On Participatory Service Provision at the Network Edge with Community Home Gateways

Amin M. Khan[a], Felix Freitag[b]

[a] *Instituto Superior Técnico, Universidade de Lisboa. INESC-ID Lisboa. Lisbon, Portugal*
amin.khan@tecnico.ulisboa.pt
[b] *Department of Computer Architecture, Universitat Politècnica de Catalunya. Barcelona, Spain*
felix@ac.upc.edu

Abstract

Edge computing is considered as a technology to enable new types of services which operate at the network edge. There are important use cases in ambient intelligence and the Internet of Things (IoT) for edge computing driven by huge business potentials. Most of today's edge computing platforms, however, consist of proprietary gateways, which are either closed or fairly restricted to deploy any third-party services. In this paper we discuss a participatory edge computing system running on home gateways to serve as an open environment to deploy local services. We present first motivating use cases and review existing approaches and design considerations for the proposed system. Then we show our platform which materializes the principles of an open and participatory edge environment, to lower the entry barriers for service deployment at the network edge. By using containers, our platform can flexibly enable third-party services, and may serve as an infrastructure to support several application domains of ambient intelligence.

Keywords: edge computing, community clouds

1. Introduction

Important use cases in ambient intelligence and the Internet of Things (IoT) drive edge computing research and deployment. Major vendors provide nowadays edge computing solutions, such as [2,11,12,13], showing the usage of edge computing in many industrial and consumer-oriented scenarios. Many of these solutions, however, are based on proprietary hardware and software platforms, and put barriers to a flexible extension with third-party services. End-users face vendor-specific solutions which are not interoperable with those of the others. Providers of ambient intelligence services may find the access to market for their services aggravated because of not having computing infrastructures available, e.g. set-top boxes or home gateways which could host their services.

The possibility of extending infrastructures with additional services, however, is needed to satisfy new user requirements. For instance in the scenario of smart homes, improved privacy and security must be included. Evolving contexts and new conditions need that edge systems can be updated in a secure way, and that they can be customized easily. Such extension to support new use cases, however, should not require to replace the whole system, but be enabled by design.

An open platform is needed where end users will benefit from being empowered to choose among several providers the most suitable services. Furthermore, end users may select an optimization criteria for these services at their choice, such as performance, security, privacy or the offered functions. Such an environment will also bring benefits for commercial providers specialized in edge services. The entry barrier for new actors to deploy services in such an open edge system will be much lower than pursing this purpose within proprietary frameworks which impose vendor lock-in.

In this paper we argue for an open edge computing environment driven by end users in which – in contrast to the above models – the platform is open for any interested stakeholder to actively participate in the service provision. Our approach combines current trends which, each by itself, have shown to be very successful and promising, namely 1) the rise of application platforms providing privacy preserving access control technologies, 2) increasing availability of powerful and energy-efficient hardware at users' premises, permanently operated, such as the Raspberry Pi, mini-PCs and enhanced home gateways, 3) containerization of services, as exemplified by Docker.

A first instance of this vision, where we run experiments of the presented work, has already become operational by a real deployment implemented as a community network cloud in Guifi.net[22]. It is formed by computing and communication resources provided by end users to host local services. The software platform which is installed on the devices is open and can be administered by the participants. User collaboration is part of the model of an open edge computing environment.

In this paper, based on our experience, we elaborate further on different aspects to be considered to foster participatory service provision, and position it as an approach to enable important use cases of IoT scenarios.

2. Motivating application scenarios

Once IoT devices and personalized mobile services become fully rolled out, the amount of data created at the user premises will soon increase in magnitudes. New requirements for ambient intelligence and IoT services at the network edge will appear. These services will need to address:

1. Privacy of the IoT data: Currently much IoT data is analyzed remotely at cloud data centers. End users are often unaware of the amount of data sent through their gateway to the service provider. Obtaining consent for this procedure from the user is often part of the service's installation process and part of the steps clicked by the user. Much of this data contains private information and if it is linked with the users' identities, personal information may become public or misused. The value of the user data is another issue to be considered. User data may help to create value for third-parties.
2. Tailored edge data analytics: An increased number of ambient intelligence services may benefit from data analytics support services running at the edge layer to take more informed decisions, become context-aware, and optimize resource usage. The access to shared information among different edge devices may in addition help to identify new meanings from the data, which may not be possible remotely where this access to shared data is not possible.
3. Resilience: There might be an increase in safety critical applications in home environments enabled by the technological capabilities. Such services may be found for instance in the domain of e-Health or surveillance and security. Relying on centralized remote cloud services, however, creates risks to such services becoming disrupted[14]. Distributed edge services may add resilience if they are able to take over temporarily critical functions in case of failures in centralized remote services.

Participatory service provision at the network edge could be enabled by decentralized cloud services hosted on edge devices. Such edge devices on which the services are deployed can be materialized by enhanced home gateways[4] and mini-PCs located at end user premises. Such an approach is different to that used in most of today's IoT data management services, which are mainly operated in large data centers.

3. Considerations for sustainability

Participatory edge cloud services will leverage on community efforts. The community is formed by users which need to find value in these services in order to support them. The type of services provided, e.g. the content, the availability of the service, and the performance of the service are issues which are important to raise users' interest. In the following, we discuss a set of dimensions to consider to make such services sustainable.

3.1. Technical dimension

Architectural issues: Being compatible with current and future architectural standards is crucial for the adoption of participatory edge services by its stakeholders. The network and service architecture for such decentralized community services should be specified as extensions to existing and emerging architectures, rather than proposing conflicting or non-standard building blocks.

Distributed services: These services are a key enabler for a decentralized platform system. Such services face the challenges such as 1) they need to be designed to be able to adapt to scale, resistant to churn, and have self* properties in order to maintain themselves operational under changing conditions in load, traffic and application patterns, and 2) they should achieve a resource-efficient usage given that edge devices are shared among several services and applications operating concurrently.

Performance characterization: Being able to understand the quality of service and experience – QoS and QoE – of decentralized community services is crucial for the wider take-up and broader applicability. It is needed to understand the services' performance and that of the applications that use them, from evaluations in real environments. Being able to predict and assess their suitability for different application scenarios is key to engage the relevant stakeholders and identify users.

3.2. Social dimension

Providing value for end users is an important element which participatory services need to fulfill in order to engage end users and volunteers. Such value may be found in the differentiating characteristic of the edge services which are enabled, such as improving privacy for the end users or enable new services which may exploit the specific access to shared edge data.

Value may also be seen from users by enabling voluntary computing tasks, which has already been successfully performed in the past for supporting research and progress[21]. Recent initiatives have proposed to exploit the computing capabilities of smartphones[6]. Some applications such as Samsung Power Sleep[7] can be run at users mobile devices when otherwise idle, which may also be applied to fixed devices such as home gateways and mini-PCs. Donating the usage of edge cloud devices for such type of services may be motivating for some users to participate in an edge cloud.

3.3. Economic dimension

Many use cases of ambient intelligence and IoT leverage on gateways[10]. Enabling commercial services is an important part of an ecosystem around the collaborative edge environment. The availability of edge computing infrastructure may be of interest for commercial service providers, since it will provide access to the resources and data needed to deploy new services. Since the proposed approach builds on commodity hardware and an open Linux platform, the business opportunity shifts from proprietary IoT gateways to the service provision level. As a consequence, a competitive market for edge service providers may appear. Volunteer efforts should co-exist with commercial exploitations. It was observed, however, that a careful regulation should be in place to govern the ecosystem[8].

4. Edge computing approaches

In this section we review several edge computing approaches and analyze their relation with the approach we propose.

The *fog*[9,23] and *edge computing* paradigms[10] disrupt the classic two-tier cloud architecture with large data centers, and consider edge computing services, such as pervasive data analytics, to run on edge devices. In edge computing, an additional infrastructure layer is built as part of a three-tier architecture by placing devices at the network edge, which take over performing specific local services[1,3].

