



Wiesner, P. and Thamsen, L. (2021) LEAF: Simulating Large Energy-Aware Fog Computing Environments. In: 2021 5th IEEE International Conference on Fog and Edge Computing (ICFEC), Melbourne, Australia, 10-13 May 2021, pp. 29-36. ISBN 9781665402910

(doi: [10.1109/ICFEC51620.2021.00012](https://doi.org/10.1109/ICFEC51620.2021.00012))

This is the Author Accepted Manuscript.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/268156/>

Deposited on: 10 June 2022

LEAF: Simulating Large Energy-Aware Fog Computing Environments

Philipp Wiesner and Lauritz Thamsen
Technische Universität Berlin, Germany
{wiesner, lauritz.thamsen}@tu-berlin.de

Abstract—Despite constant improvements in efficiency, today’s data centers and networks consume enormous amounts of energy and this demand is expected to rise even further. An important research question is whether and how fog computing can curb this trend. As real-life deployments of fog infrastructure are still rare, a significant part of research relies on simulations. However, existing power models usually only target particular components such as compute nodes or battery-constrained edge devices.

Combining analytical and discrete-event modeling, we develop a holistic but granular energy consumption model that can determine the power usage of compute nodes as well as network traffic and applications over time. Simulations can incorporate thousands of devices that execute complex application graphs on a distributed, heterogeneous, and resource-constrained infrastructure. We evaluated our publicly available prototype LEAF within a smart city traffic scenario, demonstrating that it enables research on energy-conserving fog computing architectures and can be used to assess dynamic task placement strategies and other energy-saving mechanisms.

Index Terms—Simulation, Modeling, Fog Computing, Edge Computing, Energy Consumption

I. INTRODUCTION

With the rise of the internet of things and the beginning of the fourth industrial revolution, many sectors and environments are expected to become interspersed with connected sensors. As these sensors produce an ever-increasing amount of data, the classical cloud computing approach of sending all information directly to a few powerful data centers, often far away from the data source, is not feasible anymore. One attempt to solve these emerging problems is to deploy a distributed layer of compute nodes close to the edge of the network that filters, aggregates, and preprocesses the raw data generated by end devices and sensors. This new paradigm, called fog computing, promises to save large amounts of bandwidth, reduce latencies for time-sensitive applications, and increase end-user privacy as the fog nodes can take care of aggregating and anonymizing data [1]. However, due to their distributed and heterogeneous nature, fog computing environments are much more challenging to manage than centralized data centers [2].

At the same time, energy consumption of information and communications technology is already accounting for more than 10 % of global energy consumption and is expected to exceed the 20 % mark by 2030 [3]. An open research question is whether fog computing will be able to reduce the overall energy consumption of future computing infrastructure or if it may even be responsible for a further increase [1], [4], [5].

Although the general trend of concentrating computing resources into hyper-scale data centers continues [6], research has shown that distributed architectures could save between 14 % to more than 80 % of total energy consumption compared to fully centralized architectures [7], [8].

A major part of today’s fog computing research is centered around modeling and simulations. Nevertheless, to the best of our knowledge, there exists no tool that is suitable to explore the overall power usage of large-scale fog deployments and enables research on energy-conserving architectures and algorithms in this domain. Major existing fog computing simulators provide no energy modeling at all [9], [10] or only focus on specific components of fog infrastructures [11], [12]. Furthermore, many simulators fail to realistically model fog environments by leaving out substantial aspects like mobility [10], [11]. Existing analytical approaches on the other hand provide only a high level of abstraction and cannot represent change over time, prohibiting the modeling of mobility and evaluation of online decision making algorithms [4], [7], [13].

We present LEAF¹, a new simulator for modeling Large Energy-Aware Fog computing environments. In contrast to the related work, our model was designed with the following four requirements in mind:

- (1) *Realistic Fog Computing Environment*: Enable the simulation of applications with complex task graphs in distributed, heterogeneous, and resource-constrained fog computing environments. Nodes can be mobile and they can join or leave the topology at any time.
- (2) *Holistic Energy Consumption Model*: Enable the separate assessment of power usage of compute nodes, network traffic, and applications, allowing for the creation of load profiles for individual parts of fog computing environments.
- (3) *Energy-Aware Online Decision Making*: Enable research on task placement strategies, routing policies, and energy-saving mechanisms that dynamically adapt their behavior based on the power usage of infrastructure and applications.
- (4) *Performance and Scalability*: Enable the simulation of complex scenarios with thousands of devices and applications at least two orders of magnitude faster than real-time on commodity hardware.

¹Available at <https://github.com/birnbaum/leaf>

Our model combines analytical and discrete-event modeling to enable easy-to-analyze large-scale experiments. Specifically, we model infrastructures and applications as variable graphs that can change throughout the simulation and do not rely on representing individual packets. This way, the model enables research on novel task placement strategies or energy-saving mechanisms that must be evaluated in large settings. We expect LEAF to help both researchers and practitioners assess the power usage characteristics of different cloud, fog, and edge computing architectures, which can lead to more well-founded decisions when planning future infrastructures.

The remainder of the paper is structured as follows. Section II discusses the related work. Section III introduces our infrastructure and application model for fog computing environments. Section IV explains the different kinds of power models that can be applied. Section V evaluates our model on a smart city scenario. Section VI concludes the paper.

