GLENDa: Green Label towards Energy proportionality for IaaS Data centers

David Guyon
University of Rennes 1
Rennes, France
david.guyon@irisa.fr

Anne-Cécile Orgerie
CNRS, IRISA
Rennes, France
anne-cecile.orgerie@irisa.fr

Christine Morin
Inria
Rennes, France
christine.morin@inria.fr

ABSTRACT

As cloud services multiply rapidly, so does the computing centers dedicated to them, and consequently their power consumption. Although this consumption is hampering data centers' expansion, these infrastructures have not yet reached energy proportionality, thus wasting significant amounts of energy. Numerous energy metrics have been proposed as incentives towards greener infrastructures, but none of them currently gives direct insights about the energy proportionality and green energy usage of data centers. In this paper, we propose GLENDa: a Green Label towards Energy proportionality for IaaS Data centers. We validate our metric by using traces from real infrastructures, and show that our label gets a better grade when increasing energy efficiency, increasing utilization rates, and using distributed renewable generation. We expect this new metric to become a useful reference for Cloud providers towards green data centers.

CCS CONCEPTS

• **Computer systems organization** → **Cloud computing**; • **General and reference** → *Metrics*; • **Information systems** → *Retrieval efficiency*; • **Hardware** → *Power and energy*;

KEYWORDS

Cloud computing; green computing; green label; data centers

1 INTRODUCTION

The current decade is facing an unprecedented explosion of data and services. This expansion has been powered by Cloud computing which relies on data centers for providing computing and storage facilities that are accessed through the Internet. As an example, in 2014, Facebook generated 4 new petabytes of data and ran 600,000 queries and 1 million map-reduce jobs per day [17].

Expansion of data and services means exponential growth in needed storage capacity [12]. Consequently, the energy consumed by Cloud computing raises worrying concerns. In 2014, U.S. data centers consume approximately 2% of U.S. electricity [16]. These data centers are typically deployed by

Infrastructure-as-a-Service (IaaS) Cloud providers that sell to clients virtual machines hosted on these infrastructures.

With escalating demand and rising energy prices, it becomes essential to improve data center energy and environmental performance. This can be done through three main means: increasing energy efficiency, increasing utilization rates, and using distributed renewable generation [12]. Yet, for the owners and operators of these mission-critical facilities, assessing substantial improvements requires energy efficiency and green energy metrics [9].

The PUE (Power Usage Effectiveness) metric is one of the most used metrics for measuring the energy efficiency of data center facilities [3]. Especially, it leads to good practices for cooling data centers. Its wide adoption by Cloud major companies, like Google which publishes it quarterly since 2008 [8], makes it an interesting incentive towards less consuming infrastructures.

However, it presents major drawbacks as it does not consider the energy efficiency of IT equipment, nor the nature of electricity used – renewable or not. These shortcomings make the PUE useless for assessing power-proportionality of data center infrastructures. Yet, energy-proportional designs would enable large energy savings in IT equipment, potentially doubling their efficiency in real-life use [2].

Furthermore, large data centers are reaching the limits of the PUE: Google, for instance, is achieving a PUE of 1.12 [8], close to the ideal case (i.e. 1). It then becomes necessary to define new metrics, going beyond PUE, to challenge Cloud providers that still consume huge amounts of electricity despite low PUE [16].

In this paper, we propose a new metric which assesses the energy-proportionality and green energy usage of Cloud's data centers. This metric is called GLENDa: Green Label towards Energy proportionality for IaaS Data centers. We believe that, like for the PUE, the adoption of such a metric by Cloud providers could be an incentive towards greener and more energy-proportional data centers. To facilitate its adoption by real world Cloud providers we designed the metric in order to be easily implementable. We provide experimental results exploiting real infrastructure traces showing the suitability of GLENDa for assessing energy-proportionality and renewable energy usage.

The paper is organized as follows. Section 2 presents the related work and the energy-related metrics previously introduced by literature. Our metric GLENDAs is introduced in Section 3. The validation setup is described in Section 4, and results are provided in Section 5. Section 6 discusses the drawbacks and limits of GLENDAs. Section 7 concludes and presents future work.

