

A Large-Scale Wired Network Energy Model for Flow-Level Simulations

Loic Guegan¹, Betsegaw Lemma Amersho, Anne-Cécile Orgerie²,
Martin Quinson³

Univ Rennes, Inria, CNRS, IRISA, Rennes, France
Emails: loic.guegan@ens-rennes.fr¹, anne-cecile.orgerie@irisa.fr²,
martin.quinson@ens-rennes.fr³

Abstract. The use of simulators to predict network energy consumption is a good way for scientists to improve and develop new algorithms and also to assess them. However, the average size of a network platforms is continuously increasing with the emergence of new technologies like the Internet Of Things and Fog Computing. Packet-level simulators start to reach their limits in terms of performance and this calls for newer solutions in the domain of large-scale platform energy models. In this paper, we propose two energy models for wired networks adapted to flow level simulators in order to estimate the energy consumption of large platforms. An evaluation of these models is proposed and it demonstrates their applicability in practice and their accuracy. Indeed, we obtain simulation results with a relative error lower than 4% compared to an ns-3-based solution, and our flow-based simulation is 120 times faster.

1 Introduction

Since the beginning of Internet, distributed systems are widely used and researchers put large efforts in this field in order to improve performance models and provide new solutions. Since their usage tends to grow with the arrival of new technologies like Internet of Things (IoT) and Fog computing [1], their energy consumption is an increasingly important metric to optimize. Information and communication technology (ICT) devices, that are part of distributed systems, can be categorized in three classes [12]: end-user devices, communication networks, and data centers. Researchers estimate that the overall energy consumption growth rate of each of these classes is respectively 5%, 10% and 4%. With the highest expected growth, communication networks constitute a key player in the ICT's global electricity consumption. Energy-efficient solutions are required for curbing this trend.

Several methods are available to study energy saving techniques. First, test-beds can be used to perform experiments in order to be as close to the reality as possible. However, test-beds face three main drawbacks: **1)** Deploying equipment and building a real network is expensive especially at large scale. **2)** Such a deployment requires a lot of setting-up time. **3)** Experiments are hardly reproducible due to numerous uncontrollable parameters. Another option to study

energy saving techniques consists in using simulators. They allow to conduct diverse experiments on any types of platforms in a convenient way. Moreover, they permit to save time compared to real deployments, and experiments can be reproduced easily. On the other hand, network simulators need to be carefully calibrated and assessed to produce valid results.

Several types of simulators can estimate the energy consumption of networks. However, current network platforms and traffic volumes push their scalability to their limits. Indeed, with the arrival of new technologies like the Internet Of Things and fog computing, the number of connected devices increases. Besides, generated and exchanged data volumes are augmented consequently in these increasingly distributed systems, leading to an unprecedented rise of the global energy consumption attributable to ICT.

This calls for scalable yet accurate models and simulators to estimate network energy consumption. Existing simulators that model the energy consumption work at fine grain – simulating each packet exchange – and thus are not scalable for modern research purposes. Their over-approximation is not compatible with efficient simulations. Nevertheless, scalable simulators exist but to the best of our knowledge, none of them is modeling network energy consumption or they are limited to study specific environments or protocols. Flow-based simulators operate at a flow level without simulating each packet, thus saving memory and computations. In practice, they constitute a nice compromise between network abstraction and scalability, in terms of memory usage and computation time.

In this work, we propose to extend an existing flow-based simulator with network energy models. We show that modeling the network energy consumption into a flow-level simulator leads to predictions that are in line with the predictions of packet-level simulators. We also show that simple linear models are sufficient on realistic settings to obtain accurate values, while more complex models were proposed in literature. Indeed, this model exhibits a relative error lower than 4% compared to an ns-3-based solution and provides the simulation results 120 times faster with using more than 17 times less memory.

The rest of this paper is structured as follows. Section 2 details the different simulators and the classical network energy models. Section 3 presents energy modeling in flow-level simulators. Section 4 presents the proposed energy models. Section 5 details our evaluation methodology. An evaluation of the energy models in terms of validity, scalability and applicability is proposed in Section 6. Finally, we conclude in Section 7.

2 Related Work

2.1 Network Simulators

In the literature several packet-level simulators were extended in order to simulate the energy consumption of network devices.