II. RELATED WORK

Energy-efficient data centers have been a research area for years as electricity costs account for a significant part of a data center’s operating expenditure. In fog or edge computing, however, most research on power usage only concerns specific components such as battery-constrained edge devices [12] or fog nodes [11]. The bigger picture, namely the potential impact of fog computing on the overall energy consumption of information and communications technology, has retrieved comparably little attention.

One approach to assess the estimated consumption of fog computing environments is analytical modeling. These models are usually suitable for examining large-scale scenarios since they are fast to execute and convenient to experiment with. Sarkar et al. [13] compare service latency and energy consumption of fog and cloud computing by performing a high-level examination of different architectures. The authors focus on infrastructure only and incorporate no concept for applications. Furthermore, they do not distinguish between static and dynamic energy consumption, which limits the informative value of the analyses. Jalali et al. [4] model energy usage by using a flow-based model for highly shared equipment such as core routers and switches, and a time-based model for end-user equipment that is usually not shared among many clients. Their findings show that the power impact of fog computing depends on many factors such as the type of attached access network. However their scenario is a relatively specific use case and not suitable as a general purpose model. Ahvar et al. [7] propose a model that differentiates between static and dynamic power usage and incorporates the consumption of network devices and cooling. Even though their model is relatively comprehensive, it incorporates several critical simplifications such as an even distribution of VMs among data centers and the assumption that all fog nodes connect to the same number of devices. Moreover, there exists no concept of applications or tasks, nor is there a way to calculate the relative power requirements of VMs. All presented analytical approaches suffer from the fact that they model at a high level

of abstraction and, for example, make the invalid assumption that fog nodes are homogeneous [7], [13]. Moreover, they cannot represent change over time, prohibiting the assessment of mobility and online decision making algorithms.

Besides analytical models, a large number of discrete-event simulators has been developed in recent years. Although there exist cloud computing simulators that feature comprehensive power models like GreenCloud [14], these tools are not designed to model highly heterogeneous and variable environments like those encountered in fog computing.

Many popular simulators in the domain of fog and edge computing such as EdgeCloudSim [9] or IOTSim [10] comprise no power modeling at all. The simulators that do are either overly simplistic or focus only on specific components of the infrastructure. For example, iFogSim [11] inherits some power modeling capabilities from CloudSim and can model the power consumption of data centers and fog nodes. Nevertheless, there is currently no way to simulate the power usage of edge devices, networks, or applications. PureEdgeSim [12] features a comprehensive energy consumption model for edge devices, including Wi-Fi traffic, to simulate battery-constrained devices. However, the consumption of data centers, fog nodes or wide area network (WAN) traffic cannot be modeled. FogNetSim++ [15] is the only fog computing simulator to our knowledge that features an energy model for applications. Albeit the focus on performance, the detailed modeling of network traffic on a per packet basis inevitably results in worse scalability than our model. Furthermore, due to our analytical approach, it is a lot easier to set up experiments in LEAF, explore different configurations, and analyze results.

III. FOG ENVIRONMENT MODEL

The way LEAF models fog computing environments is inspired by analytical modeling. Infrastructure and applications are represented as graphs and power models as functions that operate on these graphs. Together, these components define the state of the simulation at a certain time step, which we call a *configuration*. In contrast to purely analytical models, the state of network and applications can change over time

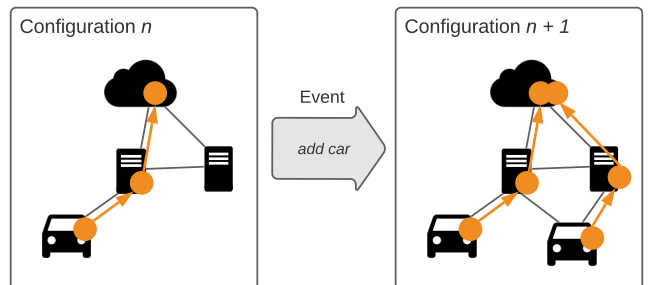


Fig. 1: LEAF uses discrete events to update the state (configuration) of the simulation. In this example, an additional compute node (a car) was added to the infrastructure graph and a new application was placed on this infrastructure.

through the use of events, as illustrated in Figure 1. LEAF uses events only in two cases: To read and to update configurations. Read events can, for example, be emitted by an algorithm that performs power measurements at a certain time step to adapt its behavior. Examples of update events are changes to the current application placements or the infrastructure graph like adapting the latency of a network link or adding a new compute node. LEAF does not use events to simulate individual network packets.

Every configuration contains all information necessary for a deeper analysis and can be inspected without any knowledge about future or past events. Modeling fog computing environments in this way enables easy-to-analyze experiments that can execute scenarios with thousands of devices and applications. Yet, our approach also has decisive advantages over purely analytical models: First, by allowing the model to change over time, heterogeneous environments with dynamically changing network topologies, mobility of edge devices, and varying workloads can be represented. Second, LEAF allows for online decision making, enabling research on dynamic energy-aware task placement strategies, scheduling algorithms, or traffic routing policies.

A. Infrastructure Model

As depicted in Figure 2, the network and computing infrastructure are represented as a weighted, directed multigraph $I = (N, L)$, where

N is the set of compute nodes. A compute node $N^i \in N$ may be resource-constrained, can have a location, and may be mobile. It can represent an entire cloud data center as well as mobile sensors with limited or no computing capacity.

