



HAL
open science

A Wi-Fi Energy Model for Scalable Simulation

Clément Courageux-Sudan, Anne-Cécile Orgerie, Martin Quinson

► **To cite this version:**

Clément Courageux-Sudan, Anne-Cécile Orgerie, Martin Quinson. A Wi-Fi Energy Model for Scalable Simulation. WoWMoM 2023 - 24th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, Jun 2023, Boston (MA), United States. pp.1-10. hal-04055720

HAL Id: hal-04055720

<https://hal.science/hal-04055720>

Submitted on 3 Apr 2023

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.

A Wi-Fi Energy Model for Scalable Simulation

Clément Courageux-Sudan
Univ. Rennes, Inria, CNRS, IRISA
Rennes, France
clement.courageux-sudan@irisa.fr

Anne-Cécile Orgerie
Univ. Rennes, Inria, CNRS, IRISA
Rennes, France
anne-cecile.orkerie@irisa.fr

Martin Quinson
Univ. Rennes, Inria, CNRS, IRISA
Rennes, France
martin.quinson@irisa.fr

Abstract—Wi-Fi devices are ubiquitous, thus they have been extensively studied to understand, for example, the impact of different channel conditions and network properties over network performance. However, improving network performance without considering energy consumption can lead to critical issues: battery depletion, higher costs, and increased latency. Existing works provide algorithms and techniques for more efficient use of energy for Wi-Fi communication, especially in the case of IoT networks, limited by battery capacity. But the ever-growing number of Wi-Fi devices along with the increase in traffic and heterogeneity of current networks make measuring the energy footprint of Wi-Fi communication particularly complex, especially at a large scale. Existing simulation models to study the energy consumption of Wi-Fi devices either suffer from scalability issues due to their fine granularity, or lack realism hindering their usage in practice. In this paper, we propose a power model tackling these scalability and accuracy issues through the use of a flow-based simulation model. By comparing the accuracy and performance of our model to state-of-the-art solution, we show that our approach achieves accurate energy predictions on large-scale and heterogeneous network infrastructures. Our flow-level model allows us to simulate the energy consumption of 800 nodes in a few seconds compared to more fine-grained simulators such as ns-3 that require more than 8 hours under the same scenario, with similar accuracy.

Index Terms—Wi-Fi Simulation, Power Model, Energy Consumption

I. INTRODUCTION

Recent moves in Information and Communications Technologies (ICT), along with paradigms such as fog computing lead to an ever-increased amount of communication, and a higher number of devices at the edge of networks. These devices heavily use wireless communication. In 2022, wireless communications are estimated to account for approximately 71% of the overall IP traffic according to a report from Cisco [10]. Among the available wireless technologies, 802.11 also known as Wi-Fi, is one of the most widespread solutions since this protocol can be used for both energy greedy high-bandwidth devices like televisions, and battery-limited devices such as smartphones.

While the energy use of communication technologies is estimated to reach approximately 3.6% of the world's total energy consumption [14], studying the energy footprint of devices and in particular of Wi-Fi NICs is a crucial and complex task. Once powered up, devices equipped with a 802.11 Network Interface Card (NIC) can send and receive data through their antennas. But powering the device, sensing

the channel, and receiving or sending data frames all consume energy.

Different approaches can be adopted to study the energy consumption of Wi-Fi devices. On one side, one can perform real-world experiments and perform energy measurements using experimental testbeds. Then, experimental results can be used to extrapolate the energy consumption of larger networks. However, this approach suffers from scalability and reproducibility issues, limiting the flexibility of the studied use cases. On the other side, it is possible to use simulation models to estimate the energy usage of existing and future network infrastructures. The use of energy models leverages some issues of real experiments, making it possible to study a large set of scenarios and heterogeneous devices in a reproducible manner. Once implemented in simulators, models can be used to study energy-saving techniques and help in the design of future networks. However, designing realistic simulation models remains challenging as it requires balancing carefully between the results' accuracy and the scalability of the simulation tools. This trade-off depends on the objective of the model: studying low-level phenomena requires more precise models, while the observation of large-scale phenomena favors scalability at the cost of a loss in realism.

While several tools have already been developed to study the energy consumption of Wi-Fi devices [20] [19] [11], existing solutions end up being either too fine-grained to study realistic large-scale infrastructures, or too coarse-grained to fully understand the consequences of application and protocol design choices. There is a need for a model enabling the simulation of realistic applications and their network constraints at scale while achieving accurate forecasts of their energy consumption.

In this work, we extend a previously designed, implemented, and validated flow-based Wi-Fi performance model [6] to enable the study of the energy consumption of the modeled devices. Through the use of a linear energy model, we manage to obtain similar accuracy compared to the energy model of the packet-based simulator ns-3 [20], while significantly improving the simulation performance both in time and memory. Our approach enables performing large-scale experiments that were previously available only using analytical models [19].

This paper is organized as follows. Section II briefly reviews the existing Wi-Fi energy models and their limits. Section III introduces our model based on flow-level simulation. Section IV assesses the validity of our model on a set of mi-

crobenchmarks comparing the output of our model to another state-of-the-art model available in ns-3. Then, we evaluate the scalability of our approach in Section V by simulating a large-scale, heterogeneous infrastructure. Section VI describes the limits of our contribution before concluding this work in Section VII.

II. CONTEXT AND STATE OF THE ART

A. Energy modeling

In literature [7]–[9], the energy consumption of ICT devices is computed as the sum of two factors, that add up:

- *Static power consumption*: The power used by a device when idle, without any active workload and is called P_{stat} in the rest of this paper.
- *Dynamic power consumption*: The additional power required to perform tasks with the device, and is called P_{dyn} in the rest of this paper.

