

# Information-Centric Connectivity

Konstantinos V. Katsaros, Vasilis Sourlas, Ioannis Psaras, Sergi Reñé, and George Pavlou

Mobile devices are often presented with multiple connectivity options, usually making a selection either randomly or based on the wireless medium's conditions, as is the case for current offloading schemes. The authors propose that link-layer connectivity can be associated with information availability, and in this respect connectivity decisions should be information-aware. To achieve information awareness at the link layer, we leverage on the information-centric networking paradigm and introduce the concept of information-centric connectivity.

## ABSTRACT

Mobile devices are often presented with multiple connectivity options, usually making a selection either randomly or based on the wireless medium's conditions, as is the case for current offloading schemes. In this article we propose that *link-layer connectivity can be associated with information availability and in this respect connectivity decisions should be information-aware*. To achieve information awareness at the link layer, we leverage on the information-centric networking paradigm and introduce the concept of information-centric connectivity (ICCON). We elaborate on different types of information availability and connectivity decisions in the context of ICCON, present specific use cases, and discuss emerging opportunities, challenges, and technical approaches. We illustrate the potential benefits of ICCON through preliminary simulation and numerical results in an example use case.

## INTRODUCTION

Motivated by the proliferation of content-centric applications in the Internet, information-centric networking (ICN) promises a shift in the operation of the network, enabling routing and forwarding based on identifiers/names of content rather than network locations [1]. As such, ICN research efforts have so far primarily focused on adapting the IP protocol stack, but also on incorporating appealing features, such as in-network caching, multicast, and mobility, to the main architecture.

At the same time, user demand for information is increasingly expressed through handheld devices (user equipment [UE]), which are normally presented with multiple connectivity options. Users are frequently in the vicinity of multiple IEEE 802.11 (WiFi) networks [2], usually offloading Internet traffic [3], but also providing access to locally stored information (e.g., major airlines offer WiFi-based multimedia services), or even enabling device-to-device (D2D) opportunistic communication between users [4] through WiFi Direct (e.g., FireChat<sup>1</sup>). Even within cellular networks, the emerging availability of variable size overlapping cells in heterogeneous networks (HetNets) [5] yields multiple options for cell selection.

As a result, UEs are presented with a multitude of connectivity opportunities to access information available in their networking vicinity in the form of opportunistically cached, intentionally pre-fetched, or locally generated (e.g., UE

content and/or services/applications). Depending on the networking environment, information can reside at a wide range of accessible network locations, including UEs, WiFi access points (APs) or cellular base stations with storage/processing capabilities, in-network content-centric routers, in-network caches/middle-boxes, and micro-clouds [6]. Diverse user mobility patterns and service interests inevitably formulate a unique and dynamic information and service availability map that is, different content/services get(s) cached/pre-fetched/placed at different network locations. In turn, awareness of information availability in the networking vicinity can lead to connectivity decisions that reflect user interests, improving users' quality of experience and facilitating or even enabling the otherwise impossible access to the desired information.

However, connectivity decisions are currently information agnostic. Offloading mechanisms are primarily focused on load and performance metrics such as downlink/uplink throughput and signal-to-noise ratio of wireless connections [3], disregarding information availability. Moreover, searching for information currently builds on the assumption that a UE is already associated with a particular network device (e.g., WiFi AP, eNodeB, or another UE). As a result, UEs need to engage in an iterative, time- and energy-consuming process comprising the network association [7] and the subsequent search for information at the application level.

In this article we argue that the ICN paradigm can and should be extended beyond the scope of the network layer (and above), enabling information-centric connectivity (ICCON). In ICCON, *information-centrism is further expressed in the connectivity decisions made by UEs*, which aim to discover networks that enable or facilitate access to the desired information itself. In essence, ICCON introduces information-awareness at the link layer.<sup>2</sup> In a characteristic ICCON use case, the selection of a WiFi AP for offloading cellular traffic can be driven by the matching between the user content interests and the content currently cached (or proactively pushed) at the AP and/or the broadband remote access router (BRAS). In another example case, a WiFi service set identifier (SSID) is carefully set up and selected by UEs to enable the exchange of information for a local event (e.g., photos taken during a concert).

ICCON is expected to enhance user experience as information-centric connectivity decisions bring the user closer to the desired

<sup>1</sup> <https://en.wikipedia.org/wiki/FireChat>

<sup>2</sup> ICCON introduces information awareness as a novel criterion in UE connectivity management; as such, it is orthogonal and fully compatible with efforts focusing on either the exclusive selection of the most appropriate access technology (i.e., vertical handovers) or the concurrent use of multiple network interfaces.

information, reducing latencies, along with network traffic and server load. Note that in the context of ICCON, these benefits come without the currently imposed need to search for information through the time- and energy-consuming network association process. This comes in sharp contrast to a substantial body of work on service discovery, which, in most cases, assumes the establishment of connectivity between participating devices, before any service discovery protocol is employed (e.g., Jini, UPnP). At the same time, the ability of UEs to intelligently discover information in their networking vicinity enables new opportunities for network operators, content and service providers, or even users themselves, to provide access to information with dedicated low-cost equipment (e.g., APs, UEs), decoupling information provisioning from Internet access.

