

On the Sustainability of Virtual Platforms: A Behavioral Intervention

BORJA MARTINEZ¹, (Senior Member, IEEE), AND

XAVIER VILAJOSANA¹, (Senior Member, IEEE)

IN3, Universitat Oberta de Catalunya, 08018 Barcelona, Spain

Corresponding author: Borja Martinez (bmartinezh@uoc.edu)

This work was supported in part by the Digital Future Society, a Program of Mobile World Capital Barcelona Foundation, in collaboration with Barcelona City Council; in part by the Spanish Ministry of Science, Innovation and Universities under Grant RTI2018-095438-A-I00; and in part by the Generalitat de Catalunya under Grant 2017-SGR-60.

ABSTRACT The energy required to supply data centers today is estimated to be around 1% of the global demand of electricity, with the cloud computing paradigm being the main driver of computing demand. Leading cloud providers are already making efforts to reduce energy expenditure of their data centers. However, the role that online platform operators and end-users can play towards a more sustainable cloud is still unclear. In this article, we raise the question whether making end users aware of their impact on cloud energy expenditure leads to more efficient use of the platforms. Focusing on non-retail platforms, we have run an A/B test in the Virtual Campus of the Universitat Oberta de Catalunya (UOC), one of the biggest online universities worldwide. In this intervention, we show the test group real-time information about the energy consumption of the platform, as well as tips on how to reduce it. Alongside, we monitor user behavior in terms of session duration and volume of traffic generated. Our results reveal that users who received this information did not change their behavior significantly. This result encourages us to find alternative ways to reduce the energy impact associated with the platform without the active participation of the end user, such as a more intelligent session management in conjunction with auto-scaling tools.

INDEX TERMS Cloud sustainability, environmental impact, virtual platform, cloud platform.

I. INTRODUCTION

Environmental threats are, at last, a global concern. Citizens, aware of the vital need to preserve our planet, are now demanding further commitment from the companies of which they are customers. Companies that base their commercial activity on the Internet cannot stay out of this global trend. For data centers are the backbone of the Internet and, unfortunately, they are energy-intensive plants. The energy required to supply their computing hardware and facilities worldwide is estimated in 205 TWh for 2018 [1]. This represents around 1% of the global demand of electricity (23.031 TWh in 2018 [2]) and is comparable to the total demand of a medium-sized country like Spain (249 TWh in 2020 [3]).

The cloud computing paradigm—on-demand availability of computer resources, mainly data storage and computing power—lies at the core of this, often referred to as “data revolution” [4]. Beyond e-commerce, cloud services have

become essential for major industrial sectors such as manufacturing systems, logistics, telecommunications and finance. Eventually, this dependence on the cloud will extend to all sectors of society, from education to public administration. The massive demand for cloud resources from both business and consumer use has led to the development of large-scale public cloud data centers called hyperscale data centers. Together, cloud and hyperscale-cloud data centers accounted for 89% of computing instances in 2018 [1], which makes them the primary target for energy optimization.

Indeed, the leading cloud providers (Amazon, Microsoft and Alphabet Inc.) are supporting sustainable policies to mitigate the environmental footprint of their data centers. BloombergNEF recently reported [5] that these main cloud providers, headed by Google, are the corporations closing more power purchase agreements (PPAs) for renewable energy worldwide, demonstrating an aim to operate their datacenters from green energy sources. Among other initiatives, since 2017 Google matches 100% of their annual electricity use with purchases of renewable energy, and

The associate editor coordinating the review of this manuscript and approving it for publication was Alon Kuperman¹.

for 2030 they aim to run on entirely 24/7 carbon-free energy [6]. No less important, modern public cloud and hyperscale data centers are much more efficient than the smaller, traditional ones, which accounted for 79% of computing instances in 2010. As reported in [1], this shift of global computation from traditional to hyperscale data centers explains why, although overall computing instances increased by 550% from 2010 to 2018, energy increased by only 6%.

Still, there are two stakeholders in the data center power equation, both trying to optimize energy consumption with different constraints. On the one hand, cloud providers strive to offer their services in a more energy-efficient way. However, they are not interested in reducing the volume of services, as their income directly depends on it. On the other hand, cloud customers' efforts focus on reducing the volume of leased services, as their cloud bill is proportional to this volume. Note that cloud providers and customers actions to reduce power consumption result in a more sustainable operation, but also an economic gain for both parties. The difference is the target they optimize.

This work is contextualized in a representative example of the cloud customer side, the Universitat Oberta de Catalunya (UOC). As many companies today, the UOC operates most of its services from servers hosted in the cloud, outsourcing resources to the three main hyperscale cloud providers: Amazon AWS, Microsoft Azure and Google Cloud. At the UOC, teaching is conducted entirely online and its platform—the UOC's Virtual Campus—serves more than 100.000 registered users.

Throughout this study, we pose several fundamental questions related to the energy sustainability of cloud-operated non-retail platforms, such as the Virtual Campus. As cloud customers, platform operators have little or no control of the underlying hardware, which is owned and managed by the cloud provider. As service providers, they have little or no control of how their users interact with the platform. Thus, with this little room for maneuver, how can a platform operator reduce its energy expenditure while still offering quality services to its users? Is energy optimization the sole responsibility of the platform operator? Alternatively, can users, once properly informed, help to reduce the impact of online services? Even more, are they essential players in this challenge?

Interestingly, most online platform users—including the students at the UOC Campus—are not aware of the huge energy demand of cloud services. Nor they are of how the applications running in the background are contributing to the overall energy expenditure beyond their terminal devices. Most importantly, end users are to a large extent free in their interaction with the digital services, but not financially responsible for the energy costs (they don't pay the bill). To make things more complex, The main goal of our platform is to help the student's learning process. This fact is key, because it implies that the use of the platform cannot be

discouraged in order to reduce consumption, but rather the opposite, a proper use has to be promoted.

To get insight into these questions, we have carried out a pilot during the second semester of 2021 in the UOC's Virtual Campus. The pilot has been conducted following an A/B test methodology in which randomly selected users have been informed of the energy impact of their activity. Throughout the pilot, we have monitored the Campus from two different angles: usage patterns and back-end resources. Regarding the first aspect, we have anonymously recorded the activities of users in their interaction with the site, such as session time, file downloads, links clicked, etc. On the other side, we have monitored—with fine temporal resolution—the dynamic allocation of cloud resources and the load of the assigned services. This information is used to compute in real time the environmental impact associated with the platform, which is shown to the connected users. The main goal of the pilot is to assess whether displaying this information affects users behavior towards more rational use.

The rest of the article is organized as follows. In Section II we overview different behavioral strategies and justify the selection of the specific one used for our study. Section III describes the UOC's Virtual Campus and methodology related to the performed A/B test as well as the monitored metrics. Then, in Section IV we present and discuss the results of our study. Finally, in Section V some concluding remarks are given.

II. BEHAVIORAL STRATEGIES

According to the European Environment Agency [7] most environmental policy interventions can be classified into one of the following strategies, or a combination of them.

The first strategy comprises traditional regulatory approaches, sometimes referred to as prescriptive or command-and-control measures (CAC). For example, emission standards establish the legal requirements governing air pollutants released into the atmosphere. Irrigation constraints for agriculture or even domestic water rationing during severe drought periods also fall in this category. The second strategy consist of market-based instruments. For example, environmental taxes (e.g. charges on single-use plastic bags) or progressive price-based approaches to water demand fall in this category. The third strategy involves awareness raising, including for example energy efficiency labels and communication campaigns.

First of all, with respect to prescriptive policies, they make sense for utilities (water, electricity and gas), typically under the control of governments. However, it is difficult to imagine how to regulate online platforms. Eventually, a regulatory policy could apply to leisure platforms, but it would hardly affect an education-oriented platform which, by contrast, should generally be encouraged.