While the static energy consumption remains constant, the dynamic energy consumption is proportional to the usage intensity of the device [8]. In the case of a network interface, the more data it receives and sends, the larger is the dynamic energy consumption because the device remains active for a longer timespan to receive and emit data.

The ratio between dynamic and static power consumption differs depending on the communication technology [7]. In the case of high throughput networks (as with Wi-Fi), the literature shows that more than 90% of the energy consumption of a network interface can be attributed to the static energy consumption alone, less than 10% depending on the activity performed by the device [9]. Consequently, the accuracy of energy models highly depends on power values considered for static and dynamic power consumption, but also on their underlying performance models, i.e. the models they employ to predict the duration of communications.

B. Wi-Fi power consumption

An example of the power usage of a Wi-Fi interface is given in Figure 1 which plots Power (P) as a function of Time (T). The interface switches between different states, and each state is dedicated to a specific task (sending, receiving data, channel sensing). Since each state turns on or off different parts of the interface’s circuitry, they consume different amounts of power. The different states to consider for a Wi-Fi network interface are:

- *IDLE*: The NIC does not perform any operation, the power used only consists of the idle power consumption.
- *Tx*: The NIC actively sends data, requiring to power the antennas of the device.
- *Rx*: The NIC actively receives data, which requires listening to the wireless channel.
- *Sleep*: Deactivates parts of the circuitry to reduce the power usage, but also deactivates some functionalities.

Each state consumes a different amount of power as shown in Figure 1. The literature provides power values for each state for different devices such as laptops in [9], or smartphones and




Fig. 1: A Wi-Fi NIC will switch between the different states depending on the current task performed by the device

routers in [16]. In [20], authors perform measurements for *Clear Channel Assessment* (CCA, i.e. sensing if the channel is busy and if a station can send data) and later use in their models the same power values for the *IDLE* and *CCA* states. We do the same in the rest of this paper.

To save energy, one can configure the device wisely and use specific policies to switch between the different states. The authors of [9] evaluate for example the efficiency of an energy-saving mechanism called *race-to-sleep*, switching the device in the *Sleep* state as fast as possible.

Accurately estimating the overall energy consumption of a Wi-Fi NIC requires estimating the time the device spends in each of the previously given states. Thus, performance metrics need to be taken into account to compute energy consumption:

A) Channel and node configurations will have an impact on the throughput of the wireless channel, increasing or decreasing the time spent in active states.

B) Poor channel conditions or phenomena such as hidden nodes (when out-of-range stations try to communicate simultaneously) can lead to important amounts of retransmissions. Each retransmission multiplies the time spent by the device in the *Tx* state, leading to additional dynamic energy consumption [13].

C) Using several antennas to either receive or send data in parallel can increase data throughput or allow for redundancy (known as Multiple-Input-Multiple-Output - MIMO). More energy is required to power each antenna when using MIMO in the *Tx* and *Rx* states [9].

D) The operating system is in charge of decoding and processing part of network packets. The implementation of the network stack in the operating system has an impact on the energy consumption of Wi-Fi frames. The authors of [16] have shown the impact of crossing the protocol stacks of operating systems, named the *cross-factor*.

While in the rest of this paper the use of dynamic MIMO configurations is possible but not validated, we ignore the impact of the cross-factor since we focus on the energy of the network card, excluding the consumption of the other devices’ components.

C. Existing Wi-Fi power models

Models have already been developed to study the energy of Wi-Fi networks at different levels of granularity.

Network simulators such as ns-3 [20] propose detailed performance models along with energy models based on the approach known as packet-level simulation [12]. These models have been extensively validated. Packet-level simulators model communications between network nodes at the scale of individual network packets, allowing for the detailed computation of the time spent by a Wi-Fi NIC in each state. Such models allow for detailed results, but the amount of computations involved at this granularity seems to limit the scalability of the simulations [8]. Simulations of over a few hundred stations cannot be performed without an extensive simulation time and a large memory footprint as shown in Section V-B.

A radically different approach is to rely on analytical models, as they do not suffer from this scalability issue. The authors of [19] and [11] model the energy consumption of Wi-Fi devices (and other communication technologies) as a simple linear equation. This equation accounts for the static power consumption, as well as a constant factor for the dynamic power consumption. Using these models, the authors of [19] can simulate the network traffic generated by 465 000 taxis in approximately 10 minutes. While such models enable extremely scalable studies, the granularity of this approach prevents any study of the impact of network events such as network congestion, realistic dynamic applications, and heterogeneity.

Flow-level simulation constitutes an intermediate approach between the packet-level and analytical approaches. It consists of modeling network communication as continuous flows of data, reducing the complexity of simulations and increasing performance. Flow-based models can be used to simulate different network technologies such as Ethernet [2], [17], [18] and Wi-Fi [6], producing throughput predictions that are closely related compared to the ones of packet models. However, compared to packet-level simulation, some information such as the exact location of a network packet at a given timestamp cannot be obtained using a flow-based approach since communications are not represented as packets going from one node to another but as a single continuous flow of data going through a set of links. This requires adapting existing models to fit the requirements of flow-level simulation. We now propose to extend an existing Wi-Fi flow-level performance model to include power consumption metrics.

III. CONTRIBUTION

In this section, we introduce our energy model for Wi-Fi devices. This model is built on top of a flow-based performance model in charge of computing the duration of flows [6], because it heavily depends on the duration of communication flows.

A. Hypotheses and conditions

Estimating the energy consumption of Wi-Fi nodes depends on the evolution of the NICs' states during simulations. We assume that the computation of the Wi-Fi communication's duration is accurate. This is important because these communications modify the state of the NIC from *IDLE* to *Rx* or *Tx*, which is the main reason for power variation.