So far, a series of promising features has been identified to promote the adoption of the ICN paradigm in wireless and mobile environments, including the inherent support of consumers' mobility, the retrieval of requested content without the need for global knowledge of the content origin identity, the native support of multicast, and the ability to cope well with intermittent short-lived connections and dynamic network topologies. Amadeo *et al.* [8] and Tyson *et al.* [9] survey the feasibility and benefits of applying ICN in various wireless scenarios such as mobile ad hoc networks (MANETs), vehicular ad hoc networks (VANETs), and wireless sensor networks (WSNs) as well as the various research prospects in extending ICN to wireless networks. However, none of the works referenced in the aforementioned surveys considered information centrism during the connectivity decision process.

In the following, we elaborate on the ICCON concept, thoroughly discussing the various connectivity options and decisions UEs can encounter, subject to different types of information available in their vicinity. We elaborate on the mechanisms required to realize ICCON and further present two example use cases. We identify a series of challenges to be met for the realization of the ICCON concept, and we delve into the details of a particular offloading use case presenting preliminary results that demonstrate the potential benefits of ICCON.

## INFORMATION AWARENESS, CONNECTIVITY OPTIONS, AND DECISIONS

### CONNECTIVITY OPTIONS AND INFORMATION AVAILABILITY

ICCON is inspired by the increasing multitude of connectivity options for UEs and the corresponding information availability. The proliferation of WiFi technology is manifested by the often dense deployments of WiFi hotspots in urban environments, even up to several tens of APs with overlapping coverage [2]. The emergence of HetNets [5] is similarly expected to yield multiple connectivity options to UEs, allowing the introduction of information awareness in the cell selection process. In such environments, information can be available in the form of opportunistically cached or pre-fetched content, as well as services hosted at the AP (e.g., set-top box with storage) or eNodeB, at in-network middle-

boxes (e.g., WAN accelerators) or content-centric routers [6].

ICCON is also an enabler of local communications where access to the Internet is often not available (e.g., underground metro), rather limited (e.g., flash crowds deplete mobile network resources), or simply not needed (e.g., sharing slides in a conference room). The formation of WiFi Direct groups (i.e., D2D communication) appears as an appealing alternative option in the quest for information; especially when this information is semantically linked to the concentration of the users (e.g., photos during a concert). Such information can have the form of previously downloaded/retrieved, locally (user) generated, or expected content.

### CONNECTIVITY DECISIONS AND INFORMATION NAMING

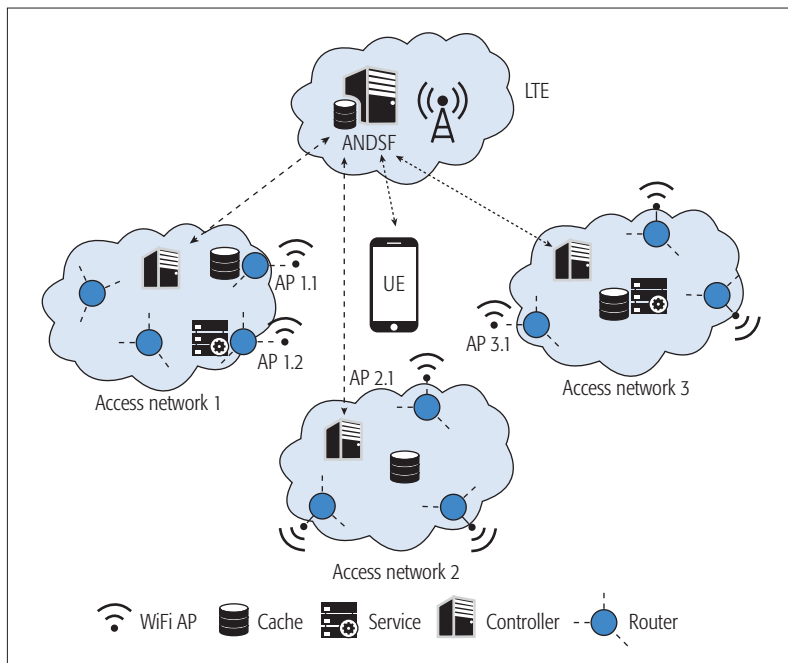
Given a wide range of connectivity opportunities, connectivity decision criteria in ICCON can vary, subject to the type of the available/expected information, as well as its representation. In the case of cached content, the objective of ICCON is to identify the connectivity option that maximizes the matching with the users' interests, and in turn the likelihood of a cache hit. To this end, UEs keep track of the requested content, building *UE profiles* through time. Content is identified following the naming conventions of the caches, for example, URL or content-centric networking/named data networking (CCN/NDN) chunk name. On the other side, a cache index provides a description of the content available, forming a corresponding *network profile*. Similar profiling mechanisms can be envisioned for pre-fetched content.

