

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

D2D Communications for Large-Scale Fog Platforms: Enabling Direct M2M Interactions

A.Virdis, C.Vallati, G.Nardini, G.Tanganelli, G.Stea, E.Mingozi (University of Pisa)

Fog computing is envisaged as the evolution of the current centralized cloud to support the forthcoming Internet of Things (IoT) revolution. While IoT devices will still communicate with applications running in the cloud, localized fog clusters, with IoT devices communicating with application logic running on a proximate fog node, will also appear. This will add proximity-based Machine-to-Machine (M2M) communications to standard cloud-computing ones, and call for efficient mobility management for entire fog clusters and energy-efficient communication within them. In this context, the LTE-Advanced (LTE-A) technology is expected to play a major role as a communication infrastructure to guarantee low deployment costs, native mobility support, and plug-and-play seamless configuration. In this work, we investigate the role of LTE-A in future large-scale IoT systems. In particular, we analyze how the recently standardized Device-to-Device (D2D) communication mode can be exploited to effectively enable direct M2M interactions within fog clusters, and we assess the expected benefits in term of network resources

and energy consumption. Moreover, we show how the fog-cluster architecture, and – in particular – its localized-communication paradigm, can be leveraged to devise enhanced mobility management, building on what LTE-A already has to offer.

Introduction

Future large-scale IoT platforms will be implemented through a multi-layered architecture [1], [14], [15]. IoT devices deployed pervasively within physical systems, such as sensors and actuators, will be accessed by IoT applications implemented in a distributed manner at different levels (see Figure 1, left). On one hand, the core application logic will largely run in a *cloud layer* implemented through powerful data centers, which, however, are placed far from IoT devices. On the other hand, simpler functions that do not require large computation and storage capabilities will run in a *fog layer* implemented by enhanced existing devices, such as network equipment, system control units, or even smartphones, that are deployed at the network edge, thus




Figure 1: Fog-based IoT Platform Architecture and Fog Clusters.

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

much closer to IoT devices. By bringing the execution of application logic to *fog nodes* closer to IoT devices, the fog layer enables low-latency communication, facilitates automatic resource discovery, and preserves context-awareness. The latter requirements are particularly relevant for Machine-to-Machine (M2M) applications, which involve closed-loop continuous interaction of the application logic with sensors and actuators.

In many future IoT/M2M scenarios, such as smart homes, factory-automation systems, a truck carrying sensorized boxes, or a patient provided with sensor and actuator wearables, fog nodes will be often deployed in direct range of communication with sensors and actuators, thus forming so-called *fog clusters* with IoT devices (see Figure 1, right). Fog clusters are (mobile) subsystems characterized by a preponderance of localized direct interactions between physically proximate devices – typically, having the fog node as an endpoint and an IoT device as the other.

One example is the IoT platform of a worldwide logistics company, distributing its goods from several warehouses. Goods are shipped in boxes equipped with sensors that monitor their internal status (e.g. temperature, for food or drugs) and position [10]. Global connectivity allows sensors to be reached worldwide by applications running anywhere in the cloud, which can track goods movements and status all the time. The company's fog clusters may be of several types: for instance, company warehouses, where a fog node implements automatic inventory of goods, or transport (e.g., trucks or ships), where boxes will interact with a fog node that runs closed-loop control logic to dynamically control the temperature read from sensors. Another example is a smart-health application [3]: in this case, patients wear sensors to monitor biometric data, and these sensors can transmit data towards control systems running into the cloud all the time and anywhere. However, IoT devices on patients residing within a hospital may be associated to fog clusters, with a fog node implementing fast-reaction control logic (e.g. to alert medical personnel in case of emergency).

In both these examples, scalability in terms of number of devices and fog clusters, mobility (possibly of entire fog clusters, as in the case of logistics) and energy-efficiency for battery-powered, constrained devices are key requirements. Large-scale IoT systems, in fact, will potentially include many such fog clusters (possibly counting thousands of devices themselves). In this context, the LTE-Advanced (LTE-A) cellular technology can play a major role, as it is able to provide seamless ubiquitous connectivity to IoT devices and fog nodes, whether implemented on devices, such as smartphones, mobile gateways (e.g. installed on public transport), or dedicated IoT devices [9]. In addition,

the LTE-A network can be exploited to enable direct communication between IoT devices and fog nodes in the same cluster.

