Università degli Studi dell'Aquila

Vrije Universiteit Amsterdam

Universiteit van Amsterdam

Master Thesis

Impact of Users' Behavior on Digital Service Energy Consumption.

Author: Bassey, Nsidibe Onoyom (2803716)

Supervisor: Prof. Benny Akesson
Supervisor: Prof. Ana Lucia Varbanescu
2nd reader: Dr. Daniele Bonetta

A thesis submitted in fulfillment of the requirements for the joint VU-UvA and Univaq Master of Science degree in Computer Science.

August 28, 2024

"Motivation is not enough, discipline, grit, perseverance, and consistency are needed. Always show up."

from unknown,

Abstract

Energy consumption in the realm of ICT is raising concerns, challenging the prospect of achieving net-zero emissions for the field. ICT solutions are often seen as efficient and low-cost, yet their impacts leave notable footprints behind. The high demand for digital services, such as online shopping, is contributing to the increased energy consumption of digital infrastructure. Energy consumption in the digital domain is mostly attributed to the hardware and software capabilities and that of the infrastructure itself, while users and how their behavior influences this consumption are highly overlooked. However, as users are those interacting with and driving these services, it is users' behavior and usage patterns that directly translate to the energy used by both the software and the hardware components.

In this context, we research the impact of user behavior on the energy consumption of digital services in the computing continuum, as effective and efficient technology alone might not provide sufficient change in achieving the desired net-zero emissions. To explore the impact of user behavior on the energy consumption of digital services, we use online shopping as a case study. The objective is to understand how diverse user interactions within the online shopping domain influence client-side and server-side energy usage. To this end, we deployed a mixed research methodology, combining theoretical and empirical analysis, and found significant differences in the energy impact of different users' actions on the server-side. However, isolating the specific impact of user behavior is difficult to implement on the client-side. Overall, our results show that product browsing is the user action with the highest impact on energy.

Based on the insights from our research, we also propose a model that can estimate the energy impact due to users' behavioral patterns on the server-side. As such, our research contributes to achieving a sustainable digital ecosystem for digital services, which is a collective responsibility where users are at the center of the technology.

Acknowledgements

My special appreciation goes to my supervisors, Prof. Benny Akesson, and Prof. Ana Lucia Varbanescu for their guidance, dedication, and support throughout this research. They guided me through technical discussions and every uncertainty encountered throughout this research. I appreciate all your efforts from device provision to intellectual discussions. I express my deep appreciation for all the weekly meetings, tasks, suggestions, discussions, and follow-ups that put me on my toes.

To my family, for having so much faith in me even when mine was failing. Thank you. To my friends, I appreciate your support, thank you for believing in me.

Contents

List of Figures

List of Tables

LIST OF TABLES

1

Introduction

Our society is increasingly digital, marked by inter-connectivity and technological advancement. Its sustainability, however, is a concern. We expect the society will require users to be mindful about their digital services consumption, the increase in digital-services demand leads to an unsustainable rise in the energy consumption within the computing continuum, i.e. compute resources that cut across the cloud data center, edge computing systems, and end devices (3). Given the complexity of this digital age, where all our actions carry environmental implications, it becomes necessary to address the reality of our collective impact. But understanding this collective impact requires transparency, to effectively inform consumers about the implications of their actions on the energy consumption of digital infrastructure.

To achieve the much-talked-about net zero CO_2 emission (17), (46), service users might even need to review their habits. Since the final decisions on which service to consume and how to interact lie with the users, their behaviors and decisions are likely to determine the industry's impact.

1.1 Problem Statement

The surge in digital services has led to extensive research into their environmental consequences, particularly concerning the energy consumption of electronic communication networks, devices, data centers, and ICT, as highlighted in the work of Preist et al 2014 (34), Li et al 2018 (28), Guegan et al 2019 (21), Ahvar et al 2019 (2), Ramboll et al 2023 (18) and Kamiya et al 2024, (27). Acrep 2020 (5) emphasized the importance of measuring energy consumption accurately, as accurate measurement serves as the initial step towards reduction. In 2024, International Energy Agency (IEA) (26) investigated the magnitude of energy consumption associated with digital technology, and reported an estimated rise in global electricity to 3.4% expected annual growth through 2026, and the global data center is estimated to reach 1000 TWh by 2026 (equivalent to Japan's yearly energy consumption). This accelerated rise in energy consumption can be attributed to the proliferation of digital services, and other data-intensive applications. These digital services rely on energy-intensive infrastructure to support their operations. Within this context, online shopping may emerge as a key driver of energy demand, since its operations span all the components of the computing ecosystem.

Despite the ubiquity of digital services and their profound impact on daily life, a gap exists in understanding how users behavior influences energy demand in the computing continuum. Different digital services have emerged over the years, with an estimation of at least 9 connected devices per person by 2025, as indicated by Safaei et al. (38). Each time these digital services are accessed, a series of processes and infrastructure components are triggered behind the scenes, from the data-center servers to the end-user devices. Each step of the process consumes energy. As the number of digital services increases, so does the number of users accessing these services: both these factors contribute to a visible increase in the overall energy consumption.

As users navigate this digital landscape, their behaviors play a vital role in shaping the trajectory of energy consumption and environmental impacts.

The COVID-19 pandemic catalyzed a shift in consumer behaviors, with online shopping experiencing unprecedented growth during the lockdown and social distancing era, compelling individuals to embrace e-commerce platforms Arcep, 2021 (1). While this digital migration offers convenience and accessibility, it also amplifies concerns regarding energy consumption and carbon emissions associated with digital transactions.

In the landscape of online shopping, where convenience converges with consumerism, sustainability implications loom. Research by McKinsey, 2021 (29)highlights the exponential growth of e-commerce, with global online sales soaring to $4.28 trillion in 2020, representing a 28% increase from the previous year, while "Statistica" 2024 (41) estimated global retail e-commerce sales to exceed 6.3 trillion U.S. dollars, with a projected 39% increase within the coming years. The surge in digital transactions underscores the profound influence of user behaviors on energy consumption and environmental footprint.

Collective action and systemic change are needed to realize the full potential of digital innovations in fostering sustainability. The United Nations Sustainable Development Solutions Network emphasizes the need for cross-sector collaboration and policy interventions to motivate sustainable practices across the digital ecosystem UN SDSN, 2022 (37).

While, European Commission, 2023 (16) proposed digital literacy to empower users in making informed choices and advocating for environmentally responsible solutions.

1.2 Objectives of the Study

Motivated by the need to address the sustainability of the ICT sector, this study aims to determine how user behaviors and interactions affect energy consumption within the digital infrastructure in the context of digital services. Using online shopping as an example of a digital service, deployed within the computing continuum, we aim to analyze the relationship between users' behavior and energy impact, identify opportunities for optimizing energy use, reducing carbon emissions, and fostering sustainable digital practices. We intend to achieve these objectives through theoretical and empirical research, and data-driven analysis.

Specifically, to understand the relationship between online shopping user behavior and the energy consumption of online shopping, we formulate the following research questions:

RQ1: What are the existing methodologies for measuring the end-to-end energy consumption of digital services across the computing continuum?

To answer this question, we first identify feasible ways to measure digital services energy consumption across the continuum. Furthermore, we analyse and briefly compare the most relevant options, and select the methods/tools to be used further in this project.

RQ2: Does the energy consumption of digital services vary across the different components of the computing continuum?

To answer this question, we propose an empirical study to assess which layer of the continuum has the highest impact on energy within the context of online shopping. Identifying the layers with the highest impact will help us to direct optimization strategies to reduce this impact.

RQ3: To what extent does the user's behavior impact energy usage in online shopping?

In answering the question, we use our case study to measure the energy impact associated with online shopping behavior. We define "shopping behavior" based on basic shopping operations, such as browsing or adding/removing items to/from the cart. We formulate 3 sub-questions to gain insight into the energy impact of these primitive actions, which allowed us to trace interaction areas with high-energy impact.

- To what extent does browsing product pages impact energy usage?

- To what extent do cart updates impact energy usage?
- To what extent does product filtering impact energy usage?

RQ4: What are the environmental impacts of online shopping?

These research questions will guide us in understanding the environmental impact left behind by online shoppers. The insights from these questions will aid us in educating digital service users of their impacts, and to inform the developers of digital services of possible areas of optimization.