In other cases, the objective of ICCON may be the discovery of particular information in the (wireless) networking environment (e.g., photos of a local event). Taking into account the enormous size of the information space, searching for individual items would likely result in significant and impractical control plane overheads. Overcoming this limitation, we envision the structuring of the information space around the notion of services/applications enabling the aggregation of individual information item names within their semantic scope (e.g., aggregating articles in a news service). Similar to publish/subscribe systems, services/applications may further define individual topics to support finer-grained semantic scopes (e.g., a World Cup Final topic created by a news service). ICCON then targets discovering the desired services/applications (or topics therein) within each connectivity option, leaving finer-grained information discovery and retrieval for the application layer (i.e., upon a connectivity decision). This avoids the need for a universal naming scheme of information items, promises better scalability, as only a few million applications are currently available at the most popular application marketplaces, and builds on the currently prevailing application-centric usage model of handheld devices.

### ICCON MECHANISMS

ICCON requires a control plane matching mechanism between user interest (user profile) and information availability. Depending on the environment, we consider a series of design options for this mechanism.

In the case of cached content, the objective of ICCON is to identify the connectivity option that maximizes the matching with the users' interests, and in turn the likelihood of a cache hit. To this end, UEs keep track of the requested content, building UE profiles through time.



**Figure 1.** Mobile network operator-supported offloading. Cached content or services may be collocated with the APs or reside further inside the network (e.g., a centralized cache in access network 2). A local per-access-network controller reports information availability to the ANDSF component of the mobile network. A UE connects to one of the APs in its vicinity (i.e., AP 1.1, AP 1.2, AP 2.1, or AP 3.1) based on recommendations from the ANDSF.

**Push vs. Pull:** Candidate connectivity points (e.g., AP, eNodeB, UE), as well as client UEs can proactively (*push*) or reactively (*pull*) advertise their network and/or UE profiles. The selection of the exact mode depends on aspects related to privacy concerns (i.e., exposure of user interests), energy availability and incentive schemes (e.g., for advertising available information), overheads (e.g., network vs. UE profile size), and so on.

**Decision Maker:** The comparison between the desired and available/expected information can be made by different entities depending on the network environment. In the case of mobile network offloading, the access network discovery and selection function (ANDSF) [10] at the mobile network operator side can take this role, collecting the necessary information and reducing the associated overheads for UEs. Individual APs can also take this role. However, in both cases, privacy concerns motivate the placement of this decision process at the UE side.

**Protocol Layer and Technical Enablers:** The exchange of the aforementioned representations can take place at the link layer or above. In the case of mobile network offloading, TCP/IP interfaces can be employed by the ANDSF for the (proactive) collection of network and UE profiles, as well as for communicating connectivity suggestions to UEs [10]. When UEs autonomously select their network of preference, a link layer mechanism is required for the exchange of the representations. UEs can retrieve network profiles (or service/application/topic identifiers) through the Access Network Query Protocol (ANQP) of IEEE 802.11u (e.g., the *ANQP vendor-specific list* element) [11]. The recently announced WiFi Neighbor Awareness Networking (NAN) proto-

col [12] also supports a low-energy-consumption device discovery mechanism enhanced with publish/subscribe primitives that can serve the same purpose. Carefully selected non-human-readable SSID values can also be used in the particular case of service/application/topic identifiers.

**Temporal Granularity:** Connectivity decisions can be made per content/service request, or once per UE, for a particular spatiotemporal context, such as a set of available APs at a certain location and time. The selected granularity affects the associated control plane communication overheads and depends on the rate of user requests, as well as the expected data volume per request (e.g., a long YouTube video vs. a tweet).

## USE CASES

As revealed in the previous section, a series of design choices result in a corresponding multitude of potential ICCON mechanisms. In this section we put together the available pieces, providing a description of two example use cases to further illustrate ICCON functionality.

### MOBILE NETWORK OPERATOR SUPPORTED OFFLOADING

We first focus on the abundance of WiFi networks [2] and the ICCON-enabled WiFi AP selection for offloading purposes (Fig. 1). AP selection is guided by the availability of the desired content in the caches of each available access network. UEs locally build their profiles by inspecting the requested URLs in the generated HTTP requests (or the requested content names in Interest packets, in case of CCN/NDN networks [1]). A local *virtual* cache is used to track item popularity in a least frequently used (LFU) fashion, for instance, without locally caching the content itself. On the access network side, the cache indexes are used as the default network profiles.

