

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

Loic Guegan, Anne-Cécile Orgerie

► **To cite this version:**

Loic Guegan, Anne-Cécile Orgerie. Estimating the end-to-end energy consumption of low-bandwidth IoT applications for WiFi devices. CloudCom 2019 - 11th IEEE International Conference on Cloud Computing Technology and Science, Dec 2019, Sydney, Australia. hal-02352637

HAL Id: hal-02352637

<https://hal.archives-ouvertes.fr/hal-02352637>

Submitted on 6 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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
Email: {loic.guegan, anne-cecile.orgetie}@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.

Index Terms—IoT devices, energy consumption, clouds, end-to-end model

I. 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 applications and new ICT devices. These devices supply services for smart building, smart factories and smart cities for instance. Through connected sensors producing data, actuators interacting with their environment and communication means – all being parts of the Internet of Things (IoT) – they provide optimized decisions.

This increase in number of devices implies an increase in the energy needed to manufacture and utilize them. 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 employ 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 in 2022, 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 requirements, 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 devices 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 global IoT energy footprint. While few studies address the energy consumption of high-bandwidth IoT applications [4], 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. Indeed, shifting computations to the cloud is often used to reduce the consumption of IoT devices [5], without studying the additional cost for the cloud infrastructure. Consequently, such an energy-saving technique, from the IoT device point of view, can result on a 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 WiFi 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 relying on WiFi devices.

The paper is organized as follows. Section II presents the state of the art. The low-bandwidth IoT application is characterized in Section III, and details on its architecture are provided in Section IV. Section V provides our experimental results combining real measurements and simulations. Section VI discusses the key findings on the end-to-end energy model. Finally, Section VII concludes this work and presents future work.

II. RELATED WORK

A. Energy consumption of IoT devices

The multiplication of smart devices and smart applications pushes the limits of Internet: IoT is now used everywhere for home automation, smart agriculture, e-health, smart cities, logistics, smart grids, smart buildings, etc. [6], [7], [8]. IoT devices are typically used to optimize processes and the envisioned application domains include the energy distribution and management. It can for instance help the energy management of product life-cycle [9]. Yet, few studies address the impact of IoT itself on global energy consumption [10], [4] or CO₂ emissions [11].

The underlying architecture of these smart applications usually includes sensing devices, cloud servers, user applications and telecommunication networks. Concerning the computing part, the cloud servers can either be located on Cloud data centers, on Fog infrastructures located at the network edge, or on home gateways [6]. Various network technologies are employed by IoT devices to communicate with their nearby gateway; either wired networks with Ethernet or wireless networks: WiFi, Bluetooth, Near Field Communication (NFC), ZigBee, cellular network (like 3G, LTE, 4G), Low Power Wide Area Network (LPWAN), etc. [12], [13]. The chosen technology depends on the smart device characteristics and the targeted communication performance. The Google Nest Thermostat can for instance use WiFi, 802.15.4 and Bluetooth [14]. In this paper, we focus on WiFi as it is broadly available and employed by IoT devices [12], [15].

Several works aim at reducing the energy consumption of the device transmission [16] or improving the energy efficiency of the access network technologies [13]. An extensive literature exists on increasing the lifetime of battery-based wireless sensor networks [5], [6]. Yet, IoT devices present more diversity than typical wireless sensors in terms of hardware characteristics, communication means and data production patterns.

Based on real measurements, previous studies have proposed energy models for IoT devices. Yet, these models are specific to a given kind of IoT device or a given transmission technology. Martinez et al. provide energy consumption measurements for wireless sensor networks using SIGFOX transmissions and employed for smart-parking systems [17].

Wu et al. implement an energy model for WiFi devices in the well-known ns3 network simulator [15].

