MOBILE DATA OFFLOADING FOR GREEN WIRELESS NETWORKS

HAORAN YU, MAN HON CHEUNG, GEORGE IOSIFIDIS, LIN GAO, LEANDROS TASSIULAS, AND JIANWEI HUANG

ABSTRACT

The explosive growth in mobile data traffic has resulted in unprecedented energy consumption in cellular networks, and has also induced huge operational expenditure for mobile network operators. A promising solution to address this problem in 5G systems is to use complementary technologies, such as Wi-Fi, to offload the traffic originally targeted toward cellular networks. In this article, we discuss the recent advances in the technologies and economics of two types of mobile data offloading: *operator-initiated offloading* and *user-initiated offloading*. In operator-initiated offloading, the mobile operators offload cellular traffic to Wi-Fi networks, which may belong to the mobile operators, mobile operators' residential subscribers, and third party Wi-Fi owners. In user-initiated offloading, the users decide how to offload traffic with or without the mobile operators' coordination. We present a taxonomy of various data offloading models, discuss various technical and economic challenges, and summarize the algorithms and mechanisms that we design to address these challenges. Finally, we outline some open problems that require further investigation.

INTRODUCTION

Mobile operators worldwide are facing enormous pressure on their cellular networks due to the explosive growth in mobile data traffic. According to Cisco's forecast [1], global mobile data traffic will grow to 30.6 EB per month by 2020, which corresponds to a nearly eight-fold increase between 2015 and 2020. Such a huge amount of traffic not only leads to a significant increase in the energy consumption and CO_2 emissions in communication networks, but also burdens the mobile operators with large operational expenditure (OPEX) [2]. Thus, it is widely anticipated that the fifth generation (5G) standardization will incorporate new communication technologies that result in significant improvements in both energy and cost efficiency [3]. With the vast number of Wi-Fi access points (APs) that have already been deployed worldwide and the availability of the unlicensed industrial, scientific, and medical (ISM) frequency bands, an energy-efficient approach to cope with the huge amount of data traffic in a timely manner is to use Wi-Fi networks to offload the traffic originally targeted toward cellular networks.

Mobile data offloading enables sustainable green networking in a number of aspects. First, it reduces the energy consumption during data transmission. Due to the shorter communication distances between mobile devices and Wi-Fi APs, Wi-Fi transmission can consume significantly less transmit power than macrocell transmission [4]. Second, mobile data offloading reduces the energy consumption in operating the macrocell base stations. By offloading some data traffic to Wi-Fi networks, it is possible to switch off some macrocell base stations or put them into sleep mode to conserve energy. This will reduce the OPEX for mobile operators. Therefore, mobile operators are increasingly interested in systematically exploiting the data offloading solutions to save energy and reduce OPEX.

Based on the initiator of the data offloading, we classify the data offloading approaches into two categories, namely, *user-initiated offloading* and *operator-initiated offloading*. In the early days of data offloading, *user-initiated offloading* is the typical option, where the users decide the network selections manually based on their local network conditions. However, with the advancement in communications standards, such as access network discovery and selection function (ANDSF) by the Third Generation Partnership Project (3GPP) and Hotspot 2.0 by the Wi-Fi Alliance (WFA), cellular and Wi-Fi networks are becoming more tightly integrated in terms of radio frequency coordination and network management. This enables *operator-initiated offloading*, where the connection manager in the mobile device connects to the operator's server, retrieves the operator's policy, and initiates the offloading. For these two offloading approaches, we discuss their corresponding technical and economic challenges, and present our recently proposed algorithms and market mechanisms that address these challenges. We also identify open problems that require further investigation.

The rest of the article is organized as follows. We first present a taxonomy of the data offloading models and discuss the major challenges. Then we show several solutions to the operator-initiated offloading and user-initiated offloading, respec-

Haoran Yu, Man Hon Cheung, and Jianwei Huang are with the Chinese University of Hong Kong.

George Iosifidis is with Trinity College Dublin and CONNECT Centre.

Lin Gao is with Harbin Institute of Technology.

Leandros Tassiulas is with Yale University.

This work was supported in part by the General Research Funds (Project Number CUHK 14202814 and 14219016) established under the University Grant Committee of the Hong Kong Special Administrative Region, China, and in part by the National Science Foundation under Grant CNS 1527090.