2 RELATED WORK

The Green Grid consortium [3] defines the Power Usage Effectiveness (PUE) metric as the ratio of total data center energy (E_{total}) over the energy of IT ($E_{ITtotal}$): $PUE = \frac{E_{total}}{E_{ITtotal}}$

The gap between the data center power and IT power is attributed to the cooling system and surrounding electrical equipment. The metric can only be greater or equal to 1 and is better when close to 1 because it means nearly all the energy is used to power the IT facility.

The Green Grid also proposes the Green Energy Coefficient (GEC) metric as the ratio of energy from renewable sources to the total amount of energy [9]. A renewable energy source is defined as an energy source that is consumed at a rate smaller than its regeneration frequency. The GEC varies from 0 to 1 and when it reaches 1, it means that the produced energy is coming at 100% from renewable sources, which is the best case. In the same study the authors present the Energy Reuse Factor (ERF) metric. By dividing the energy which is exported outside of the DC by its total energy consumption, the metric identifies the portion of energy that is exported for reuse. Thus, when the value reaches its maximum of 1 it means all of the energy brought into the DC is reused outside of the DC.

TUE [14] stands for Total-power Usage Effectiveness and aims at giving the total efficiency picture of a HPC DC. It is the ratio of the total energy consumed by a DC over the total energy consumed by the compute components. While the metric is similar to the PUE, it differs by focusing on the energy consumed by the executed computations instead of the whole IT facility power consumption. Where the value of the metric would improve with a higher idle power consumption of the IT facility in the case of the PUE, here the value gets better only if the infrastructure is doing computations.

In parallel, studies have already been done on giving ecological labels to DCs [5] and HPC centers [7]. In their report [5], Greenpeace attributes letters to Internet service providers in order to express their ecological awareness. Many factors are taken into account such as the energy transparency, their position with renewable energy sources, and the carbon and energy intensity of their infrastructure. On the other side, Green Destiny uses an energy efficiency ratio (MFLOPS/Watt) to rank the supercomputers [7].

In order to motivate DC operators to increase the energy efficiency of their DC infrastructure, the European Commission launched the Code of Conduct for Energy Efficiency in Data Centres [4]. This voluntary initiative, managed by the Joint Research Centre (JRC), aims to inform and stimulate DC operators to reduce energy consumption with respect with the mission critical function of DC. Parties signing up will be expected to follow the intent of this Code of Conduct and abide by a set of agreed commitments.

The above studies bring important efforts in showing the energy efficiency and ecological impact of DCs. However the metrics do not offer a simple way to express the global ecological awareness of the infrastructure providers. Studies focusing on giving a label or value according to the energy efficiency of computational systems have been presented. Our work is comparable as it aims at giving an easy-to-understand energy-related metric to the users. Yet, it differs in a way that we aim at expressing the energy-proportionality of DCs and their use of renewable energy sources.

3 OUR METRIC

The PUE success in driving energy efficiency of data center infrastructures can be explained by its simplicity, both mathematical and conceptual [14]. Yet, it does not faithfully account for all the energy wasted in these infrastructures. Indeed, all the IT equipment's consumption is considered as useful and effective despite the non-power proportionality of such equipment. Moreover, depending on the employed electricity sources, this energy consumption can highly impact the planet. We argue for a metric exhibiting the overall energy waste of Cloud data centers and the nature (renewable or not) of their energy supply sources in order to raise providers' awareness on this energy misuse.

A typical server consumes power depending on its resource usage (generally CPU load) [10]. Similar behavior can be observed for network devices [6]. The idle power consumption of an IT device represents its static consumption when turned on but idle, so not performing any useful work. Ideally, this consumption should be null. However, practically, it is not the case and for instance, the servers used in Section 4 for the validation experiments exhibit an idle consumption around 100 Watts. This idle power consumption represents a static consumption, while workload is inducing a dynamic energy consumption.

By definition of the PUE, the overall energy consumption of a data center E_{total} can be expressed as: $E_{total} = E_{ITtotal} \times PUE$ where $E_{ITtotal}$ represents the total energy consumption of all IT equipment (static and dynamic parts). Then, we define our metric GLENDAs as follows:

$$GLENDAs = \frac{E_{ITdynamic}}{E_{ITtotal} \times PUE} \times GEC$$

where $E_{ITdynamic}$ represents the dynamic energy consumption of all IT equipment, and PUE and GEC are respectively the PUE and GEC of the data center as defined in Section 2.