These models can be used to evaluate the energy efficiency of communication protocols or computation offloading techniques [5]. However, they do not provide an overall view of the energy consumption of the entire system architecture: from the IoT device to the cloud server. To the best of our knowledge, one previous work targets an end-to-end energy model for IoT devices [4]. However, this work focus on high-bandwidth IoT devices with data streaming-oriented applications. This study shows that, in this case (high-bandwidth IoT applications), the cloud server hosting the application consumes more energy per IoT devices than the device itself (an IP camera in the case study) [4]. In our context of low-bandwidth devices, conclusions could be the opposite as the IoT devices' consumption is optimized since they are often powered through batteries.

B. Energy consumption of network and cloud infrastructures

IoT architecture rely on telecommunication networks and Cloud infrastructures to provide services. The data produced by IoT devices are stored and exploited by servers located either in Cloud data centers or Fog edge sites. While studies exist on the energy consumption of network and cloud infrastructures in general [18], they do not consider the specific case of IoT devices. To the best of our knowledge, no study estimates the direct impact of IoT applications on the energy consumption of these infrastructures.

Most work focusing on energy consumption, Cloud architecture and IoT applications tries to answer the question: where to locate data processing in order to save energy [10], to reduce the CO₂ impact [11], or to optimize renewable energy consumption [4].

In both cases, the network and cloud infrastructures, attributing the energy consumption to a given user or application is a challenging task. The complexity comes from the shared nature of these infrastructures: a given Ethernet port in the core of the network processes many packets coming from a high number of sources [10]. Moreover, the employed equipment is not power proportional: servers and routers consume consequent amounts of energy while being idle [19], [4]. The power consumed by a device is divided into two parts: a dynamic part that varies with traffic or amount of computation to process, and a static part that is constant and dissipated even while being idle [18]. This static part implies that a consequent energy cost of running an application on a server is due to the device being simply powered on. Consequently, sharing these static energy costs among all the concerned users requires an end-to-end model [4].

In this paper, we focus on IoT devices using WiFi transmission and generating low data volumes. Our model, extracted from real measurements and simulations, can be adapted to other kinds of devices and transmission technologies.

III. CHARACTERIZATION OF LOW-BANDWIDTH IOT APPLICATIONS

In this section, we detail the characteristics of the considered IoT application. While the derived model is more generic, we

focus on a given application to obtain a precise use-case with accurate power consumption measurements.

A. IoT device side

The Google Nest Thermostat relies on five sensors: temperature, humidity, near-field activity, far-field activity and ambient light [14]. Periodical measurements, sent through wireless communications on the Internet, are stored on Google data centers and processed to learn the home inhabitants habits. The learned behavior is employed to automatically adjust the home temperature managed by heating and cooling systems.




Fig. 1. Overview of IoT devices.

Each IoT device senses periodically its environment. Then, it sends the produced data through WiFi (in our context) to its gateway or Access Point (AP). The AP is in charge of transmitting the data to the cloud using the Internet. Figure 1 illustrates this architecture. Several IoT devices can share the same AP in a home. We consider low-bandwidth applications where devices produces several network packets during each sensing period. The transmitting frequency can vary from one to several packet sent per minute [2].

B. Cloud server side

We consider that the link between the AP and the Cloud is composed of several network switches and routers using Ethernet as shown in Figure 2. The number of routers on the path depends on the location of the server, either in a Cloud data center or in a Fog site at the edge of the network.

We assume that the server hosting the application data for the users belongs to a shared cloud facility with classical service level agreement (SLA). The facility provides redundant storage and computing means as virtual machines (VMs). A server can host several VMs at the same time.




Fig. 2. Overview of the IoT architecture.

The Cloud part of the application gathers the data sent by the IoT devices. These data are treated either on the fly (e.g.

threshold detection) or periodically, and action commands are sent back to the device if required. For instance, if the user has set a targeted temperature, the connected thermostat sends the measured temperature regularly, and once the target is reached, the Cloud server detects it, and sends back to the IoT device the command to pause the heater.