Regarding the second strategy, a usage-based pricing policy may be difficult to accommodate for a virtual platform. It is worth noting that, while a household pays directly for its electricity consumption, the user/customer of a virtual

platform does not pay for the energy. And again, in the case of education, the goal should be to increase usage—encourage study—, only reducing misuse.

These constraints, both the general framework of our experiment (a virtual platform) and the specific characteristics (users undergoing university education), suggest that the best way to address the problem is through a behavioral strategy, that is, raising awareness seeking users engagement.

Behavioral influencing tactics range from reflective to automatic responses [8]. Reflective responses appeal to the conscious processing of information, in which decisions are made on the basis of rational arguments. They are analytical and dominated by reason. Complementary, but not necessarily exclusive, there is another decision mechanism based on intuition. Responses here are fast, automatic, effortless, associative, and often emotionally charged; they are also governed by habit [9].

Covering the whole range, different methods have proven effective in generating environmental awareness, but also in many other areas, from finance to commerce. [8] identifies the following categories of decision-making and choice behavior:

A. REFLECTIVE RESPONSES

- *Knowledge transfer*: provide factual information to increase awareness, without requiring any specific action. For example, in utilities, simply disclosing the environmental impact to consumers has often been effective.
- *Increasing self-efficacy*: seeks to convince people of their ability to achieve results by providing tips, advice and concrete examples. A suggestion for reusing towels in a hotel is a simple example of this category.

B. SEMI-REFLECTIVE RESPONSES

- *Social norms*: people tend to feel uncomfortable when they do not feel integrated into their social environment. Informing an individual about “normal behavior” (e.g. comparisons of their home’s energy use to similar households in the neighborhood) has proven to be quite effective in shaping behavior, a process that may occur without conscious intent and awareness.
- *Framing*: the framing effect is a cognitive bias where people decide on options based on whether they are presented with positive or negative connotations. For example, emphasizing the negative effects of climate change, such as drought or species extinction.
- *Tailoring*: personalized messages, exposing an individual to the consequences of his/her particular usage of the resource. This category includes real-time feedback tools, such a shower meter or a household energy consumption smart-phone app.

C. AUTOMATIC RESPONSES

- *Emotional shortcuts*: emotional biases occur spontaneously based on the personal feelings of an individual

at the time a decision is made. Therefore, evoking an emotion may affect the decision. For example, humor tends to break resistance.

- *Priming*: exposure to a specific stimulus that potentially activates different cognitive constructs and influences subsequent individual actions. A key point here is that such “primes” may not be directly related to the topic. For example, making someone feel ashamed (for whatever reason) has been shown to induce water conservation.
- *Nudging*: design a choice architecture [10], making it more likely that an individual will make a particular choice, without forbidding options, nor limiting freedom of choice. A simple example that has proven highly effective for utility companies is just making a “green” tariff the default choice.

In our study we have chosen an approach oriented to a rational individual decision. Our experiment has been designed to provide unbiased information about the overall consumption of the platform (“knowledge transfer” category), as well as simple and concrete actions that indicate users what they can do to reduce the energy impact of their navigation (“self-efficacy” category).

We have several reasons for this choice. First of all, being an educational institution, we consider reflective responses much more appropriate, as they appeal to rational arguments and so aim to inform and empower subjects to make rational, informed decisions. On this basis, we discard those categories in which the emotional decision-making process has a predominant weight (“framing”, “tailoring”, “emotional shortcuts” and “priming”)

Among the remaining categories, two of them deserve special attention (“nudging” and “social norms”), because they have been discarded in our case, despite having given excellent results in other types of interventions. The following two subsections go into detail on these issues. We conclude the section by reviewing some key interventions with a similar approach to ours.

D. SOCIAL NORMS

Social norms refer to common standards for behavior, set by and for members of a social group. Behavior interventions based on social norms—contrasting individual performance with the group standard—have proven quite successful in promoting environmental and social sustainability in multiple contexts, such as water conservation [11], household energy [12], sustainable transportation [13], acceptance of electric vehicles [14], among others.

The comparison of individual performance with the group may have negative unintended consequences. In our particular case, exposing a student to information regarding the study time spent by other students (indirectly through the session time, which is used to calculate energy) can have a demoralizing effect, particularly for less gifted students. It is a decision that undoubtedly conditions the strategy and

therefore the results of the pilot, but it has been taken in accordance with University policy.

E. NUDGES AND DEFAULT CHOICES

Nudging, and particularly default rules, which consist of presenting as the default the most environmentally friendly option, has proven the most efficacious means at promoting more sustainable choices. For example, [15] show the dramatic impact of automatic enrolling in savings plans. For energy, [16] and more recently [17] and [18], report substantial benefits when offering 'green' energy (a more expensive tariff) as the default option for domestic energy contracts.

Compared to other types of interventions, default choices, when applicable, very often achieve overwhelming results. However, it is not always obvious how to apply this method. For example, it is by no means clear how to define a green default for domestic water consumption. In these cases technology can be a great help. For example, motion detectors that turn out the lights when people do not appear to be in the room, create the equivalent of an 'off' default. We will discuss the applicability in our particular case later on.

F. INFORMATIVE INTERVENTIONS AND REAL-TIME FEEDBACK

As mentioned above, we have adopted an informative approach for our pilot. The end-user's response, once informed of the resources associated with their activity, and in the absence of incentives, has been extensively studied in other application areas.

For example, [19] investigates whether information about the environmental health effects of energy consumption could impact conservation behavior in household electricity. The authors conclude that environment and health-based information generated important energy savings, particularly in families with children. In the case of water, there is strong evidence that domestic consumption can be reduced by using different strategies to influence behavior, without the need for economic incentives, policy instruments or regulations [8]. Analogous success stories have been reported for recycling [20] and sustainable transportation [21].

Sometimes it is possible to show users relevant information in real time through the use of digital technologies. [22] provides an extensive review of studies in multiple domains in which digital technologies have been used to change habitual behavior. These include environmental interventions, but also self-regulation for exercise, diets and health in general. For example, focusing on environmental awareness, [23] studies the effect of information feedback about energy consumption using in-home displays. [24] analyzes the feedback through digital technologies for water conservation. [25] studies the effect of real-time feedback on resource consumption during showering, concluding that information induces substantial water and energy conservation. To the best of our knowledge, there is no published work evaluating the effect of real-time feedback on virtual platforms.

III. METHODOLOGY

In this section, we overview the methodology used in this study to assess our hypothesis. Namely, whether showing cloud consumption information affects user behavior. We first describe the UOC Virtual Campus, i.e., the system in which the intervention is deployed, and its specific constraints. We then detail the A/B test we have conducted. Finally, we present the specific information and how that information is presented to the variation group, as well as the metrics that are monitored regarding the user behavior.

A. UOC VIRTUAL CAMPUS

We have performed our experiment in the Virtual Campus of the Universitat Oberta de Catalunya (UOC). Since the academic activity at UOC is fully online, users need to interact with the teachers, colleagues and learning resources online. This interaction is mainly performed via a web application: the Virtual Campus. The teaching methodology at UOC is based on providing a guided learning plan in which the students are provided with: *i*) learning resources such as documents and videos that can be accessed online or downloaded, *ii*) interaction channels such as text and video forums, hosted in the platform, and *iii*) teacher support that can take the form of group communication in the forums or individually via e-mail. In addition, continuous assessment activities are one of the main tools for evaluation at UOC, these imply either solving online quizzes or uploading documents to the Virtual Campus. The platform has therefore a crucial role in the learning process and is heavily used by students during the academic semester. With more than 100.000 active users and 3.000.000 visits/monthly on the site, the Virtual Campus is a perfect experimentation platform for our purposes.