L is the set of network links between nodes. A network link $L^i \in L$ can represent individual wired or wireless connections as well as entire network routes such as WAN connections. Network links can be constrained by bandwidth and may have additional properties like latency that users can utilize for implementing routing policies. Edges are directed to allow modeling different constraints and power models for uplink and downlink. Multiple edges between nodes are allowed and can represent different available link types such as Wi-Fi and Bluetooth.

The infrastructure graph can change during the simulation, as nodes are free to move around and can join or leave the network. Furthermore, any property, like the latency of a link, can change.

B. Application Model

LEAF considers all applications to be streaming applications and, hence, regards all communication between different application tasks as continuous data flows. Applications are represented by directed acyclic graphs $A = (T, F)$, where T is the set of tasks in the application, and F is the set of data flows between tasks. Tasks, as well as data flows, have certain resource requirements. Applications can be placed on the fog by mapping their tasks to compute nodes, and their data flows

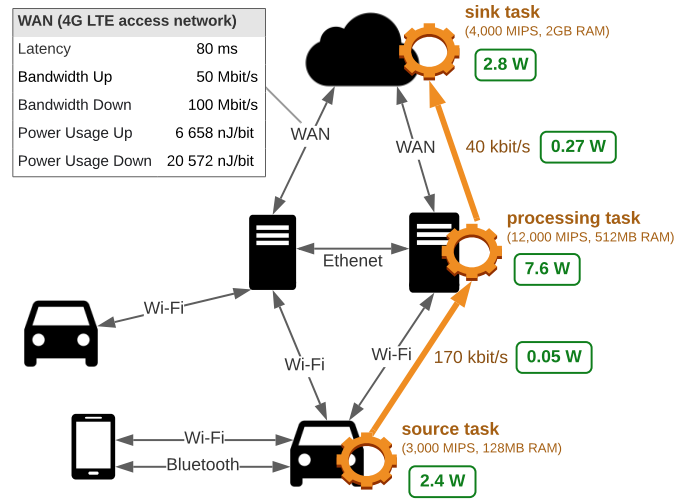


Fig. 2: Infrastructure graph with resource constraints and power usage characteristics (black), a placed application with resource requirements (orange), and its resulting power usage on the infrastructure components (green). In this configuration, the application’s combined power usage is 13.12 W.

to network links. Application placements are described in more detail below.

A task $T^i \in T$ demands certain resources from its compute node. A compute node can host as many tasks as its resources allow. There exist three different kinds of tasks that are responsible for emitting, processing, and retrieving data:

- 1) *Source tasks* are bound to a specific compute node, for example, a sensor-equipped or mobile device. They are constantly emitting data and have at least one outgoing edge and no incoming edges.
- 2) *Processing tasks* can be freely placed on any compute node $N^i \in N$ that fulfills their resource requirements. Processing tasks have at least one incoming edge and at least one outgoing edge.
- 3) *Sink tasks* are again bound to a specific compute node and have at least one incoming edge and no outgoing edges. Examples of data sinks are cloud storages or end-user terminals.

A data flow $F^i \in F$ represents a continuous stream of data between two tasks, stated in bits/s. Like any parameter, the data rate can change during the simulation. A network link can “host” as many data streams as its bandwidth allows. Due to its abstract network model, LEAF is not suitable to explicitly model complex effects resulting from network congestion like queuing delay or packet drops.

When placing an application, its tasks have to be mapped onto compute nodes, which is illustrated in orange in Figure 2. Since the assignment of source and sink tasks is fixed, task placement strategies have to determine the placement of processing tasks. After the placement, a user-defined routing policy searches for an optimal path in the network between every pair of connected tasks. If there is no node that meets

the resource requirements or if no sufficient path can be found, the placement failed.

Typically, processing tasks aggregate and filter the data stream in a way that outgoing data flows require less bandwidth than incoming data flows. Consequently, the placement of processing tasks close to the source tasks generally reduces network usage and, thus, energy consumption. Finding optimal placements in dynamic, resource-constrained, distributed, and heterogeneous environments is a complicated problem and a focus of research [16]. LEAF enables application placement strategies that take the infrastructure’s energy consumption into account.

IV. POWER MODELS

We model infrastructure power usage by assigning each compute node and network link its own power model. This enables users to determine the power usage of data centers, edge devices, and network links individually. A power model $P(t)$ is a function that returns the current power usage of its corresponding entity at any point in time t during the simulation. Power usage is reported as the sum $P_{static} + P_{dyn}$, where P_{static} is the static (load-independent) power usage and P_{dyn} the dynamic (load-dependent) fraction. Power models may maintain a state which can be used to implement energy-saving mechanisms (see Section IV-D).

To create load profiles of infrastructure components, their power models can be called through periodic events. Additionally, other simulated components can call power models at runtime to adapt their behavior or update their state. Examples of such components are energy-aware task placement strategies or hardware like batteries that update the state of charge.