As shown in Fig. 1, the ANDSF component of the mobile network operator is responsible for collecting and matching the UE and network profiles. To this end, a logically centralized controller established at each access network is responsible for regularly collecting cache indexes and *pushing* them to the ANDSF. On the user side, Bob is in a café in a central square of the city, turns on the WiFi interface of his device, and opens a video/music application (e.g., YouTube, Spotify, etc.) to listen again to the latest album of his favorite singer. The UE pushes its profile to the ANDSF, along with its coordinates or set of identified networks. The ANDSF identifies two APs/SSIDs (e.g., AP 1.1 and AP 3.1 in Fig. 1) leading to cached copies of the album. Taking into account additional information related to the current load at each AP, the ANDSF provides the UE with an ordered list of suggested APs/SSIDs. The UE proceeds with the association process (e.g., with AP 1.1), also taking into account the locally observed wireless conditions, such as signal-to-noise ratio (SNR).

### AUTONOMOUS CONNECTIVITY DECISIONS

In a second use case, we consider users looking for the latest news information, which is provided in the form of a service, that is, a news service named MyNews. The content for this service is provided by a news agency, which pushes the

news feeds to the online users of the corresponding smartphone application and certain ICCON-enabled WiFi APs around the city. The purpose of the UE-side application is to pull content when Internet access is available, and to guide the connectivity decisions of UEs in their quest for MyNews service feeds in ICCON-enabled environments.

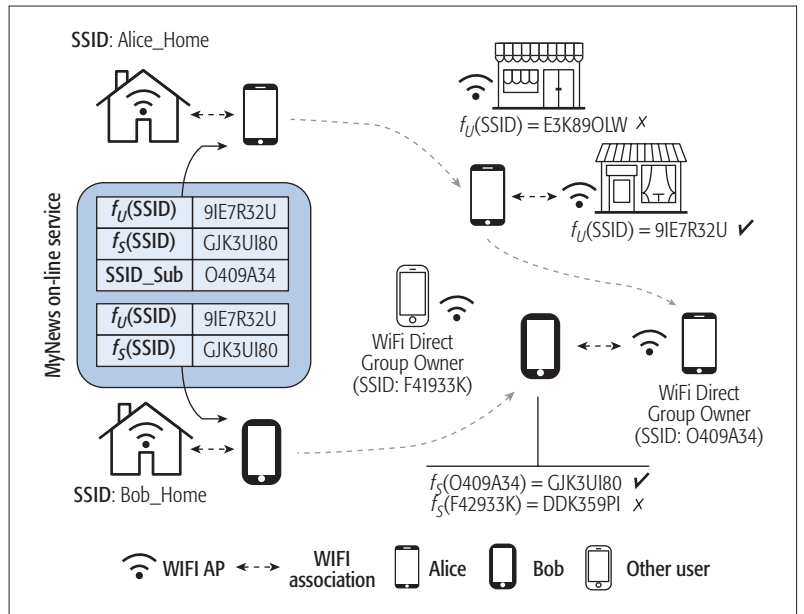
Alice checks the news while at home (Fig. 2). Later on, she roams through the city. As she has reached the quota of her cellular data plan, she turns on the WiFi interface of her device and opens the *MyNews* application to get any updates. Her application has previously downloaded a set of SSID values (or their algorithmic properties, e.g., functions  $f_U$  and  $f_S$  in Fig. 2) corresponding to WiFi-enabled *MyNews* repositories around the city. The ICCON-enabled WiFi manager of her device identifies one of the available SSIDs in one of the many beacons received from nearby APs and automatically connects using Alice's *MyNews* application user credentials. Alice selects the *Politics* category generating a request toward the AP, which responds with a data message.

At the same time, a set of SSIDs has been created to enable the sharing of user generated content between *MyNews* users participating a big music festival. Alice goes to the festival and is looking for any related information. Her UE enters the *autonomous* group formation WiFi Direct mode [7], beaconing an SSID (e.g.,  $SSID = O409A34$  in Fig. 2) that can be recognized by other *MyNews* UEs. Bob is in the area and has previously downloaded the festival program using his *MyNews* application. His UE senses Alice's SSID (e.g., since  $f_S(O409A34) = GJK3UI80$  in Fig. 2) and asks Bob if he is interested in sharing any information. Bob allows his UE to connect to Alice's group and provide the festival program. The application credits Bob's user account, enabling him in turn to search for and retrieve information during a football match the following week, for example, by providing him with a suitable SSID value matching a new  $f_S$  function defined for that particular event (Fig. 2).

## CHALLENGES

The description of the ICCON concept and the presented use cases implicitly reveal a wide range of challenges and open issues to be addressed. Here, we identify and discuss the most important of those.

**Profiling.** In the mobile offloading case, the selection of an access network is made based on the match between the UE and the network information profiles. Since network profiles correspond to caches or services in the network, it is essential to consider the scope of the profiling as expressed by the number of caches (or service instances) taken into account, as well as their location and size in order to assess the potential benefits of guiding UE decisions. At the same time, UE energy limitations, as well as performance aspects, impose important size limitations on the representation of UE profiles. Profile comparison needs to be fast in order to reduce: the corresponding energy consumption (when performed by a UE) and the time required to make a connectivity decision.



