

To be neutral or not neutral? the in-network caching dilemma

Original

To be neutral or not neutral? the in-network caching dilemma / Andreoletti, D.; Giordano, S.; Rottondi, C.; Tornatore, M.; Verticale, G.. - In: IEEE INTERNET COMPUTING. - ISSN 1089-7801. - ELETTRONICO. - 22:6(2018), pp. 18-26. [10.1109/MIC.2018.2877809]

Availability:

This version is available at: 11583/2741777 since: 2019-07-17T17:05:24Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/MIC.2018.2877809

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

TO BE NEUTRAL OR NOT NEUTRAL? THE IN-NETWORK CACHING DILEMMA

Davide Andreoletti

Università Professionale della Svizzera Italiana (SUPSI)

Silvia Giordano

Università Professionale della Svizzera Italiana (SUPSI)

Cristina Rottondi

Università Professionale della Svizzera Italiana (SUPSI)

Massimo Tornatore

Politecnico di Milano

Giacomo Verticale

Politecnico di Milano

Caching allows Internet Service Providers (ISPs) to reduce network traffic and Content Providers (CPs) to increase the offered QoS. However, when contents are encrypted, effective caching is possible only if ISPs and CPs cooperate. We suggest possible forms of non-discriminatory cooperation that make caching compliant with the principles of Net-Neutrality (NN).

According to many, the Internet can foster social and economic growth only if it provides open and non-discriminatory access to information. Such view does not really describe current Internet behavior, but can be regarded as an ideal goal, which Net-Neutrality (NN) regulations try to achieve. In practice, Internet evolution resulted in a complex service-delivery chain where Internet Service Providers (ISPs) own and manage the infrastructures that Content Providers (CPs) exploit to offer services to their end-users. As providers of new and increasingly advanced services, CPs play a key role in fulfilling the promise of an economically-valuable Internet. Hence, in a scenario of interaction between these two actors, the ISPs must upgrade their infrastructure to avoid becoming the bottleneck of the entire process.

In fact, the development of an efficient Internet infrastructure would result in a virtuous circle in which CPs can offer their services with high QoS guarantees and, in this way, induce an increasing number of users to subscribe with ISPs. However, this process is hindered by some practical issues: as vendors of valuable services, the CPs take the lion's share, while the ISPs face the risk of becoming simple providers of connectivity. This unbalance in revenue distribution risks to jeopardize the ISPs, as they are downgraded to commodity providers, and it is exacerbated by the purest NN vision, according to which ISPs should have limited or no control on the traffic traversing their network. However, this "neutral" delivery paradigm is not suitable anymore in a context where service requirements are becoming heterogeneous to such a degree that they necessarily require a different treatment of traffic. Starting with an example on current networks, it is undoubtable that users desire more guarantees for the performance of Video On Demand (VoD) than for e-mail exchange. Considering next generation networks (e.g., 5G), the current trend is to accommodate, over the same physical infrastructure, several virtual networks specifically tailored to run various bandwidth-hungry services. For example, the *network slicing*¹ paradigm embodies the principle that ISPs reserve to an external entity dedicated resources that fulfill given requirements (e.g., in terms of latency). This is a form of traffic differentiation and raises questions on the effective neutrality of ISPs. Given their social and economic role in deploying Internet infrastructures, should ISPs have the right to decide how to treat traffic in their network? If yes, to what extent can this be done neutrally and consistently with the current requirements of today's services?

The debate on NN is a long-standing one. Several frameworks have been proposed with the aim of reaching a common definition of neutrality in an era where traffic differentiation is mandatory. The proponents of pure NN see ISPs as mere "pipes" that should be agnostic to the contents they carry. On the other hand, NN opponents would give ISPs the greatest control as a



legitimate action to increase their revenue and consequently foster innovation. The reality of today's legislation is more nuanced and proposes different frameworks that aim at balancing the inherent trade-off between NN and QoS. The philosophy of such frameworks, for a thorough description of which we refer the reader to reference², can be summarized with the following statement: traffic differentiation should be allowed as long as it is not discriminatory for the CPs. Under such lens, for example, traffic can be categorized in classes (e.g., Video On Demand or VoIP) and differentiated accordingly, but contents belonging to the same class cannot be discriminated based on the owner CP. A complete review of currently-available strategies to detect illicit traffic discrimination procedures (i.e., not compliant with NN) can be found at reference³.

