF) movies (with $\bar{s}_2 = 2000$ MB), and music files such as Spotify (SF) (with $\bar{s}_3 = 20$ MB). In this scenario, we consider that one CSP has a much higher virality ($v_i = 0.7$) with respect to the other two that have a virality of 0.1, considering all three different combinations to showcase the intrinsic behavior of each caching strategy. Furthermore, we consider that the most viral operator in each case also has the highest number of content files (i.e., $K_i = 800$), while the other two CSPs participate in the global catalogue with 100 files each. The two scenarios are also presented in Table I.

TABLE I:
SIMULATION PARAMETERS

| Common parameters | | |
|----------------------------------|-------------|--------------|
| End Users | M | 100 |
| Number of CSPs | N | 3 |
| Zipf skew parameter | ρ | [0.1,0.7] |
| Global catalogue files | K | 1000 |
| Scenario 1 - Same Type CSPs | | |
| Cache size | S | 100 GB |
| File size for all CSPs | \bar{s} | 500 MB |
| Virality | v_i | {0.1, 0.7} |
| Number of files per CSP | K_i | {100, 800} |
| Scenario 2 - Different Type CSPs | | |
| Cache size | S | 50 GB |
| CSPs | i | {YT, NF, SP} |
| Average file size for YT | \bar{s}_1 | 200 MB |
| Average file size for NF | \bar{s}_2 | 2000 MB |
| Average file size for SP | \bar{s}_3 | 20 MB |
| File size standard deviation | σ_i | 12.5% |
| Virality | v_i | {0.1, 0.7} |
| Number of files per CSP | K_i | {100, 800} |

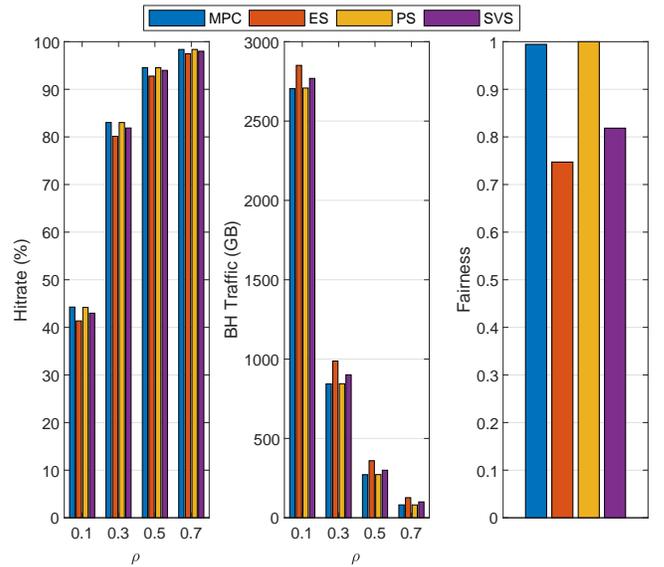


Fig. 2: Performance metrics in scenario 1a: same filesize ($\bar{s} = 500$ MB), same virality ($v_1 = v_2 = v_3 = 0.1$), and different number of files per CSP ($K_1 = 800, K_2 = 100, K_3 = 100$).

B. Simulation Results

In the first performance evaluation scenario, we start by considering the case of $N = 3$ CSPs with the same content type and the same virality ($v_1 = v_2 = v_3 = 0.1$). However, the first CSP is considered to have a much larger content catalogue than the other two providers ($K_1 = 800, K_2 = 100, K_3 = 100$). As shown in Fig. 2, the MPC and PS achieve the best (and almost identical) performance, with MPC offering slightly higher hit rate and almost optimal fairness (slightly lower than the PS which yields a Jain index equal to 1). This can be explained by the fact that the three CSPs have the same virality, meaning that their contents have the same

appeal to the end users. However, since CSP_1 has a much longer catalogue, it will receive a higher number of requests. Hence the MPC will eventually allocate each CSP storage resources that are proportional to their catalogue, obtaining a similar performance as the PS. Regarding to the variation of the network performance with respect to ρ , we observe that, as ρ increases, the network performance improves, since the requests for the most popular contents (that are placed in the cache) increase.