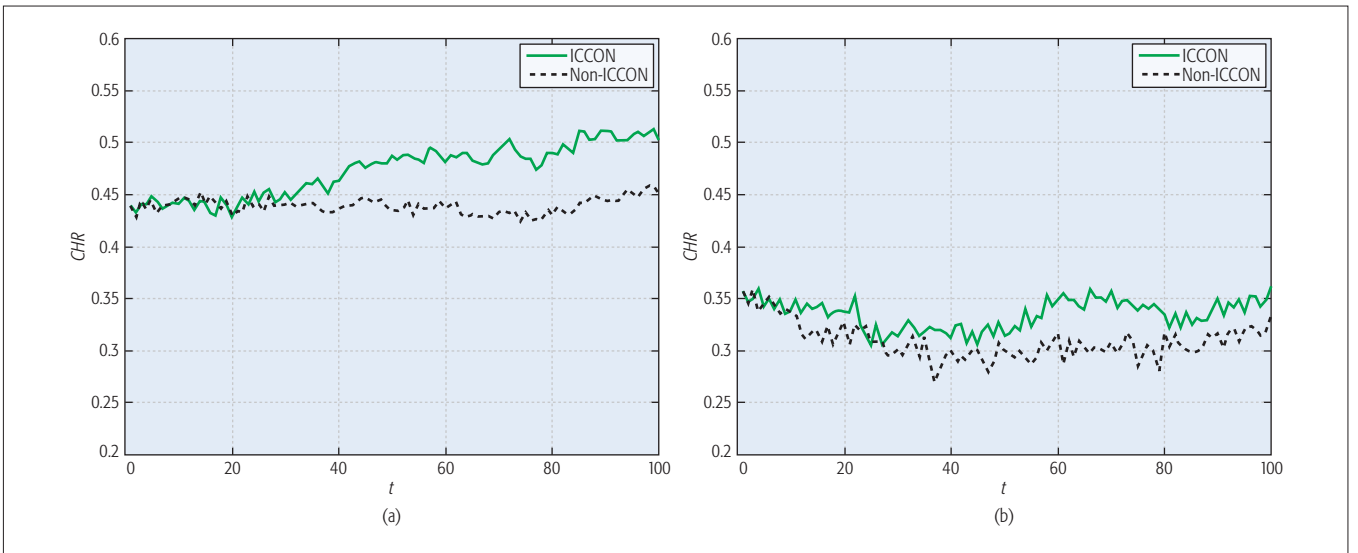
**Figure 2.** Autonomous connectivity decisions. The *MyNews* service provides Alice and Bob with functions  $f_U$  and  $f_S$ , which enable them to identify SSIDs denoting the availability of news updates and other users' interest in music festival information, respectively. Alice is also provided with an SSID value to be used for music festival information. Bob uses  $f_S$  to identify Alice's subscription to *MyNews* festival information, denoted by the corresponding WiFi Direct group. His UE joins the group to provide the festival program.

Set reconciliation methods (e.g., [13]) or the use of Bloom filters can be revisited in the ICCON context as potential approaches toward lightweight profile matching. The creation of UE profiles also poses challenges in accurately reflecting user preferences. Simple content item popularity may be augmented with contextual information, such as temporal characteristics or social information.

**Holistic connectivity management.** Information availability should be considered in the broader context of connectivity management, along with multiple other performance-related aspects, such as the wireless conditions (e.g., SNR), available data rates, and observed load on the devices providing the available information (e.g., caches). A careful fine tuning of the weight of each aspect is obviously required. To this end, using naming to expose information about the traffic type of the desired content/service (e.g., low-latency interactive traffic vs. bulk download) can also assist in decision making.

**Content/service placement.** UEs may provide indications of the desired information, implicitly allowing listening devices (e.g., APs) to collect information about the medium-/long-term demand for content/services in certain areas. This information can be subsequently used in combination with mobility patterns or social relations to make decisions on proactive content/service placement.

**Security and incentives.** Exposing UE profile information can raise privacy concerns. Careful encoding of user preferences is required, possibly calling for cooperation between information providers and users (i.e., agreeing on the naming scheme and the encoding method). Enabling



**Figure 3.** Impact of ICCON supported AP selection on CHR [ $N = 150, M = 10, C = 10^4, c = 5\%C, s = 0.8, \lambda_c = 0.01$  req/s,  $\lambda_v = 0.003$  users/s,  $U = 50, u = 10\%C, w = 0.65$ ]: a) AP selection remains the same across all UE content requests; time is measured in total number of arrival/departures; all APs are visible to all users; b) similar to a), but each user can associate with APs within a range of  $R = 30$  distance units.

access to content/services through corresponding applications that ensure common namespaces provides such a means. Moreover, truthful advertising/reporting of the available or desired information must be ensured so as to guarantee the overall system utility. Cryptographic signatures, pre-shared private/public key pairs, as well as reputation-based mechanisms may contribute in this direction. Furthermore, appropriate countermeasures are required for denial of service (DoS) attacks targeting the overwhelming of UEs with advertised information. Finally, in D2D scenarios, users must be incentivized to offer their resources. Enabling communication through smart applications also presents the potential for access control: reputation/credit management schemes linked to user application accounts and supported by applications' cloud-based back-ends.