1.3 Significance of the Study

The significance of this research lies in the limited attention given to user behaviors and their implication on the energy consumption of the digital infrastructure. While studies have explored the impact of user behaviors in the context of smart housing such as Maghsoudi et al, 2022, Huckebrink et al, 2023, (25), (30), there remains a notable gap in understanding how user behaviors shape energy demand across the computing continuum.

This research will provide insights into how user behaviors in the online shopping environment influence energy consumption in the computing continuum. The insight gained from this research will educate users, and help developers to optimize the energy utilized by online shopping platforms. We believe that by understanding which of the user's interactions impacts energy more, developers can design or redesign applications to reduce energy associated with such interactions thus fostering digital sustainability.

1.4 Organization of Study

The remainder of this thesis is structured into chapters as follows: In Chapter 1, we set the motion and give context on the importance of this research. In Chapter 2, we reviewed existing literature on energy consumption in the computing continuum, digital services energy consumption, and the influence of user behaviors on energy consumption. In Chapter 3, we give insight into understanding the energy consumption of Online shopping, the experimental design, and the research hypothesis are discussed. In Chapter 4, we discussed the energy impact associated with the client side, the challenges in measuring this impact, and the insights gained from the analysis. In Chapter 5, we discussed the server-side energy impact. In Chapter 6, we provide an overview of the energy impact associated with user behavior on both the client and the server-side, the insights from the research and their implications to stakeholders are discussed. In Chapter 7, we conclude our research by highlighting the findings, contributions, limitations, and propose direction for future work.

2

Related Work

This literature review explores existing research related to energy consumption in the computing continuum, energy consumption of digital services, and the influence of user behaviors on energy consumption, seeking to identify research gaps and propose the contributions of this empirical research to the existing body of knowledge.

2.1 Energy Consumption in the Continuum

Here, we review research papers on energy consumption in the computing continuum relating to the cloud, the edge, and the user's device. Each layer of the computing continuum consumes energy and contributes to the carbon footprint of digital services. Cloud data centers consume vast amounts of energy due to the high demand for fast data processing and storage needs. Existing studies have shown efforts made to improve their energy efficiency, from server utilization optimization, and renewable energy usage to advanced cooling technologies. Research by Beloglazov et al., 2012 (8) extensively examined energy consumption across the cloud layer including potential optimization strategies. Their research identifies that high performance and fulfillment of service level agreements have been the sole aim of data centers. This service level agreement has mostly been fulfilled without considering its impact on energy consumption. While noting the rise in energy cost due to fulfilling the high-performance agreement and the decline in energy availability they proposed a shift in optimizing the data center for high performance to optimizing for energy efficiency. This shift shows the salient need to understand the energy consumption of this infrastructure. Aldossary 2021 (4) identifies the energy overhead associated with resource provisioning in cloud data centers. They propose predictive resource management techniques that leverage dynamic resource allocation to reduce the energy expense by these infrastructures. However, these studies often focus on a specific component of the computing continuum overlooking the end-to-end energy consumption in the continuum. While, Orgerie 2019 (21) research on end-to-end energy consumption of the Internet of Things (IoT) devices, noted that IoT devices do not consume a significant amount of energy on their own, which implies that their impact on the continuum of energy consumption may not be significant. However, their impact on continuum energy consumption lies at the intersection of their pervasiveness and the infrastructures they utilize to function. Similar, to the work done by Orgerie (21) on estimating the end-to-end energy consumption of low-bandwidth IoT applications, our research focuses on the end-to-end energy consumption of digital services with a special focus on the online shopping domain. We aim to analyze how digital services energy consumption impacts the computing continuum in the context of online shopping while accounting for the impact of user behavior on this consumption. Baneshi et al 2024 (7), analyzed per-application energy consumption in a multi-application in edge/cloud layer, proposing iFogSim a simulation technique for fine-grained characterization of application energy, noting that this characterization can help in device selection and placement policies thus reducing latency and improving energy efficiency. They approached energy consumption from the system level, resource allocation, and application mapping perspective. The research underscores the importance of tracking communication energy consumption. While our research focuses on end-user devices and server-side energy, seeking to understand how user behavior impacts energy consumption, their research provides a template for tracking end-to-end energy consumption at various granularity through different resource allocation, and application mapping techniques.

Studies have shown that edge computing reduces latency and energy consumption by processing data closer to the source. Research by Shi et al. 2016 (40) and Li et al. 2018(28) shows the energy savings capabilities of edge computing in terms of reduced data transmission rate. Pérez et al. 2021 (33) and Xu et al. 2017 (47) have researched optimization strategies for energy-efficient edge computing. Notably, research on how user behavior impacts energy usage is missing in all the existing work.

At the user device layer, a medium via which digital services are being accessed is known to consume a significant amount of energy, contributing to the carbon footprint of this ecosystem. Priest et al. 2014 (34), and Roth et al. 2017 (35), in their research on the energy consumption of consumer electronics found that end-user's devices take up a greater portion of household energy use. This discovery highlights the need for energy efficiency behavior and increased user awareness. While analyzing energy consumption in mobile devices, Carroll and Heiser 2010 (11) and Preist et al. 2014 (34), recommend energy-efficient hardware design as a technique for reducing the energy consumption of these devices. Whereas, Cuervo et al. 2010 (12) and Ullah et al. 2019 (43), discussed the use of effective application off-loading techniques as a means of saving energy in end-user's devices. Aside, from optimizing hardware and effective offloading techniques, there is a need to understand how user behavior impacts the energy usage of these devices. This behavioral impact is what we intend to understand in the cause of this research.

2.2 End-to-end Energy Measurement in the Computing Continuum

In this section we review work that attempts end-to-end energy measurement in the computing continuum, as well as identifying challenges faced by this studies.

End-to-end energy measurement tracks energy usage from user devices and the network infrastructure to the servers. This measurement provides an overview of the energy impact across the continuum. However, this is not a simple endeavor as many interconnected devices with different hardware and software specifications interact in the continuum during operation.

Research by Hinton et al.2011 (24) and Baliga et al. 2010 (6) states the importance of considering the entire service delivery chain in energy assessments. However, Wohlin et al., 2006 (45); Tiwari et al. 2020 (42)noted that measuring energy consumption across the computing continuum is challenging due to hardware and software variability. Supporting this notion, Cardoso 2020 noted that (10), quantifying the energy use of digital services like a video game is a tremendous challenge because of the involvement of different heterogeneous platforms and heavy dependence on user behavior. Preist 2014 (34), in their research on Guardian news and media, analyzes the end-to-end energy consumption of this media while deploying the life cycle assessment (LCA) methodology. They identified the three factors associated with digital service energy consumption: servers, network transport, and end-user devices. They noted that deploying LCA methodology to analyze the energy consumption of digital services requires knowing how much energy each component uses. Acknowledging that several components such as the architecture, the target device, and the individual user interaction pattern can complicate this measurement. This research was done in the context of a news delivery network. Similar to their research, we will empirically analyze the end-to-end energy consumption of digital services in the context of online shopping due to the vast amount of these services springing up. Li 2017 (28) explores energy models for edge cloud-based IoT platforms, they deployed the simulation and empirical methodology, and the mathematical model was applied to data streaming analysis from a vehicle camera. The energy data was collated from the device, network, and cloud layers. Although this work is not within the premise of digital services, they showed that arriving at a single metric for end-to-end energy measurement in the computing continuum is challenging due to the heterogeneity of devices at each layer. As such the total energy consumption within the computing continuum can be determined as the sum of all energy consumption at each layer of the continuum. The summation of this energy consumption data at each layer of the continuum will give us an insight into the total consumption across the continuum. Orgerie 2022 (2), offers a comparative analysis of energy consumption across cloud, fog, and edge infrastructures, they introduced a model to assess the energy impact of different cloud edge architectures. While our work does not consider architectural design, the research shows that different architectural patterns have varying degrees of impact on energy consumption.

Existing studies have given valuable insights into energy usage in the continuum, strategies for optimization, and how architectural design can have varying impacts on energy consumption in the continuum. However, existing research does not investigate nor offer insight into the impact of user behaviors on energy consumption in the continuum. Our research aims to fill this gap by specifically analyzing how user behaviors influence energy consumption in the computing continuum, thus providing a detailed understanding of energy-efficient behaviors at the user and developer level that can contribute to minimizing digital service energy use.