In Fig. 3, we study a similar scenario with the difference that the three CSPs have the same number of files in their catalogue (i.e., $K_i = 333$) but different virality ($v_1 = 0.7$, while $v_2 = v_3 = 0.1$). As we may see, all bankruptcy-based solutions (i.e., ES, PS and SVS) achieve optimal fairness ($J = 1$), while MPC heavily prioritizes the content of the most viral operator, thus violating the network neutrality rules and achieving a very low fairness value ($J = 0.35$). Even though MPC achieves a better performance in terms of hit rate and reduced BH traffic, the difference with ES, PS and SVS is quite small, raising the question of whether it is worthwhile to make such content prioritization. It is also worth noting that the three bankruptcy-based schemes have identical performance, as the sharing takes place according to the claims, which correspond to the content catalogue size that is the same in this scenario.

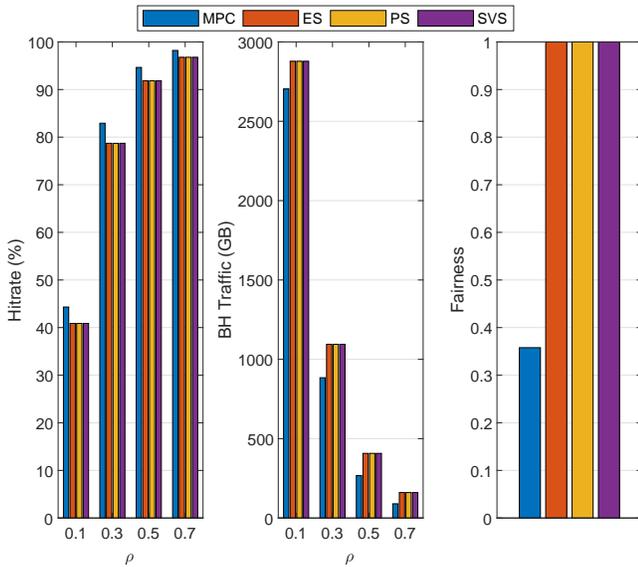


Fig. 3: Performance metrics in scenario 1b: same filesize ($\bar{s} = 500$ MB), different virality ($v_1 = 0.7, v_2 = 0.1, v_3 = 0.1$), and same number of files per CPS ($K_1 = K_2 = K_3 = 333$).

The remaining three figures refer to the second scenario where three CSPs with different content type (and therefore filesize) are considered (see Table I). In each plot, one CSP has significantly higher virality and catalogue size than the other two. Specifically, Fig. 4 presents the metrics of interest in a scenario where YT is the prevalent CSP (with $v_1 = 0.7$ and $K_1 = 800$ files in the catalogue). Under this setup, we can see that, in terms of hitrate and BH traffic, the MPC outperforms the other schemes as expected, since the most popular YT

content is prioritized. However, this gain comes with a cost on the fairness performance which is very low (i.e., Jain index $J < 0.4$), since the other two CSPs are overlooked. The ES is worse in all metrics, since sharing equally the cache storage is neither fair nor efficient when the difference in the CSPs virality and content size is so significant. On the other hand, PS and SVS manage to treat the CSPs in a fair manner (i.e., $J > 0.8$), thus being compliant with the network neutrality rules, while achieving a relatively high hit rate. In addition, we may see that, as ρ increases, the most popular contents attract much higher interest and requests, thus increasing the hit rate and reducing the BH traffic in the network.

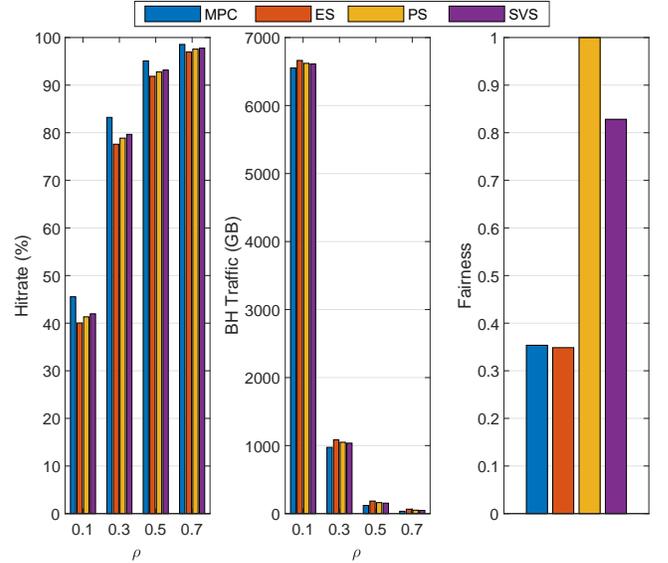


Fig. 4: Performance metrics in scenario 2a: YT has the highest virality ($v_1 = 0.7$) and catalogue size ($K_1 = 800$), while $v_2 = v_3 = 0.1$ and $K_2 = K_3 = 100$.