Once again, it is important to stress that the Virtual Campus is a crucial platform for the continuous learning process of the student. Therefore, its use cannot be discouraged in order to reduce consumption, yet there are good habits that can be promoted to reduce energy consumption. One of these is to perform manual log out once the user ends the study session. If no manual log out is performed by the user, the session is left open until a timeout expires, even when the user closes the web browser. Having a single session active marginally increments consumption, but becomes considerable when the amount of concurrent users is high and resources need to be outsourced. Another habit that unnecessarily consumes energy is repeatedly downloading some content multiple times. As resources are always available at the Virtual Campus, users may access multiple times the same content, while it would be more efficient to store a copy locally and access it offline. Other actions such reading forum entries, interacting with teacher and colleagues, solving quizzes, etc. are not only appropriate, but necessary and cannot be discouraged.

B. A/B TEST

At the core of our study is a well known method called controlled A/B experiment (or simply A/B test). This method

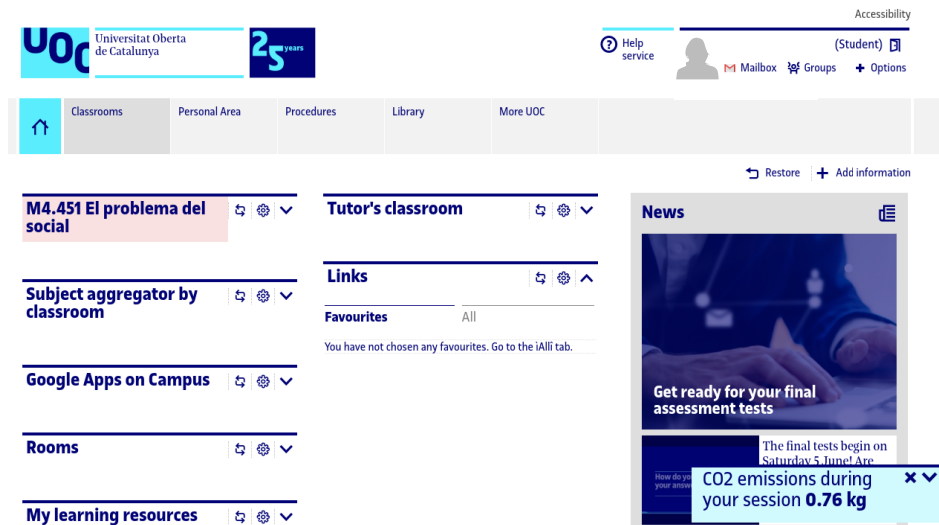


FIGURE 1. Placement of the widget in the Virtual Campus platform.

is the most reliable way for establishing a causal relationship between some change purposely introduced in the application and its influence on user-observable behavior [26]. In a controlled A/B test, users are randomly split between two groups. In the context of web A/B tests, users in the control group are presented with an unchanged version of the site, while the users in the test group are presented with a variation. The interactions with the site are instrumented, and key metrics are computed. Causality is then established when a difference with statistical significance is observed between the two groups for any of the indicators analyzed. If causality is established, the hypothesis is confirmed.

In our experiment, users were divided into two groups uniformly at random depending on whether the result of hashing¹ their Virtual Campus identification number (ID) is even or odd, effectively performing a persistent A/B test. That is, a specific user always sees the same version (unchanged or variation) of the webpage for the duration of the experiment.

Users in the test group were displayed with the widget shown in Fig. 1 (variation). The widget follows a corporate design template, which offers three levels of information as the user scrolls through the widget. In the first level, displayed by default as soon as the user logs in, the widget shows the carbon dioxide emissions of the campus while their session is active (see Figure 2.a). This information is complemented with tips on how to reduce consumption (Fig. 2.b) in the second level. More information on the aim of the test to promote awareness is placed in the third level (Fig. 2.c). In order to access the different layers of the widget, users need to interact with it. To access the tips they need to press the downward arrow (“Expand” button), while to access the information on the test, they need to press the “See More” button. In the third level, users can then click the “More Information” hyperlink

to access an external project webpage² where detailed information about the project’s goal is provided.

C. GROUP SELECTION

The UOC has a broad spectrum of users who carry out very different activities. The first distinction is obvious: faculty, administration staff and students interact with the platform very differently according to the nature of their tasks. Even within the students there is great variability – standard students, short postgraduate courses, language courses, technical vs humanistic students, etc. To isolate the effect we wanted to observe from other sources of variability, we selected the sample included in the study by creating groups that were as homogeneous as possible. To this end, all participants belonged to the students’ collective and were already enrolled at the start of the pilot (May 2nd, 2021). Within these students, we selected the 5 degree programs with the largest number of students, so that the sample was statistically representative. We also evaluated the results of these groups separately, as a way to further reduce variance. Note that segmentation is a common practice to reduce variance in A/B Tests [27]. In both groups of the A/B test, we selected 25% of the total users, as seen in Table 1.

D. METRICS

In the course of the user’s session, we collected different metrics such as the session start and end time as well as metrics on the user interaction with the learning resources. We also collected information related to the interaction of the user with the widget, such as how many users clicked on specific buttons and hyperlink. These metrics will give us insight on the interest shown by the users.

¹A 256-hash ensures the privacy and anonymity of the users.

²<https://efc.research.uoc.edu/en/what-is-all-this-about/>

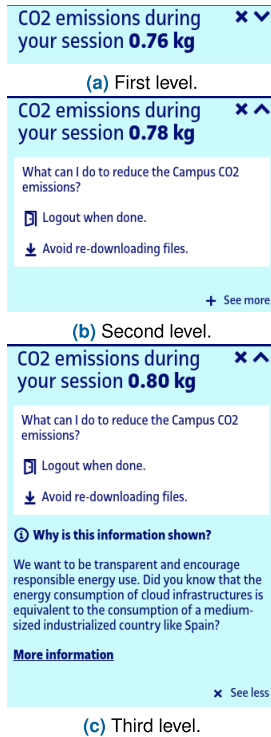


FIGURE 2. The three different expansion levels of the widget.

TABLE 1. Group segmentation, with balanced A/B samples.

Code	Group	Students	A/B Group Sizes
21	Psychology	7916	1834
28	Business Management	4600	1082
20	Law Studies	4110	935
31	Computer Science	3694	838
23	Social Education	2940	717
Total		23260	5504

1) INTERACTION WITH THE WIDGET

In order to quantify the reach of the experiment, we tracked the number of users in the test groups that, when accessing the campus, saw the widget for the first time (n_{first}). This measure has helped us to identify the test users’ population and to understand the rate at which they were incorporated into the study.

We have also tracked the number of users that click the “Expand” (n_{expand}) and “See More” (n_{see}) buttons, as well as the number of users that click the “More Info” hyperlink and go to the informational webpage (n_{info}).

2) TRACKING ACTIVE SESSIONS AND SESSION DURATION

One of the recommendations explicitly indicated by the widget was to actively close the session, i.e., logout the Campus when the activities were finalized, so back-end resources could be released. In order to assess any possible variation of the navigation patterns of the groups under study, the variables such as the number of simultaneous active

TABLE 2. Metrics collected in the A/B experiment.

Number of users that see the widget for the first time	n_{first}
Number of users that click the "Expand" button	n_{expand}
Number of users that click the "See More" button	n_{see}
Number of users that click the "More Info" hyperlink	n_{info}
Number of simultaneous active sessions	$k_{sessions}$
Accumulated session duration	$t_{sessions}$
Percentage of active disconnections/total (short sessions)	$s_{sessions}$
Average number of times an activity resource is accessed	$a_{activity}$
Average number of times an external resource is accessed	$a_{external}$
Average number of times the resources’ page is viewed	$v_{resource}$
Average number of times a tool (e.g., forum) is viewed	v_{tool}
Average number of times an activity is viewed	$v_{activity}$
Average number of times the teaching plan is viewed	v_{plan}

sessions ($k_{sessions}$) and the accumulated session time during the study ($t_{sessions}$) have been monitored.