Digital Object Identifier: 10.1109/MWC.2017.1600323

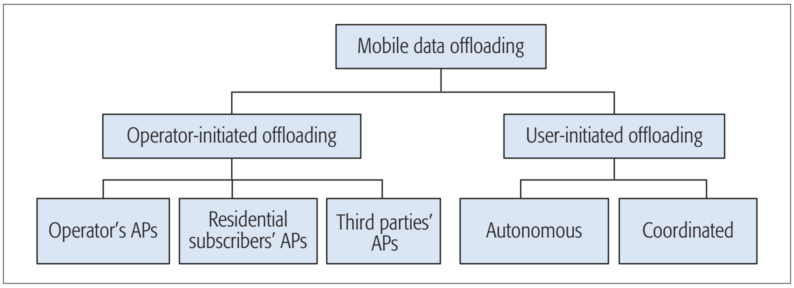

FIGURE 1. Taxonomy of data offloading models. The mobile data offloading includes operator-initiated offloading and user-initiated offloading, differentiated by the data offloading initiators. For operator-initiated offloading, we classify the models into three types based on which APs the operator uses for offloading data. For user-initiated offloading, we classify the models into two types based on whether the users can make their decisions with the operator's coordination.

tively. Finally, we outline the open problems and conclude the article.

Overview of Mobile Data Offloading Architectures

In this section, we introduce a taxonomy of data offloading models and the corresponding challenges. In order to realize the full potential of mobile data offloading, the initiator of offloading (i.e., an operator or a mobile user) needs to answer the following questions:

- When should traffice be offloaded?
- Which Wi-Fi AP should be used?
- How much traffic should be offloaded?

The answers to the above questions involve both technical and economic issues. From the technical perspective, the initiator of offloading should develop an intelligent data offloading algorithm by considering various factors, including the energy consumption, throughput, traffic delay, and service price. From the economic perspective, different entities have to reach consensus about the offloading policy. Specifically, when the operators do not own enough Wi-Fi APs, they need to lease the Wi-Fi resources (e.g., bandwidth and cache space) from their residential subscribers or third party Wi-Fi owners through effective economic mechanisms. Specifically, residential subscribers are residents who subscribe to the operators and own residential Wi-Fi APs. Different from residential subscribers, third party Wi-Fi owners are entities who own Wi-Fi networks but do not have prior relationships with the operators. Fully addressing these technical and economic issues brings the following benefits:

- The operators can efficiently save energy via the data offloading.
- The mobile users can achieve high data rates.
- The residential subscribers and third parties can generate revenues from their underutilized Wi-Fi APs.

We illustrate a taxonomy of data offloading models in Fig. 1, and discuss the detailed challenges for these models in the following.

Operator-Initiated Offloading

When an operator initiates data offloading, it can offload the cellular traffic to the following three types of Wi-Fi APs: the operator's own APs, residential subscribers' APs, and third parties' APs. When the operator offloads the traffic to its own APs, the operator is able to fully control the offloading process, and the key challenge is the design of an energy-efficient data offloading algorithm. However, due to the site acquisition and backhaul costs, it can be costly for the operator to deploy dense Wi-Fi networks by itself. A cost-efficient solution is to lease the residential subscribers' and third parties' APs, in which case the key challenge is the design of offloading policy and payment scheme that all the entities will agree on. Next we discuss the three offloading models in detail.

First, when an operator (e.g., AT&T) offloads the traffic to its own Wi-Fi APs, it needs to design an energy-efficient offloading algorithm with consideration of various criteria, such as traffic delay and revenue. Intuitively, offloading more cellular traffic to Wi-Fi reduces energy consumption, while it may also increase traffic delay due to users' intermittent Wi-Fi connectivities. Furthermore, the operator usually charges a lower price for the Wi-Fi usage than the cellular usage, and hence may lose revenue for the data offloading [5].

Second, an operator can motivate its own residential subscribers to open their Wi-Fi APs to have enough APs for offloading traffic. Such an approach has been implemented by network operators like Telefonica. An operator needs to determine the amount of cellular traffic offloaded to each residential AP as well as the payment to each corresponding subscriber.

Third, an operator can offload the traffic to the third parties' Wi-Fi APs. Since the third parties do not have prior relationships with the operator, the operator may need to compete with other operators to lease the third parties' APs. In general, each operator can lease APs from different third parties at different locations, and each third party can serve the cellular traffic from different operators. Some companies, such as BandwidthX, have built up platforms for collaborations between the operators and third party Wi-Fi owners.

User-Initiated Offloading

When a user initiates data offloading, it needs to determine its offloading strategy with consideration of various aspects, including the energy consumption at the user's smartphone, delay, data rate, and data price. For example, the Wi-Fi price is usually lower than the cellular price. Some operators charge users based on Wi-Fi connection time [6], and some Wi-Fi networks are free of charge. A key challenge for the user to design its offloading strategy is that the user does not know the complete information regarding the Wi-Fi networks' locations and throughputs. Furthermore, when the user accesses a Wi-Fi network, its data rate is affected by the total number of users connected to the same Wi-Fi network, which tightly couples the offloading strategies of different users. There are two major user-initiated offloading models. First, each user chooses its offloading strategy in an autonomous manner. In this model, each user has limited information on the locations, throughputs, and congestion levels of different Wi-Fi networks, and makes local and greedy decisions. Second, the operator provides the Wi-Fi network information to the users and coordinates the users' offloading strategies. With the operator's coordination, the users still make their own decisions in a distributed fashion, but can avoid using congested Wi-Fi networks and improve the utilization efficiency of Wi-Fi resources.