A flow-based model induces some additional constraints compared to packet-based energy models. Using a flow-based approach, the energy consumption of Wi-Fi devices cannot be measured on nodes that are not modeled separately, but at the granularity of links. In this work, we use the following approximation of the Wi-Fi network: the term *Wi-Fi link* describes the aggregation of an Access Point (AP), the communication channel, and the Stations (STA) attached to the AP. Figure 2 shows the difference between simulating power consumption using a packet model in Figure 2a where the power is estimated on each NIC, while Figure 2b illustrates our approach at the granularity of links.

Wi-Fi standards such as IEEE 802.11ax [21] define different Power Saving Mechanisms, such as Target Wake Time, PS-POLL, or Operation Mode Indication. These mechanisms are out of the scope of this work and are not considered in the current energy model.

MIMO configurations could be considered using the flow-level model but have not been validated.

Neither control and management frames nor beacons are simulated explicitly by the Wi-Fi channel performance model used to build our energy model [6]. These frames are responsible for advertising the network, performing probe requests, associating a STA to the AP, or synchronizing devices. Beacons are a type of Wi-Fi frame sent periodically from the AP toward STAs. The flow-level performance model does not inject them in the network load, thus they are not modifying the state of the simulated NICs contrarily to the packet-level simulation model of ns-3 [20].

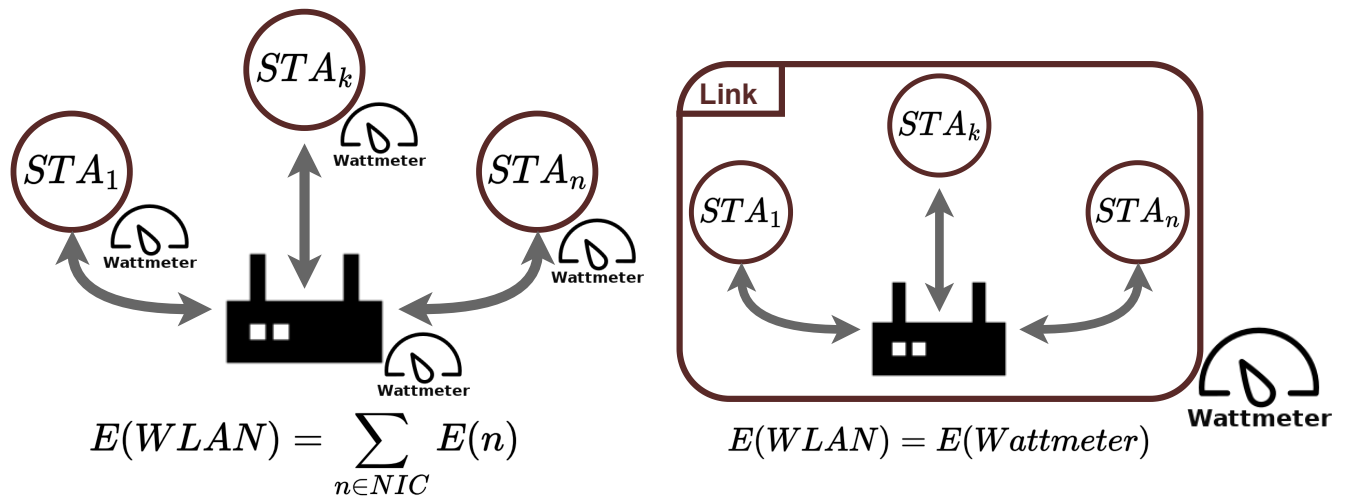
Finally, the model introduced in this section is *homogeneous*. In contrast to *heterogeneous* models where the energy consumed by each device is modeled separately, a homogeneous model considers that all the devices attached to a link consume the same amount of energy to perform the same actions. Homogeneous models have been shown to provide good accuracy in [8] while increasing scalability.

B. Model Overview

Different components need to be taken into account to compute the overall energy consumption of a Wi-Fi link over an experiment of duration T . The energy consumed by a Wi-Fi link in Joules over that period is expressed as the integral of the power of the link during that period.

$$E(T) = \int_0^T P(t)dt \quad (1)$$

The power of a Wi-Fi link corresponds to the sum of its static power consumption P_{stat} and its dynamic power



(a) Packet-based simulation, power estimated at the granularity of network interfaces

(b) Flow-based measurement, power estimated at the granularity of Wi-Fi links, aggregating the NICs of the AP and the STAs

Fig. 2: Flow vs. Packet power measurement. Packet-based measurements are done with NICs' granularity whereas flow-based measurements are done at the level of the WLAN

consumption P_{dyn} . In our case, P_{stat} and P_{dyn} correspond to the energy consumption of all devices (AP and STAs) modeled by the Wi-Fi link, expressed in Watts.

In a discrete event simulator, an event is triggered for every start or termination of a network flow. An event involving a Wi-Fi link may change the power usage of some devices, which requires an update of the overall power consumed on the link.

The energy consumed in between each event is computed using the following equation:

$$E_{tot}(link) = t * \left(\sum_{s \in STA(link)} (P_{stat}(s) + P_{dyn}(s)) + P_{stat}(AP) + P_{dyn}(AP) \right) \quad (2)$$

where t is the duration since the previous event, and $STA(link)$ is the set of all stations on the Wi-Fi link $link$. Using a network with more than one Wi-Fi link, the energy consumed by the whole network is equal to the sum of the energy consumed by each of the Wi-Fi links.