Currently, LTE-A can offer ubiquitous connectivity to IoT devices: sensors and actuators can connect to the LTE-A network as User Equipments (UEs), which can be reached by IoT applications running anywhere – and, specifically, in the cloud. On top of this, the Device-to-Device (D2D) mode has been recently introduced as an additional LTE-A feature, and is also considered a key functionality to meet energy and spectral efficiency requirements for future 5G systems. D2D enables direct data-plane communication between proximate UEs without relaying at the eNodeB. This is ideal to support intra-cluster communications efficiently in fog clusters. Using a D2D-empowered LTE-A to support intra-cluster communications would make LTE-A the only broadband technology to implement both backhaul and proximity-based connectivity, using a single interface: IoT devices can be reached by IoT applications running anywhere using the classic Device to Infrastructure (D2I) mode and, at the same time, can interact with local IoT applications running on fog nodes placed in proximity using the D2D mode.

Using a single network technology brings clear advantages, the first of which is to remove all the interoperability problems that multiple technologies would create. Moreover, LTE-A already offers plug-and-play integration, embedded security, large-scale availability, native support for mobility: thus LTE-A network operators would be able to provide plug-and-play IoT solutions to the end user, with little, if any, modifications to their infrastructure.

In this work, we discuss how a large-scale IoT system that includes fog clusters can be supported by the LTE-A technology. In particular, we show how D2D can be exploited to implement local M2M communications within fog clusters. The benefits of D2D from the network operator standpoint are assessed by means of simulations, showing that D2D entails a more efficient usage of network resources and reduces energy consumption in the infrastructure.

Integrating IoT devices and fog nodes on a large scale will also present significant challenges for the LTE-A network itself [8]. Besides the key problem of scalability, one major issue is mobility management for large groups of mobile sensors and actuators. For this, LTE-A already offers built-in mechanisms that can be used as baselines to construct more advanced solutions that meet the specific characteristics of IoT systems. We thus propose a mobility enhancement that leverages the communication pattern of fog clusters. Its potential gain is evaluated numerically and via simulation.

It is worth mentioning that the fog computing paradigm has also been proposed to implement specific LTE-A radio access network (RAN) functionalities. This is called fog-RAN, or F-

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




Figure 2 High-level illustration of large-scale IoT systems connected through the LTE-A network

RAN for short, and consists in moving caching and signal processing closer to the edge [12] with respect to a cloud-RAN solution. The goal of this paper is not to discuss how to move parts of the LTE-A RAN to the fog, instead, how to support fog clusters using RAN-specific communication mechanisms.

An Architecture for Large-scale IoT Systems using LTE-A

This section describes a large-scale layered IoT architecture in which fog computing is exploited to build a runtime environment deployed close to the IoT devices. This provides support for the execution of applications that require low-latency and context-based interactions. Specifically, we discuss how this architecture can be supported by the LTE-A technology. The overall architecture showing the structure of both the layered IoT system and the LTE-A network is illustrated in Figure 2. The LTE-A RAN, consisting of several eNodeBs each one covering a large area, provides ubiquitous connectivity to devices, which can be seamlessly reached by applications running in the cloud. As the number of connected devices and their bandwidth demand increase, novel solutions, such as femto-, micro- and pico-cells, can be locally deployed to increase the capacity per square meter and guarantee network scalability. Fog nodes will be installed close to IoT devices, to execute the simple application logic that requires direct interaction with IoT devices. Although fog nodes can be implemented on heterogeneous devices, it is expected that most will be implemented through smartphones or network equipment (e.g. LTE-A home gateways or pico-cells) that will connect to the LTE-A network as either UEs or as eNodeBs. In order to achieve direct interactions with proximate IoT devices, explicit support from the LTE-A network can be

leveraged. To this aim, the D2D communication mode appears to be a promising technology. Although the exploitation of D2D communications for IoT devices has been already envisaged [6], its usage to connect fog nodes and IoT devices has been scarcely explored and presents specific challenges [4], as highlighted in the next sections.

