Estimating the end-to-end energy consumption of low-bandwidth IoT applications for WiFi devices

Loic Guegan and Anne-Cécile Orgerie

Univ Rennes, Inria, CNRS, IRISA, Rennes, France
Emails: loic.guegan@irisa.fr, anne-cecile.orgerie@irisa.fr

Abstract. Information and Communication Technology takes a growing part in the worldwide energy consumption. One of the root causes of this increase lies in the multiplication of connected devices. Each object of the Internet-of-Things often does not consume much energy by itself. Yet, their number and the infrastructures they require to properly work have leverage. In this paper, we combine simulations and real measurements to study the energy impact of IoT devices. In particular, we analyze the energy consumption of Cloud and telecommunication infrastructures induced by the utilization of connected devices, And we propose an end-to-end energy consumption model for these devices.

1 Introduction

In 2018, Information and Communication Technology (ICT) was estimated to absorb around 3% of the global energy consumption [1]. This consumption is estimated to grow at a rate of 9% per year [1]. This alarming growth is explained by the fast emergence of numerous new applications and new ICT devices. These devices supply services for smart building, smart factories and smart cities for instance, providing optimized decisions based on data produced by smart devices. All these connected devices constitute the Internet of Things (IoT): connected devices with sensors producing data, actuators interacting with their environment and communication means.

This increase in number of devices implies an increase in the energy needed to manufacture and utilize all these devices. Yet, the overall energy bill of IoT also comprises indirect costs as it relies on computing and networking infrastructures that consume energy to enable smart services. Indeed, IoT devices communicate with Cloud computing infrastructures to store, analyze and share their data.

In February 2019, a report by Cisco stated that "IoT connections will represent more than half (14.6 billion) of all global connected devices and connections (28.5 billion) by 2022" [2]. This will represent more than 6% of global IP traffic, against 3% in 2017 [2]. This increasing impact of IoT devices on Internet connections induces a growing weight on ICT energy consumption.

The energy consumption of IoT devices themselves is only the top of the iceberg: their use induce energy costs in communication and cloud infrastructures. In this paper, we estimate the overall energy consumption of an IoT device environment by combining simulations and real measurements. We focus on a given application with low bandwidth requirement and we evaluate its overall energy consumption: from the device, through telecommunication networks, and up to the Cloud data center hosting the application. From this analysis, we derive an end-to-end energy consumption model that can be used to assess the consumption of other IoT devices.

While some IoT devices produce a lot of data, like smart vehicles for instance, many others generate only a small amount of data, like smart meters. However, the scale matters here: many small devices can end up producing big data volumes. As an example, according to a report published by Sandvine in October 2018, the Google Nest Thermostat is the most significant IoT device in terms of worldwide connections: it represents 0.16% of all connections, ranging 55th on the list of connections [3]. As a comparison, the voice assistants Alexa and Siri are respectively 97th and 102nd with 0.05% of all connections [3]. This example highlights the growing importance of low-bandwidth IoT applications on Internet infrastructures, and consequently on their energy consumption.

In this paper, we focus on IoT devices for low-bandwidth applications such as smart meters or smart sensors. These applications send few data periodically to cloud servers, either to store them or to get computing power and take decisions. This is a first step towards a comprehensive characterization of the IoT energy footprint. While few studies address the energy consumption of high-bandwidth IoT applications [7], to the best of our knowledge, none of them targets low-bandwidth applications, despite their growing importance on the Internet infrastructures.

Low-bandwidth IoT applications, such as the Nest Thermostat, often relies on sensors powered by batteries. For such sensors, reducing their energy consumption is a critical target. Yet, we argue that end-to-end energy models are required to estimate the overall impact of IoT devices and to understand how to reduce their complete energy consumption. Without such models, one could optimize the consumption of on-battery devices at a heavier cost for cloud servers and networking infrastructures, resulting on an higher overall energy consumption. Using end-to-end models could prevent these unwanted effects.

Our contributions include:

- a characterization of low-bandwidth IoT applications;
- an analysis of the energy consumption of a low-bandwidth IoT application including the energy consumption of the IoT device and the consumption induced by its utilization on the Cloud and telecommunication infrastructures;
- an end-to-end energy model for low-bandwidth IoT applications.

