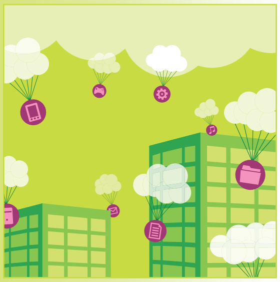

Improving Quality of Experience in Future Wireless Access Networks through Fog Computing

A novel model of Internet access networks is proposed, based on fog computing. The model hosts applications close to users by relying on virtual machines to dynamically move cloud or Web content to nodes located at the edge of access networks. Then it performs proactive caching and enforces traffic policies based on the interaction between access infrastructure and external applications. Analyzing experimental data collected from public Wi-Fi hotspots, the authors quantify this approach's benefits for optimizing bandwidth usage, reducing latency, and enhancing quality of experience. Experimental results show the fog-based approach can manage a significant portion of download data.

Nicola Iotti, Marco Picone, and Simone Cirani
Caligoo, Italy

Gianluigi Ferrari
University of Parma, Italy

Internet access networks are now facing critical challenges related to the ever-increasing number and heterogeneity of connected devices and the strict requirements of new generations of services.[1] Building networks that just provide more resources, or differentiating access networks for different types of services, isn't the right approach to handle the workload and the complexity brought about by these new challenges. For this reason, academia and industry are developing new network architectures that take into account the very nature of these changes, rather than just blindly boosting the underlying network bandwidth capabilities. Fog computing is a novel paradigm that aims at optimizing networking, computing, and storage resources and improving the quality of service brought to users (in terms of latency and throughput), by moving resources to the edge of access networks.[2] This approach is set to avoid limitations that affect traditional cloud-based solutions, such as low-bandwidth, congestion of Internet connection, and infeasibility of real-time applications, thus enabling new approaches to dynamic resource allocation, such as bandwidth usage.

Here, we analyze real data collected in the field to evaluate the benefits

Related Work on Fog Computing in Access Networks

The evolution of wireless access networks is a topic that has been deeply investigated and is rapidly evolving from research to deployment in the field. Fog computing is playing a key role in the design of future networks.

A well-defined description of fog computing architecture is introduced by Flavio Bonomi and colleagues,[1] where, even if the concept of fog computing is general, the approach is mainly focused on Internet of Things (IoT) applications. Other works introduce an architecture, similar to the one described in this article, including a smart gateway or hub deployed on the edge,[2,3] but specifically for IoT networks. Luis Vaquero and Luis Rodero-Merino[4] provide a good picture of fog computing technology and challenges. Suksant Lor and colleagues[5] provide an example of storage functions dynamically deployed in selected network sections to speed up data transfer.

Other works are focused on specific topics related to fog computing, such as security[6] and reliability,[7] but without introducing a specific architecture for users' access networks. Y. Navaneeth Krishnan and colleagues introduce a method for moving computation from the cloud to network devices,[8] focusing on the deployment of applications — basically for data preprocessing — directly on access network devices. In another work, the possibility to use fog-based access points and user equipment is envisioned,[9] providing a granular distribution and workload based on memory sharing.

Still yet, other works are mainly related to network resource management. In particular, Swati Agarwal and colleagues propose clustering techniques of sharing data and hardware resources to implement flexible management of a distributed infrastructure.[10] Beate Ottenwälder and colleagues[11] present a method for placement and migration of virtual machines for cloud and fog providers, and Takayuki Nishio and colleagues[12] introduce a system to predict future query regions for moving consumers (the method provides information to process events early).

Many of the envisioned infrastructure have common points with the one we're proposing. However, our work is focused more on leveraging storage capabilities locally provided by fog nodes, and migrating content according to users' needs and connection resources. Moreover, our work is based on the analysis of real data to estimate potential benefits of this approach, which can be applied to many other fog-based implementations.

