F2C-Aware: Enabling discovery in Wi-Fi-powered Fog-to-Cloud (F2C) systems

Zeineb Rejiba, Xavi Masip-Bruin, Alejandro Jurnet, Eva Marín-Tordera
Advanced Network Architectures Lab (CRAAX)
Universitat Politècnica de Catalunya (UPC)
Barcelona, Spain
Emails: {zeinebr, xmasip, ajurnet, eva}@ac.upc.edu

Guang-Jie Ren
IBM
Almaden Research Center
USA
Email: gren@us.ibm.com

***Abstract*—Fog-to-Cloud (F2C) is a novel architectural framework, driven by the needs imposed by emerging ubiquitous applications and relying on the coordinated use of fog and cloud computing. Within such a setup, a suitable mechanism addressing the resource discovery problem, taking into account fog-specific characteristics (i.e. mobility, heterogeneity, etc.) is yet to be defined. On the other hand, the proliferation of Wi-Fi has been remarkably growing with inherent support for an interesting feature set. One such feature is to embed vendor-specific information within the beacon frame in order to provide customized Wi-Fi-based informational services. In this paper, we propose F2C-Aware, a discovery approach that leverages this feature to allow devices in Wi-Fi-powered Fog-to-Cloud systems to become aware of each other. The proposed approach takes into account the fact that devices in the fog are not necessarily connected to the same local network and thus it provides discovery prior to network association, which could help in obtaining time and energy savings.**

***Index Terms*—fog computing, F2C computing, resource discovery, 802.11 beacon stuffing**

I. Introduction

Hosting applications on cloud infrastructures continues to be a popular trend among different businesses and industries. However, more recently, efforts are being devoted to complement this computing paradigm by leveraging the yet-unused pool of resources available at the edge of the network, a concept freshly coined as fog computing [1]. The trend to the edge is mainly driven by the rise of new applications such as the ones in Internet of Things (IoT) and Cyber Physical Systems (CPS) fields, having specific QoS constraints that, centralized, cloud-only solutions cannot guarantee. These efforts have set the ground for the advent of different research initiatives, such as the reference architecture proposed by the OpenFog Consortium [2] and the Fog-to-Cloud (F2C) [3] concept, to be designed and developed by the European H2020 mF2C project [4]. Similar in their conception and main objective, this paper focuses on the latter, which is characterized by a layered stack of computing resources from the edge up to the cloud, organized in a hierarchical structure. Within this setup, the conventional cloud sits at the topmost layer, followed by the fog, comprised of heterogeneous computing-capable devices, whereas the edge devices running IoT and CPS services lie at the lowest layer.

Recognizing the dynamicity inherent to a F2C scenario, an efficient coordination and management of the provided resources is one of its major pillars, hence undoubtedly a must in the F2C design. However, this is not possible without a suitable strategy enabling mutual discovery among the set of devices taking part in the F2C scenario, referred to as F2C-capable devices. As fog computing is rather a new research area, contributions addressing the problem of resource/service discovery in fog (and thus F2C) computing are still rare. In addition, while observing related works in similar fields (e.g. edge computing and cloudlets), we notice that the proposed solutions are generally based on the assumption that the resources to be discovered are connected to the same local network, which is not necessarily true in the F2C scenario. Hence, a non LAN-centric discovery mechanism that occurs before network association is established is envisioned as the proper strategy for devices discovery in a F2C scenario.

In the meantime, the proliferation of Wi-Fi as the most remarkable wireless technology has been considerably growing. In fact, a recent study [5] published by the Wi-Fi Alliance indicates that more than three billion Wi-Fi device shipments are expected in 2017 and more than eight billion devices are currently in use around the world. Further, according to the same study, Wi-Fi has been identified as an ideal connectivity platform for the IoT. When operating in access point (AP) mode, Wi-Fi devices make use of built-in 802.11 mechanisms to advertise their capabilities using special management frames called beacons. Interestingly, apart from their original usage, beacons can be overloaded with additional information, a concept coined by [6] as beacon stuffing. This paves the way to a new set of value-added features, ranging from announcements of specific events and services supported by the network to whatever may be designed in the future for additional applications.