The paper is organized as follows. Section 2 presents the state of the art. The low-bandwidth IoT application is characterized in Section 3, and details on its architecture are provided in Section 4. Section 5 provides our experimental results using real measurements and simulations. Section 6 discusses the key findings an the end-to-end energy model. Finally, Section 7 concludes this work and presents future work.

2 Related Work

2.1 Energy consumption of IoT devices

Smart apps and devices everywhere

Smart industry [4] : archi with sensing devices, cloud server, user applications and networks

IoT archi : devices, gateways, fog and clouds [?]

Smart cities [5]

Smart building [?]

home automation, smart agriculture, eHealth, logistics, smart grids

product life-cycle energy management [?]

focusing on access network technologies [?],

improving device transmission [?]

modeling the energy consumption of WSN devices [?] or the WiFi transmission [?]

on organizing wireless sensor communications to increase the network lifetime [4]

CO_2 impact of IoT and fog computing architectures vs Cloud [?]

Fog archi to use more renewable energy [7] or reduce communication costs [8]

2.2 Energy consumption of network and cloud infrastructures

net models server models + VM sharing

3 Characterization of low-bandwidth IoT applications

3.1 Application Characteristic

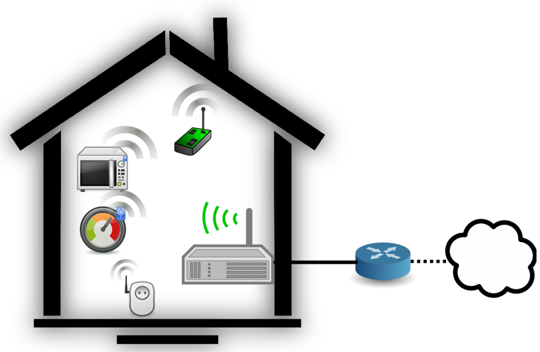

Fig. 1. Overview of IoT devices.

3.2 Cloud Infrastructure

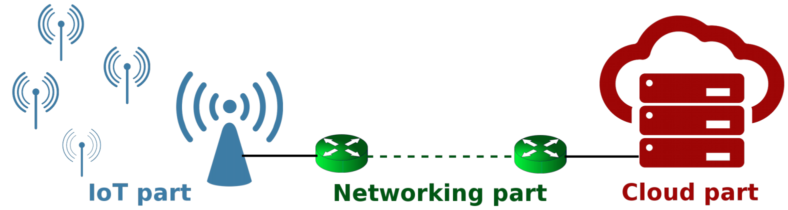

Fig. 2. Overview of the IoT architecture.

4 Experimental setup

Ajouter % de bande passante utilisé par les applis low-rate Our system model is divided in three parts. First, the IoT and the network parts are modeled through simulations. Then, the Cloud part is modeled using real servers connected to wattmeters. In this way, it is possible to evaluate the end-to-end energy consumption of the system.

4.1 IoT Part

In the first place, the IoT part is composed of several sensors connected to an Access Point (AP) which form a cell. This cell is evaluated using the ns-3 network simulator. Consequently, we setup between 5 and 15 sensors connected to the AP using WiFi 5GHz 802.11n. The node are placed randomly in a rectangle of $400m^2$ around the AP which corresponds to a typical real use case. All the cell nodes are setup with the default WIFI energy model provided by ns-3. The different energy values used by the energy model are provided on Table 1. These energy were extracted from previous work[6,7] on 802.11n. Besides, we suppose that the energy source of each nodes are unlimited and thus each of them can communicate until the end of all the simulations.

As a scenario, sensors send 192 bits packets to the AP composed of: **1)** A 128 bits sensors id **2)** A 32 bits integer representing the temperature **3)** An integer timestamp representing the temperature sensing time to store them as time series. The data are transmitted immediately at each sensing interval I varied from 1s to 60s. Finally, the AP is in charge of relaying data to the cloud via the network part.