C. P_{stat} : Static power consumption

P_{stat} corresponds to the power of network interfaces without active communications. For Wi-Fi links, computing this value requires the power consumed by NICs in the *IDLE* state. Static power consumption on the link is then equal to the sum of the static power consumption of each device connected, whether AP or STAs, as follows:

$$P_{stat}(link) = P_{idle}(AP) + \sum_{s \in STA(link)} P_{idle}(s) \quad (3)$$

In the validation of this work, only homogeneous network elements (in terms of electric consumption) are considered, so

the idle energy consumption of all the devices on a specific Wi-Fi link is the same for everyone. Thus, n STAs attached to the link (and an AP), with a homogeneous idle consumption on the link P_{idle} , we get :

$$P_{stat}(link) = (n + 1) \times P_{idle} \quad (4)$$

This equation gives the power consumed by a Wi-Fi link without any active communication. P_{stat} can be multiplied by the simulated time to obtain the static energy consumed by the link. Switching from a homogeneous to a heterogenous model only requires using Equation 3 instead of Equation 4, and considering the energy consumed by each NIC individually. However, it requires additional computation at simulation time, reducing the scalability of the model.

D. P_{dyn} : Dynamic power consumption

The dynamic power consumption P_{dyn} is the power used by the Wi-Fi devices to receive or send data actively on the communication channel. Since Wi-Fi networks use broadcasting, when one device sends data on the channel, all STAs in the range of that device receive the data at the NIC level, even if only the destination acknowledges the frame and processes it at the application level.

The computation of P_{dyn} for a Wi-Fi link is based on the power consumption of NICs in the Tx and Rx states. When a device starts a communication, it switches to the transmission state, consuming power P_{Tx} , and the other nodes in the range of the STA on the same channel (including the AP) listen for this information, consuming P_{Rx} . This way, in a network composed of n STAs and one AP, we obtain :

$$P_{dyn}(link) = U(link) \times (P_{Tx}(link) + n \times P_{Rx}(link)) \quad (5)$$

TABLE I: Power values used to calibrate the model. Obtained from ns-3's *WifiRadioEnergyModel* [20]

NIC State	Power (W)
IDLE	0.82
Rx	0.94
Tx	1.14
Sleep	0.10

where U corresponds to the usage of the link: If there is at least one active communication on the Wi-Fi link, U is equal to 1. If there is no active communication the value of U is equal to 0, consequently the dynamic power consumption is also equal to 0. Again, since the proposed model is homogeneous, the values of P_{Tx} and P_{Rx} are the same for all the interfaces within a Wi-Fi LAN.

E. Power consumption of control frames

As explained in Section III-A, the performance model used to build this power model only considers the applicative data transferred on a Wi-Fi link, ignoring all control frames.

We extend our model to account for their energy consumption using an additional factor. We only consider beacons because once STAs are connected to their AP, they represent the majority of control frames in the network. Since beacon frames are sent periodically, usually at approximately 10Hz for IEEE 802.11 (every 102.4 ms), the energy cost of sending and receiving beacons is constant over time and depends on the number of machines in the Wi-Fi link. Obtaining the overall power consumption of a link is possible by adding the cost of beacons to the static and dynamic power consumption of nodes. We note C the factor used to model the energy consumption of beacons. C corresponds to a ratio of the time spent each second by the nodes of a WLAN exchanging beacons. By introducing C into Equations 4, 5, we obtain:

$$P_{tot}(link) = P_{stat}(link) + P_{dyn}(link) + C \times (P_{Tx} + n \times P_{Rx}) \quad (6)$$

The value of C can be computed using the lowest rate of the frequency band employed in the Wi-Fi network – since it is the rate employed to transmit beacons (especially for backward compatibility reasons) – and using the beacon size (depending on the Wi-Fi version). It can also be estimated experimentally from the amount of time spent in the communication of beacons. Section IV shows how to do it on the ns-3 simulator for a single transmission rate, but this could be done for other network configurations.

Finally, applying Equation 6 to Equation 2 provides an estimation of the overall energy consumed by a Wi-Fi network.

F. Implementation

We implemented our model in the SimGrid [5] simulation framework. SimGrid is an active flow-based simulator used in many research works. It includes a Wi-Fi performance model that has been shown to accurately compute the duration of Wi-Fi flows once correctly calibrated [6].

SimGrid has a plugin system allowing easy integration of our energy model by hooking onto the network events of Wi-Fi links. Our model relies on the estimation of the power consumption of a Wi-Fi link between each arrival or termination of a Wi-Fi flow. Thus, the plugin hooks onto events linked to flow creation and destruction and updates the power consumption of the link using the previously described equations.

This implementation inherits from the limitations of the underlying Wi-Fi performance model: an inaccurate calibration of the performance model would result in wrong predictions of the flow duration, leading in turn to inaccurate energy predictions. Wi-Fi links are assumed to be either the source or the destination of network flows. It is currently not possible to have a Wi-Fi node in the middle of a network path or to create multi-hop Wi-Fi networks

The calibration of the performance model of SimGrid is made by specifying the maximum application throughput of STAs on the link instead of the theoretical throughput deduced from node configurations. [6] provides further details on this calibration process.

IV. VALIDATION


A. Experimental setup and methodology

To validate our energy model, we compare the output of our model implementation in SimGrid to the predictions of the energy model in the ns-3 simulator [20]. The comparison of the outputs of these two simulators has been done because:

- 1) Performing real experiments on Wi-Fi networks, and energy, in particular, is tedious. The reproducibility of the results requires perfectly reproducing channel conditions, which is very complex, while hardware heterogeneity can lead to biases in the measurements (both Wi-Fi NICs and power measurement hardware). Thus, results obtained through simulation are much more reproducible, easing results comparison.
- 2) ns-3 is extensively used in the literature and has a published energy model for Wi-Fi [20]. Since validation has been done for its underlying performance models [15], the predictions of the energy model are supposed realistic.
- 3) Obtaining closely related values between ns-3 and SimGrid would show that our approach, while being more coarse-grained, still allows obtaining accurate energy predictions.