Another monitored variable tracks the active disconnections (versus letting the session expire after a timeout). This is a non-default action that requires the user to click the logout button on the top right of the Virtual Campus web page. One can expect that users seeing the widget may proactively disconnect the session after finishing the activity. The variable analyzed for this metric is the percentage of active disconnections versus the total ($s_{sessions}$).

3) DOWNLOAD-RELATED METRICS

The second explicit message on the widget is a recommendation to limit the number of repeated file downloads. One could expect that users aiming to limit the energy expenditure of the Campus would reduce the number of duplicated downloads during the course. To conduct this analysis, we monitored variables related to the number of downloads on a per-user basis.

We then monitored access to bibliographic material by the students, a metric directly related to the volume of downloaded files. First, we measured the number of times a user accesses a resource associated with some activity (e.g. the problem statement for this particular activity). This variable also accounts for the access to any teaching material within the “learning resources” section in the classroom (e.g., the first chapter of the course syllabus). This variable has been called $a_{activity}$. Second, we have measured the number of times a user accesses an external resource linked with an activity (e.g a webpage with materials related to the activity to be conducted). This variable also accounts for the access to any external resource from the “learning resources” section in the classroom (e.g., a web page with examples related to the subject matter). We have named this variable $a_{external}$.

4) OTHER VARIABLES

We also monitored variables that measure the user activity but are not directly related to our explicit indications. Yet, we have analyzed them in order to observe any possible behavioral change. For example, one may think that a user conscious of the energetic impact of ICTs may minimize its activity even if there isn’t a particular message with

specific actions. Then, these variables should help us to discover other behavioral patterns during the pilot. Most importantly, we are particularly concerned about negative effects, i.e., that the widget would have resulted in disincentivizing the access to resources indeed needed for the study.

The first variable considered counts the average number of times a user accesses the “classroom resources” page ($v_{resource}$). We could think that users seeing the widget information may be more cautious when accessing the resources linked to the classroom. We also measured the average number of accesses to additional tools available in a “classroom” such as the forum and chats (v_{tool}). The Campus presents the activities in a timeline so students can have a better understanding of the course planning and timing. We also monitored the average number of times an activity is viewed ($v_{activity}$). This variable captures when users access any activity from this timeline. We also tracked the average number of times the users access the teaching plan of each subject (v_{plan}), a document detailing all the information related to planning and evaluation (the teaching plan can be directly viewed on the web or downloaded). A summary of all metrics collected is shown in Table 2.

A technical description of the tools implemented to conduct the A/B test and capture the metrics listed in this section is presented in Annex I.

IV. RESULTS

In this section we provide an analysis of the obtained results, focusing on aspects such as reach, i.e., how many users effectively saw the widget, impact on the user session patterns, and changes in patterns related to file downloads.

A. REACH

Fig. 3.a (black) shows the evolution of users as they first saw the widget (n_{first}). That is, we take into account the first time they logged-in to the Campus after launching the pilot and therefore see the widget for the first time. Therefore, at a given time this variable indicates the number of users who had viewed the widget at least once. This number stabilizes over time, as the test population is finite. During the 7 first days, 92.3% of the target users saw the widget, reaching around 5400 after a month.

One of the actions that may provide good insights about the users’ interest on the intervention is the interaction with the widget. Recall that it was designed with 3 expansion levels, as presented in Section III. In Fig. 3.a (orange) we can observe the evolution of users in the test group as they expanded the widget for the first time during the experiment (n_{expand}). In particular, we see a considerable amount of expansions during the first days till it somehow stabilizes after 5-6 days, arguably because most of the users had already seen the content. In total 2682 (49.6%) users interacted with the widget during the pilot at least once.

Fig. 3.a (blue) shows the users as they clicked on the “See more” button for the first time (n_{see}), thus arriving to the third level of information. This third level was only reached by

a minority of them, following a similar time pattern as the previous level of expansion.

To get more insight on the evolution of the users’ interest towards the experiment, Fig. 3.b, shows the daily clicks in the “Expand” button for each of the groups under study for the duration of the test. All groups show a clear fading that indicates a reduction of the “novelty effect” as time passes by. Indeed, many users interacted with the widget only once.

B. USER SESSIONS

Fig. 4 presents the number of simultaneous active sessions ($k_{sessions}$) over a 4-week period for each of the groups under study. We can clearly see the day-night patterns and we can also observe two peaks per day. During the weekdays, the main peak was around 18h and the secondary at around 12h. On the weekends the peaks were inverted, having the major incidence during the late morning. Saturdays were the days with the least activity. In overall, the number of simultaneous connections supported by the platform during the pilot was 14680 on average, reaching peaks of 34750 simultaneous sessions.

In the time span covered by Fig. 4, during the first 2 weeks the widget was not shown to any user, while after the 2nd of May 2021, the test group was exposed to the widget every-time they logged in the campus. Despite test users being informed to close the session when finished, we cannot observe any significant variation in the number of active sessions along the day between the control and test groups after the deployment of the pilot.

Furthermore, when looking at the accumulated session time for a period of 28 days after the pilot was launched ($t_{sessions}$) we cannot observe any significant variability for the control and test groups (See Fig. 5). For example, considering the Psychology studies, both test (A) and control (B) group featured a similar accumulated session time, despite the test group had seen a recommendation to close the session when done. For other studies, for example Social Education, we can observe that the test group increased the session time with respect to the control group. This happened in an inverse manner in the Computer Science studies. It must be concluded that the small variations observed are not due to the impact of the widget, but to other non-observable variables associated with the type of studies.

Fig. 6 presents the percentage of sessions for each of the groups under study that actively logged out of the campus before the timeout expiration time ($s_{sessions}$). On the left of the figure we can observe the test (A) and control (B) groups before the widget deployment, while on the right side of the figure we can observe the same test and control groups after the deployment. A first observation is that, on average, the number of sessions that were actively closed is around 9% before and after the widget deployment. In the figure, the dots represent single days (14 days before and 14 after deployment), and the crosses represent the average over each 2-week period. Quantitatively, if we take for example the Psychology studies, we can observe that 8.8% of the group

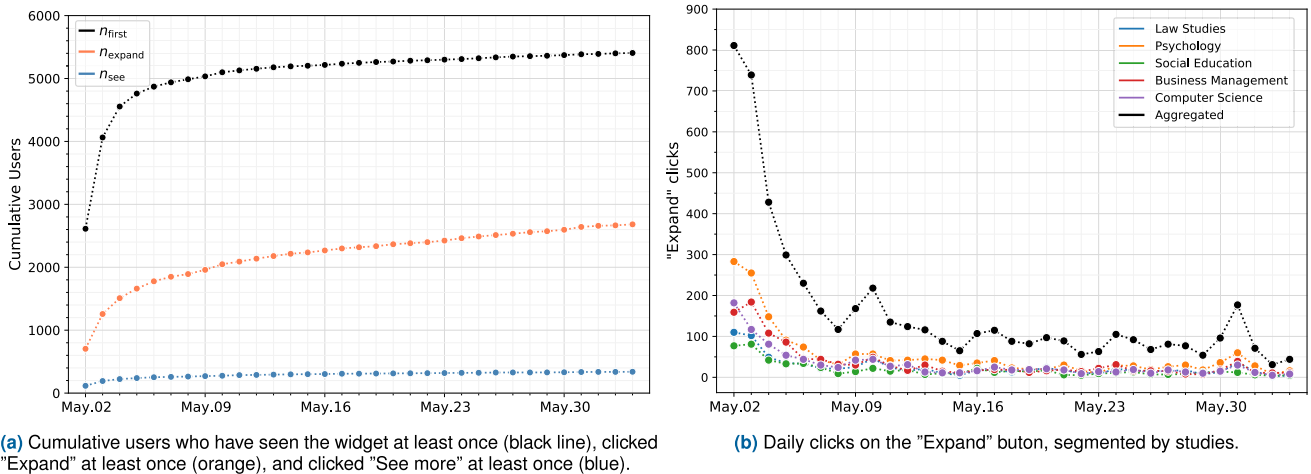


FIGURE 3. Pilot reach and interest raised indicators.