The first well-known simulator that enters in this category is ns-2. It has been extended into GreenCloud [6] which is designed for cloud simulations using

different architectures: two-tiers, three-tiers and three-tiers high-speed. However, Greencloud can be used only in cloud context and thus, it is very limited. Furthermore, as it relies on ns-2, its scalability is limited in terms of number of simulated devices and exchanges volumes of data.

As a replacement for ns-2, another packet-level simulator called ns-3 was proposed. It provides several modules including ECOFEN [8]. This module is able to simulate the energy consumption of network communications based on three energy models namely basic, linear and complete. Even if they provide accurate results [9], the packet-level nature of ns-3 prevents the usage of ECOFEN for large-scale networks.

While these simulators offer accurate predictions, they fail to scale. In fact, event rate of packet-level simulators increases very quickly on large-scale platforms with real workloads. Flow-level simulators work at a higher level of abstraction in order to have very efficient performance models. Hence, this feature makes them suitable for large-scale network simulations. Moreover, they show good prediction abilities when they are properly instantiated. This makes them good candidates for network energy consumption simulations.

CloudSim belongs flow-level simulators and is built on top of GridSim. It is a well-known event-based simulator capable of simulating cloud environments and estimating their energy consumption. It implements green strategies like Dynamic Voltage Frequency Scaling and energy-conscious VM provisioning techniques. On the other side, CloudSim is designed for cloud environment which makes it difficult to use outside of this context. Another point is that, CloudSim models are very similar to the GridSim ones which are known to contain bandwidth sharing model aberrations [13].

SimGrid is another flow-level simulator, it is very versatile and can be used to simulate High Performance Computing, Grid, Peer-to-peer etc. It is also scalable [10], but unfortunately it does not provide network energy model. Consequently, we propose to extend SimGrid by adding network energy models. However, because of the coarse granularity nature of flow-level simulators, we should carefully instantiate and validate these energy models.

2.2 Classical network energy models

Simulators are able to abstract the complexity of network devices' energy consumption. Even if a network equipment is not subject to any network traffic, it consumes energy. Indeed, the different components involved in network devices should be maintained powered on to ensure a fully functional network communication in case of packet arrival. This energy consumption part is called the idle or static part. In some cases, the idle energy cost of a network device can represent 85% of its peak power [4]. Classically, the energy consumption is estimated using the energy definition $E(t) = \int_0^t P(t)dt$. Considering a network equipment, its energy consumption varies over time according to the network traffic. Since its power is composed of two parts: one static and one dynamic, we can define its power in Equation 1 with S the set of ports on the given network device, Np_i and Nb_i respectively the number of packets and the number of bytes going

through the port i , and Pp_i and Pb_i the power costs per packet and per byte for the port i .

$$P_{netdev} = \underbrace{P_{idle}}_{\text{Static}} + \underbrace{\sum_{i \in S} Np_i \times Pp_i + Nb_i \times Pb_i}_{\text{Dynamic}} \quad (1)$$

Existing energy models in literature are based on this equation and consequently, they are not directly compatible with flow-level simulators. Indeed, flow-level simulators do not provide any notion of port in their models nor packets. These energy models should be adapted to this network representation.

3 Background on simulation

Simulation models can be built by following different approaches. One of them is to model every details of the reality. In this case, lots of parameters increase the model instantiation complexity. For the same reasons, the observations generated by these models are more difficult to interpret. On another hand, well detailed models can lead to accurate predictions, but they still should be validated in spite of their detailed nature. Another approach is to use abstract models. Thereby, no assumptions can be made on their level of accuracy. The last approach – this is the one that we chose – is to use simpler models that are still able to do accurate predictions. A good reference related to these types of models is the LogP model proposed by David Coller et al. [3]. They propose a simple model for parallel computing composed of only four parameters. Still, LogP is taking into account communication delay and CPU overhead, which makes it suitable for accurate performance predictions and algorithm optimizations. Thereby, our energy model is part of this latter approach.

As aforementioned, the energy models proposed in this paper are implemented on a flow-level simulator called SimGrid. This simulator is based on a modeling philosophy where the different models are simplified but nevertheless assessed with the aims of providing valid results. It goes against two classical believes. First, a model with more parameters is necessarily more accurate and thus gives results closer to the reality. Indeed, we showed in previous work that simple models are sometimes enough to capture complex effects [13,7]. Second, a model that implements all the parameters present in the reality are somehow correct by design (these parameters can be routing protocols, link delays, bit error rate, etc.). Yet, model validations are mandatory in all cases. It can be justified by the fact that some published simulators suffer from trivial implementation bugs [13] that have never been fixed. Hence, in this work we provide validation results that can be reproduced by everyone (our code is open source and available on the SimGrid web page).