In the scenario depicted in Fig. 5, NF is the prevalent CSP. Under this setup, we may see that PS achieves again the ideal fairness, while the difference between MPC and SVS is reduced. This can be explained taking into account the filesize of the prevalent operator, which is quite big, resulting in a prioritization of its files by schemes that do not follow a strict proportional sharing. Regarding the hit rate, MPC, PS and SVS achieve similar performance, while ES has a notable performance drop, as it provides the other two CSPs (i.e., YT and SF) with space that remains unused.

Finally, Fig. 6 illustrates the scenario where SF is the prevalent CSP (i.e., highest virality and highest number of files). Compared to Figs. 4 and 5, it is worth noting the MPC achieves better fairness than before (i.e., over 0.6), while still remaining below the performance of PS and SVS. This improvement can be explained by the small size of the files of the prevalent operator, as several files can be stored in the cache while still leaving space for the other CSPs. Another interesting observation is the reduction in the BH traffic for small values of ρ , which decreases even further as ρ increases

(i.e., BH traffic almost negligible for $\rho = 0.5$ and $\rho = 0.7$), since the music files have the most requests, generating small BH traffic in the network.

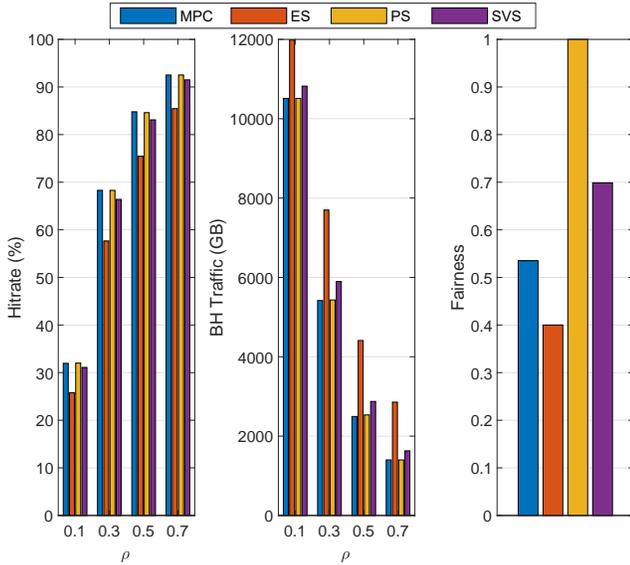


Fig. 5: Performance metrics in scenario 2b: NF has the highest virality ($v_2 = 0.7$) and catalogue size ($K_2 = 800$), while $v_1 = v_3 = 0.1$ and $K_1 = K_3 = 100$.

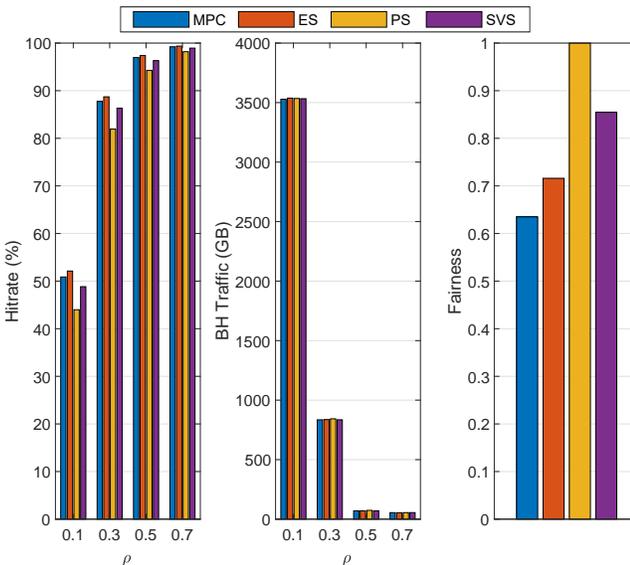


Fig. 6: Performance metrics in scenario 2c: SF has the highest virality ($v_3 = 0.7$) and catalogue size ($K_3 = 800$), while $v_1 = v_2 = 0.1$ and $K_1 = K_2 = 100$.