The ns-3 and SimGrid models are calibrated with the default energy values of ns-3 for each state of Wi-Fi NICs, as provided in Table I. The values used by default in ns-3 were obtained from [9].

The metrics compared between the two simulators are: A) the overall energy consumption of Wi-Fi interfaces, B) the dynamic energy consumption only, since it is the most challenging value to estimate, C) the memory footprint and the execution time between ns-3 and SimGrid for different experiments since the initial motivation to develop this model is to allow larger simulations.



(a) Dynamic energy consumption with different numbers of STAs and a fixed simulation time (1100 seconds) (b) Dynamic energy consumption with different simulation times and 20 STAs attached to the AP

Fig. 3: Simulating the energy cost of beacon frames in a single WLAN

All experiments have been executed using the Grid’5000 [3] experimental testbed, allowing the simulation of a large number of scenarios while using several random seeds in ns-3, which impact the underlying error models, modifying communication times.

Source code and reproducibility: The code of the model, the code of the experiments, the scripts used to generate all figures, the logs of the results used in this paper, information on how to reproduce our results, and additional results are all available online at the following address: https://github.com/klementc/wifi_energy_experiments.

B. Evaluating the cost of beacons

The first set of simulations calibrates the cost of control frames in the network, as described in Section III-E. We compute experimentally the value of the factor C to obtain accurate predictions of the energy consumed in the network without active flows by using ns-3 simulations. We run several experiments where we measure the dynamic energy consumption of ns-3 devices in a single Wi-Fi cell. Measuring only the dynamic energy consumption means setting the *idle* power consumption to 0W while leaving the other values identical to the ones of Table I. Simulations are performed with between 1 and 30 STAs connected to the AP, varying the simulated time of the experiment between 100s and 1500s, and using 100 different random seeds for ns-3 (used to compute probabilistic events). In all simulations, no flow is created between the STAs and the AP, meaning that only control frames are consuming energy on the Wi-Fi channel. Using the results, we estimate the ratio of time spent communicating beacons on the channel per second.

The results obtained are shown in Figure 3. Figure 3a shows the dynamic energy consumption of the network depending on the number of devices for 1100 seconds of simulated time, while Figure 3b shows the dynamic energy consumption for different simulated times in a cell with 20 STAs. In both cases, we first execute experiments in ns-3. We can observe that the energy consumption values obtained vary linearly with

the number of STAs in the LAN and the simulated time, which matches the description of periodic beacons given in Section III-E. Using ns-3 observations, we compute a value $C = 0.0021$ for our model, leading to the outputs observed for SimGrid in the same Figures. Overall, using this value of C gives a mean relative error (computed as $\frac{|val_{sg} - val_{ns}|}{val_{ns}}$) of 7.0% over all the simulated experiments, showing that we obtain closely related results using both models.

Consequently, in the rest of this paper, all the simulations use the value $C = 0.0021$. To match other transmission configurations, C can be modified in SimGrid before starting an experiment.

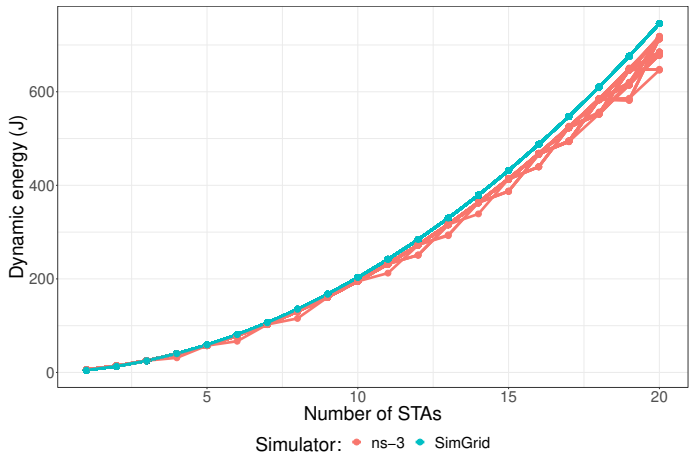
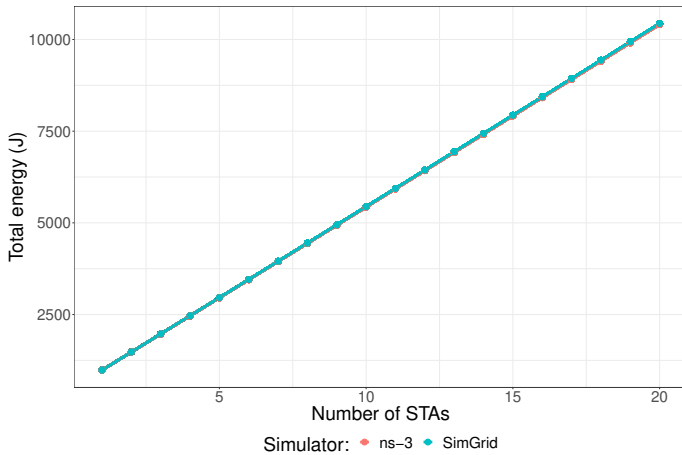
Additionally, these results allow observing that control frames have a relatively small impact on the overall energy consumption of Wi-Fi nodes. While this factor is negligible for heavily used Wi-Fi networks where large flows consume much energy, it may become important on networks with long inactive channel periods where control frames are the only source of dynamic energy consumption.