SimGrid does not include Network Interface Controllers (NICs) and routers. Instead, the paths between two hosts are represented by routes. Each route is composed of several entities, called links that represent both, the wire and the

NIC at each end. A route between two hosts can be composed of multiple links but not necessarily of multiple hosts. To this end, there is no notion of port in flow-level simulators. In spite of this, we can attach the energy consumption of ports on the links. Then, it is possible to define the power of the overall network platform using the equation $P_{total} = \sum_{i \in L} LinkPower_i \times LinkUsage_i$ with L representing the set of links on the simulated network platform.

However, in every networks, each link is connected to two different ports with potentially different energy consumption schemes. Thus, flow-level simulators abstraction introduces difficulties to transpose packet-level energy models into them. To solve this issue, we propose to define two energy models, one for homogeneous energy platform (same energy consumption for ports on a given link) and one for heterogeneous energy platform. We will see later that this distinction can in fact be simplified.

4 Contribution

4.1 Homogeneous model for flow-based simulators

We first introduce an energy model for network topologies that uses homogeneous energy consumptions. As said previously, links will hold the energy consumption of each of their ports and additionally, will ensure the bandwidth sharing. To achieve this goal, we propose to use one split-duplex link between each end-node. Then, to overcome the lack of port representation into SimGrid, we double the dynamic energy consumption of each link to account for the two ports (one at each end of the link). Since split-duplex links cause the creation of two links, the idle power will be already doubled and we only need to double the dynamic part of the power. More generally, the energy consumed by ports follows a linear model with a minimum value equals to the idle energy consumption and a maximum value define as:

$$max = idle + (BW \times ByteCons + PktCons/MTU) \times 2 \quad (2)$$

In this equation, BW represents the port rate in Bps, $ByteCons$ its energy consumption per byte in Joules, $PktCons$ represents the network device energy consumption in Joules needed for handling a packet. Finally, MTU , the Maximum Transmission Unit is used as an over-approximation of the packet size. Considering a one-way communication, the link consumes energy for the transmitting (Tx) and receiving (Rx) ports at the same time and vice-versa. If there are multiple flows, the bandwidth sharing occurs on the up or down links according to the flows' directions. Figure 1 depicts this homogeneous energy model for both kinds of simulators.

This energy model has multiple advantages. First, it is very easy to instantiate because there is only one energy cost per link. Consequently, this makes it easier to implement on flow-level simulators since these simulators use at least links for communications (but not necessarily ports). Finally, this model provides low computational overhead. However, many real topologies use heterogeneous

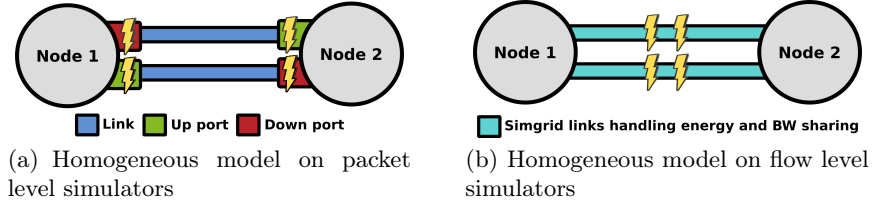


Fig. 1. Homogeneous model on packet/flow level simulators

networking devices. Thus, the energy consumption from one NIC to another varies according to its specifications. Consequently, this heterogeneity cannot be explicitly represented using the homogeneous model. Hence, we decide to introduce another model to overcome this limitation.

4.2 Heterogeneous model

In SimGrid, routes from one host to another can be composed of multiple links. Therefore, we can take advantage of this feature to build the heterogeneous model. This model is represented on Figure 2.