A sessions and 7.9% group B sessions were actively closed before the beginning of the pilot. After having deployed the widget, these numbers stayed quite similar with a 9.0% and 7.8%, respectively. If we look at the Computer Science studies we observe that before the pilot group A and B had an active disconnection ratio of 6.5% and 7.9% respectively. These ratios stayed very close during the execution of the pilot, 6.4% and 7.8% respectively.

The single dots displayed in Figure 6 help us to understand that the observed variability is indeed related to differences between studies. However, for a given study, the results for test and control groups are almost identical before and after the pilot.

C. DOWNLOADS AND TOOLS INTERACTION

Fig. 7.a presents the number of times an activity resource is accessed per user ($a_{activity}$), for the test (A) and control (B) groups segmented and averaged by studies, before and after launching the pilot. This variable is therefore an average of the number of times a file is downloaded by each user. It can be observed a decrease in the average downloads per user-file from pre- to post-deployment in almost all cases. However, we cannot attribute this small reduction to the widget effect, because the same trend is observed in both groups. Arguably, this can be attributed to the course timing, in which users download the material and activities during the early weeks of study.

Fig. 7.b shows the average number of times an external resource is accessed per user ($a_{external}$). From a users' viewpoint this is essentially the same as before, but in this case the file is downloaded from an external resource. Note that the behavior of this variable is almost identical to that observed in the previous case, with a repetition factor (i.e., average number of times a file is downloaded per user) always around 1.7.

Other variables presented in the last four entries in Table 2 include the number of times the resources' page is

visited ($v_{resource}$), access to complementary tools such as forums (v_{tool}), classroom activities views ($v_{activity}$) and accesses to the teaching plan (v_{plan}).

As can be observed in Fig. 8, the test and control groups show very similar patterns both before and after the deployment of the widget for all the monitored variables. The results of groups A and B are almost identical, corroborating the lack of impact of the widget on the user activities within the classroom.

D. RESULTS SUMMARY

We summarize in the following the results presented in this section.

The first noteworthy fact that we have observed is the quick decline in interest. The number of clicks on the widget dropped to almost zero in the first few days. This fading effect had already been observed in the literature. For instance, [28] reported a generalized loss of interest within a few weeks, even after a single interaction with the technology.

Second, we have observed that our request for active disconnections was mostly ignored. We were not able to identify any reduction in the session time of the users that show the widget compared to users in the control group. Timeout is the norm, as it is in the control group.

Finally, we asked users to download the documents only once. The results conclude that the number of downloads per document was essentially the same in the test and control groups, for all the different types of downloads monitored. The differences are most likely due to the change of activities during the semester. Thus, again in this case, we have not achieved any significant impact.

E. DISCUSSION

Our results seem conclusive: there is no significant difference in the behavior of the users who received information during the pilot and those who did not. Although non-price-based behavioral interventions have proven valuable

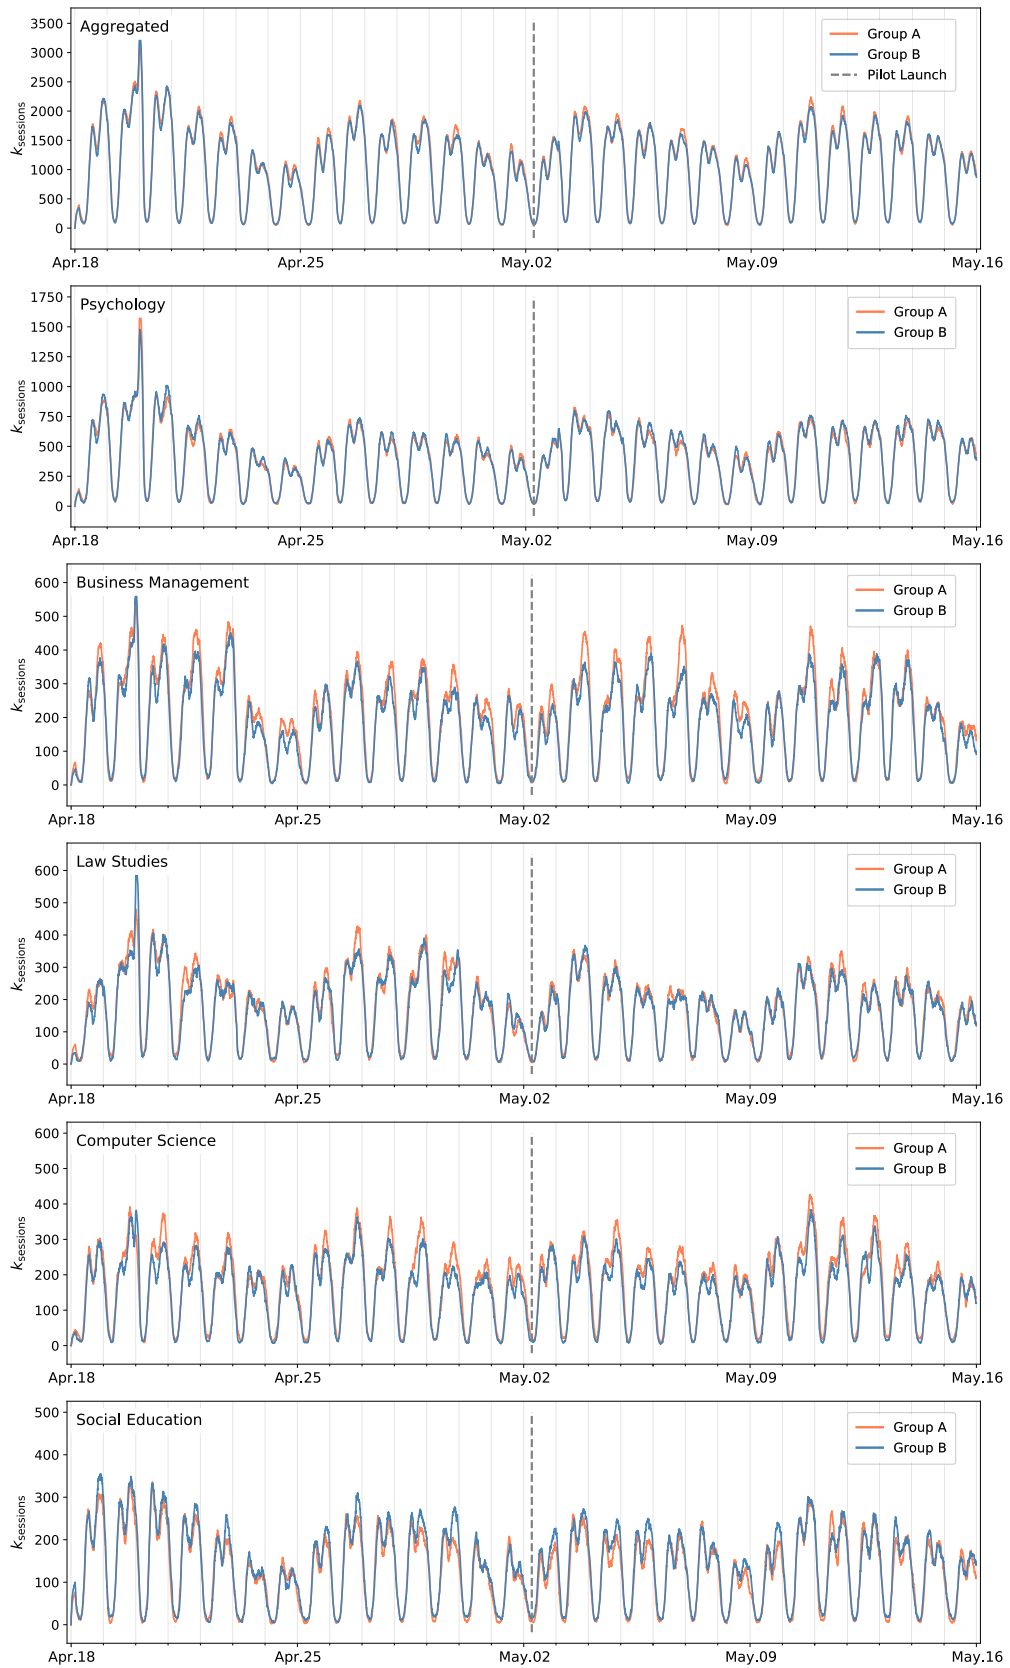


FIGURE 4. Active sessions two weeks before and after deploying the pilot.