In light of this, in this paper, we put the focus on devices discovery, particularly tailored to the needs driven by a fog or F2C scenario, and we attempt to harness the hidden potential that beacon stuffing brings in and use it as a discovery solution for Wi-Fi-enabled F2C systems. In short, we propose to embed information about F2C support inside the 802.11 beacons transmitted by F2C-capable devices, thus making them all aware of each other, hence coining the name F2C-Aware.

The remainder of this paper is organized as follows. Section II presents an overview of different discovery solutions in edge/fog computing environments. Section III describes the need for a discovery solution in the envisioned F2C architecture. The proposed discovery solution is presented in Section IV. Finally, Section V concludes the paper highlighting future research directions.

II. Related work

Resource discovery is a widely-studied topic in a variety of fields. As a result, several discovery solutions have been proposed, each driven by the inherent needs of the studied environment. In this section, we present an overview of research contributions addressing resource discovery in Fog/Edge computing, as they are of high relevance to the scenario envisioned in F2C.

Authors in [7] propose an approach based on DNS-Service Discovery (DNS-SD, RFC 6763) and Multicast DNS (RFC 6762) in order to enable cloudlet discovery. When a client submits a DNS-SD query to the cloudlet's multicast address, the cloudlet replies with a DNS-SD response indicating the IP address and port number it is using. This approach assumes that the client and the cloudlet are part of the same local network, which may not always be the case in the fog computing context. The problem of discovering the best available surrogate for computation task offloading in an edge computing scenario is addressed in [8]. In fact, authors propose a system comprised of a set of distributed brokers, implemented in conventional home routers. Discovery is enabled by advertisements made by surrogates towards the brokers, providing them with information about their capacities. In order to facilitate state and information exchange between these brokers, a Distributed-Hash-Table (DHT) overlay structure is used. However, such a solution may be complex to manage in a fog context, characterized with inherent mobility, heterogeneity and volatility of resources. In [9], a decentralized peer-to-peer approach is used to allow mobile users to discover a fog node. The mobile users are assumed to be pre-connected through a Mobile Ad hoc Social Network (MASN), using built-in communication mechanisms (Bluetooth, Wi-Fi-Direct). This enables the peers to proactively exchange knowledge about the existence of neighboring fog nodes in their respective areas. The current F2C architecture is different in that it is based on a hierarchical model, and not a P2P one, as envisioned in the MASN-based discovery. In [10], authors present Foglets as a programming infrastructure for the geo-distributed continuum comprising the fog nodes and the cloud. In Foglets, a discovery server constitutes a partitioned name server that maintains a periodically updated list of fog nodes available at different levels of the hierarchy for a given geographic area. Although authors emulate the Foglets infrastructure, including the discovery server, no information is provided on how the name server it is based on is implemented. Finally, the recently published OpenFog Reference Architecture (OFRA) [2] emphasizes the need for an autonomous and automated discovery strategy. It further states that when a new fog system is added to a clustered fog deployment, it will broadcast information about its presence and thus become available to share the workload with the rest of the cluster. However, it does not specify to what network layer this broadcast belongs nor how it could be implemented in a real world deployment.

III. Problem Analysis

An overview of the layered hierarchy of the F2C stack of resources has been briefly presented in Section I. Within this setup, the addition of a fog layer (one or many of them, according to a certain policy yet to be defined) to be managed in a coordinated way with the conventional cloud layer, poses several challenges and increases the system's complexity. Therefore, in order to facilitate a much better management and organization of the whole set of resources, we consider the following entities (see [4] and also Fig. 1):

- Agent: An agent would be any device contributing resources to the system and having a F2C-specific management software installed. We distinguish a fog agent, residing at the fog level (FA in Fig. 1) and a cloud agent, hosted within a conventional cloud infrastructure. Each has a different instance of the software installed, tailored to the inherent needs of the involved F2C layer(s).
- Area: Agents are organized into areas, based on a predefined clustering policy out of the scope of this paper. Two areas are graphically shown in Fig. 1.
- Leader: Each area is managed by a leader (L1 and L2 in Fig. 1), which is a potentially high-capacity device, selected among the set of agents within the area, according to a policy yet to be defined.