In this work, we discuss a topic that, perhaps surprisingly, has rarely been treated under the lens of NN, i.e., *in-network caching*. In-network caching is the process by which ISPs store in their network the most popular contents to reduce traffic coming from external systems, i.e., Content Providers (CPs). By using this strategy, contents are retrieved from closer servers and users experience a superior QoS. However, because of the limited storage of caches, in-network caching is an intrinsically-selective process and, as such, raises discriminatory concerns.

How can we perform caching to avoid that neither CPs nor users are discriminated? To answer this question, we first discuss some key characteristics of today's content delivery chain, namely the wide use of encryption and the raise of cooperative alliances among ISPs and CPs, whose compliance to NN principles should be carefully analyzed. Encryption is applied by CPs for security purposes, but it also prevents ISPs from effectively performing caching. In fact, due to the encryption, the ISP cannot identify the contents traversing its network and, consequently, it is not able to select the most popular ones. Cooperation is nowadays a common practice to allow CPs deploy their contents inside the domain of the ISP to better serve their users. We compare three scenarios of cooperative caching among ISPs and CPs with different degrees of NN-compliance. We then quantify the discrimination in terms of different performance experienced by the involved entities and we conclude that such process may violate the principles of NN. To overcome this problem, we finally advocate the definition of open protocols that enable a non-discriminatory cooperation towards the goal of performing a NN-compliant in-network caching.

A REVISED NET-NEUTRALITY DEBATE

The Internet is constantly evolving to adapt to the new services offered by CPs and to meet the increasing expectations of their users. Hence, the principles of NN should conform to such transformation. Today, the discussion on NN cannot neglect the following two aspects:

- *Encryption*: CPs are increasingly encrypting the traffic destined to their users to ensure privacy and security. It is estimated that the 80% of Internet traffic will be encrypted by 2019⁴, which will prevent ISPs from applying even basic forms of traffic differentiation that are allowed under a NN regime. For example, how can an ISP decide which contents to cache without being able to inspect them? In fact, requests for the same content are indistinguishable under encryption schemes, and this does not allow the ISP to infer the content popularity. A possible solution to this issue is to make CPs and ISPs jointly manage the caches, thus calling for cooperation schemes. The arbitrary selection of the counterparts of the cooperation (i.e., the CPs that the ISP decides to cooperate with) inevitably opens the doors for possible violations of the NN principles.
- *Cooperative Strategies*: nowadays ISPs commonly employ Virtualization to easily instantiate resources and offer specific services, thus widening the spectrum of the potential cooperation schemes that ISPs and CPs can implement. With ISPs that become more than simple providers of connectivity, however, new concerns are raised about their discriminatory behavior.

TOWARDS A NET-NEUTRALITY DEFINITION FOR IN-NETWORK CACHING

Without loss of generality, we consider caching to be performed on a specific class of contents, e.g., VoD. Given that ISPs should select the contents worth caching, which conditions make this process discriminatory towards the content, but not to the user? Current NN definitions, built on the principle of non-discrimination among contents belonging to the same class, do not provide rigorous guidelines to answer this question.

To shed some light on this issue, we first describe a cooperative strategy in which a CP extends its footprint by owning, maintaining and directly managing caching resources inside the domain of the ISP. Since not all the existing CPs can afford this type of cooperation, we refer to it as *partially-cooperative* caching.

Then, we propose a definition of NN-compliant caching and we compare the partially-cooperative caching with two alternative cooperative caching strategies, namely the *non-cooperative* and the *full cooperative* caching, which are NN-compliant at different degrees. In the former case, the ISPs equally divide their cache storage among all the CPs. In the latter case, a cooperative mechanism is implemented between the ISP and all the CPs to establish the right portion of cache storage that each CP is entitled to receive according to the provided definition of NN-compliant caching.

THE PARTIALLY-COOPERATIVE CACHING

Nowadays, the highest portion of traffic inside ISPs' networks is represented by VoD contents. For instance, the contents streamed by Netflix users account for an impressive percentage of the global Internet traffic (up to 40% at peak hours in the USA)⁶. Such volume of traffic inevitably requires ISPs to implement strategic countermeasures to avoid congestion, that may degrade the quality of connectivity. Being aware of the negative impact of such congestion on the QoS offered to its users, a CP may deploy its caches inside the ISPs domains to improve the streaming experience of the users and, at the same time, significantly reduce network traffic.