Exploiting D2D interactions within fog clusters

M2M communications

To ensure self-management and configuration capabilities that exempt IoT systems from human intervention, different levels of discovery services for sensors and actuators are included in all IoT systems. IoT applications running in the cloud usually rely on centralized directory services, where IoT devices register to advertise their existence and capabilities. IoT applications can interact with devices seamlessly, regardless of their location. In the logistics use-case, for example, a monitoring IoT application running in the cloud may want to connect to all the sensors to check their status. However, IoT applications running on fog nodes will need to discover devices in proximity. Fog nodes are, by definition, installed close to IoT devices to support applications that require direct access to them or to provide location-based services. This proximity discovery will be mandatory to discover local IoT devices in an opportunistic manner. In our logistics use-case, a control IoT application running on the fog node installed in a moving truck would periodically discover all the sensors installed on the boxes in its cargo.

To this aim, IoT protocols usually define a distributed procedure that relies on broadcasting probe messages to the local

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

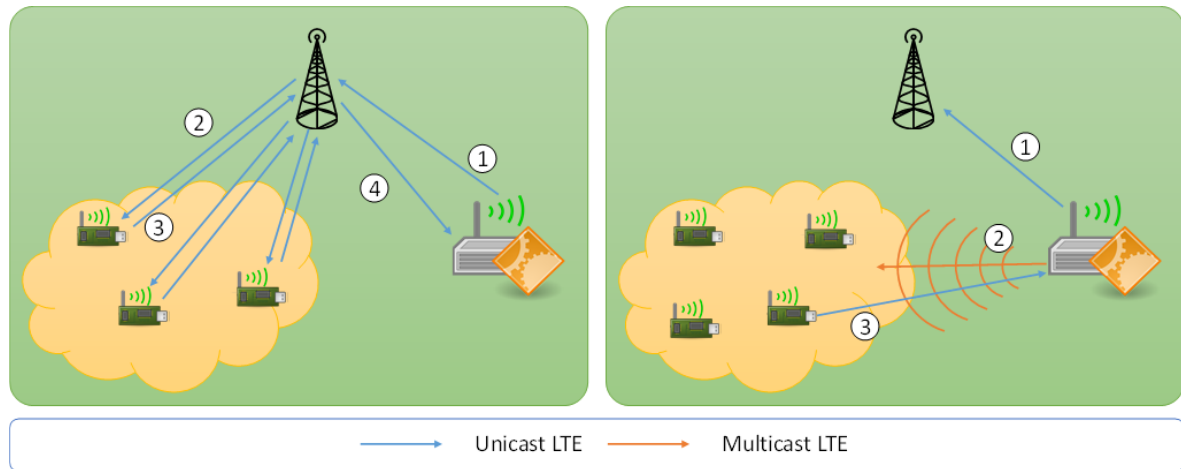


Figure 3 Legacy discovery (left) and D2D-based discovery (right).

network (or to all the devices within transmission range in case of wireless connectivity) to obtain notification messages from existing IoT devices. The Constrained Application Protocol (CoAP), for example, defines a Service Discovery procedure to discover all the sensors in its network: the IoT application (the CoAP Client) broadcasts discovery messages through the network, and every sensor (the CoAP Server) that receives the probe replies with a notification message advertising its presence. Then, the Client can run another procedure, called Resource Discovery, to obtain the list of the services that each IoT device offers.

Exploiting D2D connections

The discovery procedures described in the previous section rely on multicast communications, whose support from the underlying network is mandatory. However, the existing LTE-A standard, based on D2I mode, lacks support for proximity-based broadcast/multicast communications. The closest available feature is the Multicast Broadband single-frequency network (MBSFN), which allows an eNodeB to transmit broadcast signals over a *tracking area*, possibly consisting of several cells. However, MBSFN was envisaged to deliver services such as Mobile TV, and is unfit for proximity-based multicast for the following reasons: first, it only allows the eNodeB to send broadcasts, hence UEs would still need to use the latter as a relay to perform discovery. Second, it is inflexible, as MBSFN transmissions are scheduled over long periods, i.e. tens of seconds, and reach large areas, a tracking area easily being in the order of square kilometers.

The lack of a built-in proximity-based multicast/broadcast transmission mechanism could, in principle, be overcome using UE position information, e.g. obtained via GPS, and the eNodeB

as a relay, as illustrated in the left part of Figure 2. Position information could be exploited to obtain a proximity list for a given UE, i.e. the list of UEs in a predefined proximity area. Consider the following example, with reference to the left part of Figure 3: whenever an application at UE x needs to perform a proximity-based transmission, it will send the message to its lower layers, targeting a specific multicast group at the IP level, which will be configured by the network to be sent to the eNodeB (step 1). The latter will recognize the destination as multicast and will then forward a copy of the message (step 2) towards each UE in the proximity area of x . Each receiving UE will in turn send a response to x using again the eNodeB as relay (steps 3 and 4). This mechanism is clearly faster and more flexible than MBSFN, and allows an application to define its proximity area. However, it still requires the eNodeB to relay every transmission, both unicast and multicast, and to duplicate the original message for each UE in the proximity list, with a non-negligible cost in terms of transmission resources, which significantly limits the scalability of the system. Moreover, it requires a parallel architecture to obtain, communicate and manage UE positions.