C. Single WLAN energy prediction

Once the model is calibrated for control frames, it is possible to run a set of microbenchmarks to evaluate the accuracy of our model in the presence of network flows in a single Wi-Fi cell. We perform several experiments using two different types of flows: *a*) flows between STAs and the AP, *b*) flows between pairs of STAs (going through the AP).

The explored parameters are: the size of network flows (between 1 MB and 30 MB) and the number of STAs in the simulated Wi-Fi network (between 1 and 20 for the AP/STA flows, and between 2 and 20 for the flows between pairs of STAs). To avoid comparing our results to an outlier output of ns-3, each scenario is executed with 100 different random seeds in ns-3.

Figure 4 shows results for the simulation of flows between the STAs towards the AP with 600s of simulated time, where each flow has a size of 10MB. Figure 4a shows the overall energy consumption ($E_{tot} = E_{dyn} + E_{stat}$), while Figure 4b



(a) Total energy consumption comparison between SimGrid and ns-3 for different numbers of STAs (b) Dynamic energy consumption comparison between SimGrid and ns-3

Fig. 4: Energy consumption comparison between SimGrid and ns-3 using different numbers of STAs in the downstream flows scenario

shows only the dynamic energy consumption for the same experiments, summing both the cost of control frames and the cost of application flows in the network.

We observe that the overall energy consumption in Figure 4a consists of a quasi-linear function depending on the number of concurrent flows. This observation is understandable since static energy consumption accounts for most of the overall energy consumption in the network. In Figure 4b, we observe that more flows lead to more dynamic energy consumption, along with additional beacon receptions in both simulators.

The results of both figures show that the values computed by our model seem to match with the predictions of ns-3. The average relative error over all our experiments when computing the total energy consumption is equal to 0.15%, while the same error is equal to 5.22% regarding the dynamic energy consumption only. The error is bigger when measuring the dynamic energy consumption only since it is highly dependent on the simulated duration of network flows, which is complex to estimate (and slightly varies between different ns-3 seeds).

We also execute the flow pairs scenario, where flows start from one STA, and go towards another STA in the same WLAN, passing through the AP. Figure 5 gives an overview of the average dynamic power consumption of the Wi-Fi nodes between ns-3 and SimGrid for this scenario. Results are shown only for a flow size of 20 MB, while the results with other flow sizes are available in the online notebook. We observe that the average power in the cell depends on the size of the Wi-Fi cell, i.e. the number of STAs attached to the AP. The flow-based model provides very similar results compared to the outputs of ns-3. This shows the accuracy of our model in this scenario.

Overall, the microbenchmarks at the scale of a single Wi-Fi cell show that the outputs of our model match the energy predictions of ns-3. This is the case when comparing

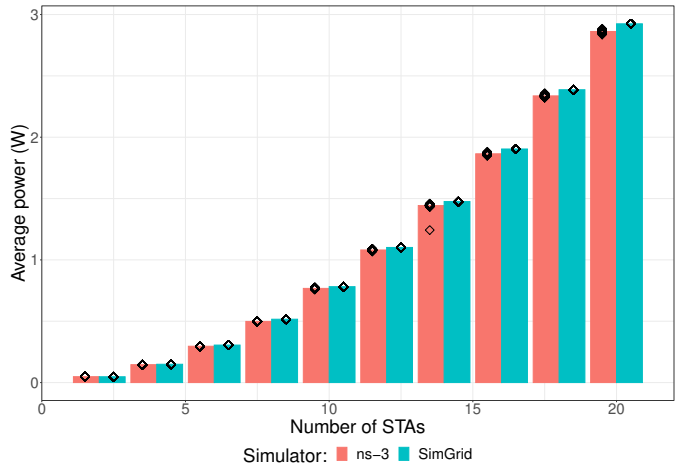


Fig. 5: Average dynamic power consumption of STAs (*IDLE* power set to 0 W) in the flow pairs scenario for different ns-3 random seeds represented as points

the overall energy consumption but also the dynamic energy consumption individually.

D. Heterogeneous network simulation

Network simulators such as ns-3 and SimGrid allow the simulation of heterogeneous networks (i.e. networks mixing different communication technologies in the same experiment). This is important to study realistic infrastructures, for example, mixing both Ethernet links and Wi-Fi cells. This set of experiments consists in simulating a small-scale heterogeneous infrastructure. It is composed of two Wi-Fi networks linked by an Ethernet link.

Each of the STAs in the first Wi-Fi network is the source of a network flow towards another STA in the second Wi-Fi network. Each network flows passes through the wired

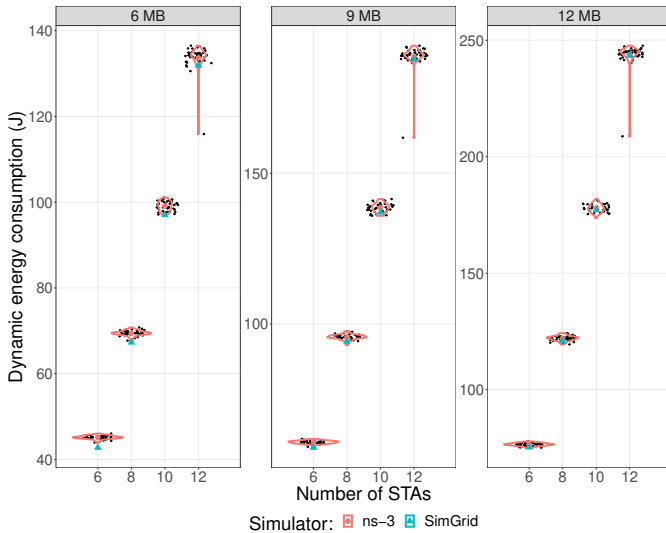


Fig. 6: Dynamic energy consumption in the heterogeneous infrastructure. Each facet corresponds to a different flow size, where colored points designate the mean dynamic energy consumption for every number of STAs. Black dots show the dispersion of ns-3 predictions depending on the random seed.

link, leading to additional simulated constraints. We use the standard *PointToPoint* model of ns-3 to simulate this link, and the default model for wired communications from SimGrid where the throughput of the link is set to 10Gbps. We run several simulations varying the size of flows between 3 MB and 20 MB, the number of STAs in each cell between 1 and 10, a simulated time between 700 and 1100 seconds, and 40 different ns-3 seeds.