This is a win-win solution that apparently satisfies all the involved entities: the contracting CP, the users, ISPs and, to some extent, also competitor CPs, which can count on a less congested network to deliver their services. However, this solution can be considered discriminatory, because contents are treated differently based on the CP ownership.

There is no clear cut on whether this type of cooperation should be considered a case of NN infringement. Following the former argument, some believe it is not. ISPs are behaving neutrally, as they seek their advantage at no expense of other players and without negatively affecting the Internet community. Conversely, some others see this strategy as a subtle form of traffic prioritization, which should therefore be prohibited in a regime of content-based indiscriminate.

We argue that the *partially-cooperative* strategy would not violate NN provided that a similar solution could be applied by every CP. However, such condition turns out to be difficult to achieve in practice. As an example, let us consider the case of two CPs operating in the same sector (e.g., VoD), but with significantly-different market power. The smaller CP is clearly disadvantaged as it does not have scale to negotiate this kind of agreement with the ISPs. The result therefore is that its users perceive a lower QoS, which confirms that cooperative caching can be discriminatory.

[NN-compliant caching](#)

The performance of a caching strategy is typically measured using the *Hit-Rate*, i.e., the percentage of requests directly served from inside the domain where the caches are located (the ISP, in our case). Hence, a high Hit-Rate is always desirable as it leads to a significant reduction

of the traffic volume across the network. The ISP's caches are expected to store far fewer contents than those made available by all the CPs. Hence, we are dealing with an intrinsically selective process that, as such, raises discriminatory concerns. Being a vital instrument for the effective management of ISPs' resources, however, caching cannot simply be prohibited in favor of a pure non-discriminatory principle. In our view, ISPs are perfectly legitimate to apply the caching strategy that maximizes the Hit-Rate, as long as this is done avoiding any discrimination among the various CPs.

Caching performance depends on the size of the caches, on the content request patterns and on the employed caching strategy. In general, content popularity is the main factor leading to high Hit-Rate values. Hence, following reference⁵, we believe that content should be cached according to their global popularity, regardless of their CP ownership. While being aware that this approach clearly favors big Content Providers (as they are expected to own most of the popular content), we also believe it is non-discriminatory. In fact, a CP with a catalogue capable of engaging its users, in our view, deserves to receive more cache space than a CP that does not own equally appealing contents. In this way, the network is neutral with respect to both the CPs and the users, which are not driven to request a content because of a strategic cooperation. Instead, contents are awarded in terms of superior Quality of Service, which results in superior Quality of Experience only by virtue of their popularity.

To avoid potentially-discriminatory solutions, ISPs may employ *transparent caching* technologies⁷ to inspect the traffic traversing their network and infer content popularity. If encryption schemes are in place, however, such techniques are ineffective and alternative strategies to manage the caches are required to avoid using single-purpose caching systems like those envisioned in the *partially-cooperative* caching paradigm. In the next Section, we provide a quantitative assessment of the traffic distortion induced by three caching strategies, which are compliant with NN at different degrees, to better understand which types of behaviors are discriminatory.

QUANTITATIVE COMPARISON OF CACHING FRAMEWORKS

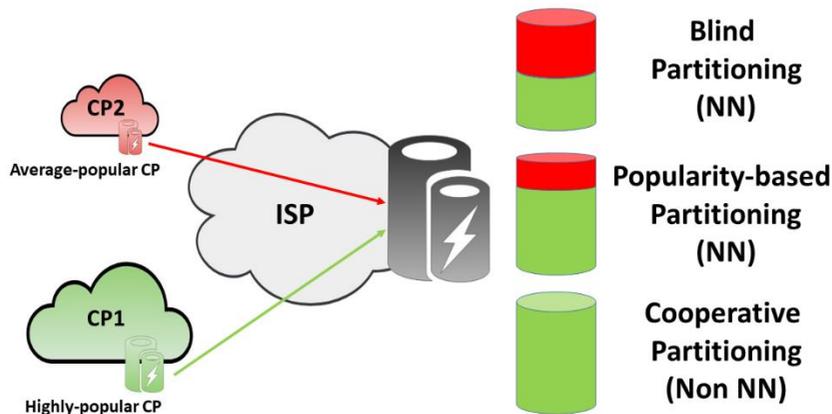


Figure 1: Representation of the considered scenario and of the employed partitioning strategies



We now consider three different scenarios of relations between two CPs and an ISP. As depicted in Fig. 1, the CPs compete to obtain a portion of the cache storage made available by the ISP in order to offer a better service to their users. CP1 and CP2 are assumed to have catalogues of different sizes and with different popularity patterns. Specifically, CP1 is considered a Content Provider with high attractiveness and most of its contents are more popular than those offered by CP2. We assume that each CP stores its most popular contents, which remain fixed during the considered time period.