2.3 User behavior and Energy Consumption

Existing research by Rusek 2022 (36) and Heikkinen 2012 (23) highlights the importance of user behaviors in designing energy-efficient digital services, noting that to understand how energy is consumed the user behaviors and activities need to be understood. Schein et al. 2013 (39) and Darby 2009 (13), while emphasizing the significant impact of user behaviors on energy consumption in resource-intensive applications like video streaming, noted that educating users and providing real-time feedback can lead to more energy-conscious behaviors. Visscher et al 2022 (30)while exploring how occupant behaviors affect the effectiveness of energy efficiency in retrofitting projects, offer insights into designing user-centric energy efficiency strategies as an option for reducing energy use. Although Dost et al.2017 (15) and Guo et al. 2012 (22), have explored energy consumption in e-commerce, their focus is primarily on data centers and the commercial sector. However,

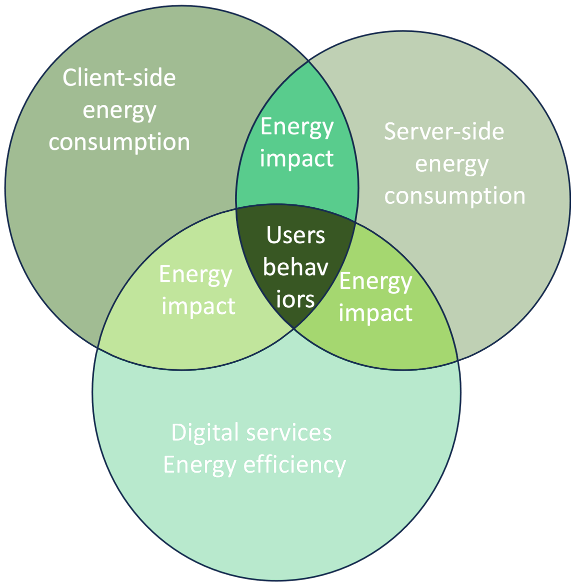

Figure 2.1: Research Venn diagram
The positioning of our research is shown in the diagram, which depicts the interaction between users' behavior and their impact on the energy impact of digital services. This Venn diagram shows that user behavior influences both client and server-side energy use. The interaction shows that user behavior has the potential to contribute to energy efficiency.

we acknowledge the research by Bertsch et al. 2023 (25) which investigates the influence of user behaviors on emissions and costs in residential energy systems, highlighting significant energy savings through simple behavioral changes. This research set the precedence on the potential gain associated with behavioral patterns as it relates to energy consumption. Also, Gram-Hanssen 2012 (20), debates the importance of efficient technologies versus user behaviors in reducing household energy consumption, highlighting the dual focus on technological solutions and user behaviors for optimizing energy efficiency, stating that even with optimized strategies without the users making a conscious effort to save energy the gains may not be visible. This literature underscores the potential for behavioral interventions in enhancing energy efficiency.

While existing studies predominantly focus on video streaming services, proposing techniques for optimizing server-side and device energy usage. Often neglected are the potential implications of user behavior on overall energy use in the continuum.

Despite acknowledging the importance of user behaviors, there is limited research and analysis of specific user behaviors in online shopping and their impact on energy consumption in the computing continuum. The only research that considered user behavior was by Preist 2014 (34), they used a parameterized model for measurement which is not disclosed making it difficult to validate. This study addresses this gap in the existing literature by providing a detailed analysis of user behaviors specific to online shopping, by isolating these behaviors and analyzing their energy impact.

Based on the knowledge gained during the analysis, we will offer insights into how user interactions influence energy usage. We will recommend sustainable behaviors that will enhance the efficient use of energy in the continuum.

2.4 Environmental Impact of Digital Services

The environmental impact of digital services is a significant concern. The Carbon Trust highlights the substantial energy consumption and environmental impact associated with digital activities. This emphasizes the need to optimize energy consumption across the computing continuum to reduce the overall environmental footprint. While Williams 2011 (44), assessed the substantial carbon footprint of ICT activities, Schien et al. 2013 (39) modeled energy consumption variability during the use of online multimedia services. Dost and Maier 2018 (15) identified significant environmental impacts from increased digital transactions in their multi-year assessment of e-commerce. Despite research on energy consumption and user behaviors, there is a critical gap in understanding how specific user behaviors in online shopping influence energy consumption and sustainability outcomes. Existing studies have focused largely on infrastructural solutions, often overlooking the nuanced ways in which user behaviors can affect energy consumption. We aim to fill this gap by investigating the influence of user behaviors on energy consumption in the online shopping domain.

2.5 Online Shopping and Users Behaviors

Lund et al.2021 (29) discuss shifts in consumer behaviors following the COVID-19 pandemic, noting an increased reliance on online shopping and digital services. This trend has significant implications for energy consumption due to greater dependence on data centers, broadband networks as well as end-user device energy consumption. Dost and Maier 2018 (15) conduct a multi-year assessment of e-commerce effects on energy consumption, highlighting significant environmental impacts from increased digital transactions. They emphasize the importance of developing sustainable e-commerce practices to balance the benefits of online shopping with a focus on reducing energy consumption and environmental impact. By examining behavioral patterns and their direct impact on digital energy use, this research will provide insights into how specific shopping behaviors impact energy usage. The insight gained will aid in developing targeted strategies to promote sustainable online shopping practices.

3

Understanding Online shopping Energy Consumption.

In this chapter, we provide detailed insight into the process of measuring the energy consumption due to online shopping, the tools, data collection process, and subject selection procedure are well-detailed herein.

3.1 Experiment Definition

This research aims to analyze how user behavior impacts the energy consumption and carbon footprint of digital services in the computing continuum from the perspective of a user, software developer, and business owner within the context of online shopping services. We used the goal, question, and metrics (GQM) proposed by Rombach 1994 (9), a framework used to establish a well-structured experimental setup. Guided by the GQM model we formulated four research questions to investigate the methodologies for measuring digital services energy consumption across the computing continuum. The impact of user's behavior on digital service energy consumption within the computing continuum in the context of online shopping, and the carbon footprint associated with these behaviors. Based on this framework the tone of this research was set and the research questions formulated.

3.2 Experiment Planning

3.2.1 Energy Monitoring Tools

To manage the energy consumption of digital services, we need to know and understand the consumption pattern. To gain this understanding energy measurement and monitoring