In the following, we describe the experimental setup, the results and the derived end-to-end model. For all these steps, we decompose the overall IoT architecture into three parts: the IoT device part, the networking part and the cloud part, as displayed on Figure 2.

IV. EXPERIMENTAL SETUP

In this section, we describe the experimental setup employed to acquire energy measurements for each of the three parts of our system model. The IoT and the network parts are modeled through simulations. 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.

A. 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 studied using the ns3 network simulator. In the experimental scenario, we setup between 5 and 15 sensors connected to the AP using WiFi 5GHz 802.11n. The sensors are placed randomly in a rectangle of $400m^2$ around the AP which corresponds to a typical use case for a home environment.

All the cell sensors employ the default WIFI energy model provided by ns3. This model comprises different power levels depending on the state of the WiFi device (i.e. idle, transmitting, receiving). The power consumption of receiving and transmitting states depends on the data rate of the device at a given time. In this paper, we consider only one data rate as the target is low-bandwidth devices in a home environment. The different energy values used by the energy model are provided in Table I. These parameters were extracted from previous work [20], [4] on IEEE 802.11n. Besides, we suppose that the energy source of each sensor is not limited during the experiments. Thus, each sensor 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 date. They are stored as time series. The data are transmitted immediately at each sensing interval I that we vary from 1s to 60s. Finally, the AP is in charge of relaying data to the cloud via the network part.

B. Network Part

The network part represents the network section starting from the AP to the Cloud excluding the server. It is also modeled into ns3. We consider the server to be 9 hops away from the AP with a typical round-trip latency of 100ms from the AP to the server [4]. Each node from the AP to the Cloud is a network switch with static and dynamic network

TABLE I
SIMULATIONS ENERGY PARAMETERS

(a) IoT part		(b) Network part	
Parameter	Value	Parameter	Value
Supply Voltage	3.3V	Idle	0.00001W
Tx	0.38A	Bytes (Tx/Rx)	3.4nJ
Rx	0.313A	Pkt (Tx/Rx)	192.0nJ
Idle	0.273A		

energy consumption. The first 8 hops are edge switches and the last one is consider to be a core router as mentioned in [10]. ECOFEN [21] is used to model the energy consumption of the network part. ECOFEN is an ns3 network energy module dedicated to wired networks. It is based on an energy-per-bit and energy-per-packet model for the dynamic energy consumption [22], [23], and it includes also a static energy consumption. The different values used to instantiate the ECOFEN energy model for the network part are shown in left part of Table I and come from previous work [24].

C. Cloud Part

Finally, to measure the energy consumed by the Cloud part, we use a real server from the large-scale test-bed Grid’5000. Grid’5000 provides clusters composed of several servers which are connected to wattmeters. The wattmeters provide 50 instantaneous power measurements per second and per server. This way, we can benefit from real energy measurements. The server used in the experiment embeds two Intel Xeon E5-2620 v4 processors 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 Debian x86_64 distribution on the Virtual Machine (VM) along with a MySQL database. We use different amounts of allocated memory for the VM namely 1024MB/2048MB/4096MB to highlight its effects on the server energy consumption. The server only hosts this VM in order to easily isolate its power consumption.




Fig. 3. Grid’5000 experimental setup.

The data sent by the IoT devices are simulated using another server from the same cluster. This server is in charge of sending the data packets to the VM hosting the application in order to fill its database. In the following, each data packet coming from an IoT device and addressed to the application VM is called a request. Consequently, it is easy to vary the different application characteristics namely: **1)** The number of requests, to virtually add/remove sensors **2)** The requests interval, to study the impact of the transmitting frequency. Figure 3 presents this simulation setup. We consider here a simple IoT application able to store the sensed values and provide them upon request. We do not include any data mining or machine learning techniques as they are highly dependent on the targeted application and quality of service. They can be added a posteriori to the derived end-to-end model if they are known, or estimated from specific energy models.