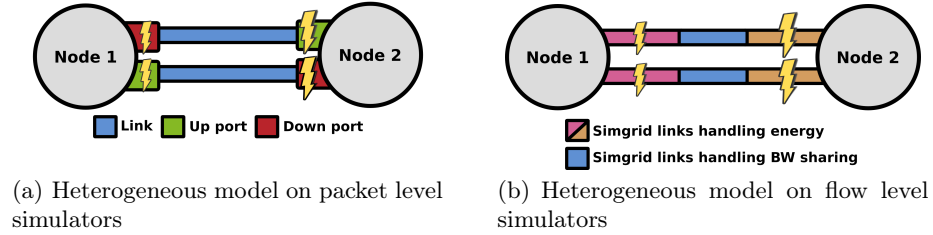


Fig. 2. Heterogeneous model on packet/flow level simulators

In order to model two ports with different energy consumption on a single link, we introduce routes made of 3 links. The first link models the energy consumption of the first port, the second link handle the bandwidth sharing and finally, the third link models the energy consumption of the second port. However, we are using split-duplex links, thus two links will be created by SimGrid and consequently the idle power will be multiplied by two. To overcome this issue, we simply have to divide by two the idle power of each energy link. Thus, this energy model follows a linear behavior with a minimum value equals to $idle/2$ and a maximum value defined as:

$$max = idle/2 + (BW \times ByteCons + PktCons/MTU) \quad (3)$$

5 Methodology and Experimental Setup

5.1 Methodology

The evaluation of the two energy models is divided in two steps. The first step is dedicated to assessing the models by means of microbenchmark simulations. Hence, the experiments are running on small platforms in order to do simulations on a simple and controlled environment. The second step of the evaluation is focused on testing the energy models on a real large-scale platform. In this way, we can demonstrate that the models are suitable for real research purposes and thus, are useful for the scientific community. In addition, it demonstrates that the energy models are scalable regardless to the simulated platform and the workload size. It is important to note that our evaluation focuses on the network energy consumption estimation and not on the network performance model. Yet, it is mandatory to have accurate time estimations to get accurate energy consumptions.

To conduct the validation experiments, we decided by lack of real large-scale platforms to do simulations. In fact, validating these energy models using test-beds requires to measure the energy consumption per byte and per packet on each network devices. However, these measurements require high-precision instruments. Nevertheless, these energy measurements results are proposed in the literature [11,2]. Therefore, we based our experiments on these works.

To validate the two energy models, the experiments are done in parallel on a packet-level simulator that acts as a trusted party. We choose to use ns-3 and ECOFEN for this purpose. It is a logical choice since ECOFEN has been assessed as accurate in the literature [9]. Moreover, ECOFEN is a packet-level simulator, thus it will be useful for the scalability experiments in order to compare its simulation performances to SimGrid. It is important to mention that all the simulations ran on Grid'5000, a large-scale test-bed for experiment-driven research in computer science. In fact, microbenchmark scenarios require 96 unique simulations for only one experiment point (4 data transfer scenarios, 3 topologies with 2 energy instantiation, 3 models are tested including one which is an optimization). Moreover, the real use case experiment was also very computationally intensive for ns-3 and ECOFEN.

All the experiments presented in this paper are available online [5]. We did our best effort to provide reproducible experiments that can be run with different parameters and require the least possible amount of user interaction for the reader desirous of reproducing our experiments.

5.2 Experimental setup

To conduct microbenchmark simulations, we designed different experimental scenarios. These scenarios are built in order to account for the effects on the energy consumption of the following parameters: **1)** The number of nodes and their connections in the network **2)** Different flow configurations, and thus different bandwidth sharing patterns **3)** The heterogeneity of the devices in terms of energy profile. For this, we define three platforms. The first one is composed of two

hosts connected by one link. The second platform is an extension of the first one where we add two hosts and two links. Finally, the last platform, called Dogbone, is composed of six hosts connected together in order to have a bottleneck on the two central hosts. Furthermore, the platforms can be set up with homogeneous or heterogeneous energy consuming devices. Thanks to these two variants, the ability of each model to predict the overall energy consumption of platforms with homogeneous and heterogeneous devices is highlighted. Then, each platform is running with different TCP flow configurations. Other classical network parameters are also defined. These parameters are equals for all the microbenchmark scenarios: the latency is set to 10ms and the bandwidth to 1Gbps for each link since the energy value for such a NIC is available in the literature [11,2]. Thus according to these papers, the energy values used in our microbenchmarks for 1Gbps links are 1.12 Watts for idle power consumption, 3.4nJ for energy consumed per byte and 197.2nJ for the energy consumption per packet. However, for the heterogeneous platform instantiation, we used energy values of 10Gbps links namely 0.53 Watts for idle power consumption, 14nJ and 1504nJ respectively for energy consumption per byte and per packet. Finally, the data flow size is varying randomly between 10MB to 100MB (x-axis) in order to build a power profile of the network platforms.