Solutions to Operator-Initiated Offloading

In this section, we show our solutions to three different types of operator-initiated offloading, where the operator offloads the data traffic to the operator's APs, residential subscribers' APs, and third parties' APs.

Offloading through an Operator's APs

In [7], we studied the scenario where an operator serves users with its cellular and Wi-Fi networks. We show an example in Fig. 2. The cellular network covers all locations, and the Wi-Fi networks have limited coverage. The users randomly move across different locations, and hence have random Wi-Fi availabilities. The operator has to dynamically select networks (cellular or Wi-Fi) for the users, and allocate cellular resources (channels and power) to the users associated with the cellular network. In order to save energy, the operator can offload some users to the Wi-Fi networks, and only allocate the channels and power to the cellular users with good channel conditions. In particular, the operator needs to design an intelligent energy-aware network selection and resource allocation algorithm to achieve a good energy-delay trade-off. For users who are not currently covered by the Wi-Fi networks, the operator can temporarily suspend their services to save energy and serve them with Wi-Fi when Wi-Fi becomes available again. However, this may lead to a large traffic delay for users who do not move around Wi-Fi very often. Therefore, the operator needs to reduce the power consumption while providing delay guarantees to all users, possibly by serving these users with the cellular network whenever needed.

We designed an Energy-Aware Network Selection and Resource Allocation (ENSRA) algorithm. In this algorithm, we monitor each user's queue backlog, the length of which denotes the amount of the user's unserved traffic. If a user has a small queue backlog, the operator will wait for the user to experience good cellular channel conditions or have access to Wi-Fi networks before serving the user. If a user has a large queue backlog, the operator will serve the user immediately using the cellular network (or Wi-Fi network if available), even if the user experiences a poor cellular channel condition such that the energy needed to serve the user is high. Therefore, the ENSRA algorithm is able to achieve a good balance between energy and delay. Moreover, the operator can achieve different energy-delay performance trade-offs by adjusting a positive control parameter V. When the operator chooses a large parameter V, it focuses more on reducing the power consumption; otherwise, the operator focuses more on reducing the traffic delay.

We compared the power-delay performance of our ENSRA algorithm to that of a heuristic algorithm. In the heuristic algorithm, the operator does not utilize the users' queue backlog information, and simply determines the network selection

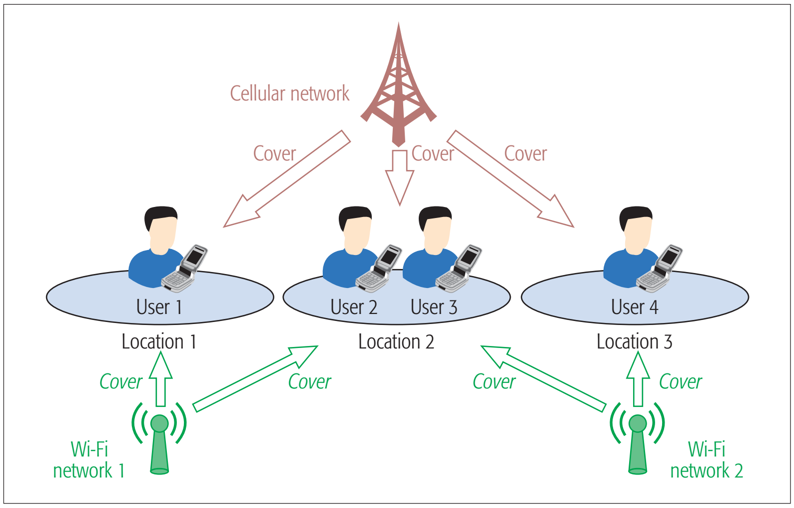

Figure 2. An example of the system model, where users 1, 2, 3, and 4 are moving among locations 1, 2, and 3. The cellular network covers all locations. Wi-Fi network 1 covers locations 1 and 2, and Wi-Fi network 2 covers locations 2 and 3.