In the latest LTE-A releases, D2D communications are introduced, in the form of proximity services, i.e. multicast communications originated at a UE, that reach other UEs in its proximity. D2D allows nearby UEs to communicate with a single hop, i.e. without the eNodeB acting as the relay in a two-hop path. The eNodeB, however, still participates in the signaling and maintains control of resource scheduling, issuing transmission grants, in a network-controlled fashion. A sender UE will request a D2D grant to the eNodeB, specifying a target ID (e.g., a group ID) within the MAC header, which allows potential receivers to filter packets at the MAC layer. Given its capabilities,

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




Figure 4 Legacy handover (left) and proxy handover (right)

this technology appears to fit the requirement of local communication from IoT applications. Consider a scenario with a fog node and some IoT devices equipped with temperature sensors. The network operator may configure a specific proximity service for temperature monitoring, allowing the interaction between these entities, and configure such service to be associated to a specific multicast IP address. The discovery procedure with D2D communications is shown in the right part of Figure 3. As a first step, the fog node requests resources to the eNodeB for transmission within the given proximity service (step 1), i.e. to the multicast IP address. The eNodeB in turn schedules resources according to its policies, and signals a grant to the fog node. At this point, the fog node performs a D2D multicast transmission (step 2), specifying the group-ID of the targeted proximity service. IoT devices in proximity receive the multicast message, infer from the group-ID that they are among the intended receivers, and decode the associated data. Finally, IoT devices reply to the fog node (step 3). Such reply messages can be transmitted using D2D, either in the same multicast manner as for the request, or in unicast, i.e. setting the fog node as a target, or, instead, it can be transmitted in D2I, using the eNodeB as a relay. Using D2D multicast for discovery also favors network scalability. First, by having UEs communicate directly, the eNodeB significantly reduces the number of transmissions that it has to perform, hence consumes less power. Second, D2D transmissions occur at a reduced power, hence generate little interference outside the proximity area. This latter characteristic in particular favors spatial reuse of frequency resources: communications that occur within two well-separated proximity areas can take place on the same frequencies. The eNodeB can leverage spatial reuse to coordinate the scheduling of proximity-based

transmissions, thus reducing the overall cell load and increasing scalability [7].

Once the discovery procedure is completed, the IoT application on the fog node can communicate with individual IoT devices using unicast D2D transmissions. These have not been standardized by 3GPP yet. However, the wide body of literature on the subject shows that they are promising for this kind of applications (see, e.g., [4]). In particular, unicast D2D transmissions are faster than D2I ones, due to the single-hop path, and they do not consume energy at the eNodeB.

It is important to highlight that, in some cases, IoT devices are required to communicate directly with the cloud, e.g. for remote monitoring or historical data collection. In these cases, in the proposed architecture the devices can still exploit traditional D2I transmissions to communicate directly with the cloud without additional overhead.

Addressing Mobility using D2D

Mobility is considered a major challenge for large-scale IoT systems [8], and the fact that fog clusters will likely insist on dense LTE-A networks, with pico- and femto-cells, exacerbates the problem. In such a scenario in fact, mobile IoT systems will trigger frequent, massive-scale handovers, which will affect the control plane, and in particular mobility management. LTE-A natively supports UE mobility. However, this feature can be enhanced and made more effective, to meet the demands of large-scale IoT systems. The key to its enhancement is to leverage the peculiar characteristics of these IoT systems, in particular, intra-fog-cluster communications, in a cross-layer approach.