6 Evaluation: validity, scalability and real use case

6.1 Validity of the proposed models

As a first step, we simulate the microbenchmark scenarios on homogeneous energy consumption platforms. The amount of data sent by the hosts are varied according to the parameters defined in Section 5.2. Due to space limitation, we only provide the results for two scenarios on Figure 3. The flow configuration used for these scenarios is the following: 3 flows overall (2 up and 1 down) between the extreme nodes of each platform.

These first results show the ability of each model to predict the energy consumption on two platforms namely Dogbone and 4 hosts and 3 links with homogeneous network devices. They clearly show a linear relation between the amount of data sent between hosts and the overall energy consumption. One interesting phenomenon to note is that both SimGrid energy models predict the exact same amount of energy. Their energy relative error are close to 0, which means they have the same energy prediction ability. In fact, it is not surprising since the first energy model is a subset of the second one. Thus, they show similar predicting abilities on homogeneous platforms. It is also interesting to note that both homogeneous and heterogeneous energy models show a very small relative error (less than 1%) compared to ns3 with ECOFEN.

Next, to demonstrate the ability of each model to predict the energy consumption on heterogeneous platforms, we did similar microbenchmark experiments using the scenarios with heterogeneous devices in terms of energy consumption. Thereby, we expect the homogeneous model to have a lower prediction ability since it is not designed to handle heterogeneous energy platforms.

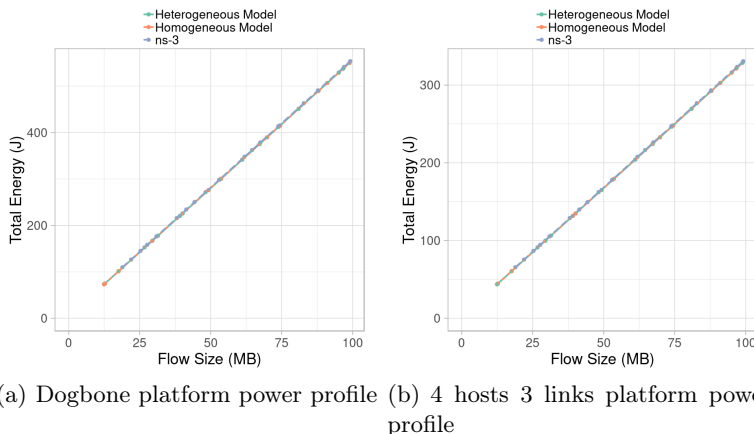


Fig. 3. Microbenchmarks energy consumption on platforms with homogeneous devices.

The simulation results are shown on Figure 4. As expected, the homogeneous model is less accurate than the heterogeneous model. In fact, we can clearly see that the heterogeneous model energy estimation corresponds to the energy predicted by ns-3 and ECOFEN. However, the homogeneous model is doing wrong approximation. This phenomenon is explain by the fact that the granularity of the first energy model is not fine enough to fully capture the dynamic energy consumption of the network devices.

Nonetheless, we decided to modify the instantiation of the homogeneous model in order to improve its prediction capability. Hence, for each link we attached an energy consumption equal to the average energy consumption of their ports. The results of these simulations are also visible on Figure 4 in pink. We can see that by using this simple optimization we surprisingly observe that the homogeneous model produces very accurate results almost equal to the heterogeneous model. This improvement can be explain by the fact that on real platforms, the difference between the energy consumption of the ports is small and thus, taking the mean has a limited impact on the overall energy consumption. Obviously, the relative error between ECOFEN and the homogeneous model will increase if this difference becomes larger. Hopefully, in real platforms this difference remains reasonable. Moreover, the heterogeneous energy model multiply by three the number of links used in the SimGrid simulation, which is not negligible on large-scale platforms. This is why, by following the Occam's razor Principle, we decide to use the homogeneous energy model with the arithmetic mean instantiation in the remaining experiments.