and resource allocation based on the current network conditions [7]. In Fig. 3a, we plot the time average power consumption of ENSRA against V. We observe that ENSRA's power consumption decreases with V. Figure 3a also illustrates the heuristic algorithm's power consumption. Since V is the control parameter of the ENSRA algorithm, the heuristic algorithm's power consumption is independent of V. We notice that ENSRA consumes less energy than the heuristic algorithm for any $V > 0.2$ Mb^2/(W · s), as a larger V forces the operator to emphasize more on the energy saving. In Fig. 3b, we plot the average traffic delay per user under ENSRA against V. We observe that the traffic delay under ENSRA increases with V, and it is smaller than the traffic delay under the heuristic algorithm for any $V < 1.1$ Mb^2/(W · s). Figure 3 implies that if the operator chooses 0.2 Mb^2/(W · s) < 1.1 Mb^2/(W · s), ENSRA outperforms the heuristic algorithm in both the power consumption and delay. For example, ENSRA with $V = 0.5$ Mb^2/(W · s) saves 40.8 percent power consumption and 47.8 percent delay over the heuristic algorithm.

Offloading through Residential Subscribers' APs

In [8], we studied the offloading of cellular traffic to Wi-Fi mesh networks that are managed collaboratively by the operator's residential subscribers. Such mesh networks emerge nowadays in various different contexts. For example, several community networks have been deployed by residents for sharing content and network resources. Similar solutions have been launched by major network operators (e.g., Telefonica) and Internet service providers (e.g., Netblazr). We investigated the integration of the cellular networks and mesh networks, and designed a framework that determines which users should be offloaded to the mesh networks and how the offloaded traffic is further routed to the cellular core.

We focused on an eNodeB that serves a set of users. The goal of the eNodeB is to deter-

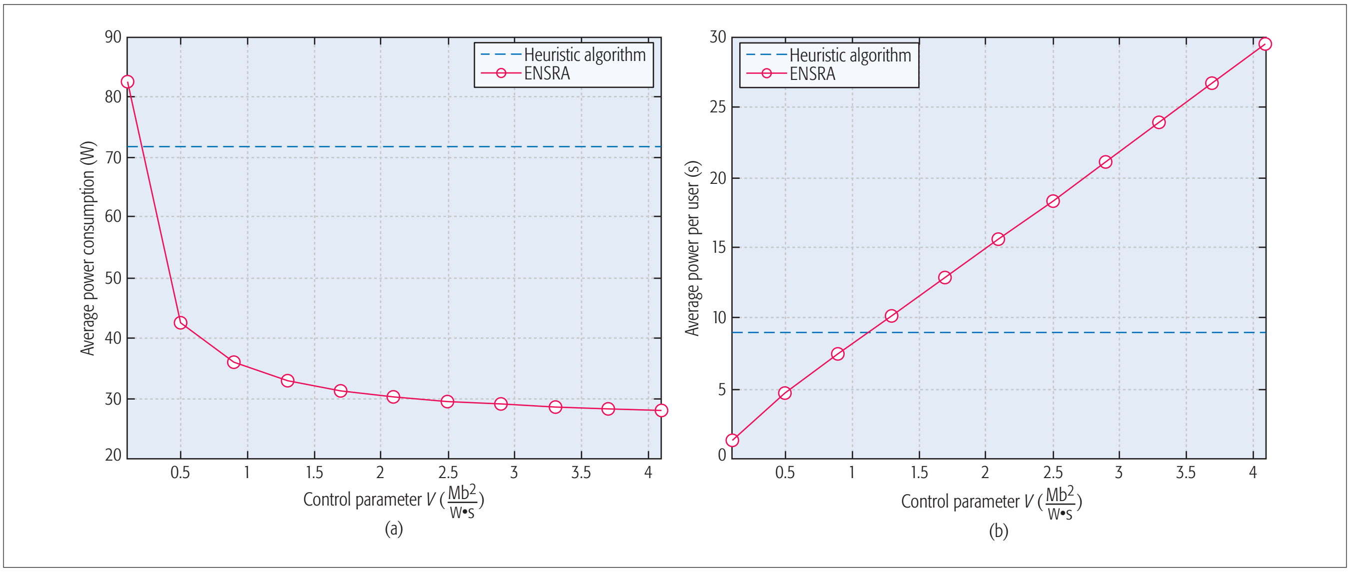

FIGURE 3. Comparison of ENSRA and heuristic algorithm: a) power consumption vs. parameter V; b) traffic delay vs. parameter V.

mine which users should be offloaded so as to maximize the energy consumption saving. Based on the consideration of energy and users' demands, we proposed an offloading algorithm that decides the set of users to be offloaded. The offloading decision is constrained by the Internet access availability of the mesh network, and the offloading algorithm runs in pseudo-polynomial time. Once the eNodeB has decided the traffic to be offloaded, the mesh network determines how the offloaded data will be further routed to/from the Internet gateways, taking into account the mesh network's available resources and servicing costs, as well as the quality of service guarantees for the users. We designed a polynomial-time routing algorithm for the mesh network to make the routing decision. We evaluated the above framework by conducting extensive experiments in a residential mesh network. We found that our framework can significantly reduce the eNodeB's energy consumption. Specifically, we considered an eNodeB that consumes 19.33 W per time slot. When the eNodeB offloads 25 percent of the mobile users to the mesh networks, the energy saving of the eNodeB ranges from 0.88 W (in the light load case) to 10.39 W (in the heavy load case).