Important initial works on fog computing came from networking hardware vendors[24], where fog computing aims to be deployed on Customer Premises Equipment (CPE), i.e. devices located in the households of the customers. Home gateways and set-top-boxes were foreseen to carry other additional services beyond the basic network services,

which telecom operators would then offer to end users[25]. In this context, the management of the system and the edge services are controlled by the telecom operator. This approach does not enable active user participation, making it different to the community home gateways we propose.

A more user-empowering approach is presented in Paradrop[26]. Paradrop is a software platform that enables applications to run on WiFi routers at user premises. It has originally been developed for OpenWRT-based home routers leveraging Linux Containers (LXC), but now also supports Ubuntu Snappy, and intgeration with Docker has also been recently reported. The resilience of Paradrop leverages on the WiFi router, which in today's homes is an always-on and always connected device. Recent open hardware routers with increased computing capabilities, such as the Turris Omnia[4], support the feasibility of Paradrop and demonstrate the timeliness of this approach. Regarding the user-centered management concept, Paradrop is similar to the proposed community home gateways. Our approach however goes one step further, by building a community infrastructure at the network edge, in which resources and services can be shared horizontally.

There are several platforms available with the goal to enable end users to easily run services in local computing devices at their homes:

Sandstorm is an open platform which significantly eases application deployment[17]. Sandstorm mimics the ease of mobile application installation by a graphical Web interface, which allows users to deploy container-based applications almost instantaneously at end users' computers with just a click. While Sandstorm is extremely efficient in usability, its primary target seems to be to run personal cloud services. Although the access to the deployed applications then can be shared, it is not the purpose to share computing resources or to support distributed services at the edge. While the number of available applications to deploy in Sandstrom is large, this scope seems not to be extensible by an end user.

Yunohost[18] is a Debian-based distribution to facilitate end users with the self-hosting of applications on local computers. User-friendliness is achieved by graphical interfaces for both the users and administrator, making this system usable by a diverse group of users. The supported applications in Yunohost consist of a list of core applications and an extensible list of third-party provided applications. The ecosystem, however, seems to be restricted to the Yunohost community, which makes a larger outreach difficult. While Yunohost facilitates application installations on local devices, its main focus seems to be personal service provision on isolated devices, rather than targeting distributed edge computing services.

Flockport[19] is a platform that runs on several Linux distributions for providing applications through Linux Containers (LXCs). Flockport offers a large list of applications through a kind of app store. The Flockport tools aim to facilitate and ease local application deployment, leveraging the features of LXC. The scope of the applications, however, seems to be limited by those available in the Flockport app store. In contrast, Docker Hub[20], the image repository of Docker-based applications, can be extended with images by any registered user.

From the review of these approaches, we can observe key limitations that may impede these works to further succeed as open participatory edge platforms, including 1) the lack of leveraging de-facto reference implementations, such as Docker, 2) the concept to conceive the host at the edge as an isolated device, rather than as a resource that can be shared to deploy distributed services, and 3) as a consequence, the lack of support services to interconnect edge devices and services.

5. Edge cloud in the Guifi community network

In this section, we describe our current implementation of an open platform for participatory edge service provision.

5.1. Concept

The Guifi.net community cloud consists of distributed heterogeneous computing devices contributed by participants. They run a set of common basic services provided by the *Cloudy* software[22]. This community cloud runs on heterogeneous hardware, since the users can contribute to the cloud with many types of devices. Often, these are inexpensive devices that can be classified as mini-PCs for home server applications. They have low energy consumption

to operate in a 24/7 mode. The infrastructure is distributed, since most of the contributed devices are located at the users' premises.

Figure 1 illustrates the Guifi.net community cloud. Different hardware devices are used as cloud nodes. The Cloudy distribution is installed on each of the devices connected to the community cloud. The cloud nodes are spread over the edge network and in different geographic locations. Desktop PCs are mainly in municipality buildings, most mini-PCs are at user premises.

Figure 1: Edge community cloud in Guifi.net, running the Cloudy distribution on heterogeneous hardware.