Basically, GLENDAs represents the ratio of the useful energy ($E_{ITdynamic}$) over the total energy consumption of the data center infrastructure ($E_{ITtotal} \times PUE$) multiplied by the ratio of green energy powering this facility (GEC). Thus, for any data center, $0 \leq GLENDAs \leq 1$ with 0 being the worst case when there is no useful energy consumed, and 1 the best case when all the energy consumed is useful and provided by renewable sources.

As for the PUE, GLENDAs is defined for a given period of time. In the validation experiments in Section 5 for instance, we will compute it daily and monthly.

In order to ease GLENDAs adoption and understanding by Cloud providers, we propose to map GLENDAs values on an alphabetical scale ranging from A++ to F as follows: A++ class from 1 to 0.875, then A+ down to 0.75, A down to 0.625, B down to 0.5, C down to 0.375, D down to 0.25, E down to 0.125 and F class from 0.125 to 0.

Various approaches exist in ecolabeling and we could have opted for a relative rating approach, for instance using fractions of the GEC as frontiers between categories or defining categories for a given PUE or a given data center type (server rooms, mid-tier data centers, enterprise data centers, etc.). But, we chose an absolute rating scale as it allows consumers to easily compare options because of its simplicity, and it does not cause misleading or confusing effects on users like relative labeling approaches [11].

4 VALIDATION SETUP

In order to evaluate the validity of the metric, we analyzed the variation of GLENDAs in different scenarios and compared the results. The considered scenarios are the following ones:

- baseline:** utilization trace and power trace taken from a real computing infrastructure used without any modification (detailed later).
- vary-on/vary-off (VO/VO):** when a server is not used, it is powered down, thus the power consumption while servers are not used has been removed from the power trace.
- power-proportional (PP):** the power consumption of the servers is considered to be proportional to the utilization ratio. Thus, servers only have a dynamic power consumption that is reaching the maximal power consumption when the server is fully used, and that is null when the server is idle.
- max power (PP with P_{max}):** this scenario expresses an infrastructure which is fully used at all time; whenever a server is utilized, its power consumption equals to the maximum server power consumption, and when unused, its consumption is null.

These four scenarios are studied under different data centers' characteristics regarding their PUE (representing the energy efficiency of the data center's infrastructure) and GEC (showing the mix of energy sources). In our study, we consider a time window of 4 weeks (28 days) and we give a per-day value of GLENDAs.

As depicted in Section 3, the calculation of GLENDAs requires the following energy information: E_{total} and $E_{ITdynamic}$. The first value is computed by summing all IT energy consumption data for a given period of time, and then multiplying it by the PUE: $E_{ITtotal} \times PUE$. The computation of the second value requires to know the idle energy consumption of the IT equipment. We compute this value by multiplying the idle power of the n IT devices (average power measured during an idle period of time) by the considered time window size (T), and then by summing of the idle power values: $E_{ITidle} = \sum_{1 \leq i \leq n} P_{idle}^i \times T$

Then we compute the following equation in order to get the dynamic consumption of the IT facility scale: $E_{ITdynamic} = E_{ITtotal} - E_{ITidle}$

Finally we divide $E_{ITdynamic}$ by E_{total} and we multiply the result by the GEC coefficient. The PUE (used to compute E_{total}) and the GEC are measured over the same time period T as the two energy values.

The Grid'5000 platform is a French experimental platform [1]. It provides access to about 1,000 servers spread on 8 sites linked through a dedicated high-speed network. The platform keeps an history of the servers' utilization (i.e. jobs) and per-second and per-device traces of the IT power consumption.

We are exploiting traces from November 2015, from the 1st to the 28th day, because this month displays a representative mix of usage and idle times. As for the location, a cluster called *Taurus* has been selected because it is equipped with fine grained wattmeters. This cluster is located in the Lyon site of Grid'5000. From the selected cluster, 10 servers were analyzed and they hosted a total of 1624 jobs over the considered period.

In order to evaluate the metric on a large-scale data center, we multiplied the power consumption of the servers by a factor of 15, thus representing 150 servers. In addition to these servers, we included in the idle energy consumption: 3 servers with a constant power of 120 Watts (representing front-end and storage servers); 4 network switches consuming 150 Watts each (one for each server rack, this value has been measured on a Grid'5000 rack switch from the Lyon site); 1 network router consuming 400W (for the outgoing network).