In [9], we proposed the leasing of wireless bandwidth and cache space in residential Wi-Fi APs for offloading purposes. Due to caching, it is possible for mobile users to fetch content directly from nearby Wi-Fi APs without using the often congested backhaul links that connect residences with the core network. Such solutions are expected to become very attractive in the near future for two reasons. First, commercial in-network storage solutions have started to proliferate, and storage is becoming an important network resource to manage and optimize.[1] Second, the storage and link capacity resources of residential APs are often underutilized. In particular, we analyzed a novel dataset about the detailed Wi-Fi usage of 167 residential subscribers over a four-month period,[2] and we found wide disparity in Wi-Fi utilization across residences and potential for offloading mobile data through the underutilized Wi-Fi APs at any time of the day.[3] The questions that we need to address for enabling such an architecture are:

- How much cache space and wireless bandwidth should the operator lease from each residential AP, and how much should be paid for this?
- Which files should the operator cache and deliver from each residential AP?

To answer these questions, we studied the joint optimization of incentive policy, caching policy, and routing policy. We proved that the joint optimization problem is NP-hard in general. For a special case of non-overlapping APs, we showed how the joint optimization problem can be optimally solved in polynomial time. For the general case, we proposed a primal-dual algorithm, which can quickly achieve a near-optimal solution to the joint optimization problem. We evaluated the performance of our solution with our dataset. We found that the operator is able to reduce the cost for serving the mobile users by up to 50 percent. Moreover, residential subscribers may see up to €9 reduction of their monthly bills if they offer to offload traffic (based on average European data prices).

In [10], we applied the Nash bargaining theory to model and analyze the negotiations between the operator and residential APs regarding the amount of offloaded traffic and payments. We formulated the entire negotiation processes between the operator and all residential APs as a one-to-many bargaining game, and analyzed the bargaining solutions systematically under two different bargaining protocols: *sequential bargaining*, where the operator bargains with APs sequentially, one AP at a time, in a given order, and *concurrent bargaining*, where the operator bargains with all APs concurrently. For clarity, we illustrate these two bargaining protocols in Fig. 4. We further studied the group bargaining scenario, where multiple APs form a group to jointly bargain with the operator, and quantified the benefits for APs when forming such a group. Our theoretical analysis, together with the numerical studies, lead to the following observations and insights:

[1] Caching at Wi-Fi APs is currently supported by commercial products such as Linksys Smart Wi-Fi and HiWiFi routers. Their advanced operating systems are capable of running various applications to customize content caching schemes.

[2] The dataset is provided by Portugal Telecom (PT), which is an Internet service provider in Portugal.

[3] Specifically, 80 percent of users utilize their Wi-Fi for less than five hours of the day.

- First, the bargaining outcomes under both bargaining protocols are efficient, in the sense that the social welfare can always be maximized.
- Second, under the sequential bargaining, an AP can obtain a higher payoff if it bargains with the operator earlier, which we call the *Early-Mover Advantage*. Under the concurrent bargaining, the payoff of each AP equals the worst-case payoff that it can achieve under the sequential bargaining, which we call the *Concurrently Moving Tragedy*.
- Third, under the sequential bargaining, group bargaining improves the payoffs of both the APs in the group and those APs bargaining before the group, while it does not affect the payoffs of those APs bargaining after the group. Under the concurrent bargaining, group bargaining only improves the payoffs of the group members.

Offloading through Third Parties' APs

In [11], we studied the data offloading from multiple operators to multiple third parties' APs, and focused on the necessary economic incentives that the operators need to provide for the APs considering the potential competition among the operators. Specifically, we designed an *offloading market*, where the operators (acting as *buyers*) compete to lease the APs (acting as *sellers*) for offloading their cellular traffic. We assumed that the marketplace is managed by a centralized broker who can be a regulator-managed clearinghouse or a private company, similar to those in the secondary spectrum market. The detailed trading process is as follows. The APs offer their services (i.e., offloading traffic for the operators) in exchange for reimbursements. The operators declare how much they are willing to pay each AP. The broker collects the operators' requests and the APs' offers, and determines how much traffic of each operator will be offloaded to each AP and at what price. The operators and the APs comply with the broker's decisions only if it is in their own interests to do so.