In the first scenario, referred to as *Blind Cache Partitioning*, ISP and CP follow the **non-cooperative** approach: the ISP equally shares the storage of its cache among the CPs, regardless of the global popularity of the contents that they own. This is the most neutral case and we consider it as the baseline for the performance comparison. In order to allow for content encryption by the CPs, we assume that CPs manage their caching portion autonomously.

In the second scenario, referred to as *Popularity-based Cache Partitioning*, the ISP selects the contents to cache only according to their global popularity, i.e., regardless of the owner CP. In case the contents are not encrypted, the ISP can infer which contents are the most popular ones. In this way, it is guaranteed that the available cache storage is divided among the involved parties proportionally to the number of the most popular contents available in their catalogues. In this way, the CP that owns more popular content is assigned a larger slice of cache storage. This approach aims at maximizing the Hit-Rate of the ISP and it is not discriminatory. However, as encryption is nowadays widely employed in the content delivery chain, ISP and CPs need to jointly compute the amount of storage that should be dedicated to each party. Hence, it is required that the involved parties implement the **full-cooperative** strategy, which might represent the prelude of NN infringement. In Section “Discussions and Future work”, we describe how two existing approaches can provide several guidelines for the development of protocols enabling a seamless and non-discriminatory cooperation.

If these protocols are not employed, however, ISPs are pushed to follow the **partially-cooperative** approach. As an example of partial cooperation, we consider a third scenario, referred to as *Cooperative Cache Partitioning*, that aims at representing a realistic situation, where the ISP cooperates only with CP1. Because in this study we consider two CPs only, the whole cache storage is assigned to CP1.

With the aim of understanding which caching solutions may lead to discrimination, we perform a simulative study considering two CPs with significantly-different degrees of attractiveness. Specifically, we assume that CP1 and CP2 own 67% and 33% of the most popular contents, respectively. This proportion can be obtained by properly choosing the sizes of the offered catalogues and the skewness of their contents’ popularity. Note that infinite combinations of such parameters can lead to the target proportion. In our simulations, we arbitrarily assume that CP1 and CP2 offer a catalogue of 50K and 400K contents, respectively, whose popularity follows the Zipf law⁸ with skewness parameter $\alpha = 0.8$ and $\alpha = 0.9$, respectively. We recall that $\alpha \in [0,1]$ and the skewness of the content popularity distribution augments with increasing α , i.e., fewer contents are highly-requested with α approaching one. The domain of the ISP is abstracted as a single cache with a total storage of 40K contents.

In Fig.2, we show the Hit-Rates experienced by the two CPs depending on the employed cooperative strategy. The Blind Partitioning approach penalizes CP1 and benefits CP2, which receives more cache storage than what it would deserve based on the global popularity of its contents. In fact, CP1 sees a significant increase of its performance if, instead, the Popularity-based Partitioning approach is used. Conversely, CP2 is penalized by such fair assignment of resources. CP1 maximally benefits from the Cooperative Partitioning approach, while CP2 experiences a null Hit-Rate because none of its contents are cached, which results in a discriminatory treatment.

Fig. 2 also shows that the ISP significantly benefits from being net-neutral. However, the presented results are not sufficient to state that the ISP would always prefer to employ a NN-compliant caching. In fact, we are aware the maximization of the Hit-Rate is just one of its possible objectives. For example, the ISP may be concerned of performing an effective Traffic Engineering, which depends on other factors beside the global popularity of contents (e.g., their size). In the next Section, we describe the characteristics of a NN-compliant in-network caching.