A. Linear and Non-Linear Power Models

Research shows that the power usage of infrastructure can often be modeled with sufficient accuracy as a linear function that is only based on the entity’s current load $C(t)$ [17], [18]. Linear power models are beneficial since they are computationally efficient and easy to comprehend. They are described as

$$P(t) = P_{static} + C(t)\sigma,$$

where $C(t)\sigma$ represents P_{dyn} . The variable σ determines the incremental energy per load. For example, for network equipment, σ represents the energy consumed per bit of transferred data (J/bit). When modeling resource-constrained compute nodes or network links that have a maximum load C_{max} with power usage P_{max} , σ can be expressed as:

$$\sigma = \frac{P_{max} - P_{static}}{C_{max}}$$

An example for more complex entities, where a simple, linear function may not model the power usage well, are compute nodes representing entire data centers. Data centers incorporate several different hosts and require additional power for operational overhead like cooling and lighting. An exemplary approach for modeling this kind of compute node is to

define a power model P_{H_i} for each host H_i and to aggregate their results to P_{dc} . To account for operational overhead, we can multiply the result with the data center’s power usage effectiveness (PUE):

$$P_{dc}(t) = \text{PUE} \cdot \sum_{H_i \in H} P_{H_i}(t)$$

Compute nodes and network links can also apply non-linear power models and power models with multiple inputs to improve the accuracy of the simulation. An example is the power model for wireless network links of [19], which, in addition to the load, also takes the distance of communicating devices into account. However, more complex power models will lead to more computationally intensive simulations, hence longer runtimes.

B. Power Models for Shared Resources

When describing the power consumption of shared resources such as cloud data centers, mobile base stations, or entire WAN links, it is not meaningful to specify static power usage. For example, it is impossible to know at which point an externally operated data center will need to switch on or off a new physical machine. Many of these resources are assumed to scale on demand, hence defining a P_{max} is not even possible. For these reasons, $P_{static} = 0$ for power models of shared infrastructure. Users must estimate a meaningful σ that directly incorporates a fraction of the static consumption.

By modeling shared resources without static power usage, we implicitly assume this infrastructure to be energy-proportional. Nevertheless, there are ways to incorporate inefficiencies into the parameters. In practice, servers as well as routers are rarely used to full capacity but are over-provisioned to ensure performance and availability. Equivalently to the energy model proposed by [4], users can estimate and incorporate an operational capacity fraction U into a power model. The effect of U is illustrated in the staircase curve depicted in Figure 3.

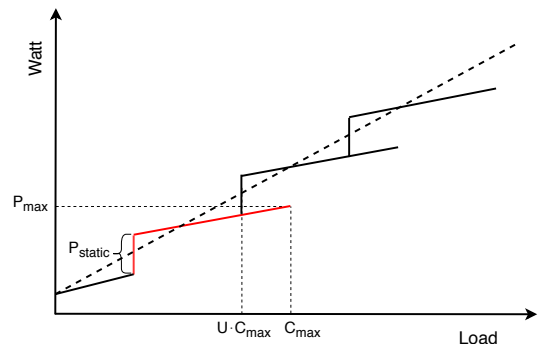


Fig. 3: Incremental performance per watt of scalable compute nodes or network links: The red section describes a sub-component such as a data center host, which is only utilized until $U \cdot C_{max}$ of its full capacity. Adapted from [4].

TABLE I: Incremental energy per bit (nJ/bit) of networking equipment adding up to an exemplary WAN link. Cisco hardware power consumption was determined with the Cisco Power Calculator².

Type	σ_{up}	σ_{down}
4G LTE module of mobile device [20]	438.4	52.0
4G LTE access network [21]	6200.0	20 500.0
Edge router (Cisco ASR 9010)	5.9	5.9
Core routers (5x Cisco CRS-3)	13.5	13.5
DC switches (2x Cisco Nexus 9500)	0.4	0.4
WAN link	6658.2	20 571.8

C. Power Models for Network Links

Although we assign power models directly to edges in the infrastructure graph, in reality, energy is not consumed by the cable or wireless connection itself but by networking equipment such as routers or wireless transmitters. Network link power models hence represent the aggregated power usage of all networking equipment involved to transfer data from one compute node to another, and σ describes the incremental energy per bit (J/bit). Static power usage should be attributed to the adjacent compute nodes, to assure it correctly drops to zero when the node is shut off. Intermediate routing devices are considered shared infrastructure and their consumption should be incorporated directly into the network link's σ . Table I depicts an exemplary parameterization of a WAN link between a mobile device and a cloud server.

For wireless connections, power usage characteristics of uplink and downlink usually differ [20], [21]. Since infrastructure is represented by a directed multigraph in LEAF, it allows modeling asymmetrical power usage characteristics by assigning different power models to the respective directions.

D. Energy-saving Mechanisms

Energy-saving mechanisms aim to reduce static power usage of hardware during periods of low utilization by throttling or powering off parts of the system. They can significantly influence overall energy consumption by pushing infrastructure towards operating energy-proportionally. Such load profiles can be modeled in LEAF by making P_{static} , which usually describes the load-independent power usage, indirectly load-dependent. For example, a common strategy in data centers is to consolidate workload on a few hosts and to power off the remaining hosts. P_{static} of such a data center could be modeled as a step function, similar to the example in Figure 3.