References

1. F. Bonomi et al., "Fog Computing and Its Role in the Internet of Things Characterization of Fog Computing," *Proc. Mobile Cloud Computing*, 2012, pp. 13–15.
2. M. Aazam and H. Eui-Nam, "Fog Computing and Smart Gateway Based Communication for Cloud of Things," *Proc. 2014 Int'l Conf. Future Internet of Things and Cloud*, 2014, pp. 464–470.
3. S. Cirani et al., "Fog Networking for 5G and IoT Workshop," *Proc. 12th Ann. IEEE Int'l Conf. Sensing, Comm. and Networking*, 2015, pp. 464–470.
4. L.M. Vaquero and L. Rodero-Merino, "Finding Your Way in the Fog," *ACM Sigcomm Computer Comm. Rev.*, vol. 44, no. 5, 2014, pp. 27–32; http://dl.acm.org/citation.cfm?id=2677046.2677052.
5. S.S. Lor, L.M. Vaquero, and P. Murray, "The Cloud in Core Networks," *IEEE Comm. Letters*, vol. 16, no. 10, 2012, pp. 1703–1706.
6. I. Stojmenovic and S. Wen, "The Fog Computing Paradigm: Scenarios and Security Issues," *Proc. 2014 Federated Conf. Computer Science and Information Systems*, 2014; doi:10.15439/2014F503.
7. H. Madsen et al., "Reliability in the Utility Computing Era: Towards Reliable Fog Computing," *Proc. 2013 20th Int'l Conf. Systems, Signals and Image Processing*, 2013, pp. 43–46.
8. Y.N. Krishnan, C.N. Bhagwat, and A.P. Utpat, "Fog Computing-Network Based Cloud Computing," *Proc. Int'l Conf. Electronics and Comm. Systems*, 2015, pp. 250–251.
9. M. Peng et al., "Fog Computing Based Radio Access Networks: Issues and Challenges," *IEEE Network*, vol. 30, no. 4, 2015; http://arxiv.org/abs/1506.04233.
10. S. Agarwal, S. Yadav, and A.K. Yadav, "An Architecture for Elastic Resource Allocation in Fog Computing," *Int'l J. Computer Science and Comm.*, vol. 6, no. 2, 2015, pp. 201–207.
11. B. Ottenwalder et al., "Migcep: Operator Migration for Mobility Driven Distributed Complex Event Processing," *Proc. 7th ACM Int'l Conf. Distributed Event-Based Systems*, 2013, pp. 183–194.
12. T. Nishio et al., "Service-Oriented Heterogeneous Resource Sharing for Optimizing Service Latency in Mobile Cloud," *Proc. 1st ACM Sigcomm Workshop on Mobile Cloud Computing*, 2013, pp. 19–26.

of fog computing on bandwidth optimization, related to functions that would be infeasible through a traditional approach. The results of this work represent an important validation of a fog-based model for Internet access networks, which is the first preliminary and essential step to legitimize the adoption of new-generation, fog-based access networks. This study brings a new meaning to fog computing, by transforming it from an effective architecture for Internet of Things (IoT) applications to a powerful approach that can have a disruptive impact on Wi-Fi Internet access services, especially in those locations where the following are relevant issues: Internet bandwidth, high user density, or shared infrastructure for Internet users and IoT nodes. Illustrative scenarios include cruise ships, trains, airplanes, hotels, and convention centers. (For others' work in this area, see the related sidebar.)

Table 1. Data related to the numbers of connected devices.					
Hotel	**Maximum connected devices (per day)**	**Average connected devices (per day)**	**Total connected devices***	**Internet download traffic (Mbytes)**	**Internet upload traffic (Mbytes)**
Hotel A	767	607.76	9,517	2,285,443	690,781
Hotel B	297	180.48	1,694	990,151	91,425
Hotel C	184	132.48	2,609	379,242	102,130

* The total is for the number of devices that activated at least one connection during the considered period (one month).

Architecture

The fog-based platform evaluated in this work, is an evolution of an infrastructure able to provide Internet access, commonly known as a hotspot, with user authentication in public locations. This novel platform is based on the deployment of local nodes able to provide networking functions, computation, and storage close to the edge of the network. For this reason, we refer to this element with the general term *fog node*. It performs authentication, authorization, and accounting (AAA) functions, like any hotspot access controller, but it also has the following three innovative capabilities:

- host applications in the form of virtual machines (VMs);
- dynamically move data to make them locally available, introducing an important intermediate point of control between end devices and the cloud platforms (this function isn't simply caching, but proactive content transfer, managed by applications running on the fog node, able to act as an extension in the fog of cloud-based services or a content repository); and
- enforce bandwidth limitation, for specific connected devices, based on the interaction between the fog node and external applications through APIs.