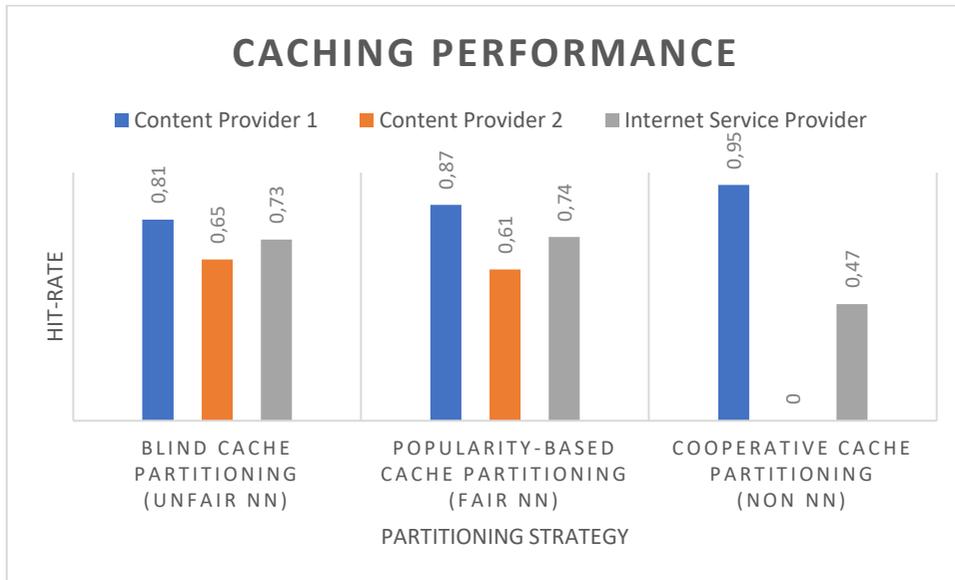


Figure 2: Hit-Rate experienced by CPs and ISP

DISCUSSIONS AND FUTURE WORK

In this final Section, we present a possible definition of NN-compliant in-network caching and we suggest few research directions toward its enforcement.

A Net-Neutrality definition of in-network caching

Since caching causes a differentiation of traffic generated by the CPs, the topic of in-network caching should be fully integrated in the NN debate. To take a first step in this direction, let us define *NN-compliant in-network caching* as the caching process that causes a traffic differentiation *only on the basis of the global popularity of contents*. Being aware that the maximization of the Hit-Rate is just one of the goals of the ISP, we also advocate a more general definition of NN-compliant caching, that is able to include other objectives based on which the ISP can legitimately perform traffic differentiation (e.g., minimizing the traffic load over their networks).

Open Questions

The fulfillment of the aforementioned principle, however, is practically hindered in a scenario of all-encrypted web, in which ISPs are incapable of inspecting the contents traversing their networks to infer the global popularity. A viable alternative would be to make CPs and ISPs cooperate with the goal of jointly managing the cache systems, which raises the following questions:

- Given that ISPs are prevented from identifying the global popularity of contents, how is it possible to implement a cooperative system that does not lead to a discrimination among the CPs?
- How can the ISP balance the available cache storage without requiring the CPs to reveal the popularity of their contents?

Research Directions

We believe that, in order to implement a caching system that is actually neutral toward all the involved parties, the available storage cannot be assigned to the CPs by employing arbitrary forms of cooperation. Instead, we advocate the definition of an open protocol enabling any potential CP to receive the amount of storage that it is entitled to manage. In our view, such protocol should meet three main requirements:

- *NN-compliance* that, in our vision, means to allow the ISP to maximize the Hit-Rate by favoring the CPs only by virtue of the attractiveness of their catalogue, i.e., only because of the global popularity of their contents;
- *Privacy* to allow the CPs not to disclose information about the popularity of their contents, which is a business-critical information;
- *Scalability* to accommodate the demands of a potentially high number of CPs.

To the best of our knowledge, existing methods that can be used to develop such protocol follow two different approaches, that we describe in the following from a high-level standpoint. In addition, we analyze the extent to which such approaches meet the three aforementioned requirements. The two approaches treat the contents of the involved CPs differently. However, under the definition of NN-compliant caching discussed throughout this paper, both are non-discriminatory.