Furthermore, to evaluate the impact of even more complex energy-saving mechanisms, users can make use of the fact that power models can maintain state. For example, [22] presented a precise model for dynamic voltage and frequency scaling (DVFS), which requires the current and past utilization of the CPU to implement different strategies. LEAF's power models

can store past utilization states of their assigned entity and take them into account when computing the current power usage.

E. Tracing Back Power Usage to Applications

To support the implementation of energy-conserving application placement strategies, LEAF allows tracing back the infrastructure's power usage to the responsible applications. The power consumption of an application is defined as the sum of all power usage that its tasks and data flows cause on their allocated resources, see Figure 2. Depending on the scenario users may decide to only attribute dynamic power usage, or dynamic and static power usage. For example, if the underlying resource is powered on, regardless of its utilization, the static energy consumption does not have to be attributed to any applications since they have no influence on it. However, if the resource can be powered off when it is idle, the placed resources actually *cause* the energy usage. In this case, the static energy should be attributed to all tasks that currently run on the node depending on their fraction of the total load.

V. EVALUATION

To evaluate the proposed model, we created an open source prototype and simulated different fog architectures and application placement strategies in a smart city traffic scenario. The prototype is implemented in Java, based on the CloudSim Plus [23] simulator.

A. Experimental Setup

The urban environment of the experimental setup is inspired by Manhattan, namely a rectangular street grid with blocks of 274 m by 80 m. It covers an area containing 16 crossings, each equipped with a smart traffic light (STL) that optionally has a fog node attached. Within this city center, we simulate 24 hours of taxi traffic with time steps of one second. Quantity and speed of taxis were modeled according to a dataset published for the 2015 DEBS Grand Challenge competition³, which contains information on approximately 173 million taxi

³<http://www.debs2015.org/call-grand-challenge.html>, accessed 2020-12-15

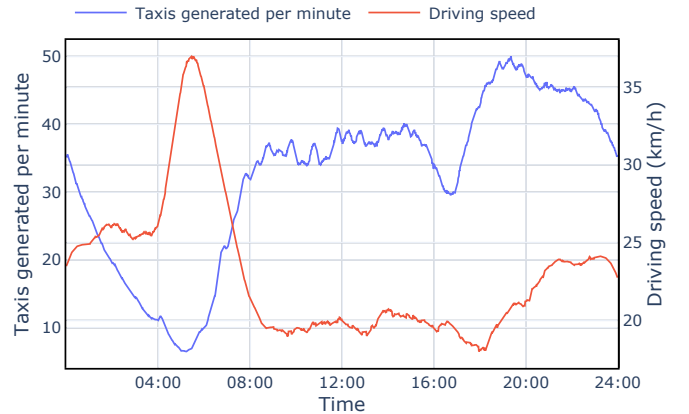


Fig. 4: Number of taxis generated per minute vs. average driving speed.

²<http://tools.cisco.com/cpc>, accessed 2020-10-17

TABLE II: Smart city infrastructure parameterization.

	Max load	P_{static}	σ
Cloud	∞	-	700 μ W/MIPS
Fog node	400 000 MIPS	100 W	350 μ W/MIPS
WAN _{STL \rightarrow Cloud}	50 Mbit/s	-	6658 nJ/bit
WAN _{Cloud \rightarrow STL}	100 Mbit/s	-	20 572 nJ/bit
Wi-Fi _{Taxi \rightarrow STL}	1.3 Gbit/s	-	300 nJ/bit
Wi-Fi _{STL \rightarrow STL}	1.3 Gbit/s	-	100 nJ/bit

trips that took place in New York in 2013. Figure 4 shows the average number and speed of taxis generated per minute.

The infrastructure graph consists of four kinds of compute nodes: The cloud, fog nodes, STLs, and taxis. Taxis are connected via Wi-Fi, e.g. *IEEE 802.11p*, to nearby STLs. The mapping of data flows to network edges is updated periodically as taxis move around. STLs communicate with other close-by STLs via Wi-Fi too, effectively forming a mesh network. Additionally, they are equipped with 4G LTE modules connecting them to the cloud node via WAN.

Table II shows the infrastructure parameterization. Only cloud and fog nodes have computing capacities stated in millions of instructions per second (MIPS). Since the load on STL and taxi nodes is the same in all experiments, their power usage is not modeled in this evaluation. All other compute nodes and network links have linear power models. The cloud is considered shared infrastructure and has no static but high dynamic power usage. Energy per bit for WAN links was configured as the example shown in Table I. Direct Wi-Fi communication between STLs consumes less energy than communication with taxis because they always have a direct line of sight and deploy energy-efficient, high-throughput access points.

We simulate a smart traffic system comprising two kinds of applications:

CCTV applications process data recorded by cameras deployed at STLs. The source task, located at the STL, sends 10 Mbit/s of video data to a processing task that requires 30 000 MIPS and is responsible for traffic monitoring, enforcing traffic laws and automatic incident detection. The 200 kbit/s of resulting data are sent to the sink task located in the cloud for further analysis and storage. 16 of these applications are running in our scenario, one for each STL.

V2I applications are vehicle-to-infrastructure, smart traffic management applications that control traffic lights to ensure maximum throughput of public transport. Each taxi on the map streams 100 kbit/s of sensor data to a processing task which requires 7000 MIPS and forwards 50 kbit/s to all STLs on the way of the taxi. The number of applications running depends on the number of taxis on the map.