4.2 Network Part

The network part represents the a network section starting from the AP to the Cloud excluding the server. It is also model into ns-3. We consider the server

Table 1. Simulations Energy Parameters

(a) Wifi

Parameter	Value
Supply Voltage	3.3V
Tx	0.38A
Rx	0.313A
Idle	0.273A

(b) Network

Parameter	Value
Idle	1W
Bytes (Tx/Rx)	3.4nJ
Pkt (Tx/Rx)	192.0nJ

to be 9 hops away from the AP with a typical round-trip latency of 100ms from the AP to the server. Each node from the AP to the Cloud is assume to be network switches with static and dynamic network energy consumption. The first 8 hop are edge switches and the last one is consider to be a core switch as mention in [8]. ECOFEN [9] is used to model the energy consumption of the network part. ECOFEN is a ns-3 network energy module dedicated to wired network. It is based on an energy-per-bit model including static energy consumption by assuming a linear relation between the amount of data sent to the network interface and its power consumption. The different energy values used to instantiate the ECOFEN energy model for the network part are shown in Table 1(b) and come from previous work [10].

4.3 Cloud Part

Finally, to measure the energy consumed by the server, we used real server from the large-scale test-beds Grid5000 (G5K). In fact, G5K has a cluster called Nova composed of several nodes which are connected to watt-meters. In this way, we can benefit from real energy measurements. The server used in the experiment include an Intel Xeon E5-2620 processor with 64 GB of RAM and 600GB of disk space on a Linux based operating system. This server is configured to use KVM as virtualization mechanism. We deploy a classical Linux x86_64 distribution on the Virtual Machine (VM) along with a MySQL database. We used different amount of allocated memory for the VM namely 1024MB/2048MB/4096MB to highlight its effects on the server energy consumption.

The sensors requests are simulated using another server. This server is in charge to send hundred of requests to the VM in order to fill the database. Consequently, it is easy to vary the different requests characteristics namely: **1)** The number request, to virtually add/remove sensors **2)** The requests interval.

5 Evaluation

5.1 IoT/Network Consumption

In a first place, we start by studying the impact of the sensors position on their energy consumption. To this end, we run several simulations in ns-3 with different sensors position. The results provided by Table 2 show that sensors position have a very low impact on the energy consumption and on the application delay. It has an impact of course, but it is very limited. This due to the fact that in such a scenario with very small number of communications spread over the time, sensors don't have to contend for accessing to the Wifi channel.

Table 2. Sensors send interval effects

Sensors Send Interval	10s	30s	50s	70s	90s
Sensors Power Consumption	13.51794W	13.51767W	13.51767W	13.51767W	13.51761W
Network Power Consumption	10.44178W	10.44167W	10.44161W	10.44161W	10.44161W
Average Appplication Delay	17.81360s	5.91265s	3.53509s	2.55086s	1.93848s

Previous work [7] on similar scenario shows that increasing application accuracy impact strongly the energy consumption in the context of data stream analysis. However, in our case, application accuracy is driven by the sensing interval and thus, the transmit frequency of the sensors. Therefore, we varied the transmission interval of the sensors from 1s to 60s. Some of these results are proposed on Table 2. In case of small and sporadic network traffic, these results show that with a reasonable transmission interval the energy consumption of the IoT/Network if almost not affected by the variation of this transmission interval. In fact, transmitted data are not large enough to leverage the energy consumed by the network.

The number of sensors is a dominant factor that leverage the energy consumption of the IoT/Network part. Therefore, we varied the number of sensors in the Wifi cell to analyze its impact. The Figure 3 represents the energy consumed by each simulated part according the the number of sensors. It is clear that the energy consumed by the network is the dominant part. However, since the number of sensors is increasing the energy consumed by the network will become negligible face to the energy consume by the sensors. In fact, deploying new sensors in the cell do not introduce much network load. To this end, sensors energy consumption is dominant.

5.2 Cloud Energy Consumption

In this End-To-End energy consumption study, cloud account for a huge part of the overall energy consumption. According a report [11] on United States data center energy usage, the average Power Usage Effectiveness (PUE) of an hyperscale data center is 1.2. Thus, in our analysis, all energy measurement on cloud server will account for this PUE.