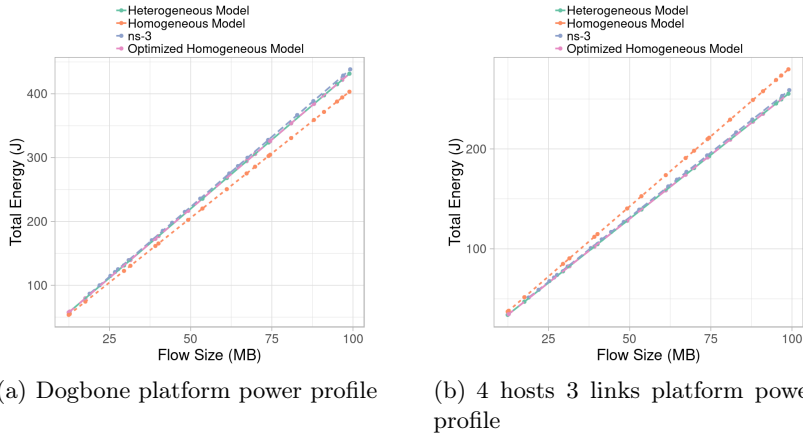


Fig. 4. Microbenchmarks energy results on heterogeneous energy platforms.

6.2 Realistic use case

To evaluate the scalability of our approach, we propose to simulate a datacenter’ network which is a classical large-scale platform widely used in research. Therefore, the final experiment has two main goals. Firstly, it demonstrates that our modified homogeneous energy model scales up in terms of execution time and memory usage. Secondly, it shows that this energy model is usable for real experiments and thus is interesting for the scientific community. The data center platform is based on a classical three-tier architecture used by Greencloud [6]. It is composed of 8 core nodes that provide access to the data center. These two core nodes are linked to 16 aggregation nodes by 10Gbps links. These aggregation nodes are linked to 512 access switches by 1Gbps links that provide access to 1,536 servers by mean of 1Gbps links. Each link latency is fixed to 0.2ms which is the average latency measured between two Grid’5000 internal nodes. The platform is set with the energy consumption of 1Gbps/10Gbps links referenced in the literature [11,2].

The simulation scenario is defined as follows. First, 10 to 400 clients requests, modeled by TCP flows, are generated and arrive simultaneously to the data center by the 4 core nodes. Next, the flows are spread among the different aggregation switches randomly and then reach the servers. Finally, the requests are handled by servers to answer the clients’ requests. The experiment runs on both simulators, SimGrid and ns-3 in order to compare their energy consumption estimation and also their performance in terms of execution time and memory usage. For each point of the simulation, we run several experiments using different random seeds in order to better estimate the accuracy of our proposition.

The energy and scalability results are shown in Figure 5. The energy consumption results show that we have similar values on both simulators. The homogeneous model provide similar results to ns-3 with ECOFEN. Even in the

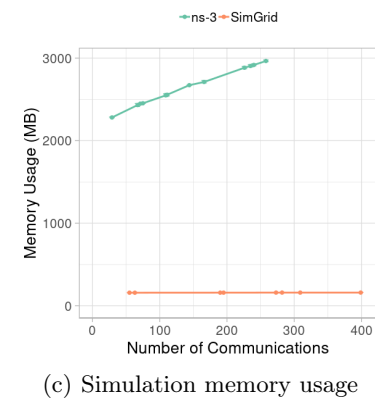
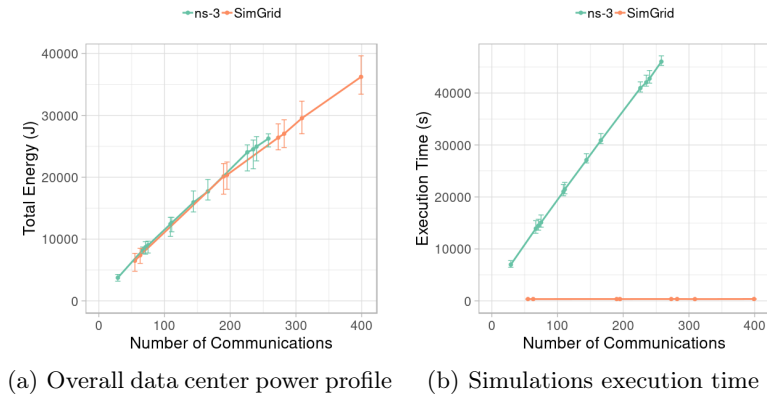