Next, we discuss the data we collected and how we analyzed it.

Validation Analysis

We analyze collected data to evaluate potential benefits introduced by the adoption of a fog-based platform on the quality of user experience. In examining this information, we're considering how much data can be moved to the fog node for local access; how it's possible to optimize Internet bandwidth use by downloading data for future local access when the connection is underused and leveraging applications running locally on the fog node; how to collect useful elements for a valid estimation of the storage needed locally in a hotel to deploy these services; and how to evaluate the impact of interactive bandwidth management.

Experiment Description and Collected Data

We deployed the fog node in three large hotels in the city of Milan, Italy, where Caligoo fully manages the Wi-Fi Internet access for hotel guests. We refer to the three hotels involved in this analysis as Hotels A, B, and C. All used data are completely anonymous to preserve users' privacy. We collected more than 50,000 connections from more than 13,800 hotel guests in February 2015. Because these hotels are mainly for business travelers, February represents a reliable perspective of normal use. All three locations and guests were unaware of the data collection to avoid biasing user behaviors with respect to usual Internet use. By analyzing the connections, we were able to collect the following data:

- connection start time: (t_0);
- connection stop time: (t_1);
- connection duration, derived as ($t_1 - t_0$);
- the amount of downlink data (bytes) during the connection;
- the amount of downlink data (bytes) during the connection;
- the actual bandwidth consumption;
- the Internet traffic generated and received by each of the three locations; and
- the number of connected users.

Data Analysis and Traffic Optimization

The most important data, for the purposes of our analysis, are related to the number of users, the amount of traffic generated by the connections, the bandwidth usage, and the type of traffic. In our analysis, we consider all connected devices as users. Table 1 reports the following data:

- the maximum number of devices connected in a single day;