In a first place, we analyze the impact of the VM allocated memory on the server energy consumption. Figure 4 depict the server energy consumption according to the VM allocated memory for 20 sensors sending data every 10s. Note that horizontal red line represent the average energy consumption for the

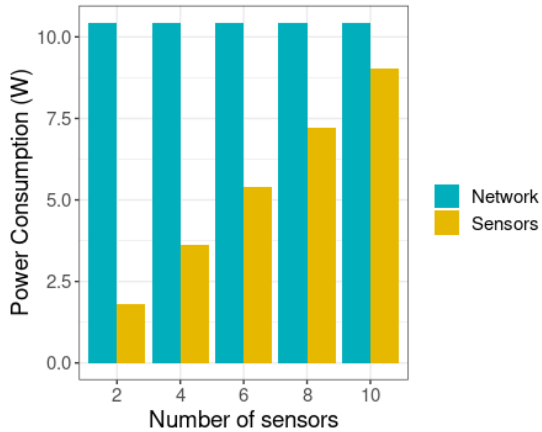

Fig. 3. Analysis of the variation of the number of sensors on the IoT/Network part energy consumption.

considered sample of energy values. We can see that at each sensing interval, server face to peaks of energy consumption. However, VM allocated memory do not influence energy consumption. In fact, simple database requests do not need any particular huge memory access and processing time. Thus, remaining experiments are based on VM with 1024MB of allocated memory.

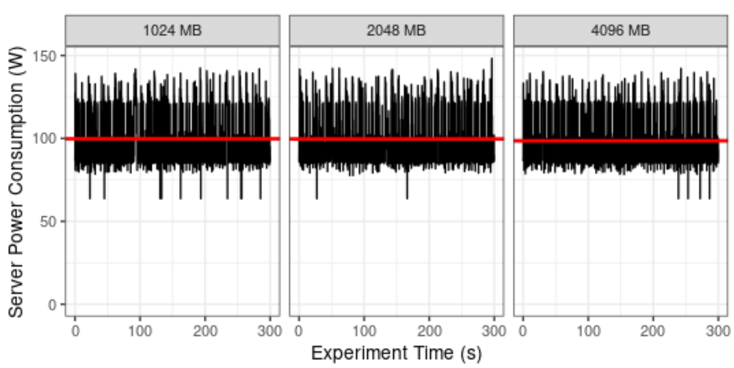

Fig. 4. VM size impact on the server energy consumption using 20 sensors sending data every 10s

Next, we study the effects of increasing the number of sensors on the server energy consumption. Figure 5(a) present the results of the average server energy consumption when varying the number of sensors from 20 to 500 while Figure 5(b) present the average server energy cost per sensors according to the number of sensors. These results show a clear linear relation between the number of sensors and the server energy consumption. Moreover, we can see that the more sensors we have per server, the more energy we can save. In fact, since the idle server energy consumption is high, it is more energy efficient to maximize the number of sensors per server. As shown on Figure 5(b), a significant amount of energy can be save when passing from 20 to 300 sensors per server.

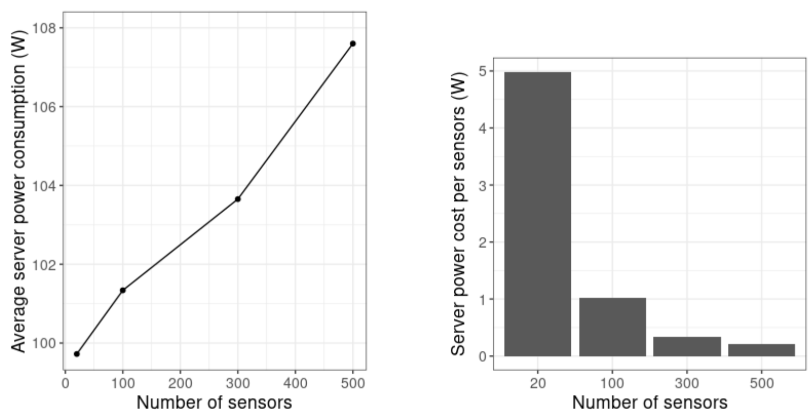

(a) Average server energy consumption

(b) Average sensors energy cost on server

Fig. 5. Server energy consumption for sensors sending data every 10s

A last parameter can leverage server energy consumption namely sensors send interval. In addition to increasing the application accuracy, sensors send interval increase network traffic and database accesses. Figure 5.2 present the impact on the server energy consumption of changing the send interval of 50 sensors to 1s, 10s and 30s. We can see that, the lower sensors send interval is, the more server energy consumption peaks occurs. Therefore, it leads to an increase of the server energy consumption.