V. EVALUATION

In this section, we analyze the experimental results. All the experiments concerning IoT devices and network parts (Table II and Figure 4) are based on simulations using ns3, while all the experiments on Cloud servers (Figures 5, 6, 7, and 8) are real measurements performed on the Grid’5000 experimental platform.

A. IoT and Network Power Consumption

In a first place, we start by studying the impact of the sensors’ transmission frequency on their energy consumption. To this end, we run several simulations in ns3 with 15 sensors using different transmission frequencies. The results provided by Table II show that the transmission frequency has a very limited impact on the energy consumption of sensor and network parts, and on the average end-to-end application delay. This is due to the fact that in such a scenario with very small number of communications spread over the time, sensors do not have to contend for accessing to the WiFi channel. Note that for the network part, we include the dynamic power consumption due to the traffic generated by the sensors themselves, and we split the static power consumption of the routers according to the utilization ratio taken by the sensors. This model is detailed in Section VI.

TABLE II
SENSORS TRANSMISSION INTERVAL EFFECTS WITH 15 SENSORS

Transm. Interval	Sensor Power	Network Power	Application Delay
10s	13.51794W	0.44188W	0.09951s
30s	13.51767W	0.44177W	0.10021s
50s	13.51767W	0.44171W	0.10100s
70s	13.51767W	0.44171W	0.10203s
90s	13.51761W	0.44171W	0.10202s

Previous work [4] on a similar scenario shows that increasing application accuracy impacts 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 transmission frequency of the sensors. In

our case with small and sporadic network traffic, these results show that with a reasonable transmission interval, the energy consumption of the IoT and the network parts are 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.

We then vary the number of sensors in the WiFi cell. Figure 4 represents the energy consumed by the sensor and the network (from the AP to the cloud) parts according to the number of sensors. Similarly to the results of Table II, the network part is almost not affected by the number of sensors as their traffic is negligible compared to the network devices capacities. Consequently, sensors energy consumption is dominant, as each sensor adds its own consumption.




Fig. 4. Analysis of the variation of the number of sensors on the IoT/Network part energy consumption for a transmission interval of 10s.

B. Cloud Energy Consumption

In this end-to-end energy consumption study, cloud accounts for a huge part of the overall energy consumption. According to a report [25] on United States data center energy usage, the average Power Usage Effectiveness (PUE) of an hyper-scale data center is 1.2. This metric accounts for indirect data center power costs, such as the cooling infrastructure and the power distribution losses.

In our analysis, we use the PUE to account for these costs and all energy measurement on cloud servers use it. It means that the power consumption of the server is multiplied by the PUE [18].

Firstly, we analyze the impact of the VM allocated memory on the server energy consumption. Figure 5 depicts the server energy consumption according to the VM allocated memory for 20 sensors sending data every 10s. Note that the horizontal red line represents the average energy consumption for the considered sample of energy values.

We can see that at each transmission interval, the server faces spikes of energy consumption. However, the amount of allocated memory to the VM does not significantly influence the server energy consumption. In fact, simple database requests do not need any particular heavy memory accesses and

processing time. Thus, remaining experiments are based on VM with 1024MB of allocated memory.

Here, for clarity's sake, we use VMs with one virtual CPU (i.e. one physical CPU core). The influence of the number of core on the server' energy consumption has been widely studied in the literature [26]. For a given application that scales smoothly when adding cores, the relation between number of cores and power consumption is linear. In other words, adding cores for the execution of a parallel application multiplies accordingly the dynamic energy consumption on the server part.

Next, we study the effects of increasing the number of sensors on the server energy consumption. Figure 6 presents the results of the average server energy consumption when varying the number of sensors from 20 to 500, while Figure 8 presents the average server energy cost per sensor 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 VM, the more energy we can save. In fact, since the server's idle power consumption is high (around 97 Watts), it is more energy efficient to maximize the number of sensors per server. As shown on Figure 8, a significant amount of energy can be save when passing from 20 to 300 sensors per VM. Note that these measurements are not the row measurements taken from the wattmeters: they include the PUE but they are not shared among all the VMs that could be hosted on this server.