One of them has been proposed in ⁹ and envisions a content delivery architecture that allows to efficiently cache, in the domain of the ISP, contents encrypted by a CP. Each request issued to the CP results in the generation of a pseudonym (which univocally identifies the requested content) that is freely readable by the ISP. Based on this pseudonym, the ISP can locate the requested content and directly serve it to the user. By employing this architecture, two main goals are achieved:

N° of Refreshes	Hit-Rate
1	0, 50
2	0, 43
3	0, 37
4	0, 30
5	0, 24

Table 1: Impact of number of refresh events on caching performance

(1) the ISP can count the occurrences of the pseudonyms to infer the popularity of contents and, based on that, apply any known caching strategy; (2) the CPs can keep their contents encrypted to ensure security and privacy to their users without preventing the application of caching. Assuming that the ISP aims at maximizing the Hit-Rate, each CP receives a fair amount of storage based on the global popularity of its contents. The privacy of the CPs is guaranteed as long as the ISP is not able to infer the actual content names from their pseudonyms.

However, the information about the most popular contents can be considered public, to some extent (e.g., the hit-parade is typically publicly known), and this may enable the ISP to associate the highly-popular contents with the most frequently occurring pseudonyms. To cope with this issue, the association between contents and their pseudonyms can be refreshed with some frequency (as already suggested by the authors of ⁹). However, since the ISP performs caching by using the pseudonyms as substitutes of the actual content names, the refreshing event nullifies the acquired knowledge of popularity. Hence, such refreshes deteriorate caching performance.

To evaluate the effectiveness of this strategy, we performed simulations considering the same scenario used to obtain the results shown in Fig. 2. We assume that, in the period between two events of refresh (referred to as round), the ISP learns the popularity of contents by analyzing the pseudonyms and, in the next round, stores the most popular ones. We summarize the obtained results in Table 1. Note that the values of the Hit-Rate are much worse than those shown in Fig. 2. In fact, those results have been obtained by considering the caches to be permanently filled with the actual most popular contents. Instead, by using the architecture described in this paragraph, the ISP decides which contents to cache based on an imprecise knowledge of popularity, as this is inferred from the analysis of a limited number of pseudonyms (because in practice the number of requests between two refresh events is limited). Concerning scalability issues, the authors of ⁹ state that the proposed solution can easily handle a large amount of requests (up to hundreds of thousand users even for a basic implementation of the architecture).

In practice, the architecture can be optimized to be deployed in real scenarios (e.g., by increasing the computational power of its main functional blocks).

In the other approach, which is proposed in¹⁰, the ISP cooperates with several CPs and reserves a portion of its cache storage to each of them. The ISP can obtain the information about the Hit-Rate experienced by each individual CP by virtue of the current resource allocation. Then, the ISP executes an iterative algorithm that takes as input the Hit-Rates of the individual CPs and converges to the optimum partitioning (i.e., the partitioning that maximizes the global Hit-Rate measured by the ISP). Note that the individual Hit-Rates are the only information that CPs are required to disclose to the ISP, which makes the approach privacy-preserving. One fundamental difference with respect to the former approach is that the ISP is only in charge of partitioning the available storage, while caching is performed by the individual CPs. Practically, this approach is a way to extend the footprint of the CPs by giving them a proper amount of cache resources inside the domain of the ISP. Scalability is analyzed in detail in¹⁰, where it is shown that the computational complexity of the algorithm is polynomial in the number of CPs, which implies a scalable application of the proposed approach in real scenarios.

The solutions proposed in reference⁹ and reference¹⁰ are presented as mechanisms to allow the ISP to perform optimal caching in a scenario of all-encrypted web. We do a step forward by showing that they can represent important building blocks toward the realization of a NN-compliant in-network caching architecture.

Final Remarks

In-network caching is the process by which ISPs store within their domains the most frequently requested content in order to reduce the traffic burden within the network. By selecting the individual contents that are worth caching, however, the ISP introduces a distortive effect on the traffic of the Content Providers, and this may raise discriminatory issues. In spite of the importance of the topic, in-network caching has not received enough attention in the Net-Neutrality debate. In this work, we partially fill such gap by analyzing how the wide use of encryption prevents ISPs from applying caching strategies without the cooperation of the Content Providers, which further increases the tendency towards violation of the Net-Neutrality principles. In order to cope with this issue, we advocate the definition of open protocols that help the involved parties to perform such cooperation seamlessly (i.e., without requiring the formation of arbitrary cooperative alliances). We also encourage the definitions of standards to foster their wide application.