B. Experiments

We conducted eight different experiments to demonstrate that our model enables research on energy-efficient fog architectures, task placement strategies and energy-saving mechanisms. Figure 5 provides an overview of the results. The experiments are explained and analyzed below.

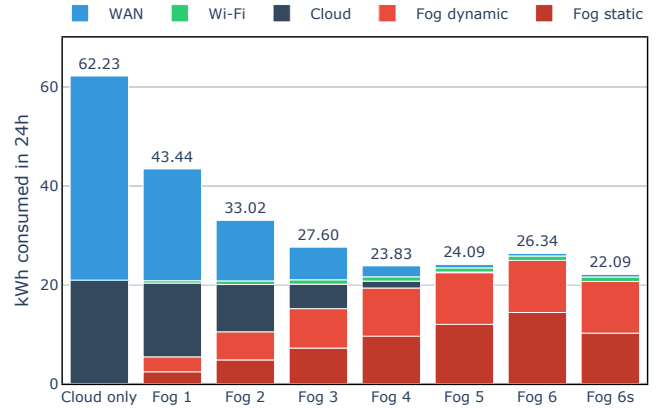


Fig. 5: Aggregated energy consumption of infrastructure components of all conducted experiments.

1) *Cloud Only*: In the first experiment, all processing tasks of both application types were placed in the cloud. There exist no fog nodes. Figures 6a and 6d depict the consumption of infrastructure components and application types over time. We can observe that the major part of the power consumption, namely 66.4%, can be attributed to the WAN as all raw sensor data have to travel to the processing tasks placed in the cloud. The accumulated power required by V2I applications correlates with the number of taxis on the map.

2) *Fog Only*: In experiment *Fog 6*, six traffic light systems were equipped with fog nodes, to avoid data transfer over the power-intensive WAN. These nodes provide sufficient computing capacity so no tasks have to be offloaded to the cloud. Processing tasks are distributed evenly across the fog nodes. As expected, the power usage of network decreased sharply. The fog nodes themselves are now the most relevant power consumers, accounting for 94.7% of the total energy consumption. A major fraction of this, namely 57.8%, is static power consumption: The six fog node require 100 W of static power each, totaling in 14.4 kW during a day.

Although the computational and network load required by the CCTV applications is constant during the entire simulation, the reported power consumption varies over time. This is because we allocate the static power consumption of fog nodes proportionally to applications running on them and fog nodes are utilized inefficiently in this experiment. Especially at night when only a few taxis are on the road, the relative power demanded by the CCTV applications rises.

3) *Fog and Cloud*: Experiments *Fog 1* to *Fog 5* contain one to five fog nodes, respectively. Tasks are still distributed evenly across the available fog nodes, but once all fog nodes are running at more than 85% capacity, tasks are being offloaded

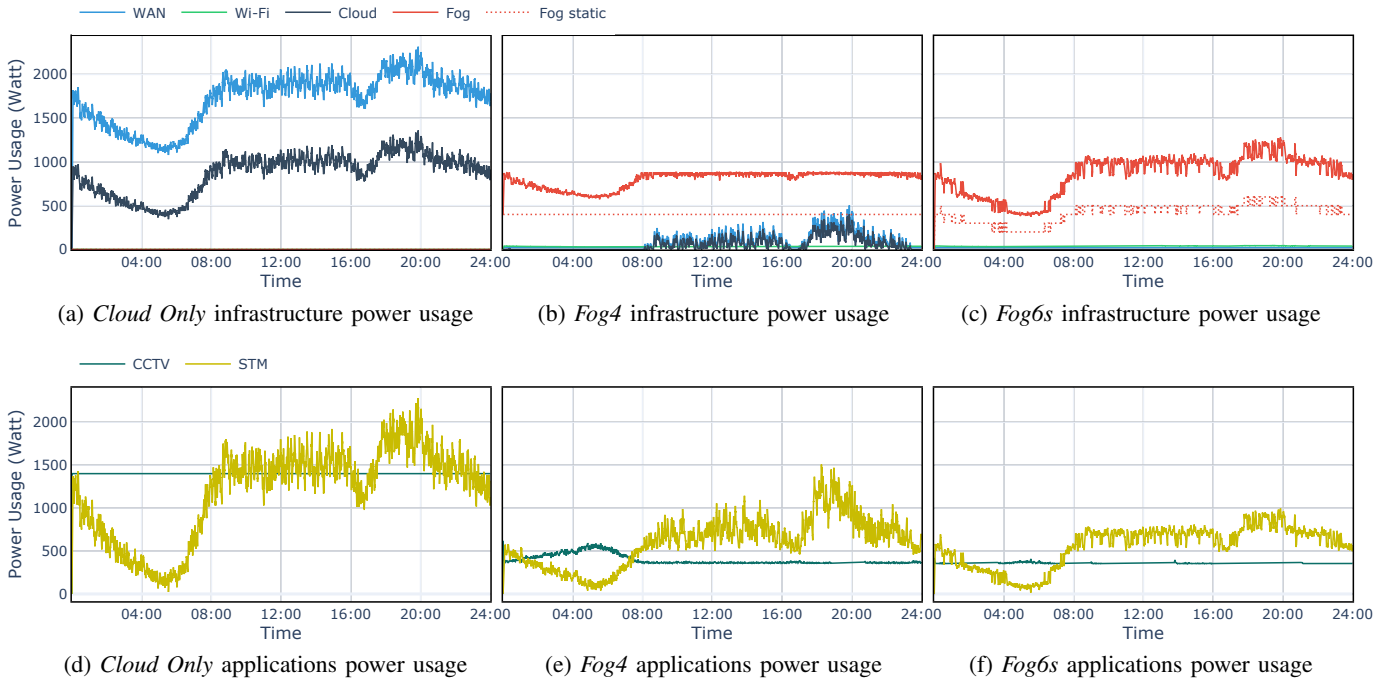


Fig. 6: Power usage of infrastructure and applications for different experiments.