For the studied server, its static power consumption (also called idle consumption) is around 83.2 Watts. We consider a PUE of 1.2, this value is taken from [25]}. Since traditionally cloud servers host several VMs at the same time, our model proportionally shares the static power consumption of the server among the VMs it can host, depending on their VM size (allocated CPU and RAM). This model is detailed in Section VI.

A last parameter can leverage server energy consumption, namely sensors transmission interval. In addition to increasing the application accuracy, sensors transmission frequency increases network traffic and database accesses. Figure 7 presents the impact on the server energy consumption when changing the transmission interval of 50 sensors to 1s, 10s and 30s. We can see that, the lower sensors transmission interval is, the more server energy consumption peaks occur. Therefore, it leads to an increase of the server energy consumption.

In the next section, we use the hints detailed here and extracted from the real and simulated experiments in order to provide an end-to-end energy model that can be used for low-bandwidth IoT applications.

VI. END-TO-END CONSUMPTION MODEL

To have an overview of the energy consumed by the overall system, it is important to consider the end-to-end energy consumption. We detail here the model used to attribute the energy consumption of our application for each part of the architecture.




Fig. 5. Server power consumption multiplied by the PUE (= 1.2) using 20 sensors sending data every 10s for various VM memory sizes

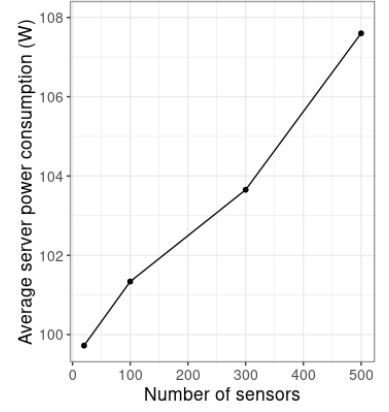


Fig. 6. Average server power consumption multiplied by the PUE (= 1.2) for sensors sending data every 10s




Fig. 7. Server power consumption multiplied by the PUE (= 1.2) for 50 sensors sending requests at different transmission interval.




Fig. 8. Average sensors power cost on the server hosting only our VM with PUE (= 1.2) for sensors sending data every 10s

For a given IoT device, we have:

- 1) For the IoT part, the entire consumption of the IoT device belongs to the system's accounted consumption.
- 2) For the network part, the data packets generated by the IoT device travel through network switches, routers and ports that are shared with other traffic.
- 3) For the cloud part, the VM hosting the data is shared with other IoT devices belonging to the same application and the server hosting the VM also hosts other VMs. Furthermore, the server belongs to a data center and takes part in the overall energy drawn to cool the server room.

Concerning the IoT part, we include the entire IoT device power consumption. Indeed, in our targeted low-bandwidth IoT application, the sensor is dedicated to this application. From Table I, one can derive that the static power consumption of one IoT sensor is around 0.9 Watts. Its dynamic part depends on the transmission frequency. So the power consumption of an IoT device:

$$\begin{aligned}
 P_{IoTdevice} &= P_{static}^{IoT} + P_{dynamic}^{IoT} \\
 &= \frac{P_{idle} \times (T - t_{RX} - t_{TX})}{T} + \frac{P_{RX} \times t_{RX} + P_{TX} \times t_{TX}}{T}
 \end{aligned}$$

where P_{static}^{IoT} and $P_{dynamic}^{IoT}$ are respectively the static and dynamic power consumption of the IoT device, t_{RX} , t_{TX} , and t_{idle} are the duration spent in each mode (receiving, transmitting and idle) and P_{RX} , P_{TX} , and P_{idle} the respective power consumption of each mode, and T is the transmission interval between two communications from the IoT device to the cloud server.