**Naming granularity and spectrum sharing.** Setting up per-application or -service SSIDs links information naming granularity to use of the wireless spectrum. Also subject to information demand, very fine-grained naming can lead to overheads related to the support of high numbers of SSIDs, and broader medium access control. This calls for a closer look at the impact of naming granularity on the utilization of the wireless spectrum, and the sharing of the wireless medium by the various applications or services.


### PRELIMINARY RESULTS

Taking a first step in addressing some of the identified challenges, we engage in a preliminary investigation of the potential benefits of ICCON in our first example use case of mobile operator supported offloading. To this end, using Matlab-based simulations, we first explore the potential benefits in terms of cache hit ratio (CHR). Then we investigate the impact of the different cache aggregation levels on the characteristic time of cache objects and its relation to the inter-arrival time of user requests.

### IMPACT OF ICCON ON CACHE HIT RATIO

In our first setup, a set of  $N$  UEs, affiliated with a single mobile operator, visit a certain location (e.g., a shopping mall) where a set of  $M = 10$  WiFi APs offer access to the Internet. We assume that the APs are deployed in an orthogonal room/area of  $100 \times 40$  distance units. The APs are deployed in two rows (5 APs per row), where the rows are 20 distance units away, and each row is 10 distance units away from the edge of the room. Each AP is also 20 distance units away from the neighboring APs of the same row. Each AP is backed up by a single cache of size  $c$ . Users enter the considered area from random directions and activate their WiFi interfaces at randomly selected locations. Upon activation users remain at their location until they switch off their WiFi interface. While activated, UEs request content from a catalog of size  $C$ . Content popularity follows a Zipf-like distribution of slope  $s$ .  $U$  unique UE profiles are derived from the Zipf-like distribution. Each profile consists of a set of  $u$  unique, uniformly randomly selected items from the content catalog.

In this context, we first consider a scenario where  $N/3$  of the users initially visit the shopping mall and associate with a randomly selected AP. This connectivity decision is ICCON-agnostic and uniformly random, given the simplifying assumption of uniform wireless conditions. Upon association, each user generates content requests at a rate  $\lambda_c$ . We let caches stabilize so as to observe the cache performance in the absence of any ICCON-related mechanism. Subsequently, an arrival-departure process starts in which, at each step, one already connected user departs from the system, and a new one enters it. Each newly arrived UE now connects to the AP leading to the cache with the best fit against its profile. Between consecutive departure/arrival events we also let caches stabilize so as to observe the cache performance. The process has a rate  $\lambda_v$  and completes once all  $2N/3$  ICCON-supported UEs enter the system.



**Figure 4.** Impact of cache size and AP selection mode on CHR achieved by ICCON, [ $N = 150, M = 10, C = 10^4, c = 5\%C, s = 0.8, \lambda_c = 0.01$  req/s,  $\lambda_v = 0.003$  users/s,  $U = 50, u = 10\%C, w = 0.65$ ]; a) CHR for different cache sizes,  $c$  is expressed as a percentage of the total catalog size; each point represents the average CHR value observed throughout the arrival/departure process, along with the minimum and maximum values; b) AP selection is enabled per content request; time is measured in time slots of 100,000 content requests each.

The selection of the AP for the ICCON-supported UEs is based on a fit function  $F$  that takes into account the weighted ( $w$ ) average of both the UE-network profile match ( $f$ ) and AP load ( $l$ ), that is,

$$F = wf + (1 - w)l$$

$f$  is calculated as the ratio of a UE profile's content items found in the LFU-index of the corresponding cache, and  $l$  is calculated as

$$l = 1 - \frac{n_i}{\sum_{i=1}^M n_i},$$

where  $n_i$  denotes the number of users currently being served at AP  $i, i \in [1..M]$ .

We first focus on the investigation of the potential benefits of ICCON. For this reason we simplify our setup assuming that all APs are visible to all users, under the same wireless conditions (e.g., SNR), and we further ignore any interference between APs. Figure 3a shows the evolution of the CHR observed across APs for the entire lifetime of the aforementioned departure/arrival process. The figure compares the observed performance against an information-agnostic scenario where the newly arrived UEs associate with the least loaded AP ("non-ICCON" plot in Fig. 3a). Interestingly enough, we observe that when half of the initially randomly assigned UEs depart from the network (i.e., 25 UEs here), the ICCON mechanism starts improving the CHR, since from that point on, it manages to cluster the UEs at different APs based on their profile. ICCON increasingly improves the observed CHR until it reaches an increase of 10 percent against the non-ICCON case where the observed CHR remains relatively stable.