As already anticipated, fog clusters will in general be mobile. However, the introduction of logic running close to IoT devices will bring a new mobility pattern. In fog clusters, fog nodes

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

<i>Parameter</i>	<i>Value</i>
Carrier frequency	2 GHz
Bandwidth	10 MHz (50 RBs)
eNodeB Tx Power	40 dB
UE Tx Power (in both D2D and D2I mode)	20 dB
eNodeB base power consumption	260 W
eNodeB power consumption per transmitted RB	3.76 W/RB

Table 1 Main simulation parameters.

are expected to move along with IoT devices directly connected, e.g. in public transport vehicles. This will entail large groups of IoT devices moving simultaneously, some of which - fog nodes - will need to keep alive their local D2D connections, in addition to the regular handover operations. This new mobility pattern is called *mobile fog* [2] and will bring new challenges to the LTE-A mobility management. Current handover operations, in particular, will require additional support to allow fog nodes to preserve D2D connections established with IoT devices in proximity along with the LTE-A backhaul connectivity.

This group mobility will represent a challenge for the regular handover operations. A large number of devices moving together, hence crossing cell borders simultaneously, will trigger many near-simultaneous handover requests, which can impair the performance. Moreover, this group of UEs will require specific support to preserve the resources allocated for D2D communications when moving from a cell to another.

To address these issues, two concurrent optimizations should be pursued. First, the network should exploit context information provided by the IoT system to implement specific solutions for IoT group mobility. For instance, the fog node could coordinate with the eNodeB, communicating the number of clients currently subscribed and/or mobility information, such as direction, expected destination of the group and the amount of resources allocated to D2D communications. The network can exploit such information to detect potential handovers of large-scale groups of UEs and take proactive action. For example, part of the handover procedures might be proactively triggered or even delayed to avoid simultaneous handover requests from many devices. Moreover, by gaining information on the group, such as its size or its expected traffic volume, the network may also start to pre-provision resources even before the actual handover procedure starts, thus potentially resulting in a smoother cell transition. Second, the fog node may act as handover relay for all its subscribers, i.e. receiving handover information common to the whole group and propagating it locally, e.g. again exploiting D2D multicast capabilities and reducing the amount of signaling traffic.

Normal handover operations are generally started by the UE as soon as it detects a *target* eNodeB having a higher signal quality than its serving one (the *source* eNodeB). This will trigger a

handover procedure, which requires several message exchanges among these three entities [13] and can be split in two main phases, as shown in Figure 4 (left). First, the source eNodeB sends a handover request to the target eNodeBs, which acknowledges it after a successful admission control. These messages are typically exchanged through the X2 interface defined by the 3GPP standard. Then, a connection reconfiguration message is sent to the UE using a downlink (DL) transmission, to notify both UE- and cell-related configuration parameters. A *proxy handover* can be envisaged to reduce the number of communications when UEs belonging to a fog cluster perform simultaneous handover, as shown in Figure 4 (right). In this case, the involved eNodeBs exchange *aggregate*, cluster-wide handover requests and acknowledgements. Then, the fog node acts as a proxy, receiving the connection reconfiguration message (2a in Figure 4, right) and transferring the information to the whole group using either a unicast or multicast D2D transmission (2b in Figure 4, right). Note that the format of the connection reconfiguration message must be modified in order to support multiple UEs, e.g. turning UE-specific fields into vectors. Both approaches are potentially efficient and can significantly contribute to the system scalability. However, they are clearly challenging and require both standardization and research efforts.

Performance evaluation

To evaluate the expected benefits of D2D multicast for local discovery and the proposed optimization for mobility, we present simulation results obtained using SimuLTE [5], an OM-NeT++-based system-level simulator. Simulation parameters are summarized in Table 1. We consider a scenario where three fog clusters are served by one eNodeB. Each fog cluster has an increasing number of IoT devices, randomly placed within a 50m-radius circle centered at the fog node. The latter is located at a distance of 250m to the eNodeB. Periodically, the fog nodes start a discovery procedure, by sending a discovery message. IoT devices that receive the message send a reply using a unicast transmission. With D2D, the discovery message is sent using a multicast D2D transmission. With D2I, the fog node sends its message to the eNodeB, which in turn relays it to all its connected users, via unicast DL transmissions. D2D transmissions are allocated using the transmission resources in the uplink (UL) part of the LTE spectrum. The latter is in fact less loaded than the DL part (due to the well-known traffic asymmetry) and allows better overall channel quality [11]. Confidence intervals at 95% level are shown only when visible.