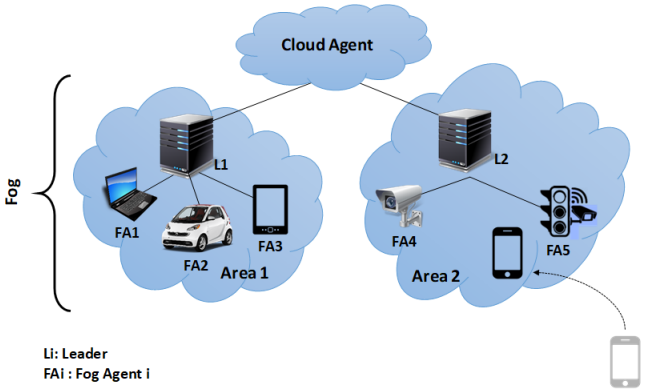

Fig. 1. F2C Architecture

One question that arises within such a context would be: "How does an agent willing to join a leader's area (thus willing to contribute resources to the F2C system) become aware of its presence, while initially not having any prior knowledge about such a presence?". To answer this question, a suitable discovery strategy should be implemented both at the leader's side and at the fog agent's side.

For this strategy to be efficient and fog-compliant, several aspects have to be taken into account, such as the inherent dynamicity and mobility of devices, resources volatility and

heterogeneity in the fog, and limited computing resources (as compared to the cloud layer). Additionally, real-time constraints imposed by certain applications and the need for core network traffic reduction are also worth considering.

We define discovery hereafter as a bidirectional mechanism where a leader makes use of special announcements to advertise its presence in the vicinity whereas fog agents [1] listen to these announcements in order to detect the leader's presence.

More specifically, in this paper we envision beacon stuffing as a suitable approach for devices discovery in the F2C scenario (certainly open to other fog-based architectures, as well), and thus we propose to use the 802.11 beacons to carry leader announcements, as it is described in Section IV.

IV. Proposed solution

The main rationale for the proposed solution leverages 802.11 beacon frames to be used to embed additional information. We first provide an overview of different beacon stuffing techniques that have been used in the literature. Next, we dig into the proposed solution, describing the technique to be used. Finally, we introduce preliminary efforts towards realizing a prototype, in order to prove feasibility of the proposed approach.

A. Beacon stuffing review

Embedding customized information into standard 802.11 beacon frames has been first described in [6] and referred to as "beacon stuffing". Authors present three possible techniques that can be used for this purpose. Each uses a different field of the original 802.11 beacon frame (Basic Service Set Identifier (BSSID),Service Set Identifier (SSID), Vendor-Specific Information Element (VSIE)), as depicted in Fig. 2.

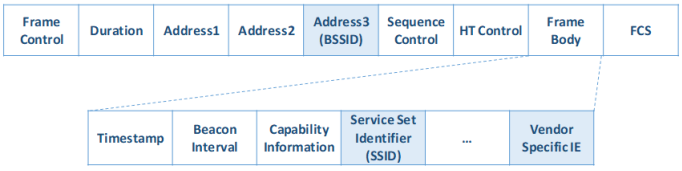

Fig. 2. 802.11 beacon frame format

Subsequent research efforts have also promoted the use of beacon stuffing for a wide set of use cases, such as broadcasting distress signals using the SSID field [11], advertising local opportunities using the VSIE field [12] and advertising privacy-awareness information [13] (also using the VSIE).

Aligned with [12] and [13], we also choose the VSIE-based information embedding technique to provide discovery in F2C systems, although it may be associated with the need for kernel space changes, unlike using SSID- or BSSID-based information embedding. In fact, using SSID-based information embedding, even users with little technical background could easily fake beacons. Besides, beacons with custom SSIDs would be seen by conventional Wi-Fi client applications as Wi-Fi networks, thus confusing legitimate Wi-Fi users. Additionally, the use of the BSSID for purposes other than address transportation is not compliant with the 802.11 standard [12]. On the other hand, the VSIE field has indeed been defined by the standard in order to carry custom vendor-specific information, which further backs our choice.