to the cloud. The lowest overall energy consumption was achieved in experiment *Fog 4*, saving 2300 Wh compared to *Fog 6*. Four fog nodes provide enough capacity to host all CCTV processing tasks and around 125 V2I processing tasks. Although during some periods of the day more than 125 taxis are on the map simultaneously, the savings on static fog node power usage outweigh the additional cloud and WAN power usage caused by task offloading with four fog nodes.

4) *Energy-Saving Mechanism*: Experiment *Fog 6s* extends experiment *Fog 6* with an adaptive task placement strategy and energy-saving mechanism. This time, if a fog node is idle for five seconds, it will be put to sleep, reducing its static power consumption of zero. In order to fully exploit this energy-saving mechanism, the task placement algorithm tries to

consolidate as much work as possible on a minimum number of nodes to maximize the number of idle fog nodes. Individual fog nodes are still only utilized up to 85% capacity. Figure 6c shows that the overall static power usage of fog nodes is now no longer constant. All six fog nodes are only active for less than 2 hours during the day, while during the night only two out of six are active. Consequently, the power usage of applications was reduced too, see Figure 6f. Resulting from the improved placement strategy, CCTV applications are now responsible for less static power usage, and, thus, have a very stable load profile. V2I applications benefit from the energy-saving mechanism by becoming more power proportional. At night, the attributed power usage of all V2I applications drops below 80 W.

Figure 7 displays how experiment *Fog 6s* undercuts the second best performing experiment *Fog 4* at different times of the day. When there is little traffic in the city, *Fog 6s* benefits from highly reduced static power usage of shut-off fog nodes. In times of high utilization, *Fog 6s* continues to provide sufficient resources to process all data in the fog layer, avoiding expensive WAN traffic for offloading. At times of average utilization, the load profile of both experiments is almost identical.

C. Analysis

So far the evaluation has shown that we have successfully met our Requirements (1), (2), and (3): LEAF was used to model a realistic fog computing scenario with different kinds of compute nodes, network links and applications. By assigning individual power models to different parts of the infrastructure, we were able to analyze the energy footprint

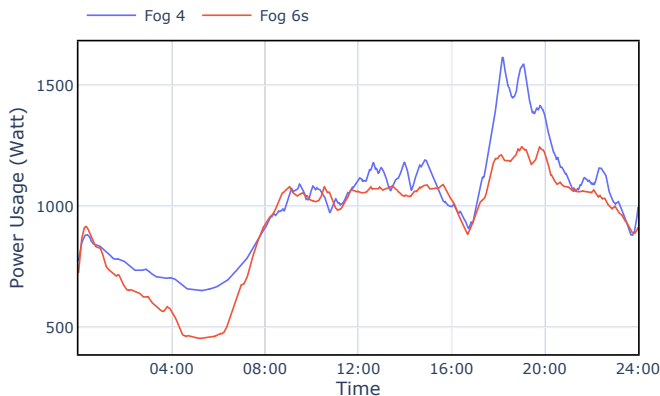


Fig. 7: Comparison of overall power usage in experiments *Fog 4* and *Fog 6s* (smoothing applied).

of different fog node deployments in order to find the most energy-conserving configuration. In the last experiment, we demonstrated how our model enables research on dynamic task placement strategies and energy-saving mechanisms.

To evaluate Requirement (4), *Performance and Scalability*, we analyzed the runtimes of simulations. Each presented experiment simulates around 46 500 taxis executing around 330 000 tasks and finishes in less than 50 seconds on a single core of a 1.4 GHz Intel Core i5. A series of further experiments show that in the outlined scenario, the only algorithmic component that shows non-linear growth is the algorithm for finding shortest paths between tasks. Even when scaling experiments by factor 10, namely 465 000 taxis and 3 300 000 tasks, they had a median runtime of 613 seconds. Consequently, the presented model is suitable to simulate large-scale experiments several magnitudes faster than real time on commodity hardware.

VI. CONCLUSION

This paper presents LEAF, a simulator for modeling large energy-aware fog computing environments. LEAF features a holistic but granular energy consumption model that covers data centers, edge devices, and network as well as applications which are running on this infrastructure. By combining analytical and discrete-event modeling, the proposed model enables the simulation of thousands of streaming applications on a distributed, heterogeneous, and dynamic infrastructure. The publicly available implementation of LEAF was evaluated within a realistic smart city traffic scenario and proved to be a valuable tool for research on energy-conserving fog computing architectures, task placement strategies, and energy-saving mechanisms.