Here, the network devices are included only in the idle part of the energy consumption (E_{ITidle}) because over the Grid'5000 2015 traces, we observed at maximum a variation of 2.5% of their power consumption, which is considered

negligible in our experiments (less than 4 Watts of dynamic power consumption). Similarly, the front-end and storage servers are considered here to be fully part of the idle part as we do not have wattmeters on these devices, and being in the idle part is the worst case for GLENDa.

In the utilization trace, each job has start and end timestamp values and a server location which are utilized to synchronize the jobs with the power consumption traces. These values are then used to calculate the $T_{ITdynamic}$ and $T_{ITtotal}$. The idle power consumption is measured each week. It is indeed stable over such a period of time.

The RTE France website provides an hourly detail of their electricity production [15]. It contains the average power in MW produced from all of the power sources: biomass, gas, coal, fuel oil, nuclear, wind turbines, incineration of wastes, Pumped Heat Electrical Storage (PHES) and Run-Of-the-River (ROR) hydroelectricity with and without storage reservoir. With such a detailed information, it is possible to calculate the hourly ratio of renewable energy versus the total energy received, thus the GEC.

The sources which are considered as renewable are: biomass, PHES, ROR (with and without storage reservoir), solar, and wind turbine.

From the RTE trace of November 2015, the calculated average GEC is 0.14. Figure 1 shows its variation over the first 24 hours of this specific month. The daily average is of 0.14 which is the same as the weekly and monthly average.

Figure 1: Per-hour GEC variation of the 1st of November 2015

In order to see the influence of the GEC on GLENDa, we also use two other values for the GEC: 0.50 when the green energy represents 50% of the source, and 0.80 when it represents 80% of the source. These are not real values but set in order to represent an on-site production of energy with solar panels for example.

In the validation experiments, three different PUE values are considered. The first PUE is 1.53 which is the PUE measured at the Rennes site of Grid'5000 on February 2017 (no PUE measurement for the Lyon site in November 2015). The

second PUE of 1.12 is the one given by Google for its data centers at the time of the study [8]. A PUE of 1.83 is taken for the third value, because it represents the lower bound case of data centers from 2010 [13].

5 VALIDATION RESULTS

GLENDa is evaluated for different scenarios. The first experiment gives the daily variation of our metric with the baseline scenario. The subsequent experiments compare the GLENDa values obtained in all the scenarios. The last experiment shows the distribution of the GLENDa ecolabels according to the scale described in Section 3.

Figure 2 presents the results with the real utilization and power traces unmodified. In this baseline scenario, the PUE and the GEC are respectively fixed at 1.53 and 0.14. At the bottom of the figure, the usage ratio varies from 0.36 to 1.0 with an average of 0.86. This average means that the platform is utilized 86% of the time. Above this plot, the daily average power consumption is shown alongside the idle power consumption. At the top is the plot of the daily variation of GLENDa over the full month of November 2015.

Figure 2: Variation of GLENDa, power consumption and usage ratio from November 2015 in the baseline scenario

For the two first days, we can see that for a high and a low usage the power consumption is nearly equal to the idle power. While it is understandable that low usage implies low power consumption, it seems uncommon that high usage does not increase the power. The reason is that user's jobs

do not always execute intense computations. The variation of GLENDa follows the dynamic power, thus validating the high utilization rate property. GLENDa is close to 0 when there is no dynamic power and reaches 0.043 when the dynamic power is at its maximum.

In this experiment, we compare the variation of GLENDa with the different scenarios presented in Section 4.

Figure 3 compares the baseline, VO/VO, PP and max power scenarios. The PUE and the GEC are fixed at respectively 1.53 and 0.14. VO/VO allows to decrease the idle power part when usage is not at its maximum. With this scenario the value of GLENDa becomes higher than with the baseline as expected by the property designed to highlight energy efficient infrastructures. Yet, it is not far from the baseline value because of the high usage ratio (86% on average), which implies small periods of idle time.