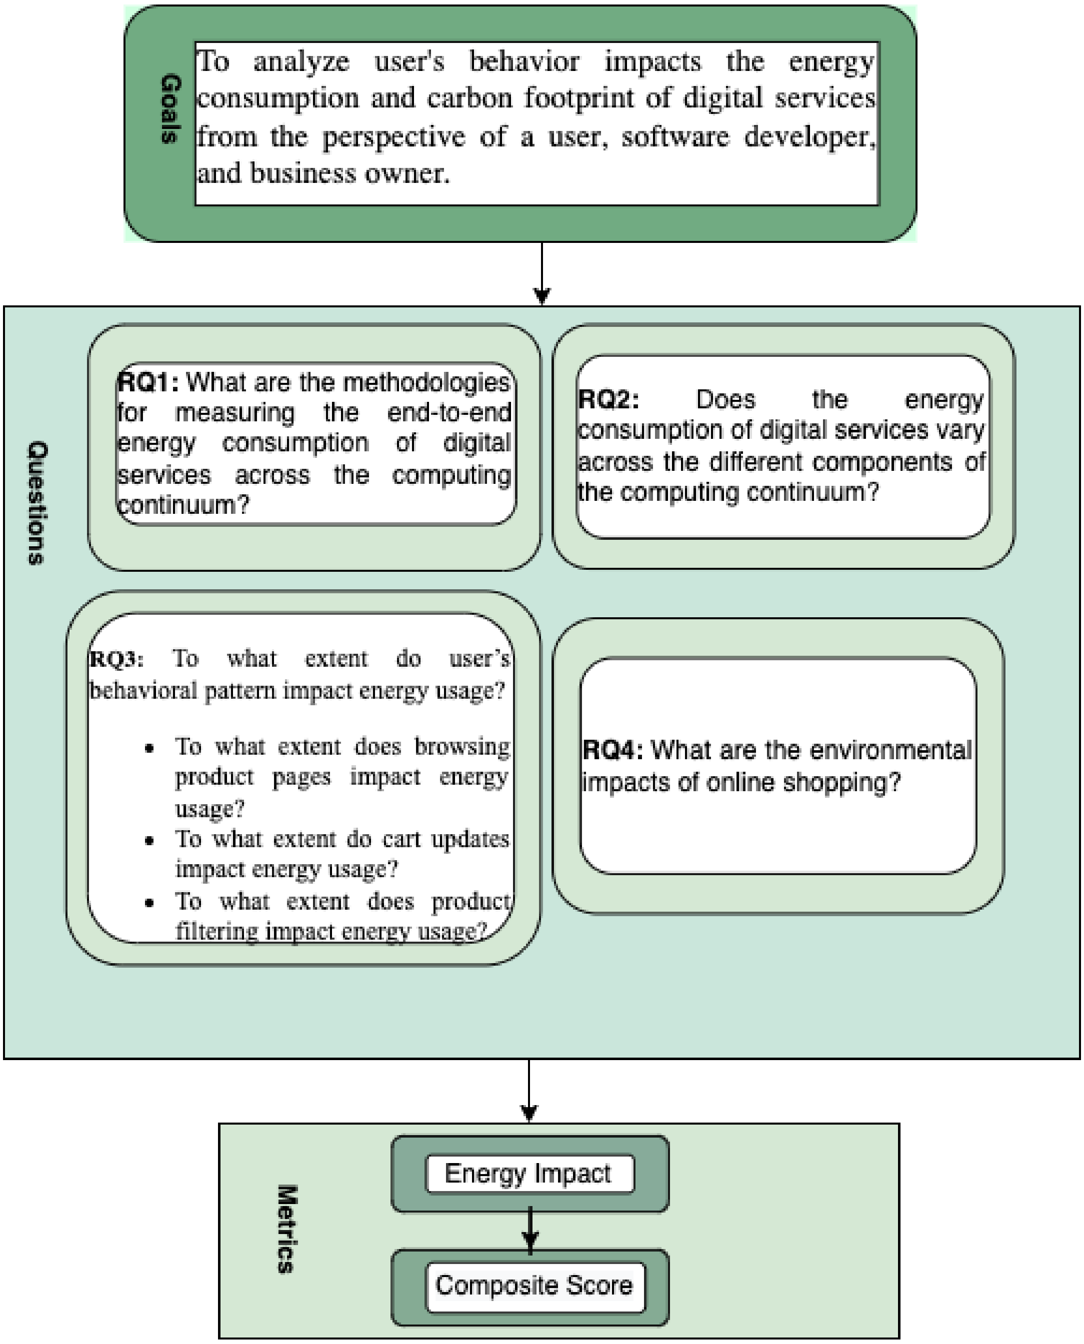

Figure 3.1: Visual description of GQM framework

Table 3.1: The goal, quality metrics framework

Goal, Quality, Metrics framework	
Analyze	the impact of user behavior on online shopping
For the purpose	Evaluating
with respect to	energy consumption and carbon footprint in the online shopping domain
from the perspective	of a user, software developer, and policy maker
in the context of	online shopping in the computing continuum

Table 3.1 describes the GQM framework, within the context of the user's behaviors impact on the computing continuum energy usage.

tools are required. These tools can be characterized by the environment they run on, software or hardware, the granularity of data they report, and the operating system they require to function. They are often limited by hardware and operating system requirements. When it comes to tools for measuring energy impact in the computing domain, there is no *one-size-fits-all* tool. When choosing a tool for measuring and monitoring the energy usage of any application or software system, consideration should be given to the operating system and the level of energy detail needed.

The tools used to monitor and measure the energy consumption of digital infrastructures are the key factors that determine energy consumption. They must provide data that are both accurate and repeatable. Accuracy and repeatability are necessary to validate energy consumption.

Based on our systematic study we highlight the following as the requirements of a good tool for measuring energy consumption in the ICT domain. The tool should ensure the precision of energy usage measurements across various system components, from the CPUs, GPUs, Network, and memory to the peripherals. The equipment should be calibrated correctly, with a detailed document specifying the calibration methods, tolerance level, and sampling frequency. A clear and concise documentation should be available to the users. This ensures that energy consumption data is valid and reliable, allowing for effective verification and understanding of energy consumption patterns. The energy cost associated with such tools should be minimal, as any overhead cost can impact the system's performance. Current energy measurement and monitoring tools are mostly platform-dependent.

A versatile tool that can effectively measure energy consumption data without hardware or operating system compatibility issues will generally be considered a good tool if the energy consumption data are accurate and verifiable. Also, these tools should allow for instrumentation, such that users can tune their parameters such as sampling intervals to suit their needs. *PowerMetrics*, a tool designed for macOS, exemplifies many of these requirements. It offers energy impact measurements at the level of granularity of interest that suits our experiment. Although, the energy impact reported is a composite score without a unit. It gives a detailed insight into how each process and service impacts energy usage.

3.2.2 Subject Selection

To effectively assess the energy impact of online shopping behaviors on the energy consumption of the computing continuum, we chose the Microsoft *"eShopOnContainers"* (32), a microservice reference architecture application. The application allows each component to be deployed as a separate service in the same container. The application represents a real-world use case of a typical online shop, with a wide range of features. The functionalities implemented by this application allow us to measure the primitive energy impact associated with specific shopping activities. The choice of this application was based on the versatility of the features implemented by the application. These features allowed us to uniquely measure the desired behaviors of the users while interacting with the application. The ability to monitor independent service energy consumption helps us to analyze how specific primitive users' shopping behavior impacts energy consumption. Also, the *"eShopOncontainers"* application is designed to handle varying levels of workloads making it easy for us to simulate different user behaviors and load patterns while measuring their impact on energy consumption. We deployed *powermetrics* (31) as the monitoring tool because it allows us to read the energy impact value at the process level. This choice was due to the compatibility with our experimental environment. With *powermetrics* we can confidently see the energy impact associated with the service we interact with in the online shop. Although, *powermetrics* provides performance indices at different granularity, the tool is limited to MacOS and the energy impact data reported is unit-less making it difficult to generalize insight from this measurement.

3.2.3 Experimental Variables

In answering the proposed research questions, we consider the energy impact associated with user behaviors as the dependent variable. The independent variables are the user interaction and time. The energy impact associated with the user's behaviors is measured while interacting with the application. The energy impact data obtained during this interaction serves the objective energy used during the online shopping sessions. This measurement is taken for each primitive shopping task the users perform. The primitive behaviors of interest in this research are cart updates, product browsing, and filtering of product catalogs. For the energy impact associated with server-side consumption, we define energy impact as the energy consumed by each of the processes/services that the user interacts with, while on the client-side energy impact is the function of the browser energy consumption plus the user activity.

3.3 Experimental Hypothesis

For research questions 1 and 4 due to the nature of the question we will theoretically assess and answer these questions as such no hypothesis is formulated for it. However, this question will be analyzed based on the insights of our systematic study.

To answer the research questions 2 and 3, the following hypotheses are formulated:

RQ2:

- H_0: There is no variation in the energy consumption of digital services across the different components of the computing continuum.
- H_1: The energy consumption of digital services varies across the different components of the computing continuum.

RQ3.1:

- H_0: Browsing product pages does not impact the energy consumption of the computing continuum.
- H_1: Browsing product pages has an impact on the energy consumption of the computing continuum.

RQ3.2

- H_0: Updating shopping carts has no impact on the energy consumption of the continuum.
- H_1: Updating the shopping cart has an impact on the energy consumption of the computing continuum.

RQ3.3

- H_0 Product search and filtering do not impact the continuum energy consumption.
- H_1: Product search and filtering have an impact on energy consumption in the continuum.

These research hypotheses guided us to empirically analyze the extent to which user interaction impacts the energy usage of the computing continuum. Through the insight gained from hypothesis testing, we can adequately advise users, developers, and business owners on the energy impact associated with different primitive shopping behaviors.

3.4 Experimental Design

To access the variability in user's behaviors impact on energy consumption in the computing continuum. We chose three primitive actions based on the most frequent tasks performed by online shoppers: product browsing, cart updates, and product filtering. Each of these actions was isolated on the server side and their energy impact was measured.

The execution duration for each experiment was set to five minutes, with a cooling time of 1 minute. The cooling time helps in reducing confound effects due to the thermal heat and environmental state of the device. The experiment is repeated 10 times for each behavior measured. This allows us to gain some level of statistical significance. The experimental subject and use case of this experiment reflect a real-world usage scenario. The experimental combination is outlined in Table 3.2.