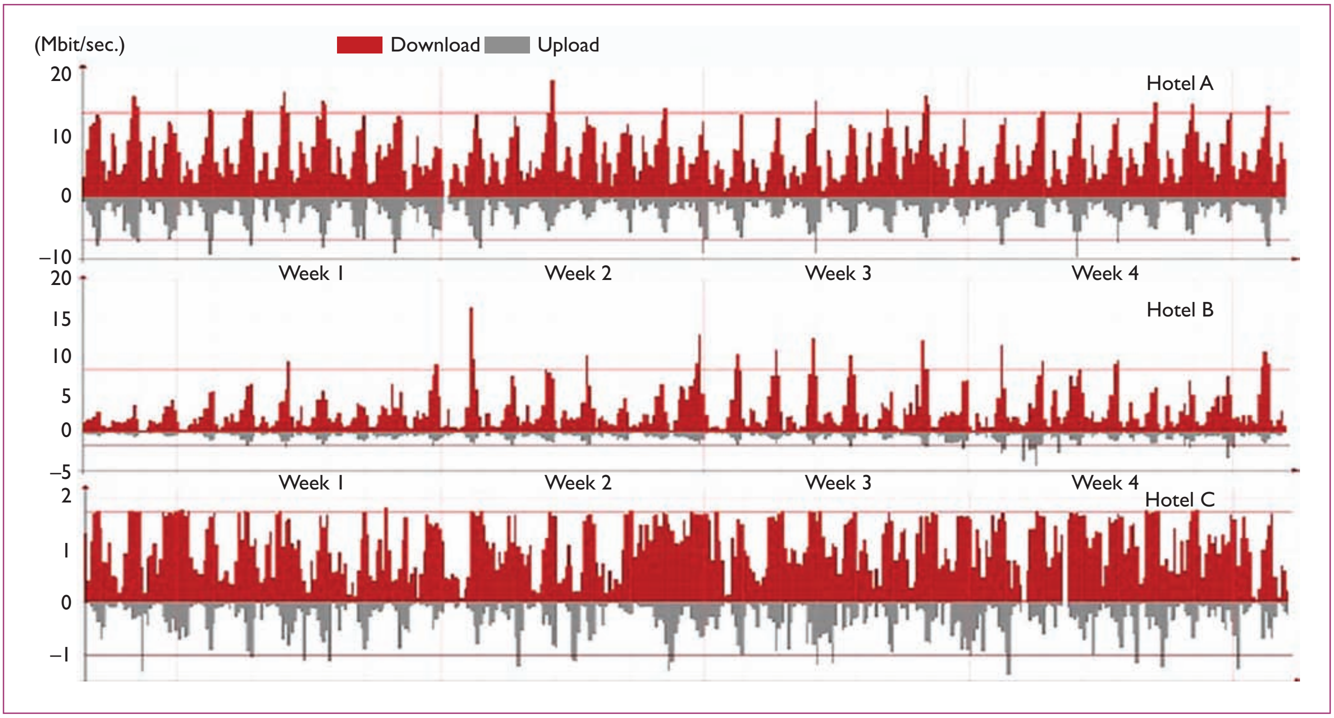

Figure 1. Average bandwidth use in a one-hour period, measured in Hotels A, B, and C. Bandwidth consumption isn't uniform but highly variable, and it's possible to observe periodic oscillations within a 24-hour period.

- the average number of devices connected in a single day;
- the total number of devices that activated at least one connection during the considered period;
- the total download traffic (Mbytes); and
- the total upload traffic (Mbytes).

We decided to measure bandwidth utilization calculating the average value, over a one-hour interval, on the number of bytes received and transmitted on the fog node interface connected to the Internet router. Figure 1 shows the results.

The collected data show clearly that the bandwidth consumption isn't uniform but highly variable, and it's possible to observe periodic oscillations within a 24-hour period. The reason of this particular bandwidth use is because during the day a large number of hotel guests aren't in the building, but in the evening, as they come back, a larger number of people in the hotel use Internet connections. This behavior leads to two relevant considerations: hotels need high bandwidth Internet connection to handle the heavy traffic peaks; and during a large portion of the day, the Internet connection is underused. This is the primary reason that this specific scenario would greatly benefit from an intelligent content-management system, able to move content proactively when the connection isn't used by guests, to reduce traffic when the connection is heavily loaded.

The main goal of this study is to evaluate the effect of the capability to run applications on the network edge. For this reason, one of the key factors is to estimate the amount of traffic that could be proactively cached (or, in general, managed) by applications running on the fog node, scheduling the download or upload, on the basis of the available bandwidth or the interaction with local or remote applications. We focus on the traffic that can't be managed by traditional caching, based on the general idea of multiple access to the same content. We identify the traffic related to content that could be proactively moved on the local node, before the first request for this content has been issued, and sometimes even before the user, interested in this content, arrives at the hotel. We also detect traffic that could be avoided, running applications on the local fog node. We refer to this traffic as *manageable traffic*. Our ability to proactively move specific content or to avoid traffic is made possible by the applications' fog-based design, such as the hotel room booking system or the Wi-Fi AAA service, which can run software modules in the cloud as well as on the fog node, and trigger content download based on specific events.

To estimate manageable traffic, we analyzed the download traffic captured during one week, in one location, and then calculated the average over 24 hours to present meaningful results for a typical day. The average number of daily captured packets on the hotels' networks was larger than 1,200,000. For our purposes, there isn't an immediate or simple method to identify manageable traffic by analyzing captured packets; we thus decided to start identifying DNS queries and replies. On the basis of the DNS resolution, we classified all the traffic from the corresponding IP address as follows.

- *Well-known*. This is traffic from well-known sources that are basically known in advance to be requested by users, such as popular non-real-time multimedia content platforms.
- *Frequent*. We detected specific IP source-generating traffic to a large number of different users. In this group, we identify non-real-time traffic related to the authentication system that could be moved to the fog node, or to specific events occurring in the hotel. We can also envision systems able to automatically detect frequently downloaded content and automatically move it to the fog node.
- *User-specific*. This is a large amount of non-real-time traffic from a specific source to a specific user, such as multimedia or content related to entertainment. This traffic could be moved in advance to the fog node for a local access — for example, as part of the interaction between the fog-based infrastructure and the room booking system.
- *Other*. The first three categories are considered manageable traffic, whereas this last one is unmanageable. In the case of traffic from IP addresses not corresponding to any DNS resolution, we directly contacted the IP, whenever possible, to classify the traffic as described before. This analysis is complicated because of the presence of traffic from content delivery network (CDN) platforms, such as Akamai or Amazon, that expose their domain names hiding the real content and making it impossible to classify the traffic. To avoid an overestimation of the manageable traffic, in this analysis we consider all the unidentified traffic as unmanageable.