The cloud infrastructure is not bound to a certain number of devices, but grows organically by every new node that community members add to the infrastructure. A similar situation holds for services. While there is a pre-configured set of services integrated in the Cloudy distribution, this community cloud is open to provide new services. The Guifi community cloud currently runs on around thirty devices. Instantaneous values on services and devices can be obtained at any moment from a publicly available Cloudy instance[1].

5.2. Cloudy software platform

For enabling the community network cloud, the *Cloudy* distribution has been developed[2]. Cloudy is based on Debian GNU/Linux. It can be freely downloaded from public repositories[3].

Users of the Guifi.net community cloud are expected to install this distribution on the contributed nodes. Making Cloudy the default system for community network clouds ensures homogeneity in terms of a basic set of common services, which are needed for every participant to join and interact in the community cloud. Once the Cloudy distribution is installed, it provides a set of service categories, grouped as Search, Guifi.net, Community cloud, Personal cloud and Enterprise cloud. Figure 2 shows the Cloudy web user interface after installation at the user's device.

---

[1] http://demo.cloudy.community User: *guest*, Password: *guest*
[2] http://cloudy.community/
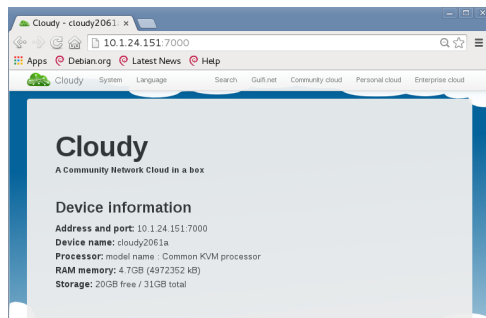[3] https://github.com/Clommunity/

Figure 2: Cloudy distribution Web user interface.

Figure 3: Example of a home server in the edge community cloud.

5.3. Hardware for the community home gateways

Cloudy is meant to be installable on any kind of on-premise computing device, which can then become part of the community network cloud. Figure 3 shows a typical node which has been recently deployed in the Guifi community cloud. The device from Minix [4] comes with a low energy consuming Intel Z3735F (64-bit) processor, 2 GB of RAM and 32 GB of internal storage. Over the USB port, additional storage capacity can be added by the user.

5.4. Service deployment with Docker containers

Our current on-going work addresses service personalization. Service personalization would allow the Cloudy service offer to adapt to the individual needs of the users. It is likely that a user will contribute a device as Cloudy host to the community cloud if she is able to run community *and* personal services [5]. Given that Cloudy is conceived as an open platform, this openness should support users to install their own applications, have the flexibility to choose which kind of services they want to run, and which ones of them to run for personal use, and which ones to share with the community.

We have chosen Docker as the technology to enable in Cloudy the support for service personalization and customized service deployments in the *Enterprise cloud* service category. In the *Docker* sub-menu of the Cloudy web interface the user can find a list of pre-configured applications available that can be started as Docker containers.

5.5. Experimentation

We consider the scenario when a Cloudy user is willing to install and run a new application in its Cloudy node and publish it to the other Cloudy users. The steps which the user has to conduct are installing and deploying the Docker container of the application in Cloudy and publishing this new application, thus it will be known to the other Cloudy users.

Figure 4 shows how a Cloudy user using the *Search service* finds several applications deployed as Docker containers in the cloud, which are shared by their owners with the community. We can see four different applications that are found, i.e. ownCloud, Kanban, Redis, BitTorrent tracker, installed as Docker containers on other Cloudy instances, of a total of six shared applications. It can be seen that for two of these applications, i.e. ownCloud and Kanban, there are two instances published with different IPs, corresponding to instances of different owners. In this case, a user interested in using this application could chose between the two providers. Finally, four different IPs are in the list of found applications, meaning that there are four Cloudy instances acting as services providers.