3.5 Experiment Setup

The experiment environment was set up on a Macbook Pro, 14-inch 2021. The specification of the experimental devices is shown in Table 3.3. The experiment subject is installed and hosted locally on the experiment device. The experiment subject was cloned from the official GitHub repository. It is worth noting that the *eShopOncontainer* requires docker installation for a seamless run. Interaction with the front end of the application was performed manually. However, this can also be done using GUI interaction tools, we could not explore this option due to limited time. For interaction with the back end of the application, we used a Python script utilizing *Locust.io* (Locust.io) to simulate workload and different numbers of users to measure the energy impact of varying workloads and different numbers of users on the energy consumption of the computing energy consumption.

Table 3.2: Shopping behavior and the combination of users tested

shopping behavior	Number of users	Duration(mins)
Browsing	1	5
Browsing	5	5
Browsing	10	5
Browsing	100	5
Browsing	300	5
Browsing	600	5
Browsing	1000	5
cart updates	1	5
cart updates	5	5
cart updates	10	5
cart updates	100	5
carts updates	300	5
carts updates	600	5
carts updates	1000	5
product filtering	1	5
product filtering	5	5
product filtering	10	5
product filtering	100	5
product filtering	300	5
product filtering	600	5
product filtering	1000	5

Table 3.3: Technical specification of the experimental device

System Specification	
Parameters	**Specification**
Chip	Apple M1 pro
Memory	32GB
Display	14 inch (3024 *1964)
OS	macOS Sonoma 14.5
Flash storage	994,66 GB
No. of cores	10 (8 performance and 2 efficiency

3.6 Experiment Preparation

To assess the feasibility and accuracy of the experiment, we pre-tested the application with different workloads and usage scenarios, as well as on different browser engines. This preliminary test helped us identify whether all behaviors of interest were captured by the application. During this test, we identified the most energy-efficient browser engine for the *"eShopOnContainer"* application, to be Safari. To clarify whether this browser's energy efficiency was application and activity-dependent, we assessed three other sites: the CNN news site, live sports scores, and YouTube, all performing the same task. We discovered that Safari and Google Chrome were more energy-efficient in all cases than Firefox. Based on the results of the preliminary test, we chose the Safari browser as the web engine for interacting with the client-side of the web application. We concluded on this because the client-side energy consumption is a function of the browser's energy usage and interaction performed by the users. Since *"PowerMetrics"* reports energy consumption per process, we did not account for any system processes. However, we ensured that only the terminal and the experimental subjects were active while taking the measurement. This precautionary measure guarantees that the device operates solely within a controlled environment dedicated to the experiment. When considering the server-side energy consumption, we ensured that *"locust.io"*, which simulates the workload, did not contribute to the energy consumption by running it from a separate browser. This allowed for the separation of concerns and ensured that *"locust.io"* did not contribute to the actual energy consumption of the web browser we were interacting with. The entire experimental setup for our experiment is described in Figure 3.2, where all components and their interactions are shown.

3.7 Measuring Digital Service Energy Impact

Reducing the energy impact of digital services entails understanding user's needs, and business demands. Reducing this environmental impact entails measuring the energy consumed by these services. In this section, we detail the strategy to measure the energy consumption associated with user behavior in online shopping.

We deployed a software tool to measure the energy impact of user interaction on online shops. We chose the software tool because of the granularity in which energy measurements are given. Although hardware tools are known to be more accurate than software tools, they provide global energy consumption measurement i.e. the entire system against per process or application. We cannot use hardware measurement when targeting fine-grain

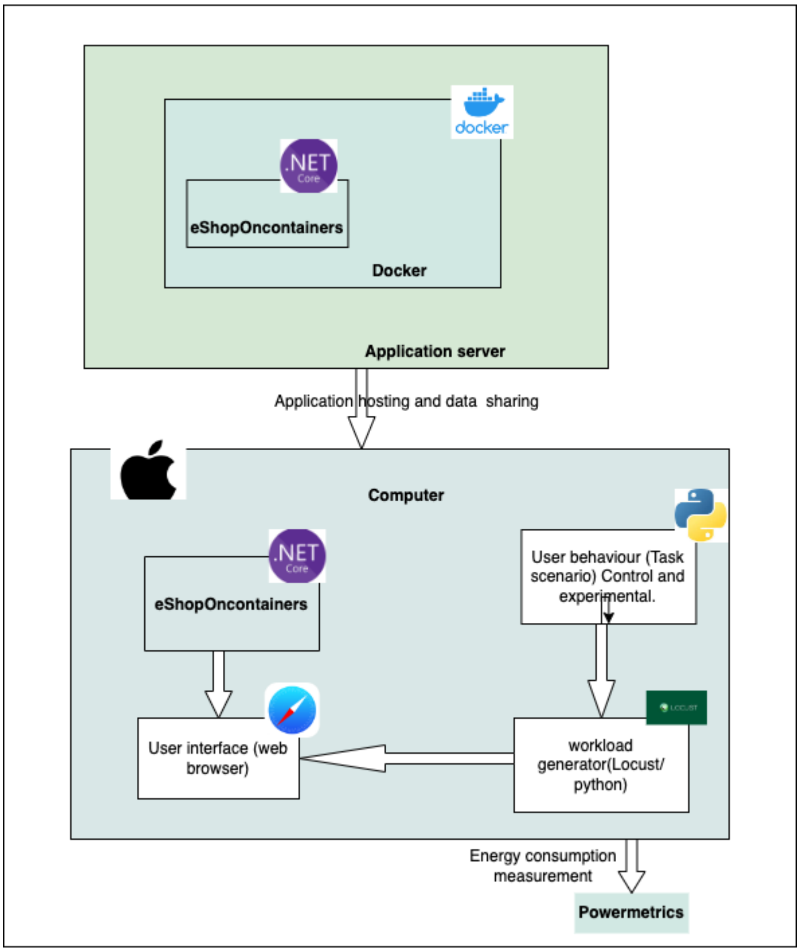

Figure 3.2: Experimental setup

energy consumption as in the case of user behavior impact on energy consumption. Since we are interested in accessing the energy impact associated with specific user interactions. It is important to choose tools that allow see the direct energy impact of these interactions. On the client-side the direct energy impact associated with a process is being accessed not the total energy of the access device (i.e. browser energy per interaction). Measuring the total energy of the access device will not reflect the impacts associated with user interaction. Software tools are known to give the granularity of energy measurement up to process, threads, and methods level, we chose the software tool because of this granularity.

Since software measurement tools are built on empirical estimation the energy consumption may vary across different systems. This implies that the accuracy and steadiness obtained running the same job/ application in the same environment and setup may result in different measures in the energy impact due to factors such as processor thermal effect. As noted by Fahad 2019 (19) in their comparison studies, software tools can report significant differences in energy measurement even with the same setup and environment.

3.8 Data Collection and Analysis

The data for this research was collected using *powermetrics*, while the *"locust.io"* was interacting with the specific service that represents our user behavior of interest. For each user behavior, we repeated the experiment 10 times, this was to gain statistical significance and observe if the readings from the measurements widely vary from each other.

To statistically analysis the data collected during the experiment, we cleaned the data collected. We perform basic descriptive statistics to gain insight into the energy impact of each user behavior. We also perform a preliminary exploratory analysis that shows the trends and patterns in the energy impact of different user behaviors.

We proceed to perform the Kruskal Wallis's test, a non-parametric statistical test, used when the assumption for normality in data distribution fails. It shows the difference in ranks of a dependent variable (Energy impact) across all the independent variables (product browsing, product filtering, and cart updates) that made up our user behavior. Our hypotheses are tested to help us in answering the formulated research questions.

4

Client-Side Energy Impact of User's Behaviors: Online shopping

Client-side energy consumption is the energy consumed by end-users when engaging in online shopping, this energy entails the device's energy, the network device, and every other energy that the user consumes when accessing the application. In the context, of this research, we focus on clients-side energy consumption solely on the user device and the energy expense to access an online shop. The energy used by other components such as the WiFi access and the network devices are not considered here. To understand how user behavior impacts this consumption we analyze various user interactions on the end-user device to gain insights on the contributions of these interactions to the overall energy usage of these services.