To manage automatically the highly variable amount of captured traffic, we decided to organize the obtained data in chunks, including the same amount of captured traffic. Every block includes about 1.3 Gbytes and has different time extensions, depending on the time needed to collect the 1.3 Gbytes.

Table 2 shows the percentage of total traffic, identified as manageable, for every category and total. This analysis shows that an average amount corresponding to 28.89 percent of the total traffic could be managed in a smart way by a fog-based infrastructure to optimize the available resources. We also observed a considerable amount of traffic related to specific well-known cloud services or platforms such as Facebook, Skype, or Microsoft Office 365. We didn't consider this traffic manageable, because part of these applications are real-time, encrypted, or out of control in general. But this traffic could be managed in collaboration with the service provider, leveraging the nature of fog nodes that could host third-party VMs. With this approach, all the entities that want to provide services could produce software agents able to run on the fog node to optimize the resources usage, mainly in terms of available bandwidth. We estimated this traffic (potentially manageable) as 29.74 percent of the traffic categorized as non-manageable by the previous analysis. The manageable traffic thus could rise up to be around 50 percent of the total traffic.

Figure 2 shows the non-manageable and the manageable traffic, divided in classes, as part of the total traffic, distributed during the day, with the same approach as Table 2. Figure 3 shows a possible redistribution of the manageable traffic, to improve bandwidth use compared to the actual total traffic distribution during the day.

Resources Evaluation

Our data analysis provides interesting information about the local storage capability needed on the fog node to perform content management. Considering an optimal bandwidth usage — that is, using all the available bandwidth to download data to be cached on the node, we can set a theoretical upper limit for the local storage capacity on a fog node of 687 Gbytes at Hotel A, 817 Gbytes at Hotel B, and 47 Gbytes at Hotel C. Analyzing the captured traffic, we can estimate a more realistic value of 5.1 Gbytes of manageable traffic that must be cached per day.

Table 2. Percentages of manageable traffic in an average day.

Part of the day	Well-known traffic	Frequent traffic	User-specific traffic	Total
Part 0	16.63	10.81	3.52	30.97
Part 1	27.24	7.35	0.647	35.24
Part 2	11.41	11.21	0.77	23.41
Part 3	5.05	12.48	8.76	26.30
Part 4	7.97	7.20	7.44	22.63
Part 5	19.61	7.95	0.96	28.53
Part 6	14.02	7.76	7.77	29.57
Part 7	11.39	6.27	7.827	25.49
Part 8	14.30	6.48	9.59	30.38
Part 9	10.46	5.12	7.91	23.50
Part 10	17.17	4.55	9.21	30.94
Part 11	16.43	3.97	9.82	30.23
Part 12	12.38	3.57	9.16	25.12
Part 13	15.29	3.54	9.32	28.16
Part 14	15.67	3.45	13.36	32.49
Part 15	17.82	6.00	9.10	32.94
Part 16	11.29	5.26	12.54	29.11
Part 17	18.21	7.09	9.76	35.07
Day total	14.62	6.67	7.59	28.89