From the experiments we could observe that a Cloudy device in the community cloud is able to deploy, on behalf of third-parties, additional services that are provided as Docker images. This is an important features for enabling new cloud-based services at the network edge. Cloudy users could support with their devices the deployment of commercial containerized services. The Docker integration in Cloudy is still evolving and may benefit in the future

---

[4]http://www.minix.com.hk/Products/MINIX-NEO-Z64-ANDROID-TV.html
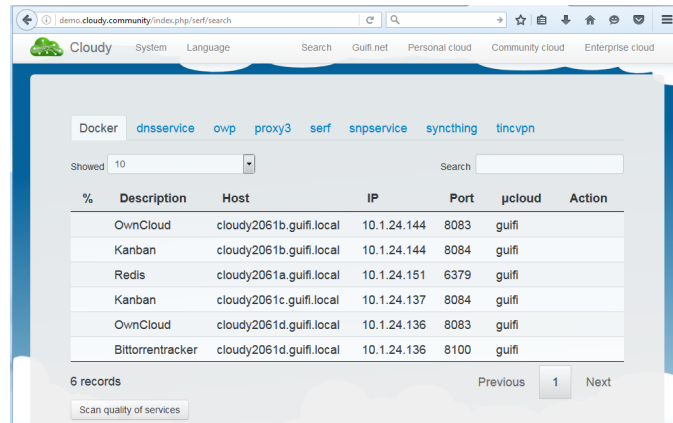
Figure 4: A Cloudy user finds four service providers publishing Docker-based applications.

from additional tools becoming available, as result from the standardization efforts recently started around the Docker ecosystem[5].

6. Conclusions and Outlook

Participatory service provision, presented in this paper through community network clouds, was proposed as a base for an open flexible platform for distributed edge service provision. An important potential for application can be seen for the deployment of tailored services at the network edge.

The presented solution takes advantage of the recent trend towards container-based service provision, and integrates technology provided by Docker into a platform for edge devices. By providing an implementation to make this platform user-friendly, the average user may install services through the platform's graphical interface.

Container management tools are evolving very fast, and next steps should include further integration of Docker tools to enable complex application deployments. Another dimension of future work should address building an ecosystem that brings together volunteer service contributions and commercial exploitation on these community home gateways, as well as exploiting its potential to provide security solutions.