5.3 End-To-End Consumption

To have an overview of the energy consume by the system, it is important to consider the end-to-end energy consumption. The Figure 5.3 represents the end-to-end system energy consumption while varying the number of sensors. It is important to see that, for small-scale systems, the server energy consumption is dominant face to the energy consumed by the sensors. However, since we are using a single server, large-scale sensors deployment lead to an increasing consumption of energy in the IoT part. On the other side, network energy consumption is stable regarding to the number of sensors since the system use case do not required large data transfer. Thus, it is important to remember that, to

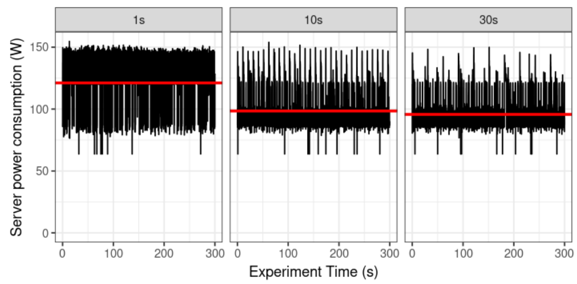

Fig. 6. Server energy consumption for 50 sensors sending request at different interval.

save energy, we should maximize the number of sensors handle by each cloud server while keeping reasonable sensors request intervals.

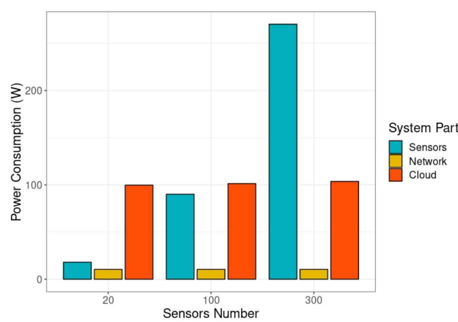

Fig. 7. End-to-end network energy consumption using sensors interval of 10s

6 Discussion

7 Conclusion

References

1. T. S. Project, "Lean ICT, Pour une sobriété numérique," https://theshiftproject.org/article/pour-une-sobriete-numerique-rapport-shift/, Oct. 2018.
2. Cisco, "Cisco Visual Networking Index: Forecast and Trends, 2017–2022, White paper," https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html, Feb. 2019.
3. Sandvine, "The Global Internet Phenomena Report," https://www.sandvine.com/phenomena, Oct. 2018.

4. K. Wang, Y. Wang, Y. Sun, S. Guo, and J. Wu, "Green Industrial Internet of Things Architecture: An Energy-Efficient Perspective," *IEEE Communications Magazine*, vol. 54, no. 12, pp. 48–54, 2016.
5. W. Ejaz, M. Naeem, A. Shahid, A. Anpalagan, and M. Jo, "Efficient energy management for the internet of things in smart cities," *IEEE Communications Magazine*, vol. 55, no. 1, pp. 84–91, 2017.
6. D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, "Demystifying 802.11n Power Consumption," p. 5.
7. Y. Li, A.-C. Orgerie, I. Rodero, B. L. Amersho, M. Parashar, and J.-M. Menaud, "End-to-end energy models for Edge Cloud-based IoT platforms: Application to data stream analysis in IoT," *Future Generation Computer Systems*, vol. 87, pp. 667–678, Oct. 2018. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0167739X17314309
8. F. Jalali, K. Hinton, R. Ayre, T. Alpcan, and R. S. Tucker, "Fog Computing May Help to Save Energy in Cloud Computing," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 5, pp. 1728–1739, May 2016. [Online]. Available: http://ieeexplore.ieee.org/document/7439752/
9. A. C. Orgerie, L. Lefèvre, I. Guérin-Lassous, and D. M. L. Pacheco, "ECOFEN: An End-to-end energy Cost mOdel and simulator For Evaluating power consumption in large-scale Networks," in *2011 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, Jun. 2011, pp. 1–6.
10. B. F. Cornea, A. C. Orgerie, and L. Lefèvre, "Studying the energy consumption of data transfers in Clouds: the Ecofen approach," in *2014 IEEE 3rd International Conference on Cloud Networking (CloudNet)*, Oct. 2014, pp. 143–148.
11. A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United States Data Center Energy Usage Report," Tech. Rep. LBNL–1005775, 1372902, Jun. 2016. [Online]. Available: http://www.osti.gov/servlets/purl/1372902/