[1]From this point onwards, agent refers to the fog agent, and not the one in the cloud layer.

B. F2C-Aware: A Device Discovery approach for F2C systems

In this section, we describe the proposed discovery solution, referred to as F2C-Aware, both from the leader's perspective and from the agent's perspective.

1) Leader side: In order to make itself discoverable, a device serving as leader in the F2C architecture makes use of 802.11 beacons to advertise F2C-specific information. To achieve this, we propose to use the Vendor-Specific Information Element (VSIE) of the beacon frame, resulting in the structure depicted in Fig. 3.

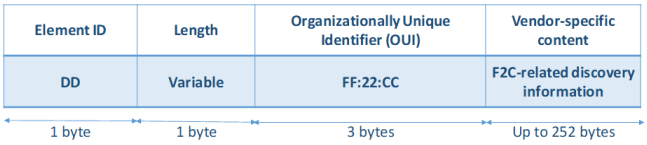

Fig. 3. F2C-specific information element (IE)

As defined by the 802.11 standard, the first byte of this information element (IE) represents the VSIE ID, which is 221 (DD in hexadecimal), followed by one byte for the length field. The Organizationally Unique Identifier (OUI) is included next. For the purposes of this work, we use the "FF:22:CC" OUI, since it has not been assigned to any organization yet. Although the OUI itself should be a sufficient indicator allowing an agent to become aware of the presence of F2C in its proximity, additional F2C-related information can be included in the subsequent vendor-specific content field. The following items could constitute information that may be included in this field:

- The leader ID, generated using the identification scheme in use within the F2C system.
- The types of tasks usually managed by the leader. This will allow the agent receiving the message, to be aware of how its resources are likely to be used. This should be represented in a format that an agent is able to interpret.
- Information indicating how urgently resources are to be needed, tentatively on a scale ranging from Low, Medium to High. For instance, high urgency would mean the resources the leader is already managing are overloaded. Therefore, if an agent takes this information into account and associates with the leader, thus contributing its resources to it, the leader's load may be alleviated.

These items will be provided as a list of consecutive F2C attributes. A first intuition regarding the representation of such attributes is to use the Type-Length-Value (TLV) encoding scheme, widely used in data communication protocols. "Type"

is one byte representing the attribute types listed above. "Length" is the length in bytes of the following field. "Value" is the actual attribute value.

All in all, once the necessary information has been included, the leader starts broadcasting F2C-enhanced beacons according to an interval, which is dynamically determined by the F2C system policies in place.

2) Agent side: In order to discover potential leaders in its proximity, an agent would start listening for F2C-enhanced beacons. Once such a beacon is detected (through the F2C OUI it contains) and knowing the F2C-specific encoding format, the agent will be able to decode the subsequent information in order to extract the advertised characteristics. If the characteristics match its preferences, then it will connect to the leader.

C. Preliminary efforts toward a prototype

In order to show the feasibility of our proposed discovery strategy in a real setup, initial efforts have been recently devoted to build a prototype using two Linux-based tools. The first one is the open source access point implementation *hostapd* which provides an option to configure static VSIEs to be carried within beacons and probe response frames. We use this option at the leader side to embed F2C-specific information after it has been encoded into the proper structure. The second tool we use at the agent side is the wireless configuration utility *iw*. It returns the list of IEs found in beacons and probe responses along with other information in the scan results. We use a simple python script that parses the *iw* scan results looking for F2C IEs and then decodes their content.

Using the proposed prototype, a set of experiments would be set in order to evaluate the efficiency of the proposed solution. From a leader perspective, we would be interested in measuring the overhead that beaconing will cause in terms of additional wireless channel usage in addition to evaluating whether beaconing is likely to cause a performance degradation to other leader functionalities. As for the agent side, the energy consumption and the processing usage caused by the beacon discovery will be evaluated under different scenarios (e.g. when an agent is continually exposed to beacon broadcasts, when there are no broadcasts and when there are alternate broadcasting and non-broadcasting periods) and compared to other state-of-the-art proposals.