Concerning the sharing of the network costs, for each router, we consider its aggregate bandwidth (on all the ports), its average link utilization and the share taken by our IoT application. For a given network device, we compute our share of the static energy part as follows:

$$P_{netdevice}^{static} = \frac{P_{static}^{device} \times Bandwidth^{application}}{AggregateBandwidth^{device} \times LinkUtilization^{device}}$$

where P_{static}^{device} is the static power consumption of the network device (switch fabrics for instance or gateway), $Bandwidth^{application}$ is the bandwidth used by our IoT application, $AggregateBandwidth^{device}$ is the overall aggregated bandwidth of the network device on all its ports, and $LinkUtilization^{device}$ is the effective link utilization percentage. The $Bandwidth^{application}$ depends on the transmission frequency in our use-case. The formula includes the link utilization in order to charge for the effective energy cost per traffic and not for the theoretical upper bound which is the link bandwidth. Indeed, using such an upper bound leads to greatly underestimate our energy part, since link utilization typically varies between 5 to 40% [27], [28].

Similarly, for each network port, we take the share attributable to our application: the ratio of our bandwidth utilization over the port bandwidth multiplied by the link utilization and the overall static power consumption of the port. Table III summarizes the parameters used in our model, they are taken from [19], [27]. These are the parameters used in our formula to compute the values that we used in the simulations and that are presented in left part of Table I.

TABLE III
NETWORK DEVICES PARAMETERS

Device	Parameters
Gateway	Static power = 8.3 Watts Bandwidth = 54Mbps Utilization = 10%
Core router	Static power = 555 Watts 48 ports of 1 Gbps Utilization = 25%
Edge switch	Static power = 150 Watts 48 ports of 1 Gbps Utilization = 25%

The dynamic consumption of the network part includes a cost per packet and a cost per byte for each network device as detailed in [21]:

$$P_{dynamic}^{netdevice} = \frac{P_{byte}^{device} \times NbBytes + P_{pkt}^{device} \times NbPkts}{T}$$

with $NbBytes$ and $NbPkts$ respectively the number of bytes and packets sent by the application during one transmission interval and P_{byte}^{device} and P_{pkt}^{device} the power consumption per network device for each byte and each packet respectively.

For the sharing of the Cloud costs, we take into account the number of VMs that a server can host, the CPU utilization of a VM and the PUE. For a given Cloud server hosting our IoT application, we compute our share of the static energy part as follows:

$$P_{static}^{Cloudserver} = \frac{P_{static}^{server} \times PUE^{DataCenter}}{HostedVMs^{server}}$$

Where P_{static}^{server} is the static power consumption of the server, $PUE^{DataCenter}$ is the data center PUE, and $HostedVMs^{server}$ is the number of VMs a server can host. This last parameter should be adjusted in the case of VMs

with multiple virtual CPUs. We do not consider here over-commitment of Cloud servers. Yet, the dynamic energy part is computed with the real dynamic measurements, so it accounts for VM over-provisioning and resource under-utilization.

In our case, the Cloud server has 14 cores, which corresponds to the potential hosting of 14 small VMs with one virtual CPU each, and each vCPU is pinned to a server core. We consider that for fault-tolerance purpose, the IoT application has a replication factor of 2, meaning that two cloud servers store its database.

The Figure VI represents the end-to-end system energy consumption using the model described above while varying the number of sensors for a transmission interval of 10 seconds. The values are extracted from the experiments presented in the previous section.