Figure 5 shows the percentage of UL resources (i.e., Resource Blocks in the UL subframe) saved when D2D multicast is employed. The saving increases with device density: in fact, proximity to the fog node allows devices to transmit with higher

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

modulations, thus occupying fewer resources. Additional savings can be achieved by enabling spatial reuse among non-interfering transmissions, as shown in the figure. In particular, this scenario enforces frequency reuse among devices belonging to different fog clusters. However, such benefits are limited due to the low traffic rate generated by the devices.

Figure 6, instead, compares the power consumed by the network (i.e. the eNodeB) with D2I and D2D transmissions. The power consumption is computed according to the models provided in the EARTH EU project, where power is an affine function of the transmitted Resource Blocks (RBs) in the DL leg. The D2I scenario requires more RBs as the device density increases, resulting in more consumed power. On the other hand, when D2D is exploited, the power consumed by the network stays constant, as no DL transmission occurs. Thus, supporting localized IoT communications through D2D is beneficial for the operator, and allows its network to scale to higher numbers of devices through reuse.

As for the devices themselves, the fact that D2D allows higher modulations due to proximity (or the same modulation at a lower transmission power) increases energy efficiency, hence – indirectly – battery lifetime. Another point to be considered is that, in mobile fog clusters, a fog node and a communicating IoT device are likely to be stationary with respect to each other: thus the quality of their D2D channel will be more stable over time than that of D2I channels. Let us consider a scenario where one fog node and one IoT device are moving along a road, at 25 m/s. Different eNodeBs are deployed along the trajectory. Figure 7 shows the channel quality indicator (CQI) reported by the IoT device and measured both in D2D and in D2I with the best serving cell. We observe that the channel quality of the D2I communication fluctuates, deteriorating when the device approaches the cell border. Instead, the CQI of the D2D communication stays at 15, since the relative distance between the fog node and the device is constant over the entire simulation.

As far as mobility optimization is concerned, we consider a fog cluster composed of one fog node and N UEs, which perform handover from one cell to the neighboring one. Table 2 reports the comparison between legacy and proxy handover in terms of number of signaling messages. With legacy handover, each UE (including the fog node) performs the handover autonomously, hence $N+1$ messages are required for each phase of the procedure. On the other hand, proxy handover requires only two messages to be sent along the X2. Also, the DL subframe of the radio interface is offloaded using proxy handover, since only one DL message (from the eNodeB to the fog node) is needed, independently of the number of the number of UEs. Relaying from the fog node to the cluster’s members can be performed using either N unicast D2D messages or one multicast D2D message.

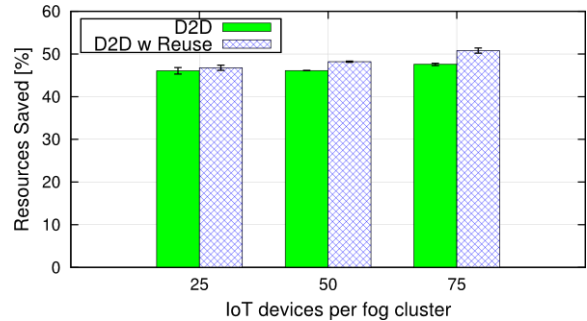


Figure 5 Resource saving using multicast D2D, with and without spatial reuse.




Figure 6 Power depleted by the network with D2I and D2D-based discovery.




Figure 7 Reported CQI using D2I and D2D in a mobile fog cluster




Figure 8 Number of control messages per minute using legacy and proxy handover

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

Type of handover	Phase 1		Phase 2
	Handover request msgs	Handover acknowledgement msgs	Connection reconfiguration msgs
Legacy	$N+1$ (over X2)	$N+1$ (over X2)	$N+1$ (DL)
Proxy, unicast D2D	1 (over X2)	1 (over X2)	1 (DL) + N (Unicast D2D)
Proxy, multicast D2D	1 (over X2)	1 (over X2)	1 (DL) + 1 (Multicast D2D)

Table 2 Number of messages required in a legacy and proxy handover

Despite the higher number of messages, the former approach might benefit from more reliable unicast D2D transmissions. In order to provide a quantitative analysis of the advantage of the proposed handover mechanism, let us consider a highway scenario where eNodeBs are deployed along the road to provide vehicles with continuous connectivity. A number of trucks travel along the highway, each one carrying a fog cluster consisting of goods equipped with sensors for tracking purposes and a device acting as fog node. We consider an increasing number of trucks per minute traveling along the highway, with different number of devices per truck. Figure 8 reports a comparison between legacy and proxy handover, in terms of control messages per minute per cell. Results show that the required overhead depends on the road traffic and, using legacy handover, it may become non-negligible. On the other hand, the proposed handover mechanism with unicast D2D messages reduces the traffic load on the X2 and DL connections, whereas a significant number of D2D messages still needs to be sent. However, D2D transmissions occur on the UL subframe, which is likely to be less loaded than the DL one. Proxy handover with multicast messages can potentially reduce the number of messages by up to 95%.