In order to further take into consideration the characteristics of the network topology in the association process, we next consider a scenar-

io where a UE can only associate with APs that are within a range of  $R$  distance units. Figure 3b depicts the evolution of the CHR for  $R = 30$  distance units. Due to the fact that each user has fewer connectivity options, both the ICCON and non-ICCON approaches perform worse compared to the previous scenario. However, the ICCON case still improves the observed CHR by approximately 17 percent against the non-ICCON case.

Figure 4a further illustrates the impact of the cache size on the perceived CHR (for the simple scenario of Fig. 3a). We see that ICCON increases CHR values compared to an information-agnostic AP selection scheme for the entire range of cache capacity values considered.

We proceed to further examine the impact of ICCON when connectivity decisions are taken per content request (e.g., long YouTube videos). We exploit the fit function  $F$  for each generated content request against each AP with the UEs eventually associating with the AP yielding the best fit. Figure 4b shows the observed CHR for 50 time slots of 100,000 requests each, for two scenarios:

- The "ICCON" scenario where all associations are driven by the fit function  $F$
  - The "non-ICCON" scenario where a UE selects the least loaded AP to associate with
- We see a significant improvement against the non-ICCON scenario, reaching a CHR of around 78 percent.

#### IMPACT OF CACHE AGGREGATION LEVEL

Although our preliminary results show an improvement of the achieved CHR, the overall effectiveness of ICCON is subject to a wide range of aspects also related to the dimensioning of the caching infrastructure. In this context, it is important to observe that connectivity decisions based on UE-cache profile matches become inefficient when the lifetime of items




Figure 5. Impact of aggregation level  $\alpha$  on  $r$  (Fig. 4a) and the CHR (Fig. 4b). The results correspond to  $C = 10^6$  and  $\lambda_c = 0.01$  requests/s. Cache sizes ( $c$ ) are expressed as a ratio of the catalog size.

in the cache is such that items are evicted before being requested. Here, we investigate the impact of the cache size, as well as the size of the aggregate request load, on the resulting effectiveness of an ICCON mechanism, as expressed by the relation between cache content lifetimes and user request rates. We express the aggregate request load with the *aggregation level* parameter  $\alpha$ , which denotes the number of UEs sharing a single cache. Intuitively, large  $\alpha$  values contribute to contention in the cache and correspondingly shorter lifetime values. Assuming a fixed request rate for all users ( $\lambda_c$ ), we employ Che's approximation [14] to compare the *characteristic time* ( $\tau$ ) of cached objects (i.e., the time spent in the cache until evicted in a least recently used manner) against the inter-arrival time of UE requests ( $1/\lambda_c$ ). For this purpose we define  $r = \tau\lambda_c$ . Obviously, ICCON becomes meaningful for  $r$  values greater than 1, that is, when the lifetime of objects in the cache is greater than the inter-arrival time of requests and the cached items are more likely to not have been evicted before they are requested.

Our setup involves a single cache, which is shared by a total user population size  $\alpha$ . Large  $\alpha$  values correspond to centralized caches shared by multiple end users connected to the corresponding access network. We focus on the behavior of the system at steady state once a connectivity decision has been made (i.e., in this scenario we do not consider user arrivals/departures). Figure 5 shows the  $r$  and CHR values for different aggregation levels ( $\alpha$ ) and cache sizes ( $c$ ). We see that large caches and low UE populations significantly increase  $r$  (note the logarithmic scale on both axes). We notice that an order of magnitude larger cache size has the same effect as an order of magnitude lower  $\alpha$ . Highly centralized caches result in low  $r$  values, lowering the expectations from a good profile match. However, as expected, Fig. 5b shows considerably lower CHR for small cache sizes, resulting in small UE benefits even if content is actually found in a cache, as indicated by the profile matching.

## SUMMARY AND CONCLUSIONS

In this article, we propose an extension to the ICN paradigm to further encompass connectivity decisions, so as to enable the efficient discovery and access of the desired information in the networking vicinity of mobile devices. The resulting ICCON concept can be applied in several different environments ranging from cellular offloading decisions to D2D communications. Although several technical enablers have been identified, a series of research challenges need to be faced in order to reap the benefits of ICCON. Our preliminary investigations demonstrate such benefits in the example case of cellular offloading, motivating further research efforts.

## ACKNOWLEDGMENTS

This work has been supported by the EC H2020 UMOBILE project (GA no. 645124) and the EPSRC INSP Early Career Fellowship (no. EP/M003787/1). V. Sourlas is supported by the European Commission through the FP7-PEOPLE-IEF INTENT project, (GA no. 628360).