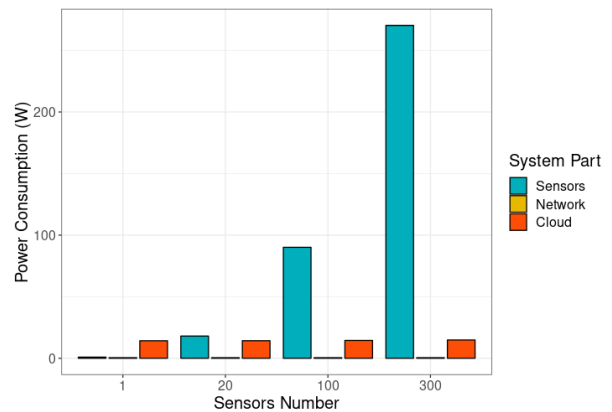


Fig. 9. End-to-end network energy consumption using sensors interval of 10s

Note that, for small-scale systems, with WiFi IoT devices, the IoT sensor part is dominant in the overall energy consumption. Indeed, the IoT application induces a very small cost on Cloud and network infrastructures compared to the IoT device cost. But, our model assumes that a single VM is handling multiple users (up to 300 sensors), thus being energy-efficient. Conclusions would be different with one VM per user in the case of no over-commitment on the Cloud side. For the network infrastructure, in our case of low-bandwidth utilization (one data packet every 10 seconds), the impact is almost negligible.

Another way of looking at these results is to observe that only for a high number of sensors (more than 300), the power consumption of Cloud and network parts start to be negligible (few percent). It means that, if IoT applications handle clients one by one (i.e. one VM per client), the impact is high on cloud and network part if they have only few sensors. The energy efficiency is really poor for only few devices: with 20 IoT sensors, the overall energy cost to handle these devices is almost doubled compared to the energy consumption of the IoT devices themselves. Instead of increasing the number of sensors, which would result in a higher overall energy consumption, one should focus on reducing the energy consumption of IoT devices, especially WiFi devices which

are common due to WiFi availability everywhere. One could also focus on improving the energy cost of Cloud and network infrastructure for low-bandwidth applications and few devices.

VII. CONCLUSION

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. Through the fine-grain analysis of a given low-bandwidth IoT device periodically sending data to a Cloud server using WiFi, we propose an end-to-end energy consumption model. This model provides insights on the hidden part of the iceberg: the impact of IoT devices on the energy consumption of Cloud and network infrastructures.

On our use-case, we show that for a given sensor, its larger energy consumption is on the sensor part. But the impact on the Cloud and network part is huge when using only few sensors with low-bandwidth applications. Consequently, with the IoT exploding growth, it becomes necessary to improve the energy efficiency of applications hosted on Cloud infrastructures and of IoT devices.

Our future work includes studying other types of IoT wireless transmission techniques that would be more energy-efficient. We also plan to study other IoT applications in order to increase the applicability of our model and provide advice for increasing the energy-efficiency of IoT infrastructures.