In the future, we plan to extend our model to provide time-based and location-based calculations of the carbon emissions and electricity costs, thus enabling research on infrastructures and algorithms optimized for a low environmental footprint and operating expenditure. Moreover, we plan to integrate LEAF with relevant simulators for network, urban traffic, or electric grids to co-simulate and evaluate more realistic scenarios.

ACKNOWLEDGMENT

This work was supported by the German Ministry for Education and Research (BMBF) as BIFOLD (01IS18025A) and the German Academic Exchange Service (DAAD) as ide3a.

REFERENCES

- [1] A. V. Dastjerdi and R. Buyya, "Fog computing: Helping the internet of things realize its potential," *Computer*, vol. 49, no. 8, p. 112–116, 2016.
- [2] A. Yousefpour, C. Fung, T. Nguyen, K. Kadiyala, F. Jalali, A. Niakanlahiji, J. Kong, and J. P. Jue, "All one needs to know about fog computing and related edge computing paradigms: A complete survey," *Journal of Systems Architecture*, vol. 98, pp. 289–330, 2019.
- [3] N. Jones, "How to stop data centres from gobbling up the world's electricity," *Nature*, vol. 561, pp. 163–166, 2018.
- [4] 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, 2016.

- [5] T. H. Noor, S. Zeadally, A. Alfazi, and Q. Z. Sheng, "Mobile cloud computing: Challenges and future research directions," *Journal of Network and Computer Applications*, vol. 115, p. 70–85, 2018.
- [6] R. Miller, "Hyperscale data centers: A data center frontier special report," Data Center Frontier, Tech. Rep., 2019.
- [7] E. Ahvar, A.-C. Orgerie, and A. Lèbre, "Estimating energy consumption of cloud, fog and edge computing infrastructures," *IEEE Transactions on Sustainable Computing*, 2019.
- [8] M. Yan, C. Chan, A. Gygax, J. Yan, L. Campbell, A. Nirmalathas, and C. Leckie, "Modeling the total energy consumption of mobile network services and applications," *Energies*, vol. 12, p. 184, 2019.
- [9] C. Sonmez, A. Ozgovde, and C. Ersoy, "EdgeCloudSim: An environment for performance evaluation of edge computing systems," *Transactions on Emerging Telecommunications Technologies*, vol. 29, no. 11, pp. 39–44, 2018.
- [10] X. Zeng, S. K. Garg, P. Strazdins, P. P. Jayaraman, D. Georgakopoulos, and R. Ranjan, "IOTSim: A simulator for analysing iot applications," *Journal of Systems Architecture*, vol. 72, pp. 93–107, 2017.
- [11] H. Gupta, A. Vahid Dastjerdi, S. K. Ghosh, and R. Buyya, "iFogSim: A toolkit for modeling and simulation of resource management techniques in the internet of things, edge and fog computing environments," *Software: Practice and Experience*, vol. 47, no. 9, p. 1275–1296, 2017.
- [12] C. Mechalikh, H. Taktak, and F. Moussa, "PureEdgeSim: A Simulation Toolkit for Performance Evaluation of Cloud, Fog, and Pure Edge Computing Environments," in *2019 International Conference on High Performance Computing & Simulation (HPCS)*, 2019, pp. 700–707.
- [13] S. Sarkar and S. Misra, "Theoretical modelling of fog computing: a green computing paradigm to support iot applications," *IET Networks*, vol. 5, no. 2, pp. 23–29, 2016.
- [14] D. Kliazovich, P. Bouvry, Y. Audzevich, and S. Khan, "GreenCloud: A packet-level simulator of energy-aware cloud computing data centers," in *The Journal of Supercomputing*, vol. 62, 2010, pp. 1–5.
- [15] T. Qayyum, A. Malik, M. Khan, O. Khalid, and S. Khan, "FogNet-Sim++: A toolkit for modeling and simulation of distributed fog environment," *IEEE Access*, vol. 6, pp. 63 570–63 583, 2018.
- [16] A. Brogi, S. Forti, C. Guerrero, and I. Lera, "How to place your apps in the fog: State of the art and open challenges," *Software: Practice and Experience*, vol. 50, pp. 719 – 740, 2020.
- [17] L. A. Barroso and U. Hözl, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, pp. 33–37, 2008.
- [18] K. Hinton, F. Jalali, and A. Matin, "Energy consumption modelling of optical networks," *Photonic Network Communications*, vol. 30, no. 1, p. 4–16, 2015.
- [19] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*, vol. 8, 2000, p. 8020.
- [20] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4G LTE networks," in *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2012, pp. 225–238.
- [21] A. Vishwanath, F. Jalali, R. Ayre, T. Alpcan, K. Hinton, and R. Tucker, "Energy consumption of interactive cloud-based document processing applications," in *2013 IEEE International Conference on Communications (ICC)*, 2013, pp. 4212–4216.
- [22] T. Guérout, T. Monteil, G. D. Costa, R. N. Calheiros, R. Buyya, and M. Alexandru, "Energy-aware simulation with DVFS," *Simulation Modelling Practice and Theory*, vol. 39, pp. 76–91, 2013.
- [23] M. C. S. Filho, R. L. Oliveira, C. C. Monteiro, P. R. M. Inácio, and M. M. Freire, "CloudSim Plus: A cloud computing simulation framework pursuing software engineering principles for improved modularity, extensibility and correctness," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017, pp. 400–406.