During this process, we identify browser efficiency as one of the factors that can affect client-side energy consumption for shopping applications accessed via the web. Previous studies by Saraiva 2020 (14) have shown that browser engines like Safari (version 17.6) and Google Chrome (version 128.0.6613.114 arm64) are more energy efficient than their counterpart like Firefox (version 127.0.2). Based on the previous studies and our preliminary experiment we chose Safari as the browser engine through which our online store will be accessed. Other influencing factors of client-side energy used are embedded in user interactions such as page scrolling, and clicking.

The amount of time plus the level of interaction during online shopping activities directly translates to energy consumption on the client side. Shopping applications that are multimedia rich with high-resolution images cost more energy, however, this was not within the context of our research.

On the client-side our research shows that the more interactive a user gets with the online shopping application the more energy it consumes. Another, influencing factor is the duration of interaction, as energy seems to progress linearly with the duration spent until the resource reaches saturation.

User behavior plays a key role in determining client-side energy consumption during online shopping, as each click and scroll costs energy.

4.1 Challenges in measuring the Client-side Energy Consumption of Online Shopping

Challenges encountered during the client-side measurement stem from the energy efficiency of the browser, operating system variation, and the type of device used by the online shopper to access the platform, each of which constitutes to variation in energy impact values reported.

We also identified that the battery efficiency of the access device and the processor capability influenced the client-side energy usage. To manage this effect, we ensured that the battery life of our experimental device was between 100 - 90 percent throughout our experiment.

Different operating systems manage energy consumption differently. This further complicates the energy consumption process making it difficult to generalize on different platforms. As such the results obtained in this experiment are based on the setup defined earlier in section 3.

4.2 Measurement Procedure

We utilize *powermetrics* a software measurement tool, that uses established power models built on hardware parameters. This tool reports the energy impact of all running processes. It can however be instrumented to measure specific process energy impact. We took a baseline measurement of Safari without the online shop application running. We simulate user interaction: product browsing, filtering products, and cart updates and measure the energy impact value.

On the Client-side isolating this behavior to measure independently was challenging. We took the client-side measurement while randomly performing different tasks for 5 minutes, after each cycle a cooling time of 1 minute was observed, this was repeated 10 times to gain statistical significance.

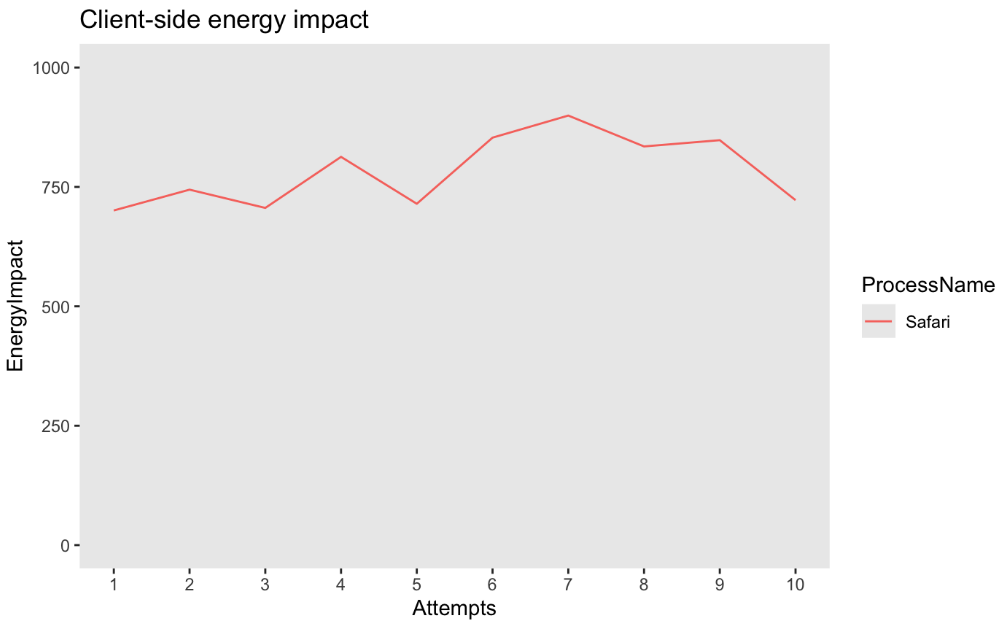

Figure 4.1: Client-side energy impact

Figure 4.1 shows the client-side energy impact of a single user interacting with the online shopping platform for 5 minutes. The variation in the energy impact value from the ten experimental runs may be attributed to the fact that this interaction was done manually. The only explanation we can give from this could be that for every click and scroll different energy impact is associated with them. since this was manually performed, maintaining the same scroll and clicks for each experimental run is difficult.

Also, running the same job and process in the same environment and setup may result in different measurements. Fahad 2019 (19), supports this observation, noting that software tools can report significant differences in energy measurement even with the same setup and environment.

4.3 Different Behaviors Energy Consumption on the Client-side

We could isolate specific energy impacts associated with different user's behavior on the client-side. The difficulty in consistently performing a single task without triggering the other behaviors deprived us of gaining insight into how this behavior impacts energy on the client-side.

However, from our measurement, we discovered that for a typical shopping section by a single user for 5 minutes, the energy impact ranges from about 700 - 900 which is a composite score that shows that the cost associated with the client-side impact is on the high side. This variation is widely attributed to the difficulty in replicating the exact interaction manually for all replicas. Worth noting is that the impact of online shopping on the client-Side is dependent on many factors aside from the influence of the users.

5

Server-side Energy Impact of User Behavior: Online Shopping

To analyze and estimate energy usage associated with online shopping behavior on the server-side. We measured and analyzed how each user shopping behavior performed by online shoppers on the client-side impacts energy usage on the server-side. We conducted different experiments to account for different numbers of users and different behaviors. On the Server-side, isolating different behaviors and measuring the energy impact associated with them was easy against what we noticed on the client-side. This ease came with the utilization of Python script that leveraged *locust.io* to simulate different workloads. This allows for direct interaction with specific behaviors of interest that we measured at any given time. We used *PowerMetrics* to collect the energy impact data associated with the service we are interacting with at that point. However, modeling the server-side energy impact of user behavior at scale may be challenging due to the interconnection of different devices, their design, and the underlying architecture. In the context of this research, most of this complexity was controlled because of the simplicity of our experimental setup. As described in the experimental design 3.2, our experimental setup runs on a single platform.

5.1 Findings and Discussion

During the exploratory data analysis phase, we discovered that different primitive user behaviors have varying degrees of energy impact on the server. This exploration shows that the energy impacts increase linearly with an increase in the number of concurrent users. However, this linearity gets to an equilibrium point after which the energy impact begins to drop. In our experiment, the energy impact process linearly upward up to

300 concurrent users. After, this point the energy impact begins to drop, based on our observation during the experiment this drop in energy impact resulted from the fact that the number of requests processed and delivered dropped. We deduced that the average number of requests processed and delivered has a direct influence on the energy impact of servers. We also observed that the response time increases with the number of concurrent users, but this rise in the response time did not directly translate to a higher energy impact. What we could derive from this, points to the fact that the system was overloaded leading to system resources being saturated.

Based on the queuing theory principle, when a system is overloaded jobs waiting to be processed join the queue leading to a longer time in the CPU, this also explains why the percentage of CPU usage did not directly translate to higher energy impact. This can simply be explained that although the CPU operates at a high capacity the number of successful jobs processed and delivered is less. So, the CPU is less efficient at this point. It might also be that at peak usage the CPU enters context-switching mode instead of actual task performance leading to inefficiencies in the amount of workload successfully processed and delivered. The queuing leads to tasks spending more time in the CPU. This delay and queuing causes some requests to drop thus reducing the actual work performed, as such when requests are no longer processed effectively, energy impact reduces because less computation and resources are being engaged.

From our analysis on the server-side, as shown in Figure 5.2, product browsing had the highest energy impact on the server while shopping. Cart update energy impact is minimal compared to product browsing and filtering for different numbers of users. While the result shows different levels of energy impact on the server due to user behaviors we acknowledge that other parameters such as the frequency of user interaction, and the energy efficiency of the server itself have an impact on server energy usage.

Also, Figure 5.1 shows the energy impact of different numbers of online shoppers, simultaneously performing different shopping behaviors: product browsing, product filtering, and cart updates. As shown in the plots the energy impact increases with the number of users, until resources are saturated.