VI. CONCLUSIONS

In this paper, we dealt with the issue of cache sharing in multi-CSP scenarios. In particular, taking into account the limited cache resources and the importance of network neutrality, we formulated the cache sharing as a bankruptcy

problem and we proposed three different strategies (i.e., Equal Sharing, Proportional Sharing and Shapley Value Sharing) for its solution. We conducted extensive simulation experiments to evaluate the performance of the proposed strategies in terms of hit rate, backhaul traffic and fairness and we compared them with the widely adopted Most Popular Content caching policy. Our results revealed interesting insights, highlighting the unfair treatment of the content in various scenarios by the MPC policy, showing at the same time that bankruptcy-based solutions may achieve close-to-optimal fairness without compromising the network performance.

ACKNOWLEDGMENTS

This work has been partially supported by the EC-funded (Horizon 2020) research project Affordable5G under Grant Agreement n° 957317.

REFERENCES

- [1] Cisco, “Cisco visual networking index: Forecast and trends, 2017–2022 white paper,” Feb 2019.
- [2] X. Wang, M. Chen, T. Taleb, A. Ksentini, and V. C. M. Leung, “Cache in the air: Exploiting content caching and delivery techniques for 5G systems,” *IEEE Communications Magazine*, vol. 52, no. 2, pp. 131–139, 2014.
- [3] I. Parvez, A. Rahmati, I. Guvenc, A. I. Sarwat, and H. Dai, “A survey on low latency towards 5G: RAN, core network and caching solutions,” *IEEE Communications Surveys Tutorials*, vol. 20, no. 4, pp. 3098–3130, Fourthquarter 2018.
- [4] F. De Pellegrini, A. Massaro, L. Goratti, and R. El-Azouzi, “Competitive caching of contents in 5g edge cloud networks,” in *2017 15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, May 2017, pp. 1–8.
- [5] Z. Xiong, S. Feng, D. Niyato, P. Wang, A. Leshem, and Z. Han, “Joint sponsored and edge caching content service market: A game-theoretic approach,” *IEEE Transactions on Wireless Communications*, vol. 18, no. 2, pp. 1166–1181, Feb 2019.
- [6] C. Fang, H. Yao, Z. Wang, P. Si, Y. Chen, X. Wang, and F. R. Yu, “Edge cache-based ISP-CP collaboration scheme for content delivery services,” *IEEE Access*, vol. 7, pp. 5277–5284, 2019.
- [7] D. Mitra and A. Sridhar, “Consortiums of ISP-content providers formed by nash bargaining for Internet content delivery,” in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, April 2019, pp. 631–639.
- [8] M. Dehghan, W. Chu, P. Nain, D. Towsley, and Z. Zhang, “Sharing cache resources among content providers: A utility-based approach,” *IEEE/ACM Transactions on Networking*, vol. 27, no. 2, pp. 477–490, April 2019.
- [9] A. Antonopoulos, E. Kartsakli, C. Perillo, and C. Verikoukis, “Shedding light on the Internet: Stakeholders and network neutrality,” *IEEE Communications Magazine*, vol. 55, no. 7, pp. 216–223, July 2017.
- [10] J. Kraemer, L. Wiewiorra, and C. Weinhardt, “Net neutrality: A progress report,” *Telecommunications Policy*, vol. 37, no. 9, p. 794–813, October 2013.
- [11] S. Hoteit, M. El Chamie, D. Saucez, and S. Secci, “On fair network cache allocation to content providers,” *Computer Networks*, vol. 103, pp. 477–490, 2016.
- [12] D. Andreoletti, S. Giordano, C. Rottondi, M. Tornatore, and G. Verticale, “To be neutral or not neutral? The in-network caching dilemma,” *IEEE Internet Computing*, vol. 22, no. 6, pp. 18–26, Nov 2018.
- [13] R. Chai, Y. Li, and Q. Chen, “Joint cache partitioning, content placement, and user association for D2D-enabled heterogeneous cellular networks,” *IEEE Access*, vol. 7, pp. 56 642–56 655, 2019.
- [14] A. Antonopoulos, “Bankruptcy problem in network sharing: Fundamentals, applications and challenges,” *IEEE Wireless Communications*, vol. 27, no. 4, pp. 81–87, 2020.