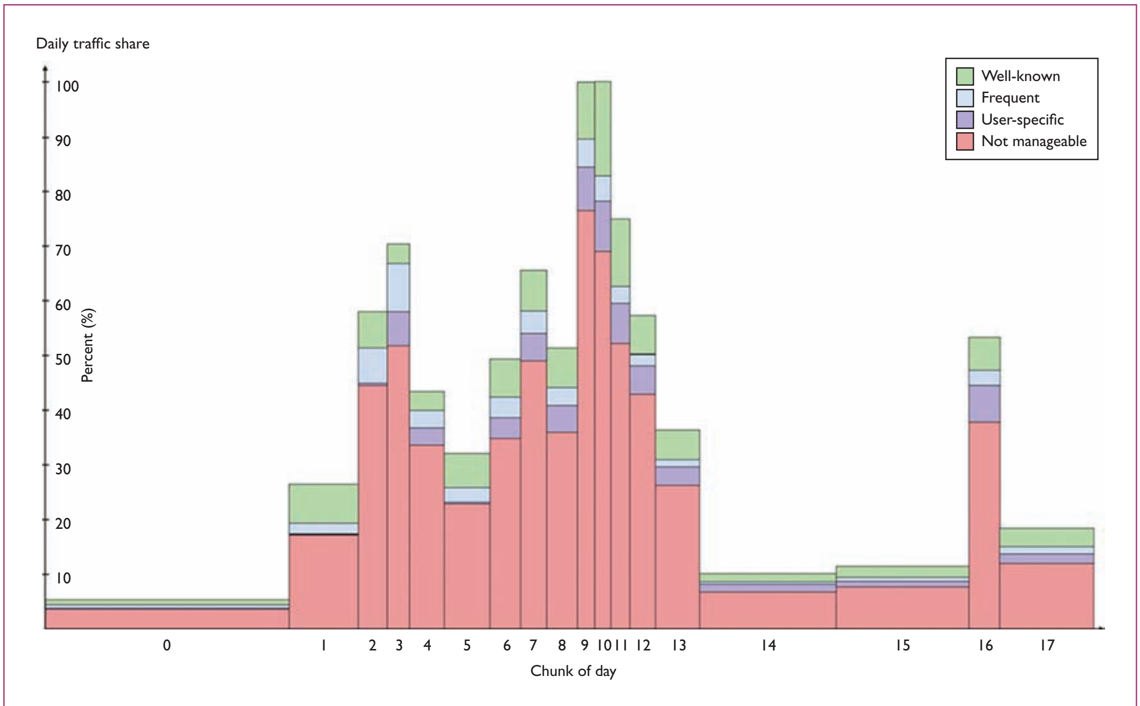

Figure 2. Manageable and non-manageable traffic downloaded during the day. The manageable traffic is classified as well-known, frequent, or user-specific.

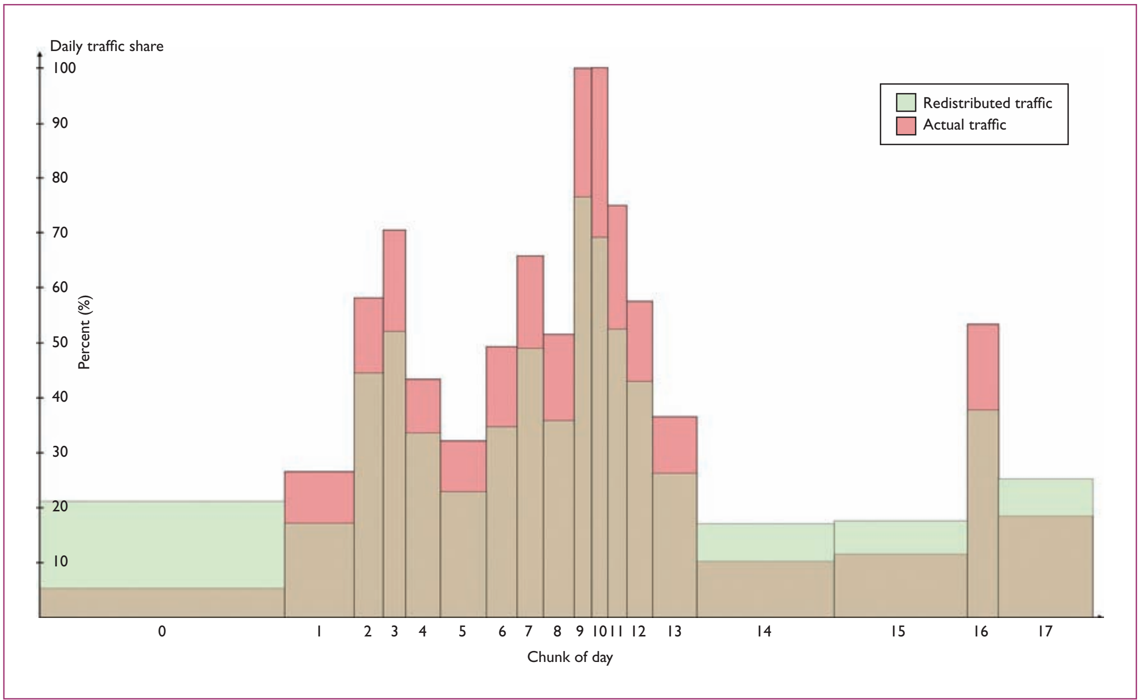

Figure 3. Possible downloaded traffic redistribution during the day, compared to the actual total traffic distribution. This could optimize bandwidth usage.