Figure 6 shows the results obtained with a simulation time of 900 seconds, between 3 and 6 STAs in each cell, and flow sizes of 6, 9, and 12 MB respectively. In each case, we provide a violin plot of the dynamic energy consumption estimated in ns-3 using different seeds against the energy predictions of SimGrid (SimGrid predictions do not depend on random computations, leading to a single point). The violin and the black dots allow observing the dispersion of ns-3 results. Depending on the random seed, the Wi-Fi channel is shared differently in ns-3, leading to different communication times, hence different dynamic energy consumption. Even if most values are relatively close, there are outliers such as in the 12 STAs setup where we can observe a difference of up to 20J between different executions of ns-3.

Overall, our results remain close to the average of ns-3, with a mean relative error of 4.4% for the dynamic energy consumption. In this setup, the error is more important for small flow sizes. As in the previous experiments, looking at the overall energy consumption of the network further reduces this error given the important proportion of energy consumed in the *IDLE* state, which is easier to predict.

This set of experiments validates our model at a small scale in the case of a heterogeneous network.




Fig. 7: Simulated architecture for the large-scale experiment

V. LARGE SCALE SIMULATIONS


The experiments presented in this Section explore the ability of a flow-based model to efficiently simulate large-scale networks. To that extent, we simulate an infrastructure reproducing a metropolitan public Wi-Fi network. Such networks are extensively used in wide areas such as university campus [13], commercial centers [4], or to cover entire cities as in the case of Google Wi-Fi [1]. These networks are typically used to allow users in a large area to connect to an AP and access services over the internet. The network is usually composed of two sections: **a**) a set of Wi-Fi APs to which clients connect using their devices, and **b**) a wired gateway between the APs and the core network to access the Internet.

Figure 7 provides an overview of the simulated infrastructure. We perform several simulations in both ns-3 and SimGrid using different parameters: the number APs available for STAs to connect to (between 1 and 50), the number of STAs connected to each AP (17 on average with a standard deviation of 2), the size of the network flows, the timestamp at which network flows are created, and 30 ns-3 seeds. In every experiment, each STA in the network sends 25 messages on average towards the destination node (*DST*) connected to the APs via wired network links.


A. Energy predictions

Figure 8 shows the energy consumption of the overall infrastructure over time for one execution of the experiment with 15 Wi-Fi cells and a simulated time of 1100 seconds. Figure 8a plots the overall energy consumption while Figure 8b only shows the dynamic energy. In each case, we measure by steps of 10 seconds the energy consumed by the nodes, and compute the relative error between ns-3 and SimGrid, visible as black points. Similar figures using other parameters are available in our online notebook.

We observe that, even at a larger scale, ns-3 and SimGrid compute similar energy consumption values. The relative error doesn't exceed 0.3% for the overall energy consumption,



(a) Total energy consumption



(b) Dynamic energy consumption

Fig. 8: Evolution over time of the overall and dynamic energy consumption values between ns-3 and SimGrid for the large-scale experiment

while it is almost always below 10% for the dynamic energy consumption. The error in dynamic energy consumption is more important during the first few time slots since only a small amount of data is going through the link, while our model is calibrated for large data throughput.

During the phase when most messages are sent (i.e. between 200 and 550 seconds), a rapid increase in the dynamic energy consumption can be observed, but this increase seems to have a limited impact on the overall energy consumption of the network that remains quasi-linear.

Even at a much larger scale than the previous microbenchmarks, a flow based-model provides accurate energy prediction values compared to ns-3.

B. Performance comparison

We now compare the memory footprint of ns-3 and SimGrid simulations in Figure 9 while Figure 10 shows the time required to simulate each scenario with the two energy models. For both metrics, we can observe that using a flow-




Fig. 9: Memory footprint comparison of SimGrid and ns-3 simulations




Fig. 10: Simulation time comparison of SimGrid against ns-3

based model increases the performance by several orders of magnitude compared to the packet-based model of ns-3.

Regarding the memory footprint, Figure 9 shows that simulating up to 55 Wi-Fi cells required up to 80 GB of memory with ns-3, while 1.2 GB was needed at maximum to study the same scenarios using SimGrid. Regarding runtime, simulations take up to 8 hours to complete using ns-3, while the runtime of SimGrid simulations does not exceed a few seconds.

Comparing the performance gains between a packet-based energy model and our approach, we can see the opportunity brought by our solution to perform large-scale simulations with much better performance than the packet model of ns-3, while keeping similar prediction values.

VI. THREATS TO VALIDITY

Section IV and V show the validity of the flow-level energy model in different use cases. However, the model could be limited under certain conditions. Below is a list of issues not directly taken into account by our model.

The Wi-Fi performance model used [6] only models the successful application throughput on the link while bad conditions may result in an important amount of errors and retransmissions. One solution to account for data loss due to retransmissions in energy predictions would be to multiply the energy values of Table I by the average number of retransmissions in the channel, but this remains to be done and carefully validated.