Conclusions and future works

In this paper we showed how D2D communications can be exploited to support fog computing nodes deployed in LTE-A networks close to IoT devices. Specifically, multicast local communication, generally exploited by IoT applications to discover devices in proximity, can be implemented through D2D interactions, and this allows an operator to save power and reuse resources. Moreover, we have proposed an enhancement of the handover signaling of LTE, which still leverages D2D interactions, to increase its scalability with large-scale mobile fog clusters.

Future work will entail investigating other key issues of the presented LTE-based architecture. For instance, LTE allows UE to save power through *discontinuous reception (DRX)*. Whether and how this technique can be used to mitigate the problem of *sleepy nodes* in IoT is subject of ongoing investigation. Moreover, cross-layer synergies between network- and application-layer mechanisms, such as e.g. CoAP and LWM2M, are being investigated.

References

- [1] F. Bonomi, R. Milito, J. Zhu, S. Addepalli. 2012. "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). ACM, New York, NY, USA, 13-16.
- [2] K. Hong, D. Lillethun, U. Ramachandran, B. Ottenwalder, and B. Koldehofe. "Mobile fog: a programming model for large-scale applications on the internet of things". In Proceedings of the second ACM SIGCOMM workshop on Mobile cloud computing (MCC '13). ACM, New York, NY, USA, 15-20.
- [3] M. Ahmad, M. Bilal Amin, S. Hussain, B. Ho Kang, T. Cheong, S. Lee, "Health Fog: a novel framework for health and wellness applications", The Journal of Supercomputing, Springer, 1–19, 2016
- [4] C. Vallati, A. Viridis, E. Mingozzi and G. Stea, "Exploiting LTE D2D communications in M2M Fog platforms: Deployment and practical issues," Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on, Milan, 2015, pp. 585-590.
- [5] A. Viridis, G. Stea, G. Nardini, "Simulating LTE/LTE-Advanced Networks with SimuLTE", DOI 10.1007/978-3-319-26470-7_5, in: M.S. Obaidat, J. Kacprzyk, T. ˆOren, J. Filipe, (eds.) "Simulation and Modeling Methodologies, Technologies and Applications", Volume 402 of the series Advances in Intelligent Systems and Computing, pp. 83-105, Springer, ISBN 978-3-319-26469-1, 15 January 2016
- [6] O. Bello; S. Zeadally, "Intelligent Device-to-Device Communication in the Internet of Things," in IEEE Systems Journal, vol.PP, no.99, pp.1-11
- [7] A. Asadi, Q. Wang, V. Mancuso, "A survey on device-to-device communication in cellular networks," IEEE Comm. Surveys and Tutorials., vol.16, no.4, pp.1801,1819, Fourth quarter 2014
- [8] F. Ghavimi and H. H. Chen, "M2M Communications in 3GPP LTE/LTE-A Networks: Architectures, Service Requirements, Challenges, and Applications," in IEEE Communications Surveys & Tutorials, vol. 17, no. 2, pp. 525-549, Second quarter 2015.
- [9] "LTE evolution for IoT connectivity", Nokia white paper, 2016
- [10] M. Witte, N. Hossain, "smaRT ULD - a low power approach for smart systems integration enabling the air freight container of the future", European Technology Platform on Smart Systems Integration (EPoSS Annual Forum), http://publica.fraunhofer.de/eprints/urn_nbn_de_0011-n-1433356.pdf
- [11] Yu, C. H., Tirkkonen, O., Doppler, K., & Ribeiro, C. (2009). On the performance of device-to-device underlay communication with simple power control. In Proceedings of IEEE VTC Spring, (pp. 1–5). April 26–29, 2009.
- [12] M. Peng and K. Zhang, "Recent Advances in Fog Radio Access Networks: Performance Analysis and Radio Resource Allocation," in IEEE Access, vol. 4, no. , pp. 5003-5009, 2016.