V. Conclusion and future work

Exploiting the Wi-Fi capabilities provided out of the box in most of the modern smart devices brings endless possibilities for value-added services. Within this context, we show how the vendor-specific information element already present in 802.11 beacons can be enhanced with customized information to enable mutual discovery of devices in the recently proposed Fog-to-Cloud (F2C) computing paradigm, as well as other fog-based deployments. Apart from the novel use of an existing technology for the purpose of discovery in F2C, the proposed solution differs from other proposals in the literature in the fact that it provides a pre-association discovery. Since the beacon is received regardless of the existence of a network connection, no manual intervention from the agent side is required to discover whether a leader exists in its network, thus resulting in both time and energy savings. In the future, a performance evaluation of the proposed approach is planned in order to further prove its suitability to a F2C system. Additionally, we envision to investigate context-aware broadcasting and scanning frequencies, properly-adapted to a fog-context, in order to increase our proposal's efficiency.

Acknowledgments

This work was partially supported by the H2020 EU mF2C Project ref. 730929, and for UPC authors also by the Spanish Ministry of Economy and Competitiveness and by the European Regional Development Fund, under contract TEC2015-66220-R (MINECO/FEDER)

References

[1] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing - MCC '12*, p. 13, ACM, 2012.

[2] "OpenFog Reference Architecture : OpenFog Consortium." https://www.openfogconsortium.org/ra/. Accessed: 2018-01-02.

[3] X. Masip-Bruin, E. Marín-Tordera, G. Tashakor, A. Jukan, and G. J. Ren, "Foggy clouds and cloudy fogs: A real need for coordinated management of fog-to-cloud computing systems," *IEEE Wireless Communications*, vol. 23, no. 5, pp. 120–128, 2016.

[4] "mF2C project." http://www.mf2c-project.eu/. Accessed: 2018-01-02.

[5] "Wi-Fi Alliance® publishes 7 for '17 Wi-Fi® predictions — Wi-Fi Alliance." https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-publishes-7-for-17-wi-fi-predictions. Accessed: 2018-01-02.

[6] R. Chandra, J. Padhye, L. Ravindranath, and A. Wolman, "Beacon-stuffing: Wi-Fi without associations," *Proceedings - 8th IEEE Workshop on Mobile Computing Systems and Applications, HOTMOBILE 2007*, pp. 53–57, 2007.

[7] G. Lewis, S. Echeverría, S. Simanta, B. Bradshaw, and J. Root, "Tactical cloudlets: Moving cloud computing to the edge," in *Proceedings - IEEE Military Communications Conference MILCOM*, pp. 1440–1446, IEEE, 2014.

[8] J. Gedeon, C. Meurisch, D. Bhat, M. Stein, L. Wang, and M. Mühlhäuser, "Router-Based Brokering for Surrogate Discovery in Edge Computing," in *Proceedings - IEEE 37th International Conference on Distributed Computing Systems Workshops, ICDCSW 2017*, pp. 145–150, IEEE, 2017.

[9] S. Soo, C. Chang, and S. N. Srirama, "Proactive Service Discovery in Fog Computing Using Mobile Ad Hoc Social Network in Proximity," in *2016 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pp. 561–566, IEEE, 2016.

[10] E. Saurez, K. Hong, D. Lillethun, U. Ramachandran, and B. Ottenwälder, "Incremental deployment and migration of geo-distributed situation awareness applications in the fog," in *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems - DEBS '16*, pp. 258–269, ACM, 2016.

[11] A. Al-Akkad, L. Ramirez, A. Boden, D. Randall, and A. Zimmermann, "Help beacons," *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, pp. 1485–1494, 2014.

[12] S. Zehl, N. Karowski, A. Zubow, and A. Wolisz, "LoWS: A complete Open Source solution for Wi-Fi beacon stuffing based Location-based Services," *2016 9th IFIP Wireless and Mobile Networking Conference, WMNC 2016*, pp. 25–32, 2016.

[13] B. Könings, F. Schaub, and M. Weber, "Prifi beacons," *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication - UbiComp '13 Adjunct*, pp. 83–86, 2013.