## REFERENCES

- [1] G. Xylomenos et al., "A Survey of Information-Centric Networking Research," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 2, 2014, pp. 1024–49.
- [2] A. Akella et al., "Self-Management in Chaotic Wireless Deployments," *Proc. 11th Annual Int'l. Conf. Mobile Computing and Networking*, 2005, pp. 185–99.
- [3] A. Aijaz et al., "A Survey on Mobile Data Offloading: Technical and Business Perspectives," *IEEE Wireless Commun.*, vol. 20, no. 2, Apr. 2013, pp. 104–12.
- [4] S. Liu et al., "Exploring the Potential in Practice for Opportunistic Networks Amongst Smart Mobile Devices," *Proc. 19th Annual Int'l. Conf. Mobile Computing & Networking*, 2013, pp. 315–26.
- [5] J. G. Andrews, "Seven Ways That HetNets are a Cellular Paradigm Shift," *IEEE Commun. Mag.*, vol. 51, no. 3, Mar. 2013, pp. 136–44.
- [6] F. Bonomi et al., "Fog Computing and Its Role in the Internet of Things," *Proc. 1st Ed. MCC Wksp. Mobile Cloud Computing*, 2012, pp. 13–16.
- [7] D. Camps-Mur et al., "Device-to-Device Communications with Wi-Fi Direct: Overview and Experimentation," *IEEE Wireless Commun.*, vol. 20, no. 3, June 2013, pp. 96–104.
- [8] M. Amadeo et al., "Content-Centric Wireless Networking: A Survey," *Elsevier Computer Networks*, vol. 72, Oct. 2014, pp. 1–13.
- [9] G. Tyson et al., "A Survey of Mobility in Information-Centric Networks: Challenges and Research Directions," *Proc. 1st ACM Wksp. Emerging Name-Oriented Mobile Networking Design – Architecture, Algorithms, and Applications*, 2012, pp. 1–6.

- [10] 3GPP, "TS 23.402 V13.3.0, Technical Specification Group Services and System Aspects; Architecture Enhancements for Non-3GPP Accesses (Release 13)," Sept. 2015.
- [11] IEEE Std 802.11u, "IEEE Standard for Information Technology-Telecommunications and Information Exchange between Systems-Local and Metropolitan Networks-Specific Requirements-Part II: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 9: Interworking with External Networks," 2011.
- [12] D. Camps-Mur *et al.*, "Enabling Always on Service Discovery: WiFi Neighbor Awareness Networking," *IEEE Wireless Commun.*, vol. 22, no. 2, Apr. 2015, pp. 118–25.
- [13] A. Z. Broder, "On the Resemblance and Containment of Documents," *Proc. Compression and Complexity of Sequences*, June 1997.
- [14] H. Che *et al.*, "Hierarchical Web Caching Systems: Modeling, Design and Experimental Results," *IEEE JSAC*, Sept 2002.

## BIOGRAPHIES

KONSTANTINOS V. KATSAROS (k.katsaros@ucl.ac.uk) received his B.Sc. in informatics (2003), and his M.Sc. (Honours, 2005) and Ph.D. (2010) degrees in computer science from Athens University of Economics and Business (AUEB), Greece. His research interests focus on information-centric networking, smart grid communications, network function virtualization for 5G networks, and IoT interoperability. In 2013 he joined University College London (UCL), United Kingdom, as a research associate. Before that, he held positions at Télécom ParisTech, France, and AUEB, Greece.

VASILIS SOURLAS (v.sourlas@ucl.ac.uk) received his Diploma and M.Sc. degree in computer science from the Computer Engineering Department, University

of Patras, Greece, in 2004 and 2006, respectively. In 2013 he received his Ph.D. from the Electrical and Computer Engineering Department, University of Thessaly, Greece. In 2015 he joined UCL to pursue his two-year Marie Curie IEF fellowship. His interests are in the area of information-centric networks and future Internet.

IOANNIS PSARAS (i.pсарas@ucl.ac.uk) is an EPSRC Fellow at the Electrical and Electronic Engineering Department of UCL. He is interested in resource management techniques for current and future networking architectures with particular focus on routing, replication, caching, and congestion control. Before joining UCL in 2010, he held positions at the University of Surrey, United Kingdom, and Democritus University of Thrace, Greece, where he also obtained his Ph.D. in 2008. More information can be found at <http://www.ee.ucl.ac.uk/~ipsaras/>

Sergi René (s.rene@ucl.ac.uk) received his B.S. in electrical engineering and M.S. in network engineering from the Technical University of Catalonia (UPC), Barcelona, Spain, in 2006 and 2010, respectively. In 2015, he also received his Ph.D. from the Network Engineering Department at UPC. In 2015, he joined UCL, carrying out research on information-centric networks, congestion control, and mobility.

GEORGE PAVLOU (g.pavlou@ucl.ac.uk) holds the Chair of Communication Networks at the Department of Electronic Engineering, UCL, where he coordinates networks and services research activities. His research interests focus on networking and network management, including traffic engineering, content-based networking, autonomic networking, and software-defined networks. He has been instrumental in a number of research projects that have produced significant results with real-world uptake, and has contributed to standardization activities in ISO, ITU-T, and IETF.