Bandwidth Management

Most hotels have conference or meeting rooms. This leads to deviations from the 24-hour periodic pattern of bandwidth consumption (see Figure 1) because of specific events hosted during the considered period of observation. Conferences and meetings are particularly critical because of the high density of connected users and the specific needs of the speakers who require reserved bandwidth for presentations or product demos. This situation can't be managed easily using preconfigured policies or rules based on traffic recognition. Our fog node dynamically enforces selective bandwidth limitations, to specific connected devices, as a consequence of the interaction with external applications through APIs. With this feature, a speaker, simply using an application running on his/her device (such as a smartphone) can ask for reserved bandwidth, calling an API exposed by the fog node. The same function can be used by a smart object, such as a fire alarm during an emergency, on a shared network. APIs are open only to authorized parties and if a request can't be served, it can be queued, keeping the application informed, and served as soon as resources are available, rather than just relying on a traditional best-effort approach. This approach makes fog-based access networks able to locally manage network policies, along with computation and storage, as a flexible and shared infrastructure.

The interaction between the fog node and local applications lets us dynamically and flexibly manage the available bandwidth. This is expedient to deal with complex situations such as: conferences, with a high density of connected users with different needs; and emergencies, when selected smart objects (such as fire alarms) might need to have priority in shared networks. The collected experimental results show that this new approach to the design of access networks has great potential benefits, in terms of resources optimization and performance. Future development of this approach will be related to a more detailed analysis, considering different types of locations and studying algorithms for moving content and optimizing resource usage.

Acknowledgments

This work was sponsored and supported by Caligoo.

References

1. E. Adler, "Here's Why 'The Internet of Things' Will Be Huge, and Drive Tremendous Value for People and Businesses," *Business Insider*, 7 Dec. 2013; www.businessinsider.com/growth-in-the-internet-of-things-2013-10?IR=T.
2. F. Bonomi et al., "Fog Computing and Its Role in the Internet of Things Characterization of Fog Computing," *Proc. Mobile Cloud Computing*, 2012, pp. 13–15.

Nicola Iotti is a founding member and CTO of Guglielmo and Caligoo. His research focuses on advanced platforms for public Internet access networks. Iotti has an MSc in electronic engineering, with an emphasis in networking, from the University of Parma, Italy. He's an editorial board member of the *International Journal for Engineering, Information and Communications*. Contact him at nicola.iotti@caligoo.com.

Marco Picone is a co-founder and the head of the Mobile Computing group at Caligoo. He also is an adjunct professor in the Department of Information Engineering at the University of Parma, and a visiting researcher in the NetOS group at the Computer Laboratory for the University of Cambridge. His research interests include distributed systems, Internet of Things (IoT), fog computing, vehicular networks, and pervasive and mobile computing. Picone has a PhD in information technologies from the University of Parma. Contact him at marco.picone@caligoo.com.

Simone Cirani is a co-founder and head of the IoT group at Caligoo. He also is an adjunct professor in the Department of Information Engineering at the University of Parma. His research interests are the IoT, distributed systems, pervasive and mobile computing, and network security. Cirani has a PhD in information technologies from the University of Parma, Italy. He's a member of the IEEE Computer Society and IEEE Communications Society. Contact him at simone.cirani@caligoo.com.

Gianluigi Ferrari is an associate professor of telecommunications in the Department of Information Engineering and Architecture, and coordinates the IoT Lab (http://iotlab.unipr.it) at the University of Parma, Italy. He is a co-founder and president of Things2i, a spin-off company of the University of Parma, dedicated to IoT and smart systems. His research interests include signal processing, advanced networking/communication, and IoT and smart systems. Ferrari has a PhD in information technologies from the University of Parma, Italy. He's an IEEE senior member. Contact him at gianluigi.ferrari@unipr.it.