REFERENCES

1. Rost, Peter, et al. «Network slicing to enable scalability and flexibility in 5G mobile networks.» *IEEE Communications magazine* (2017): 72-79. 55.5.
2. Van Schewick, Barbara. «Network neutrality and quality of service: what a nondiscrimination rule should look like.» *Stan. L. Rev.* 67 (2015): 1.
3. Garrett, Thiago, et al. "Monitoring Network Neutrality: A Survey on Traffic Differentiation Detection." *IEEE Communications Surveys & Tutorials* (2018).
4. Liu, Jason. «Cisco Blogs.» 4 October 2017. <https://blogs.cisco.com/enterprise/a-guide-for-encrypted-traffic-analytics>. 28 March 2018
5. Maillé, Patrick, Gwendal Simon, and Bruno Tuffin. «Toward a net neutrality debate that conforms to the 2010s.» *IEEE communications magazine* 54.3 (2016): 94-99.
6. Pariag, David, and Tim Brecht. «Application Bandwidth and Flow Rates from 3 Trillion Flows Across 45 Carrier Networks.» *International Conference on Passive and Active Network Measurement*. Springer, Cham (2017): 129-141. 5 April 2018.
7. Jia, Qingmin, et al. «The Collaboration for Content Delivery and Network Infrastructures: A Survey.» *IEEE Access* 5 (2017): 18088-18106.
8. Adamic, Lada A., and Bernardo A. Huberman. «Zipf's law and the Internet .» *Glottometrics* (2002): 143-150. 3.

9. Yuan, Xingliang, et al. «Enabling secure and efficient video delivery through encrypted in-network caching.» *IEEE Journal on Selected Areas in Communications* 34.8 (2016): 2077-2090.
10. Araldo, Andrea, Gyorgy Dan, and Dario Rossi. «Caching Encrypted Content Via Stochastic Cache Partitioning.» *IEEE/ACM Transactions on Networking (TON)* 26.1 (2018): 548-561.

ABOUT THE AUTHORS

Davide Andreoletti received Bachelor and Master Degrees (“cum laude”) in Telecommunications Engineering from Politecnico di Milano in 2012 and 2015, respectively. He is currently pursuing a Ph.D in Information Engineering at Politecnico di Milano and he is employed as researcher at the Dipartimento di Tecnologie Innovative (DTI) at SUPSI University in Lugano, Switzerland. His research mainly focuses on Network Neutrality and privacy-preserving strategies within the context of content delivery in Internet. Contact him at davide.andreoletti@supsi.ch.

Silvia Giordano, Ph.D. from EPFL, is currently the head of the NetLab and direction member of the Institute of System for Informatics and Networking, at SUPSI University in Lugano, Switzerland. She is an associate researcher at CNR, and Distinguished professor of the University of Tianjin. Her main research interests include Complex and Social Networking, Human Mobility, Pervasive Computing and Networking, Wireless and Mobile Ad Hoc Networks, Quality of Service and Traffic Control. She is ACM Distinguished Scientist 2014 and ACM Stars in Computer Networking and Communications 2017. Contact her at silvia.giordano@supsi.ch.

Cristina Rottondi received both Bachelor and Master Degrees “cum laude” in Telecommunications Engineering and a PhD in Information Engineering from Politecnico di Milano in 2008, 2010 and 2014 respectively. She is currently employed as researcher by the Dalle Molle Institute for Artificial Intelligence (IDSIA) in Lugano, Switzerland. Her research interests include data privacy and security in Smart Grids and optical networks planning. Contact her at cristina.rottondi@supsi.ch.

Massimo Tornatore received a Ph.D. degree in information engineering from Politecnico di Milano, Italy, in 2006. He is currently an Associate Professor with the Department of Electronics, Information, and Bioengineering, Politecnico di Milano. He also holds an appointment as Adjunct Professor with the Department of Computer Science, University of California, Davis, USA. He is the author of more than 300 peer-reviewed conference and journal papers. His research interests include performance evaluation, optimization and design of communication networks, cloud computing and privacy. XX He was the co-recipient of eleven best-paper awards. Contact him at massimo.tornatore@polimi.it.

Giacomo Verticale received the Ph.D. degree in telecommunications engineering from the Politecnico di Milano in 2003. He is currently an Assistant Professor with the Politecnico di Milano, Italy. His Ph.D. dissertation was on the performance of packet transmission in 3G mobile systems. Previously, he was with the Research Center CEFRIEL, where he worked on DSL high-speed access technologies. He was involved in several European research projects advancing the Internet technology. His current interests focus on the security issues of the Smart Grid, on Network Function Virtualization, and on Edge Computing in 5G. Contact him at giacomo.verticale@polimi.it.