Figure 3: GLENDa comparison with scenarios where the infrastructure uses PP servers and full resource usage users' applications

In the case of power-proportional servers, GLENDa is significantly higher than with VO/VO. This is because E_{ITidle} is then very low and the energy consumption of the servers is driven by the usage and the computations caused by the jobs. However, a high usage without computation (idle job) gives a low GLENDa value as shown in the first day. In the fourth scenario, GLENDa remains stable around 0.087. This behavior is due to the absence of idle job. Indeed, in this scenario, each job reaches a maximal power consumption and the idle power of servers is null. It explains why the first day does not have a low value of GLENDa. Compared to the other scenarios, $E_{ITdynamic}$ is larger, and E_{ITidle} energy is the same as the PP scenario. This experiment shows that GLENDa meets the expected properties for increased energy efficiency and increased utilization rate. Indeed, the metric increases when the energy efficiency improves (VO/VO scenario), when servers are power-proportional (PP scenario) and when the utilization rate is at its maximum (PP with P_{max} scenario).

In Figure 4, we compare the GLENDa variation with 3 different values of PUE for the baseline and the PP scenarios.

Figure 4: GLENDa comparison of a varying PUE with the baseline and PP scenarios

The GEC is fixed to 0.14. In comparison with the previously used PUE of 1.53, GLENDa is globally higher when the PUE is 1.12. This is because the total energy consumed by the DC is lower when the PUE gets closer to 1. As expected GLENDa decreases when the PUE increases (1.83). Indeed, the data center is wasting more energy in surrounding equipment such as the cooling system.

Figure 5: GLENDa comparison of a varying GEC with the baseline and PP scenarios

For the baseline and PP scenarios and with a fixed PUE of 1.53, Figure 5 shows the variation of GLENDa with three different values of GEC: 0.14, 0.50 and 0.80. When the part of renewable energy increases to 50% GLENDa increases in both scenarios. The same behavior is observed when it reaches 80%. As expected by the last property of the metric, increasing the portion of used renewable energy increases consequently the value of GLENDa.

From the previous experiments, we compute the average value of GLENDa over the whole period and plot the values in Figure 6. The x axis follows the GLENDa ranges presented in Section 3 with a logarithmic scale to make the low values more distinguishable and a color has been given to each label.

The four markers at the top of the figure are the average GLENDa for the baseline, VO/VO, PP and max power scenarios, which daily values have been presented in Figure 3. The second line shows the six markers for the baseline (in blue)

Figure 6: The ecolabels are distributed to each scenario according to their monthly average value of GLEND A

and PP (in green) scenarios with three different values of PUE, previously detailed in Figure 4. The bottom line exposes the average value of GLEND A for the same two scenarios but with three different values of GEC, previously detailed in Figure 5.

As shown by the figure, most of the values are located in the F label area. Nonetheless the scenarios with power-proportional servers are closer to the E label than the baseline and the VO/VO scenarios. The PP scenario is almost in the E label when its PUE value falls at 1.12. On the other hand, both scenarios move away from the E label when the PUE gets worse. The attributed label changes from F to E when the baseline scenario has 80% of its energy from renewable energy. However the PP scenario surpasses the baseline scenario given that it is attributed the D label when GEC is 50%, and the C label when GEC is 80%. Such a label should then spur Cloud providers to greater efforts for operating greener and more energy-proportional data centers.

6 DISCUSSION

From a practical point of view, in order to measure the power consumption of the IT devices, wattmeters are required. In addition a dedicated server is necessary to store the data along with a software layer to process the data stream and the incoming requests. The installation represents an investment in terms of time, price and maintenance.

Measuring the idle power consumption of IT devices can be done after their first installation and after maintenance periods as it may vary over time [10]. This measurement of the idle consumption can be done for entire racks or the entire infrastructure, given that all devices are idle at the same time (like during a maintenance on a rack).

For the sake of clarity, in our experiments, we consider only the dynamic consumption and the power proportionality of computing servers. But, GLEND A's scope is more generic and includes other infrastructure devices such as

network switches, gateways and storage servers. Improving their energy efficiency also impacts GLEND A.

Considering the GEC, as explained by the Green Grid consortium, the green energy is defined as any form of renewable energy for which the data center owns the rights to the green energy certificate or renewable energy certificate, as defined by a local/regional authority [9]. This includes on-site electricity generation from solar panel or wind turbines for instance. In the extreme case where the GEC is null, GLEND A is also null. This is desirable feature, as using no renewable energy at all cannot be considered as a sustainable behavior.

GLEND A keeps the simplicity's spirit of the PUE: the energy consumption values are easy to get (if the data center has wattmeters). GLEND A's main target is the data center operators which would like to estimate their energy proportionality. Then, if Cloud providers expose the GLEND A label of their data centers, it can help energy-aware users to choose between them.