We proposed a two-sided iterative double auction that enables the broker to gradually reach the socially efficient solution, without any prior knowledge of the market. More specifically, the broker designs and announces a traffic scheduling and a payment rule, which are parameterized by a set of prices. In each round, the broker announces the prices, and then the operators and APs submit their requests and offers, respectively. Based on the received operators' requests and APs' offers, the broker updates and announces the prices in the next round. The time complexity for updating the prices is polynomial in both the number of operators and the number of APs. The iterative double auction continues until the prices converge to certain stable values.

Solutions to User-Initiated Offloading

In this section, we show our solutions to the user-initiated offloading, where users make the data offloading decisions with or without the operator's coordination.

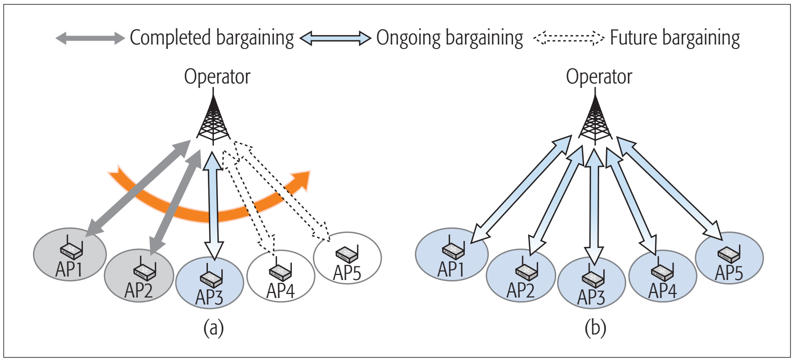

Figure 4. Illustration of bargaining protocols: a) sequential bargaining; b) concurrent bargaining.

Autonomous Offloading Decision Making

In [12], we studied a particular user's individual offloading decision without any coordination among users. In particular, we considered *delay-tolerant* applications, such as software update and movie downloading, where a user can tolerate some delays without experiencing a significant satisfaction loss. As illustrated in Fig. 5, we assumed that the cellular network has ubiquitous coverage, while the Wi-Fi availability is location-dependent. Given the statistical information on the user's mobility pattern and the deadline of the delay-tolerant application, the user needs to consider the following trade-off:

- *Data usage payment*: On one hand, the user wants to pay as little as possible, so it is preferable to wait for the cheaper Wi-Fi and offload as much data traffic to Wi-Fi as possible.
- *Quality of service requirement*: On the other hand, the user has a deadline to meet. If it cannot encounter Wi-Fi soon, it should use the cellular network immediately.

To systematically consider this trade-off, we formulated the dynamic network selection as a Markov decision process (MDP) problem, where the action is the network selection (i.e., either remaining idle, using the cellular network, or using a Wi-Fi network if available), and the state consists of the user's current location and the remaining size of the file to be transferred. The objective is to find the optimal policy that minimizes the user's expected total payment plus the penalty for incomplete file transfer. By using the finite-horizon MDP model, we computed the optimal offloading policy and designed a Delay-Aware Wi-Fi Offloading and Network Selection (DAWN) algorithm. The complexity of the DAWN algorithm is $\mathcal{O}(KLT\sigma)$, where K is the file size, L is the number of locations, T is the number of time slots, and σ is the granularity of the file size dimension in the algorithm. Furthermore, we established the sufficient conditions under which the optimal policy exhibits a nice threshold structure in both the file size and time dimensions, which led to the design of a lower complexity monotone DAWN algorithm. Contrary to some widely adopted practices, such as "Wi-Fi first," our results showed that it is not always optimal to perform Wi-Fi offloading (even when Wi-Fi is immediately available) when the deadline requirement is stringent and when the