Our validation experiments do not account for hidden nodes, which can lead to collisions when two STAs are not in range and start sending data simultaneously. The collisions occurring on STAs in range with both emitting nodes are difficult to detect for the emitters. This leads to a reduced application throughput. A homogeneous model where STAs subject to hidden node issues consume more energy due to retransmissions could improve the accuracy of predictions in this context. Again, this remains to be validated.

Regarding the energy consumption of control frames, we only considered beacons and ignored other frames such as association requests given that they are very sparse in the studied use cases. If simulating a network where such frames appear frequently, their energy consumption could be estimated using the same approach, and by recalibrating the model to account for them.

Since the hardware of APs is often different than the hardware of STAs, another possibility is to have separated static power consumption values between STAs and APs. It would be possible to model this difference by having two different values for P_{idle} in Equation 4: one for STAs and the other for the AP.

VII. CONCLUSION

This paper introduces a new flow-based model to simulate the energy consumption of Wi-Fi networks. This linear model computes the energy consumed by Wi-Fi interfaces when sending or receiving data as well as the energy consumed by control frames in the network.

The experimental results on both small and large-scale infrastructures show that the implementation of the model in SimGrid provides accurate energy predictions compared to ns-3 in a large set of scenarios while requiring much fewer resources than ns-3. While ns-3 provides more fine-grained simulation results, the flow-level approach enables large-scale simulations with similar results without extensive runtime and memory usage.

In the future, this model could be extended to take into account energy-saving modes defined in modern Wi-Fi standards such as PS-POLL, and validated under degraded channel conditions. We plan to use this model to simulate the energy consumption of large-scale and heterogeneous end-to-end networks mixing different types of communications. This would allow studying the impact of different application structures and deployment configurations on the network's energy consumption.

ACKNOWLEDGMENT

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] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren, "Analysis of a mixed-use urban wifi network: when metropolitan becomes neapolitan," in *ACM SIGCOMM conference on Internet Measurement*, 2008, pp. 85–98.
- [2] G. Anggono and T. Moors, "A Flow-Level Extension to OMNeT++ for Long Simulations of Large Networks," *IEEE Communications Letters*, vol. 21, no. 3, pp. 496–499, 2017.
- [3] D. Balouek *et al.*, "Adding virtualization capabilities to the Grid'5000 testbed," in *Cloud Computing and Services Science*, ser. Communications in Computer and Information Science. Springer, 2013, vol. 367, pp. 3–20.
- [4] V. Brik, S. Rayanchu, S. Saha, S. Sen, V. Shrivastava, and S. Banerjee, "A measurement study of a commercial-grade urban wifi mesh," in *ACM SIGCOMM conference on Internet Measurement*, 2008, pp. 111–124.
- [5] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, scalable, and accurate simulation of distributed applications and platforms," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2899–2917, Jun. 2014.
- [6] C. Courageux-Sudan, L. Guegan, A.-C. Orgerie, and M. Quinson, "A flow-level wi-fi model for large scale network simulation," in *International Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2022, pp. 111–119.
- [7] D. Feng, C. Jiang, G. Lim, L. J. Cimini, G. Feng, and G. Y. Li, "A survey of energy-efficient wireless communications," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, pp. 167–178, 2012.
- [8] L. Guegan, B. L. Amersho, A.-C. Orgerie, and M. Quinson, "A large-scale wired network energy model for flow-level simulations," in *International Conference on Advanced Information Networking and Applications*. Springer, 2019, pp. 1047–1058.
- [9] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, "Demystifying 802.11 n power consumption," in *International conference on Power aware computing and systems*, 2010.
- [10] C. V. N. Index, "Forecast and trends, 2017–2022," *Cisco Systems*, pp. 1–7, 2018.
- [11] 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.
- [12] B. Liu, D. Figueiredo, Y. Guo, J. Kurose, and D. Towsley, "A study of networks simulation efficiency: fluid simulation vs. packet-level simulation," in *IEEE Conference on Computer Communications (INFOCOM)*, vol. 3, 2001, pp. 1244–1253.
- [13] M. Maity, B. Raman, and M. Vutukuru, "TCP download performance in dense wifi scenarios," in *International Conference on Communication Systems and Networks (COMSNETS)*, 2015, pp. 1–8.
- [14] J. Malmodin and D. Lundén, "The energy and carbon footprint of the global ICT and E&M sectors 2010–2015," *Sustainability*, vol. 10, no. 9, p. 3027, 2018.
- [15] R. Patidar, S. Roy, T. R. Henderson, and A. Chandramohan, "Link-to-system mapping for ns-3 Wi-Fi OFDM error models," in *Workshop on ns-3*, 2017, pp. 31–38.
- [16] 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 Transactions on Networking (ToN)*, vol. 23, no. 4, pp. 1243–1256, 2015.
- [17] J. Sommers and Y. Mok, "Balancing Accuracy and Efficiency in TCP Flow Simulation," in *IEEE International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2014, pp. 475–478.
- [18] P. Velho, L. M. Schnorr, H. Casanova, and A. Legrand, "On the validity of flow-level tcp network models for grid and cloud simulations," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 23, no. 4, p. 23, 2013.
- [19] P. Wiesner and L. Thamsen, "Leaf: Simulating large energy-aware fog computing environments," in *IEEE International Conference on Fog and Edge Computing (ICFEC)*, 2021, pp. 29–36.
- [20] 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*, 2011, pp. 222–230.
- [21] H. Yang, D.-J. Deng, and K.-C. Chen, "On energy saving in IEEE 802.11 ax," *IEEE Access*, vol. 6, pp. 47 546–47 556, 2018.