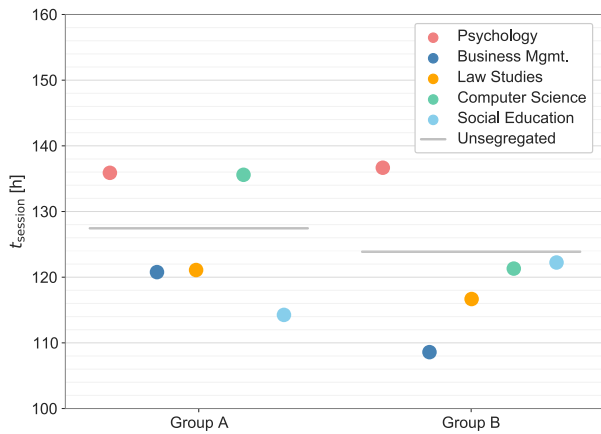


FIGURE 5. Accumulated session time after 28 days.

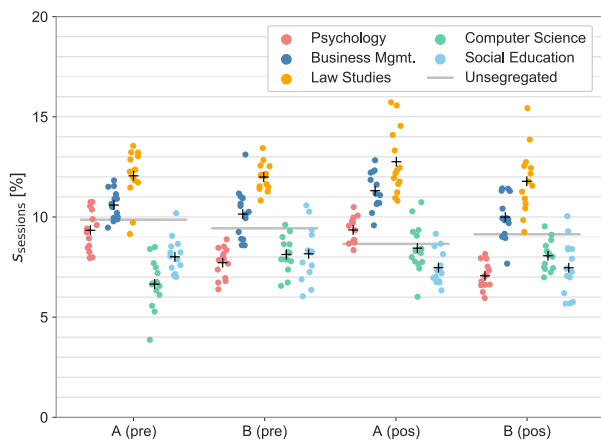
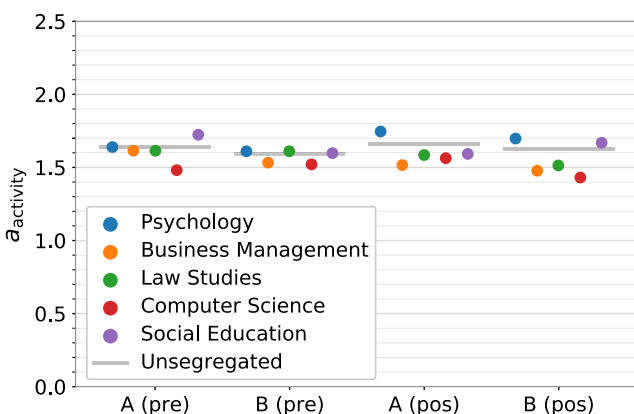


FIGURE 6. Active disconnections (% of total sessions) 2 weeks before and 2 after the widget deployment.

in improving end users efficiency in other fields such as household energy [29], we have not been able to promote similar responses among our platform’s users. There may be several reasons for this.



First, users may not perceive a direct relation between their actions and the overall benefit. There are two aspects that could be involved in this assumption. On the one hand, in other reference interventions, the end user was financially responsible for the optimized element. For example, in the case of household electricity, an energy saving has a direct impact on the bill. On the other hand, it is very difficult to get a sense of the energy consumption of a technology that is not familiar to users —the cloud. As reported in [19] for the electricity use-case, the link between individual use and the resulting impact on global benefit is elusive for most users. In our case, it is even more complex, because the cloud is an abstract entity for most of them.

Second, we must not forget that users are always looking for the shortest path to their goals, and their goal in this case is to fulfill their obligations as students. Somehow, we are asking to exchange individual time for a fuzzy global energy saving. Faced with a selfish or altruistic decision, it is very difficult to opt for the latter. Maybe this placed excessive responsibility on end users [30].

Finally, as has been reported in the literature [31], seemingly harmless variations can have a detrimental impact on the intervention results. The way information is displayed matters. Finding an optimal way for this use case (maximum impact on user awareness, with minimum impact on a platform fruitful usage) requires further study.

Overall, our results should alert us of tackling similar problems through so-called persuasive technologies, those designed to encourage behavioral change [31]. But even in the case that the potential issues described could be alleviated, there is still an open question about the persistence of behavioral interventions, as has been widely discussed in the literature. Indeed, there is a considerable scientific controversy about the long-term effects. First of all, sustained habit change is less studied than the short-term effects [22]. Some works, such as [32] (water conservation), argue that in the long-term the reduction resulting from the interventions eventually dissipated. Other studies have obtained more

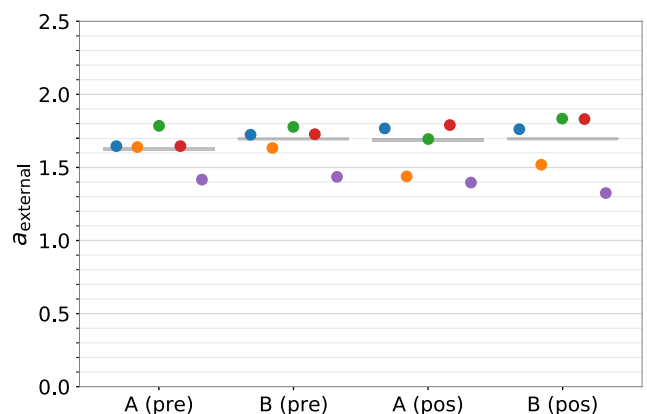


FIGURE 7. Average number of downloads per user and file, segmented by studies.