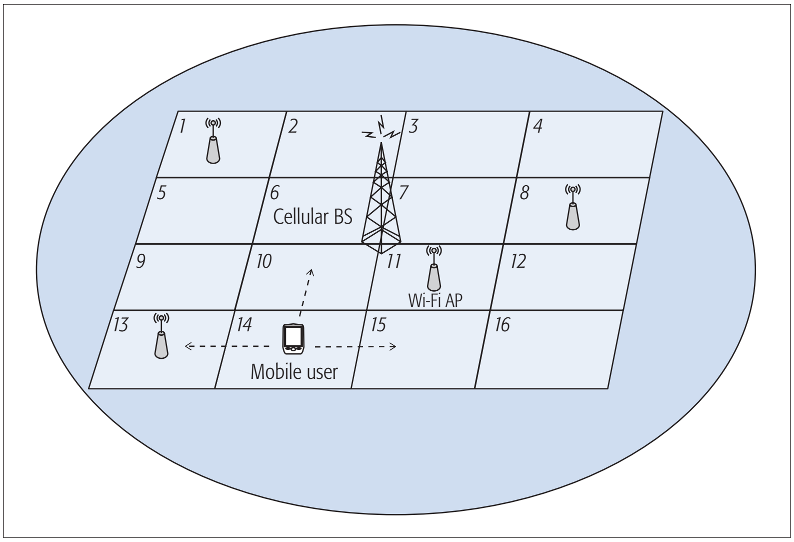

Figure 5. An example of the network setting, where a user is moving within 16 locations. The user is always under the coverage of a cellular base station, but Wi-Fi is only available at locations 1, 8, 11, and 13.

cellular network provides much higher throughput than the Wi-Fi network (e.g., a 4G LTE-Advanced cellular network vs. a congested Wi-Fi network).

Coordinated Offloading Decision Making

In [13], we studied the coordination of the users' network selection and data offloading decisions. We aim to address two practical problems in network selection:

- *Wi-Fi network congestion*: As mentioned earlier, "Wi-Fi first" is a widely adopted policy that aims to offload the data traffic to the best available Wi-Fi network. However, if all the devices in proximity adopt this same policy, they will simultaneously choose the same Wi-Fi network, which may lead to network congestion and thus low throughput.
- *Ping-pong effect*: Due to the delay and power consumption during network handoff, a user may incur the *switching penalty* when it decides to change the network connection. Without considering this penalty, a fast-moving user adopting the "Wi-Fi first" policy may experience too frequent network switching, which can lead to a throughput reduction and fast battery degradation.

In this work, we formulated the interactions among multiple users as a non-cooperative network selection game, where a user's strategy is a particular network-time route regarding his/her network selections across different time slots. His/her payoff for choosing a network-time route is the total throughput obtained minus the total network switching costs incurred. We showed that asynchronous best response updates of users' strategies are guaranteed to converge to a pure Nash equilibrium. In addition, we characterized that each best response update can be computed in $\mathcal{O}(N^3 T^3)$ time, where N is the number of networks and T is the number of time slots. As a result, we proposed a network selection algorithm to facilitate the users' coordination, which led to a better load balancing and fairness performance than some heuristic schemes.

Open Problems

Despite the research efforts that we have described above, there are still some open problems in the algorithm and market mechanism design for the data offloading.

Distributed and network-centric algorithm design: In user-initiated offloading, although the users make their network selections in a distributed manner, they seldom consider the operator's objective in terms of overall network performance and OPEX. In contrast, in operator-initiated offloading, although the network operations are under the operator's control, most of the proposed algorithms operate in a centralized fashion, where the operator makes the network selection and resource allocation decisions under the full information of the system. In order to increase the scalability and reduce the amount of information exchange, it is important to consider a distributed design in which each user makes a distributed decision based on the local network information while taking into account the operator's objectives.

Real-time network information: In order to perform the intelligent network selection as discussed in [7, 12, 13], it is necessary for the users to be able to assess the real-time network information, such as the data rates and congestion levels at different macrocell base stations or Wi-Fi APs. While the ongoing industry standards, such as ANDSF and Hotspot 2.0, are aiming to provide users with more network information, their interoperability for a comprehensive information sharing still requires further development.

Incomplete user information: In practice, the operator can design proper cellular and Wi-Fi pricing schemes to motivate the users to offload their traffic to the Wi-Fi networks [5]. However, it is challenging for the operator to accurately estimate the users' demands and predict their behavior.

Conclusion

Mobile data offloading has recently emerged as a promising approach to reduce the energy consumption in 5G systems. However, due to the highly dynamic network conditions and user mobility, the interactions among multiple operators and mobile users are particularly challenging to analyze. In this article, we present recent advances in data offloading technologies and economics. We classify the data offloading approaches, discuss their technical and economic challenges, and present our recently proposed solutions. Finally, we identify some open challenges in this fast growing research area.

References

[1] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015–2020, Feb. 2016.

[2] S. Zhang *et al.*, *Wireless Traffic Steering for Green Cellular Networks*, 1st ed., Springer, 2016.

[3] J. G. Andrews *et al.*, "What Will 5G Be?," *IEEE JSAC*, vol. 32, no. 6, June 2014, pp. 1065–82.

[4] J. Huang *et al.*, "A Close Examination of Performance and Power Characteristics of 4G LTE Networks," *Proc. ACM MobiSys*, Low Wood Bay, U.K., June 2012, pp. 225–38.

[5] K. Poularakis, G. Iosifidis, and L. Tassiulas, "Deploying Carrier-Grade WiFi: Offload Traffic, Not Money," *Proc. ACM MobiHoc*, Paderborn, Germany, July 2016, pp. 131–40.