ACKNOWLEDGMENTS

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

REFERENCES

- [1] The Shift 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, Feb. 2019.
- [3] Sandvine, "The Global Internet Phenomena Report," <https://www.sandvine.com/phenomena>, Oct. 2018.
- [4] 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.
- [5] K. Kumar and Y. Lu, "Cloud Computing for Mobile Users: Can Offloading Computation Save Energy?" *Computer*, vol. 43, no. 4, pp. 51–56, 2010.
- [6] 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.
- [7] 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.
- [8] D. Minoli, K. Sohrawy, and B. Occhiogrosso, "IoT Considerations, Requirements, and Architectures for Smart Buildings—Energy Optimization and Next-Generation Building Management Systems," *IEEE Internet of Things Journal*, vol. 4, no. 1, pp. 269–283, 2017.
- [9] F. Tao, Y. Wang, Y. Zuo, H. Yang, and M. Zhang, "Internet of Things in product life-cycle energy management," *Journal of Industrial Information Integration*, vol. 1, pp. 26–39, 2016.
- [10] F. Jalali, K. Hinton, R. Ayre, T. Alpcan, and R. S. Tucker, "Fog Computing May Help to Save Energy in Cloud Computing," *IEEE J. on Selected Areas in Communications*, vol. 34, no. 5, pp. 1728–1739, 2016.
- [11] S. Sarkar, S. Chatterjee, and S. Misra, "Assessment of the Suitability of Fog Computing in the Context of Internet of Things," *IEEE Transactions on Cloud Computing*, vol. 6, no. 1, pp. 46–59, 2018.
- [12] F. Samie, L. Bauer, and J. Henkel, "IoT technologies for embedded computing: A survey," in *IEEE/ACM/IFIP CODES*, 2016.
- [13] C. Gray, R. Ayre, K. Hinton, and R. S. Tucker, "Power consumption of IoT access network technologies," in *IEEE International Conference on Communication Workshop (ICCW)*, 2015, pp. 2818–2823.
- [14] Google, "Nest Learning Thermostat – Spec Sheet," https://nest.com/downloads/press/documents/nest-thermostat-fact-sheet_2017.pdf, 2017.
- [15] H. Wu, S. Nabar, and R. Poovendran, "An Energy Framework for the Network Simulator 3 (NS-3)," in *International ICST Conference on Simulation Tools and Techniques (SIMUTools)*, 2011, pp. 222–230.
- [16] P. Andres-Maldonado, P. Ameigeiras, J. Prados-Garzon, J. J. Ramos-Munoz, and J. M. Lopez-Soler, "Optimized LTE data transmission procedures for IoT: Device side energy consumption analysis," in *IEEE International Conference on Communications Workshops (ICC Workshops)*, 2017, pp. 540–545.
- [17] B. Martinez, M. Montón, I. Vilajosana, and J. D. Prades, "The Power of Models: Modeling Power Consumption for IoT Devices," *IEEE Sensors Journal*, vol. 15, no. 10, pp. 5777–5789, 2015.
- [18] E. Ahvar, A.-C. Orgerie, and A. Lebre, "Estimating energy consumption of cloud, fog and edge computing infrastructures," *IEEE Trans. on Sust. Comp.*, 2019.
- [19] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A Power Benchmarking Framework for Network Devices," in *NETWORKING*, ser. Lecture Notes in Computer Science, 2009, pp. 795–808.
- [20] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, "Demystifying 802.11n Power Consumption," in *International Conference on Power Aware Computing and Systems (HotPower)*, 2010, p. 5.
- [21] A.-C. Orgerie, B. L. Amersho, T. Haudebourg, M. Quinson, M. Rifai, D. L. Pacheco, and L. Lefèvre, "Simulation Toolbox for Studying Energy Consumption in Wired Networks," in *CNSM: International Conference on Network and Service Management*, 2017, pp. 1–5.
- [22] V. Sivaraman, A. Vishwanath, Z. Zhao, and C. Russell, "Profiling per-packet and per-byte energy consumption in the NetFPGA Gigabit router," in *IEEE INFOCOM Workshops*, 2011, pp. 331–336.
- [23] P. Serrano, A. Garcia-Saavedra, G. Bianchi, A. Banchs, and A. Azcorra, "Per-Frame Energy Consumption in 802.11 Devices and Its Implication on Modeling and Design," *IEEE/ACM Trans. on Net.*, vol. 23, no. 4, pp. 1243–1256, 2015.
- [24] 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.
- [25] 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," LBNL, Tech. Rep. LBNL–1005775, 1372902, Jun. 2016.
- [26] F. C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A.-C. Orgerie, and M. Quinson, "Predicting the Energy-Consumption of MPI Applications at Scale Using Only a Single Node," in *IEEE Cluster Conference*, 2017, pp. 92–102.
- [27] A. Hassidim, D. Raz, M. Segalov, and A. Shaqed, "Network utilization: The flow view," in *IEEE INFOCOM*, 2013, pp. 1429–1437.
- [28] Z. Zhang, Y. Bejerano, and S. Antonakopoulos, "Energy-Efficient IP Core Network Configuration Under General Traffic Demands," *IEEE/ACM Trans. on Networking*, vol. 24, no. 2, pp. 745–758, 2016.