Fig. 5. Energy and scalability simulations results

worst cases, it still predicts energy values in the ECOFEN confidence interval. The homogeneous model has a relative error to ECOFEN lower than 4%, which is a reasonable accuracy regarding to its level of granularity. Performance and scalability results are also shown on Figures 5(b) and 5(c). The execution time is clearly higher on ns-3 than on SimGrid. In fact, for 258 requests, ns-3 takes more than 12 hours against only 6 minutes for 897 requests on SimGrid. This is why we stopped ns-3 experiments earlier than SimGrid ones. On this example, SimGrid is more than 120 times faster than ns-3 with ECOFEN. Similarly, the memory usage is also reduced: for 258 requests ns-3 requires 3GB of memory whereas SimGrid uses at most 169.08MB. It is logical since flow level simulators do not model every packets that are transferred from one host to another. Hence, their memory footprint and their computational overhead are drastically reduced. In this case, SimGrid uses 17 times less memory than ns-3.

7 Conclusion

In this paper, we proposed a network energy model for flow based simulators. We showed that classical models have to be adapted for flow-based simulators due to their coarse grain nature. To evaluate these energy models, we performed microbenchmark experiments. They showed that with simple energy models, we can produce energy consumption estimations close to the ones provided by packet level simulators. We also demonstrated that a fine grain modeling of platforms energy heterogeneity is not required to have suitable energy estimations. Moreover, on simulations of a data center network comprising more than 1,500 servers, we obtained a precision close to packet-level simulators with less than 4% relative error, and with simulation runtimes 120 times faster on flow-level simulators. This realistic use case highlighted how it can be used by the scientific community on large-scale platforms. The implementation is open-source and already available on the website [5].

References

1. Bonomi, F., Milito, R., Zhu, J., Addepalli, S.: Fog computing and its role in the internet of things. In: MCC workshop on Mobile cloud computing (2012)
2. Cornea, B.F., Orgerie, A.C., Lefèvre, L.: Studying the energy consumption of data transfers in Clouds: the Ecofen approach. In: IEEE CloudNet, pp. 143–148 (2014)
3. Culler, D., et al.: LogP: Towards a realistic model of parallel computation. In: ACM Sigplan Notices, vol. 28 (1993)
4. Fiandrino, C., Kliazovich, D., Bouvry, P., Zomaya, A.Y.: Performance Metrics for Data Center Communication Systems. In: IEEE CLOUD (2015)
5. Guegan, L., Orgerie, A.C., Quinson, M.: Reproducible Experiments Repository. <https://gitlab.inria.fr/lguegan/flowlvlwiredenergy> (2018)
6. Kliazovich, D., Bouvry, P., Audzevich, Y., Khan, S.U.: GreenCloud: A Packet-Level Simulator of Energy-Aware Cloud Computing Data Centers. In: IEEE GLOBE-COM (2010)
7. Legrand, A., Marchal, L., Casanova, H.: Scheduling distributed applications: the SimGrid simulation framework. In: IEEE/ACM CCGrid (2003)
8. Orgerie, A.C., Lefèvre, L., Guérin-Lassous, I., Pacheco, D.M.L.: ECOFEN: An End-to-end energy Cost mOdel and simulator For Evaluating power consumption in large-scale Networks. In: Sustainet (WoWMoM Workshop) (2011)
9. Orgerie, A.C., et al.: Simulation Toolbox for Studying Energy Consumption in Wired Networks. In: Int. Conf. on Network and Service Management (2017)
10. Quinson, M., Rosa, C., Thiéry, C.: Parallel Simulation of Peer-to-Peer Systems. In: IEEE/ACM Int. Symp. on Cluster, Cloud and Grid Computing (CCGrid) (2012)
11. Sivaraman, V., Vishwanath, A., Zhao, Z., Russell, C.: Profiling per-packet and per-byte energy consumption in the NetFPGA Gigabit router. In: INFOCOM Workshops (2011)
12. Van Heddeghem, W., Lambert, S., Lannoo, B., Colle, D., Pickavet, M., Demeester, P.: Trends in worldwide ICT electricity consumption from 2007 to 2012. *Computer Communications* **50** (2014)
13. Velho, P., Schnorr, L.M., Casanova, H., Legrand, A.: On the validity of flow-level tcp network models for grid and cloud simulations. *ACM Transactions on Modeling and Computer Simulation* **23**(4) (2013)