[6] L. Duan, J. Huang, and B. Shou, "Pricing for Local and Global Wi-Fi Markets," *IEEE Trans. Mobile Comp.*, vol. 14, no. 5, May 2015, pp. 1056–70.

[7] H. Yu *et al.*, "Power-Delay Tradeoff with Predictive Scheduling in Integrated Cellular and Wi-Fi Networks," *IEEE JSAC*, vol. 34, no. 4, Apr. 2016, pp. 735–42.
[8] A. Apostolaras *et al.*, "A Mechanism for Mobile Data Offloading to Wireless Mesh Networks," *IEEE Trans. Wireless Commun.*, vol. 15, no. 9, Sept. 2016, pp. 5984–97.
[9] K. Poularakis *et al.*, "Mobile Data Offloading Through Caching in Residential 802.11 Wireless Networks," *IEEE Trans. Net. Serv. Mgmt.*, vol. 13, no. 1, Jan. 2016, pp. 71–84.
[10] L. Gao *et al.*, "Bargaining-Based Mobile Data Offloading," *IEEE JSAC*, vol. 32, no. 6, June 2014, pp. 1114–25.
[11] G. Iosifidis *et al.*, "A Double-Auction Mechanism for Mobile Data-Offloading Markets," *IEEE/ACM Trans. Net.*, vol. 23, no. 5, Oct. 2015, pp. 1634–47.
[12] M. H. Cheung and J. Huang, "DAWN: Delay-Aware Wi-Fi Offloading and Network Selection," *IEEE JSAC*, vol. 33, no. 6, June 2015, pp. 1214–23.
[13] M. H. Cheung, R. Southwell, and J. Huang, "Congestion-Aware Network Selection and Data Offloading," *Proc. IEEE CISS*, Princeton, NJ, Mar. 2014.

Biographies

Haoran Yu [S'14, M'17] (hryu@ie.cuhk.edu.hk) is a postdoctoral fellow at the Chinese University of Hong Kong, where he received his Ph.D. degree in information engineering in 2016. He was a visiting student at Yale University from 2015 to 2016. His research interests lie in the field of wireless communications and network economics, with current emphasis on cellular/Wi-Fi integration and economics of Wi-Fi networks. His paper in *IEEE INFOCOM 2016* was selected as a Best Paper Award finalist.

Man Hon Cheung (mhcheung@ie.cuhk.edu.hk) is a research assistant professor with the Department of Information Engineering at the Chinese University of Hong Kong. He received his Ph.D. degree in electrical and computer engineering from the University of British Columbia. He has served as a Technical Program Committee member at IEEE ICC, GLOBECOM, WCNC, and WiOpt. His current research focus includes mobile data offloading and mobile crowdsensing.

George Iosifidis (george.iosifidis@tcd.ie) received his Ph.D. degree in electrical engineering from the University of Thessaly, Greece, in 2012, and is currently the Ussher Assistant Professor in Future Networks with Trinity College Dublin. His research interests lie in the broad area of wireless network optimization and economics.

Lin Gao [S'08, M'10, SM'16] (gaol@hit.edu.cn) is an associate professor at Harbin Institute of Technology, Shenzhen, China. He received his Ph.D. degree in electronic engineering from Shanghai Jiao Tong University in 2010, and worked as a postdoctoral fellow at the Chinese University of Hong Kong from 2010 to 2015. He received the IEEE ComSoc Asia-Pacific Outstanding Young Researcher Award in 2016. His research interests are in the interdisciplinary area of network economics and games.

Leandros Tassiulas [S'89, M'91, SM'05, F'07] (leandros.tassiulas@yale.edu) is the John C. Malone Professor of Electrical Engineering and member of the Institute for Network Science, Yale University. He received the IEEE Koji Kobayashi Computer and Communications Award (2016), the Inaugural INFOCOM 2007 Achievement Award, the INFOCOM 1994 Best Paper Award, the National Science Foundation (NSF) Research Initiation Award (1992), the NSF CAREER Award (1995), the Office of Naval Research Young Investigator Award (1997), and the Bodossaki Foundation Award (1999).

Jianwei Huang [F'16] (jwhuang@ie.cuhk.edu.hk) is a professor in the Department of Information Engineering at the Chinese University of Hong Kong. He has received eight international Best Paper awards, including the IEEE Marconi Prize Paper Award in Wireless Communications 2011. He has served as the Chair of IEEE Technical Committee on Cognitive Networks and Technical Committee on Multimedia Communications. He is an IEEE ComSoc Distinguished Lecturer and a Thomson Reuters Highly Cited Researcher.}