7 CONCLUSION

In this paper, we introduce GLEND A, a metric to evaluate the energy-proportionality of data centers and their usage of energy from renewable sources. Based on traces from a real infrastructure and specifically designed scenarios, we evaluate how the metric behaves over a month. Our validation shows higher values of GLEND A are given whenever an infrastructure has power-proportional servers, its IT facility power consumption is near the total power consumption and when it consumes its energy from non-fossil sources. Finally we establish value ranges in order to give a label along the value of GLEND A, thus offering an easy-to-understand metric to the users. The metric is designed in order to be easily implementable by DC operators. Its wide adoption, like for the PUE metric, could be rapid if included in the best practice guidelines of the Code of Conduct [4].

ACKNOWLEDGMENTS

Experiments presented in this paper were carried out using the Grid'5000 experimental test-bed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>).

REFERENCES

- [1] Daniel Balouek and others. 2013. Adding Virtualization Capabilities to the Grid'5000 Testbed. In *Cloud Computing and Services Science*. Communications in Computer and Information Science, Vol. 367. Springer, 3–20.
- [2] L. A. Barroso and U. Hölzl. 2007. The Case for Energy-Proportional Computing. *Computer* 40, 12 (2007), 33–37.

- [3] Christian Belady, Andy Rawson, John Pflueger, and Tahir Cader. 2008. *Green grid data center power efficiency metrics: PUE and DCIE*. Technical Report. Technical report, Green Grid.
- [4] European Commission. 2008. Code of Conduct on Data Centres Energy Efficiency. <https://e3p.jrc.ec.europa.eu/communities/data-centres-code-conduct>. (2008).
- [5] Gary Cook and David Pomerantz. 2015. Clicking clean: A guide to building the green internet. (2015).
- [6] I. Cuadrado-Cordero, A.-C. Orgerie, and C. Morin. 2015. GRaNADA: A Network-Aware and Energy-Efficient PaaS Cloud Architecture. In *IEEE International Conference on Green Computing and Communications (GreenCom)*. 412–419.
- [7] Wu-Chun Feng and Kirk Cameron. 2007. The green500 list: Encouraging sustainable supercomputing. *Computer* 40, 12 (2007).
- [8] Google. 2017. Google data center PUE performance. <https://www.google.com/about/datacenters/efficiency/internal/>. (2017).
- [9] Green Grid. 2014. *Harmonizing Global Metrics for Data Center Energy Efficiency, Global Taskforce Reaches Agreement Regarding Data Center Productivity*. Technical Report. The Green Grid Consortium.
- [10] Franz C. Heinrich, Alexandra Carpen-Amarie, Augustin Degomme, Sascha Hunold, Arnaud Legrand, Anne-Cécile Orgerie, and Martin Quinson. 2017. *Predicting the Performance and the Power Consumption of MPI Applications With SimGrid*. Technical Report. Inria.
- [11] Stefanie Lena Hille, Christian Geiger, Moritz Loock, and John Pelozo. 2016. Best in Class or Simply the Best? The Impact of Absolute Versus Relative Ecolabeling Approaches. *Journal of Public Policy & Marketing* (2016), 1–50.
- [12] Nathaniel Horner and Ines Azevedo. 2016. Power usage effectiveness in data centers: overloaded and underachieving. *The Electricity Journal* 29, 4 (2016), 61–69.
- [13] Jonathan Koomey. 2011. Growth in data center electricity use 2005 to 2010. *A report by Analytical Press, completed at the request of The New York Times* 9 (2011).
- [14] Michael K. Patterson, Stephen W. Poole, Chung-Hsing Hsu, Don Maxwell, William Tschudi, Henry Coles, David J. Martinez, and Natalie Bates. 2013. TUE, a New Energy-Efficiency Metric Applied at ORNL's Jaguar. In *International Supercomputing Conference (ISC)*. 372–382.
- [15] RTE. 2017. Eco2mix. <http://www.rte-france.com/en/eco2mix/eco2mix>. (2017).
- [16] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner. 2016. *United States data center energy usage report*. Technical Report. Lawrence Berkeley National Laboratory Tech. Rep. LBNL-1005775.
- [17] J. Wiener and N. Bronson. 2014. Facebook's Top Open Data Problems. Facebook Research, <https://research.fb.com/facebook-s-top-open-data-problems/>. (2014).