Acknowledgements

---

[5]Cloud Native Computing Foundation. https://cncf.io/

References

1. Intelligent Gateway for IoT powered by Red Hat. https://www.redhat.com/en/en/about/videos/intelligent-gateway.
2. The Value of Bringing Analytics to the Edge. http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Value_of_Analytics_at_the_Edge_Final.pdf.
3. Luis M. Vaquero and Luis Rodero-Merino. Finding your way in the fog: Towards a comprehensive definition of fog computing. *SIGCOMM Comput. Commun. Rev.*, 44(5):27–32, October 2014.
4. Turris Omnia. More than just a router. The open-source center of your home. https://omnia.turris.cz/en/
5. Nuno Apolnia, Felix Freitag, and Leandro Navarro. Leveraging deployment models on low-resource devices for cloud services in community networks. *Simulation Modelling Practice and Theory*, 2016.
6. F. Bsching, S. Schildt and L. Wolf. DroidCluster: Towards Smartphone Cluster Computing. *32nd International Conference on Distributed Computing Systems Workshops*, Macau, 2012, pp. 114-117.
7. Samsung Power Sleep. http://www.samsung.com/at/microsite/powersleep/en/faq.html
8. Roger Baig, Llus Dalmau, Ramon Roca, Leandro Navarro, Felix Freitag, and Arjuna Sathiaseelan. Making Community Networks economically sustainable, the Guifi.net experience. *ACM SIGCOMM Global Access to the Internet for All (GAIA '16)*. Florianpolis, Brazil, 31-36.
9. Bonomi Flavio, Milito Rodolfo, Zhu Jiang, Addepalli Sateesh. Fog Computing and Its Role in the Internet of Things. *First Edition of the MCC Workshop on Mobile Cloud Computing (MCC '12)*, 2012.
10. W. Shi and S. Dustdar. The promise of edge computing. *Computer*, 49(5):78–81, May 2016.
11. Azure and Internet of Things. https://azure.microsoft.com/en-us/documentation/articles/iot-suite-what-is-azure-iot/.
12. Intel IoT Platform Reference Architecture. http://www.intel.com/content/www/us/en/internet-of-things/white-papers/iot-platform-reference-architecture-paper.html.
13. Watson Internet of Things Platform. http://www-03.ibm.com/software/products/en/internet-of-things-platform.
14. Centralized Web Services Are Wonderful - Until They Go Wrong. https://www.technologyreview.com/s/603738/centralized-web-services-are-wonderful-until-they-go-wrong/.
15. guifi.net – a Community Network. http://guifi.net.
16. Alexandros Marinos and Gerard Briscoe. Community Cloud Computing. *CloudCom*, 5931(December):11, July 2009.
17. Sandstorm open source operating system for personal and private clouds. https://sandstorm.io/.
18. YunoHost: A server operating system aiming to make self-hosting accessible to everyone. https://yunohost.org/.
19. Flockport: Run apps on any cloud, any server, any provider. https://www.flockport.com/.
20. Docker Hub. https://hub.docker.com/.
21. Adam L. Beberg, Daniel L. Ensign, Guha Jayachandran, Siraj Khaliq, and Vijay S. Pande. Folding@home: Lessons From Eight Years of Volunteer Distributed Computing. In *8th IEEE International Workshop on High Performance Computational Biology (HiCOMB '09), within IPDPS*, pages 1–8, Rome, Italy, May 2009. IEEE.
22. Mennan Selimi, Amin M. Khan, Emmanouil Dimogerontakis, Felix Freitag, and Roger Pueyo Centelles. Cloud services in the Guifi.net community network. *Computer Networks*, 93, Part 2:373 – 388, 2015.
23. Blesson Varghese, Nan Wang, Dimitrios S. Nikolopoulos and Rajkumar Buyya. Feasibility of Fog Computing. *arXiv preprint arXiv:1701.05451*. 2017
24. Internet of Things (IoT): Transform Data into Action at the Network Edge. http://www.cisco.com/web/solutions/trends/iot/fog-computing.html
25. Jon Whiteaker, Fabian Schneider, Renata Teixeira, Christophe Diot, Augustin Soule, Fabio Picconi, and Martin May. Expanding home services with advanced gateways. *SIGCOMM Comput. Commun. Rev.*, 42(5):37–43, September 2012.
26. Dale F. Willis, Arkodeb Dasgupta, and Suman Banerjee. Paradrop: A multi-tenant platform for dynamically installed third party services on home gateways. In *ACM SIGCOMM Workshop on Distributed Cloud Computing (DCC '14)*, pages 43–44, Chicago, USA, 2014.
27. A. Chandra, J. Weissman, and B. Heintz. Decentralized edge clouds. *IEEE Internet Computing*, 17(5):70–73, Sept 2013.
28. Pedro Garcia Lopez, Alberto Montresor, Dick Epema, Anwitaman Datta, Teruo Higashino, Adriana Iamnitchi, Marinho Barcellos, Pascal Felber, and Etienne Riviere. Edge-centric computing: Vision and challenges. *SIGCOMM Comput. Commun. Rev.*, 45(5):37–42, September 2015.
29. Simon Caton, Christian Haas, Kyle Chard, Kris Bubendorfer, and Omer F. Rana. A Social Compute Cloud: Allocating and Sharing Infrastructure Resources via Social Networks. *IEEE Transactions on Services Computing*, 7(3):359–372, 2014.
30. Kyle Chard, Kris Bubendorfer, Simon Caton, and Omer F. Rana. Social Cloud Computing: A Vision for Socially Motivated Resource Sharing. *IEEE Transactions on Services Computing*, 5(4):551–563, 2012.
31. Atta Rehman Khan, Mazliza Othman, Sajjad Ahmad Madani, and Samee Ullah Khan. A Survey of Mobile Cloud Computing Application Models. *IEEE Communications Surveys & Tutorials*, 16(1):393–413, January 2014.
32. N. Apolnia, R. Sedar, F. Freitag, and L. Navarro. Leveraging low-power devices for cloud services in community networks. In *3rd International Conference on Future Internet of Things and Cloud (FiCloud '15)*, pages 363–370, Aug 2015.
33. Clommunity - a community networking cloud in a box. http://clommunity-project.eu.