To test our hypotheses, we performed descriptive statistics to gain insight into the distribution of our data. The distribution is visualized using density plots and qqplots. These plots are known for visualizing the distribution and normality level of experimental data.

The density distribution 5.3 shows that our energy impact data is not normally distributed, as the plot did not show the normal bell-like shape that indicates normality.

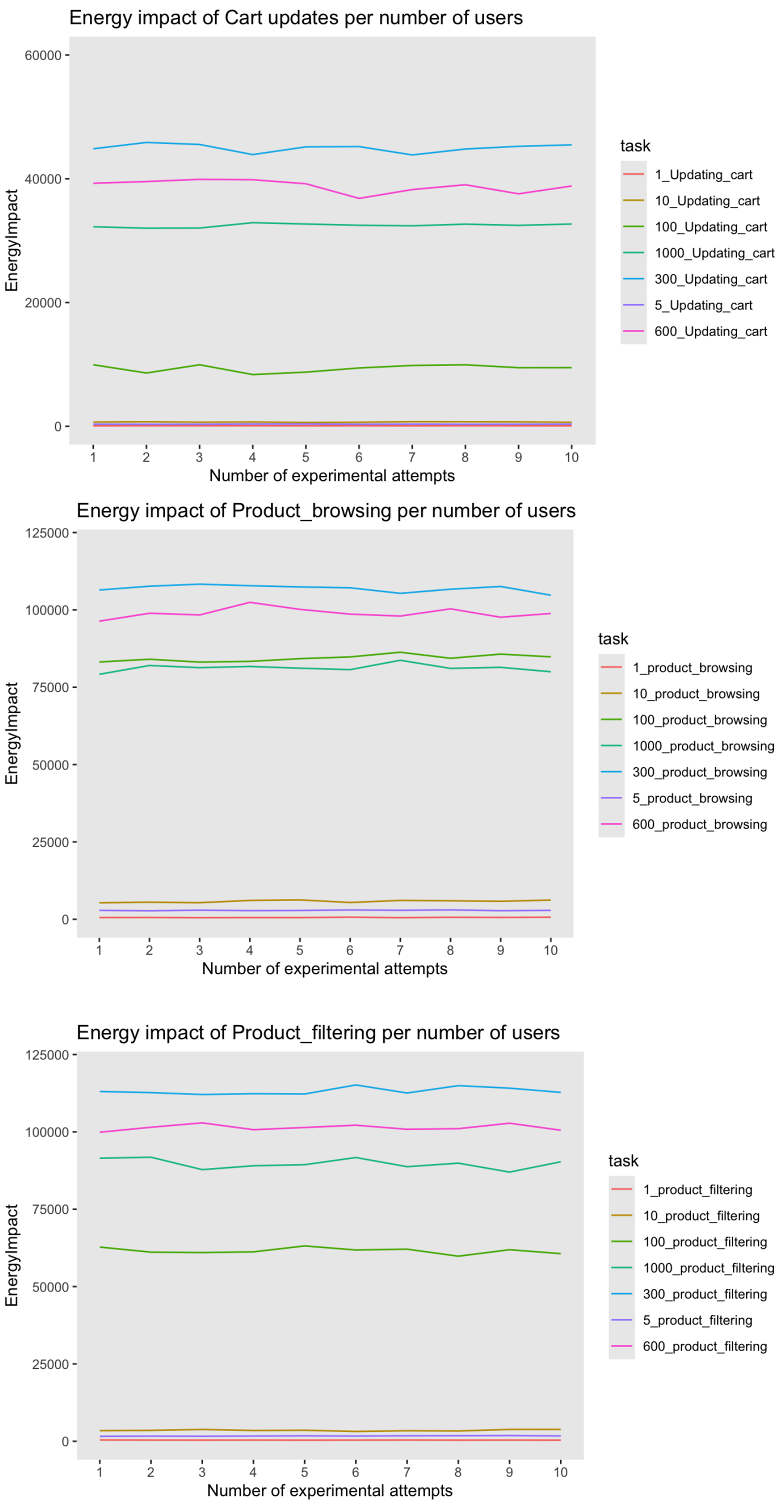

Figure 5.1: Energy impact of different online shopping behaviors across varying numbers of users.

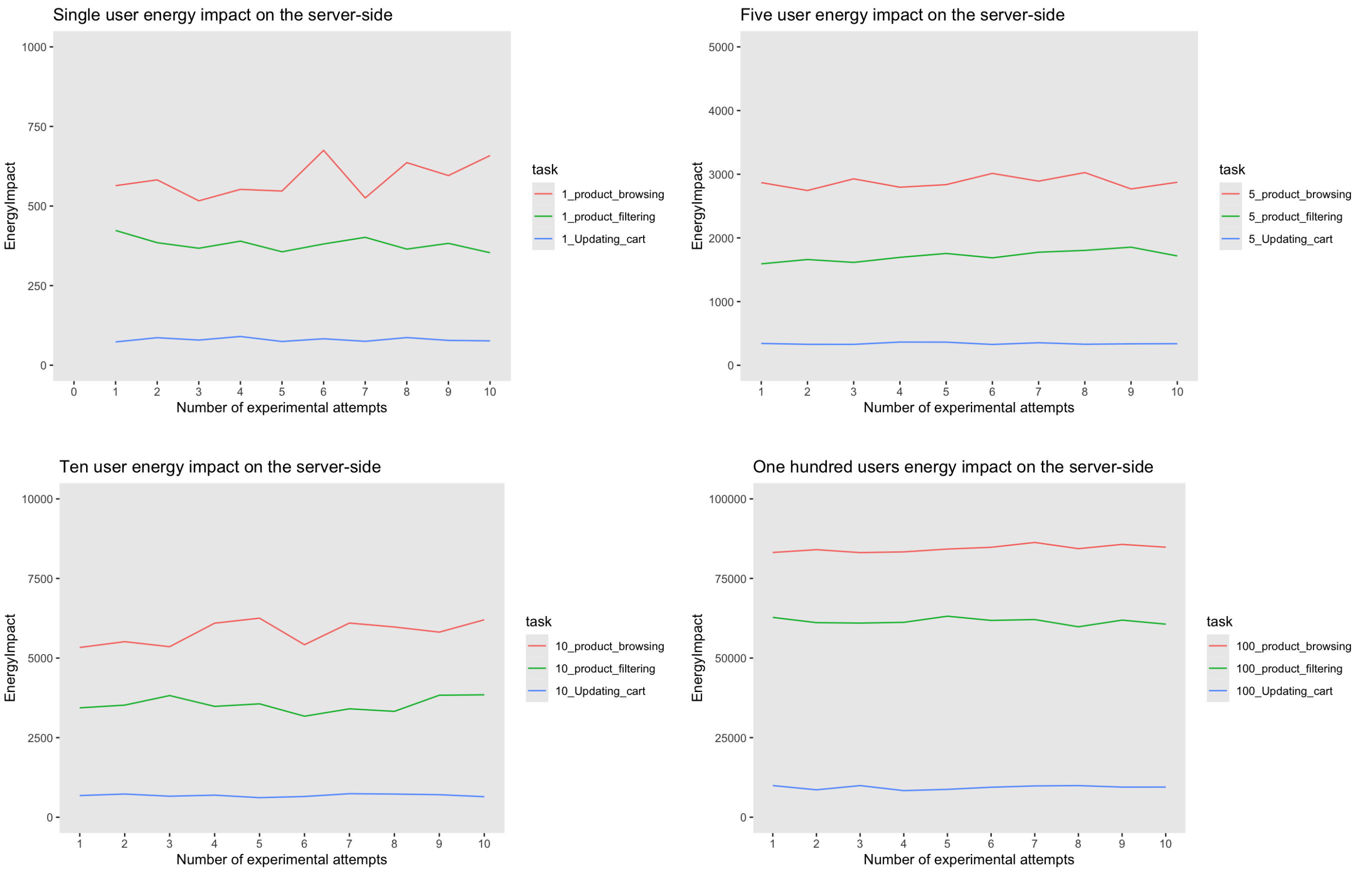

Figure 5.2: Energy impact for different user behaviors across varying numbers of users.

We confirm this distribution using qqplots 5.4 which confirms the non-normality of our distribution as the points widely vary from the source.