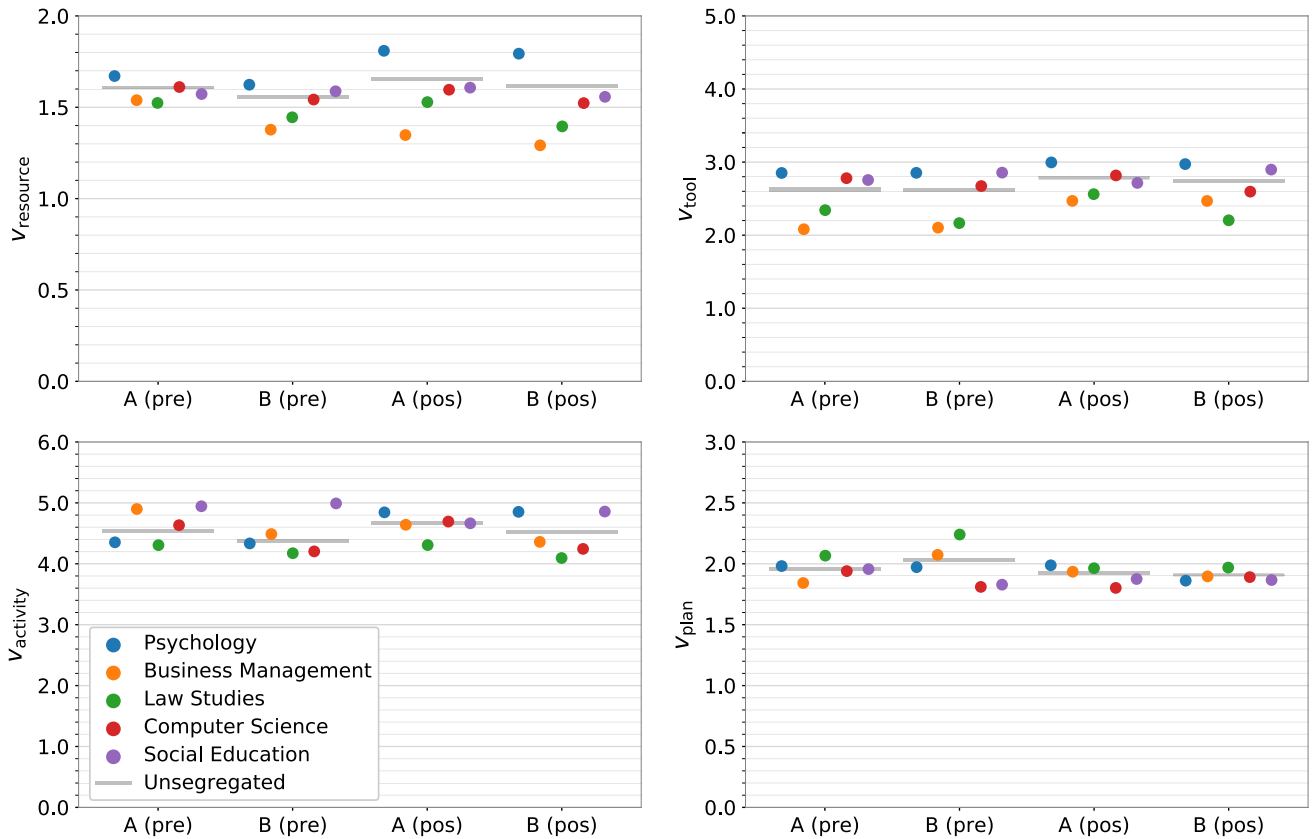


FIGURE 8. Average number of actions per user, for different actions, segmented by studies.

persistent results [33], [34]. But it is questionable whether these findings can be generalized.

In particular, feedback through digital technology has repeatedly demonstrated its ability to promote habit change, but it is not yet clear whether this disruption leads to a lasting change. Interest shown in the feedback quickly fades for the vast majority of users, even after a single interaction with the technology, as previously noted by [28] and confirmed by our results. It should be noted that reducing interest in feedback does not necessarily mean that the initial effect has been lost.

V. CONCLUSION

In this work we have conducted a large-scale behavioral intervention aimed at reducing unnecessary use –and thus the associated energy consumption– of an educational virtual platform. Through a widget embedded in the platform, a Test group of platform users were informed of the overall platform environmental impact, as well as the reasons why this information was shown. They also received very simple instructions on how to reduce their impact without disturbing the intended use of the platform. Test users received this information on a purely informational basis, without any monetary incentive or regulatory action. The widget also collected anonymous metrics about the user’s interaction with

the Virtual Campus, in particular the session time and the volume of data transferred.

Our results reveal that, despite having shown users information about the overall energy consumption of the platform, along with tips on how to reduce it, users did not significantly change their behavior. We have seen that the duration of the sessions remained unaffected, even when one of the tips to reduce consumption was to log out manually when finished instead of letting the session open to timeout. We also observed that the download habits were not significantly modified, while one of the tips was to avoid downloading the same content multiple times. We have discussed different potential causes for this lack of interest in our intervention. Among these, we believe, is the difficulty of understanding the impact of their individual actions towards the global goal of reducing energy waste of the Cloud, a technology that may seem abstract to end users and for which they are not financially responsible.

Fortunately, there is a way forward in reducing energy consumption of the cloud from the platform side. Namely, to provide end users with technical aids. Technology today offers appealing means, which are relatively easy to apply in our case. On the one hand, SaaS models with auto-scaling allow the use of only the hardware resources needed at any given time, freeing the rest for other applications.

In an orchestrated manner, this allows a datacenter to optimize hardware resources while maintaining the same quality of service. On the other hand, the platform can easily detect and interpret users' activity, performing automatic actions without requiring their explicit attention. The platform's intelligence will therefore optimize the resources requested from the SaaS. Moreover, users can be informed of what is being done for them, potentially increasing the level of satisfaction with the platform. This is, in our opinion, a much more promising approach that does not require changes in the way users interact with the platform, nor places the responsibility on them.

ANNEX I. BRIEF TECHNICAL DESCRIPTION

To carry out this experiment, we have developed a specific tool set, which has been designed to be easily ported to other platforms. The tool has two main functions. First, it monitors the physical resources hosted in the cloud as well as on-premise servers. This information is used to estimate the global power consumption of the platform in real-time, which is shown to the user through a widget. Second, it monitors users' interactions with the platform. The gathered information is used in a post-process stage, aimed to identify the potential impact of the information displayed.

The tool is composed of different modules that run both at the cloud and at the user web browser:

At the cloud side, a power analyzer tool is designed to periodically collect energy traces from any virtualized instance, container or serverless function, identifying the most relevant runtime parameters in each case. The tool aggregates this information and estimates the total real-time power consumption associated with the platform using a proprietary model [35]. The result is then exposed through an API, so any front-end instance at the user-side has access to this information.

On the terminal side, when a user visits the site an agent is injected in the navigator as part of the platform. This agent periodically queries the power analyzer tool to show the real-time power consumption computed by the cloud agent. The query period is configurable in order to find a trade-off between the user's perception of fresh information and the bandwidth supported by the API (or alternatively the cost associated with this bandwidth). In addition, the local agent collects metrics about the user's session, in particular the session time and the volume of data transferred. It also captures interaction events, such as clicks on different widget options. The information collected is then sent back to a specific cloud back-end and stored for later analysis. The tool preserves users' anonymity and privacy at all times.

Technically speaking, the agent has been implemented in plain JavaScript and injected into the UOC front-end using the Google Tag Manager. This enabled a non-intrusive deployment of the widget without requiring any modification or affecting the operation of the virtual campus. At the back-end, different tools have been used to monitor the UOC infrastructures. The Amazon Web Services hosted services in the

form of lambdas, containers and entire instances, have been monitored using the CloudWatch API from AWS. The on-premises infrastructure (private cloud) required the connection through SNMP to the different racks hosting the servers supporting the UOC campus. Via SNMP we could access the power management system in the rack structure and obtain fine-grained traces of the energy required by the system.

ACKNOWLEDGMENT

The authors would like to thank the UOC Technology Architecture and Security Department for their unconditional support to carry out this intervention. They would like to explicitly express their gratitude to the MWC Management Team who had a fundamental role in this project. Any opinions, findings and conclusions herein are those of the authors and do not necessarily represent those of UOC.