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

- [13] K. Dimou et al., "Handover within 3GPP LTE: Design Principles and Performance," 2009 IEEE 70th Vehicular Technology Conference Fall, Anchorage, AK, 2009, pp. 1-5.
- [14] S. Sarkar; S. Chatterjee; S. Misra, "Assessment of the Suitability of Fog Computing in the Context of Internet of Things," in IEEE Transactions on Cloud Computing , vol.PP, no.99, pp.1-1 doi: 10.1109/TCC.2015.2485206 .
- [15] Li, Wei, et al. "System modelling and performance evaluation of a three-tier Cloud of Things." Future Generation Computer Systems (2016).

Antonio Viridis (a.virdis@iet.unipi.it) is a post-doc researcher at the University of Pisa, where he obtained his MSc degree in Computer System Engineering in 2011, and his PhD in Information Engineering in 2015. His research interests include Quality of Service, scheduling and resource allocation in wireless networks, network simulation and performance evaluation. He has been and is currently involved in national, EU-funded and industry-funded research projects. He coauthored six patents and 20 peer-reviewed papers in the field of cellular network modeling and algorithms.

Carlo Vallati (c.vallati@iet.unipi.it) is Assistant Professor at the Department of Information Engineering of the University of Pisa. He received a Master's Degree (magna cum laude) and a PhD in Computer Systems Engineering in 2008 and 2012, respectively, from the University of Pisa. In 2010, he visited the Computer Science department of the University of California at Davis. His main research interests include next generation broadband networks, wireless mesh networks, sensor networks, M2M communications. He is co-author of +30 peer-reviewed papers in international journals and conference proceedings. He has been involved in the project BETaaS, Building the Environment for the Things as a Service, funded by the European Union under the 7th Framework Program and in several research projects supported by private industries (Nokia Siemens Networks and Fluidmesh). He has served as a member of the organization committee for the international conference VALUETOOLS 2009 and WOWMOM 2011 and as member of the technical program committee for several international conferences.

Giovanni Nardini (g.nardini@ing.unipi.it) is a post-doc researcher at the University of Pisa, where he obtained his MSc degree in Computer System Engineering in 2013, and his PhD in Information Engineering in 2017. His research interests are cellular networks, and specifically resource allocation in LTE-Advanced networks. In this fields, he has coauthored five patents and sixteen papers. He has been involved in EU-funded and industrial research projects.

Giacomo Tanganelli (g.tanganelli@ing.unipi.it) is a Postdoctoral Researcher at the Department of Information Engineering of the University of Pisa. He received a Master's Degree and a

PhD in Computer Systems Engineering in 2012 and 2016, respectively, from the University of Pisa. In 2015, he was a PhD visiting student at the Computer Science Department of the ETH Zurich. His main research area is the Internet of Things including also wireless sensor networks, M2M communications and low power WANs. He has been involved in the project BETaaS, Building the Environment for the Things as a Service, funded by the European Union under the 7th Framework Program.

Giovanni Stea (g.stea@ing.unipi.it) is an Associate Professor at the Department of Information Engineering of the University of Pisa, Italy, where he also got his PhD in 2003. His current research interests include quality of service and resource allocation in wireline and wireless networks, performance evaluation through simulation and analytical techniques, traffic engineering. In these fields he has coauthored more than 80 peer-reviewed papers and 15 patents. He has been involved in national and European research projects, and he has led joint research projects with industrial partners. He has served as a member of the technical and/or organization committees for several international conferences, including SIGCOMM, WoWMoM, and VALUETOOLS, and he is serving on the editorial board of the Springer Wireless Networks journal.

Enzo Mingozzi (e.mingozzi@ing.unipi.it) is an Associate Professor at the Department of Information Engineering of the University of Pisa, Italy. His research activities span several areas, including the design and performance evaluation of multiple access protocols for wireless (multi-hop/mesh) networks, quality of service (QoS) provisioning and service integration, architectures and protocols for the Internet of Things. He is a co-author of more than 100 peer-reviewed papers in international journal and conference proceedings, 5 book chapters and 10 patents. He currently serves on the Editorial Board of the Internet of Things Journal (IEEE), Computer Networks (Elsevier), and Computer Communications (Elsevier). He was also guest editor of a few journal special issues on topics related to wireless and mobile networks. He served in the organizing and technical program committee of several international conferences and was a co-founder of the IEEE IoT-SoS and IEEE MeshTech workshops.