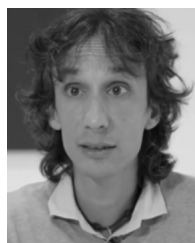
REFERENCES

- [1] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, Feb. 2020.
- [2] International Energy Agency. (2019). *World Energy Outlook 2019*. Accessed: Nov. 2021. [Online]. Available: <https://www.iea.org/reports/world-energy-outlook-2019>
- [3] Red Eléctrica de España. (Jul. 2021). *The Spanish Electricity System 2020*. Accessed: Nov. 2021. [Online]. Available: <https://www.ree.es/en/datos/publications/annual-system-report/spanish-electricity-system-2020-report>
- [4] L. Einav and J. Levin, "The data revolution and economic analysis," *Innov. Policy Economy*, vol. 14, pp. 1–24, Jan. 2014.
- [5] BloombergNEF. (Jan. 2020). *IH 2020 Corporate Energy Market Outlook*. Accessed: Aug. 2020. [Online]. Available: <https://www.bnef.com/core/shorts/5665>
- [6] Google. (Dec. 2021). *Google Environmental Report 2021*. Accessed: Dec. 2021. [Online]. Available: <https://www.gstatic.com/gumdrop/sustainability/google-2021-environmental-report.pdf>
- [7] European Environment Agency. (Feb. 2021). *Policy Instruments*. Accessed: Dec. 2021. [Online]. Available: <https://www.eea.europa.eu/themes/policy/intro>
- [8] S. H. A. Koop, A. J. Van Dorssen, and S. Brouwer, "Enhancing domestic water conservation behaviour: A review of empirical studies on influencing tactics," *J. Environ. Manage.*, vol. 247, pp. 867–876, Oct. 2019.
- [9] D. Kahneman, "Maps of bounded rationality: Psychology for behavioral economics," *Amer. Econ. Rev.*, vol. 93, no. 5, pp. 1449–1475, Nov. 2003.
- [10] R. H. Thaler, C. R. Sunstein, and J. P. Balz. (Apr. 2010). *Choice Architecture*. Accessed: Dec. 2021. [Online]. Available: <https://ssrn.com/abstract=1583509>
- [11] W. Schultz, S. Javey, and A. Sorokina, "Social comparison as a tool to promote residential water conservation," *Frontiers Water*, vol. 1, p. 2, Jul. 2019.
- [12] K. S. Wolske, K. T. Gillingham, and P. W. Schultz, "Peer influence on household energy behaviours," *Nature Energy*, vol. 5, no. 3, pp. 202–212, Mar. 2020.
- [13] C. Kormos, R. Gifford, and E. Brown, "The influence of descriptive social norm information on sustainable transportation behavior: A field experiment," *Environ. Behav.*, vol. 47, no. 5, pp. 479–501, Jun. 2015.
- [14] M. Barth, P. Jugert, and I. Fritsche, "Still underdetected—Social norms and collective efficacy predict the acceptance of electric vehicles in Germany," *Transp. Res. F, Traffic Psychol. Behav.*, vol. 37, pp. 64–77, Feb. 2016.
- [15] B. C. Madrian and D. F. Shea, "The power of suggestion: Inertia in 401(k) participation and savings behavior," *Quart. J. Econ.*, vol. 116, no. 4, pp. 1149–1187, Nov. 2001.
- [16] F. Ebeling and S. Lotz, "Domestic uptake of green energy promoted by opt-out tariffs," *Nature Climate Change*, vol. 5, no. 9, pp. 868–871, Sep. 2015.
- [17] M. Kaiser, M. Bernauer, C. R. Sunstein, and L. A. Reisch, "The power of green defaults: The impact of regional variation of opt-out tariffs on green energy demand in Germany," *Ecol. Econ.*, vol. 174, Aug. 2020, Art. no. 106685.

- [18] U. Liebe, J. Gewinner, and A. Diekmann, "Large and persistent effects of green energy defaults in the household and business sectors," *Nature Hum. Behav.*, vol. 5, no. 5, pp. 576–585, May 2021.
- [19] O. I. Asensio and M. A. Delmas, "Nonprice incentives and energy conservation," *Proc. Nat. Acad. Sci. USA*, vol. 112, no. 6, pp. E510–E515, Feb. 2015.
- [20] J. R. Hopper and J. M. Nielsen, "Recycling as altruistic behavior: Normative and behavioral strategies to expand participation in a community recycling program," *Environ. Behav.*, vol. 23, no. 2, pp. 195–220, 1991.
- [21] H. W. Warner and S. Forward, "The effectiveness of road safety interventions using three different messages: Emotional, factual or a combination of both messages," *Transp. Res. F, Traffic Psychol. Behav.*, vol. 36, pp. 25–34, Jan. 2016.
- [22] S. Hermsen, J. Frost, R. J. Renes, and P. Kerkhof, "Using feedback through digital technology to disrupt and change habitual behavior: A critical review of current literature," *Comput. Hum. Behav.*, vol. 57, pp. 61–74, Apr. 2016.
- [23] K. Jessoe and D. Rapson, "Knowledge is (less) power: Experimental evidence from residential energy use," *Amer. Econ. Rev.*, vol. 104, no. 4, pp. 1417–1438, 2014.
- [24] K. Hartley, N. S. W. Lim, and C. Tortajada, "Policy note: Digital feedback-based interventions for water conservation," *Water Econ. Policy*, vol. 7, no. 1, Jan. 2021, Art. no. 2071004.
- [25] V. Tiefenbeck, A. Wörner, S. Schöb, E. Fleisch, and T. Staake, "Real-time feedback promotes energy conservation in the absence of volunteer selection bias and monetary incentives," *Nature Energy*, vol. 4, no. 1, pp. 35–41, Jan. 2019.
- [26] R. Kohavi, R. Longbotham, D. Sommerfield, and R. M. Henne, "Controlled experiments on the web: Survey and practical guide," *Data Mining Knowl. Discovery*, vol. 18, no. 1, pp. 140–181, Feb. 2009.
- [27] C. A. Gomez-Urbe and N. Hunt, "The Netflix recommender system: Algorithms, business value, and innovation," *ACM Trans. Manage. Inf. Syst.*, vol. 6, no. 4, pp. 1–19, Jan. 2016.
- [28] S. Gözl, "Does feedback usage lead to electricity savings? Analysis of goals for usage, feedback seeking, and consumption behavior," *Energy Efficiency*, vol. 10, no. 6, pp. 1453–1473, Dec. 2017.
- [29] H. Allcott and S. Mullainathan, "Behavior and energy policy," *Science*, vol. 327, no. 5970, pp. 1204–1205, 2010.
- [30] C. Gonzalez-Arcos, A. M. Joubert, D. Scaraboto, R. Guesalaga, and J. Sandberg, "'How do I carry all this now?' understanding consumer resistance to sustainability interventions," *J. Marketing*, vol. 85, no. 3, pp. 44–61, May 2021.
- [31] L. Ableitner, V. Tiefenbeck, S. Hosseini, S. Schöb, G. Fridgen, and T. Staake, "Real-world impact of information systems: The effect of seemingly small design choices," in *Proc. Workshop Inf. Technol. Syst. (WITS)*, Seoul, South Korea, 2017, pp. 387–402.
- [32] K. S. Fielding, A. Spinks, S. Russell, R. McCrea, R. Stewart, and J. Gardner, "An experimental test of voluntary strategies to promote urban water demand management," *J. Environ. Manage.*, vol. 114, pp. 343–351, Jan. 2013.
- [33] H. Allcott and T. Rogers, "The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation," *Amer. Econ. Rev.*, vol. 104, no. 10, pp. 37–3003, Oct. 2014.
- [34] V. Tiefenbeck, V. Tasic, S. Schöb, and T. Staake, "Long-lasting effects or short-term spark? On the persistence of behaviour change induced by real-time feedback on resource consumption," in *Proc. ECIS*, no. 84, Jun. 2016, pp. 1–17.
- [35] B. Martinez and X. Vilajosana, "Exploiting the solar energy surplus for edge computing," *IEEE Trans. Sustain. Comput.*, vol. 7, no. 1, pp. 135–143, Jan. 2022.



BORJA MARTINEZ (Senior Member, IEEE) received the B.Sc. degree in physics and the Ph.D. degree in informatics from the Universidad Autónoma de Barcelona (UAB), Spain. From 2005 to 2015, he was an Assistant Professor with UAB. He is currently a Research Fellow at the IN3, UOC. His research interests include low-power wireless technologies and energy management policies.



XAVIER VILAJOSANA (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in computer science from the Universitat Politècnica de Catalunya, and the Ph.D. degree in computer science from the Universitat Oberta de Catalunya (UOC). He has been a Researcher at Orange Labs, UC Berkeley, and HP Labs. He is currently a Professor at UOC. His research interests include wireless communications and standardization.

...