Figure 5.3 shows the energy impact data for one hundred users with different shopping behaviors. The shape of the plots varies from the well-known bell-like shape that signifies normality and Figure 5.4 shows that our energy impacts data varied from the source. As all the points did not cluster around the center.

While we roll out the possibility of normality in our data. We applied the Kruskal Wallis test, a non-parametric test that the assumptions support the distribution of our data to test our hypothesis since our data set violates the assumptions of normality and homogeneity of variance.

Kruskal Wallis test reports the chi-square test statistic that shows the difference in ranks of dependent variables "EnergyImpact" across all the independent variables "product browsing, product filtering, and cart updates". A degree of freedom that shows the number of groups to be compared minus 1. A p-value that states the level of statistical significance of the results. For this experiment, we chose the 95 percent confidence interval which is the standard commonly used, it strikes a balance between accuracy and precision for accepting or refuting our null hypothesis.

Table 5.1: Kruskal-Wallis Comparison Test

Kruskal chi-square	df	p-value
30.814	2	2.037×10^{-7}

Table 5.1 shows the result obtained from our statistical test. A p-value of 2.037×10^{-7} is obtained, which is less than the significance level of 0.05. Based on this, we reject the null hypothesis and accept the alternative hypothesis. We conclude that with 95% confidence there is a significant difference in energy impact across different user behaviors.

Since we accept the alternate, hypothesis that states that the energy impact varies across different user behaviors, we further probe to know if the behaviors vary from each other. To check this variation we used the pair-wise Wilcox test, which is a statistical test used to compare the differences between groups.

This test provides a matrix that shows the p-values of the comparison between different groups.

The pair-wise test in table 5.2, shows that in the comparison between groups 1 and 2, there exists a significant difference between (product filtering vs. cart update). Between groups 1 and 3 (product filtering and product browsing), the p-value is higher than 0.05,

Table 5.2: Wilcox Pairwise Test

m	n	p-value	significance
1 (product filtering)	2 (cart update)	4.6×10^{-6}	significant
1 (product filtering)	3 (product browsing)	1	not significant
2 (cart update)	3 (product browsing)	5.5×10^{-6}	significant

shows that there exists no statistically significant difference in energy impact between product filtering and product browsing and between groups 2 and 3, the p-value is less than 0.05 shows that in this group statistically there exists a significant difference in energy impact between cart update and product browsing due to user behavior.

5.2 Modeling Server-side Energy Impact

We propose a regression model that can help predict how user behavior impacts server-side energy, taking into account the user's shopping behavior, the number of users, and the duration of interaction. This model can be used for predictive analysis, by providing insight into a shopping platform's future energy needs. The model can be used to identify areas and behaviors with high energy requirements.

Mathematically, we define the server-side consumption model as:

$$S_E = B_0 + C_1 * U_i + C_2.F_j + C_3 * T * N \tag{5.1}$$

Where,

S_E = server-side energy impact

B_0 = the baseline energy impact of the server (this represents the energy impact of an idle server when no shopping interaction is performed).

C_1, C_2, C_3 = coefficient (these are weights assigned to variables, they are multipliers that determine how much influence each variable, i.e, U_i, F_i, and T has on energy impact.

U_i = user's shopping behavior

F_i = fail request T = time per interaction N = the number of users.
[1]

[1] These coefficients are determined by analyzing real data, using regression tools, this tool adjusts the coefficients to find the best fit that the model can use to make predictions close to the real data. Theoretically, coefficients can be set manually by guessing the weight. Practically, this is not recommended, because guessing the weight might not reflect the true relationship between energy impact and user behavior.

This model captures the number of users interacting with the shop, the shopping behavior exhibited, the failed request, and the duration of interaction, which are the variables that collectively influence server-side energy impact.

5.3 Energy Cost of Reprocessing Failed Request

When the system is still efficient and processing requests and jobs at 100% without any failure the energy impact increases linearly. When the server exceeds the point of 100% request delivery the efficiency of the server reduces.

On re-sending a failed request, we reason that this might incur extra costs on energy used. This is because each request will require an extra CPU cycle, computation as well as network activity, which collectively consumes energy. When resending such requests right after a failure, the energy cost may spike up because the server becomes less efficient under high load. However, we acknowledge that efficient caching techniques, error handling, and load balancing can minimize the energy overhead of re-sending failed requests.

On the server side, the shopping behavior that has the highest energy impact is product browsing this can be explained by the high volume of requests sent to the server, product filtering follows this impact because filtering products demands a high level of processing from the servers. However, we discovered that carts update had the least energy impact on the server side. For failed requests, re-sending and re-processing such requests may incur extra costs in terms of energy usage.

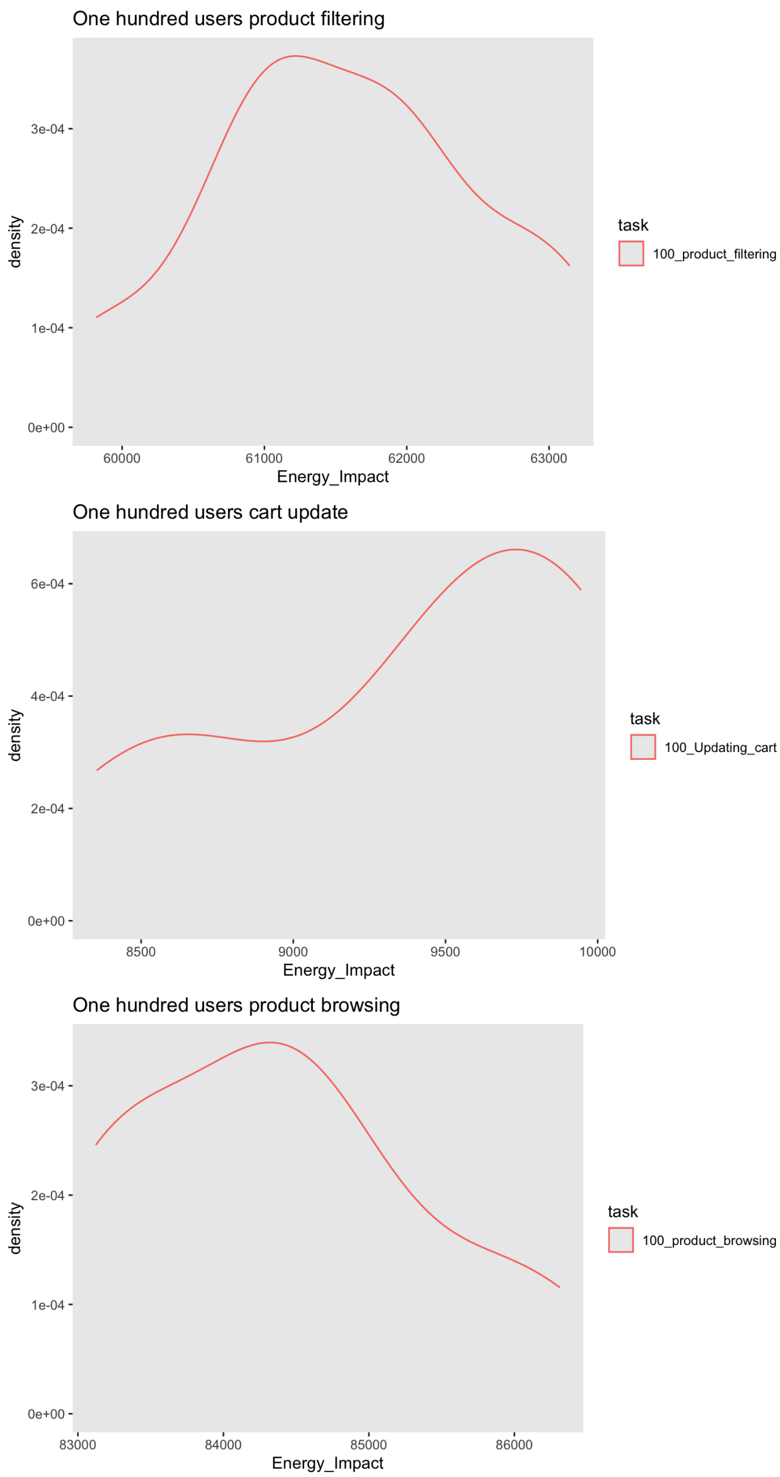

Figure 5.3: Density plots Showing distribution of different behavior for 100 users

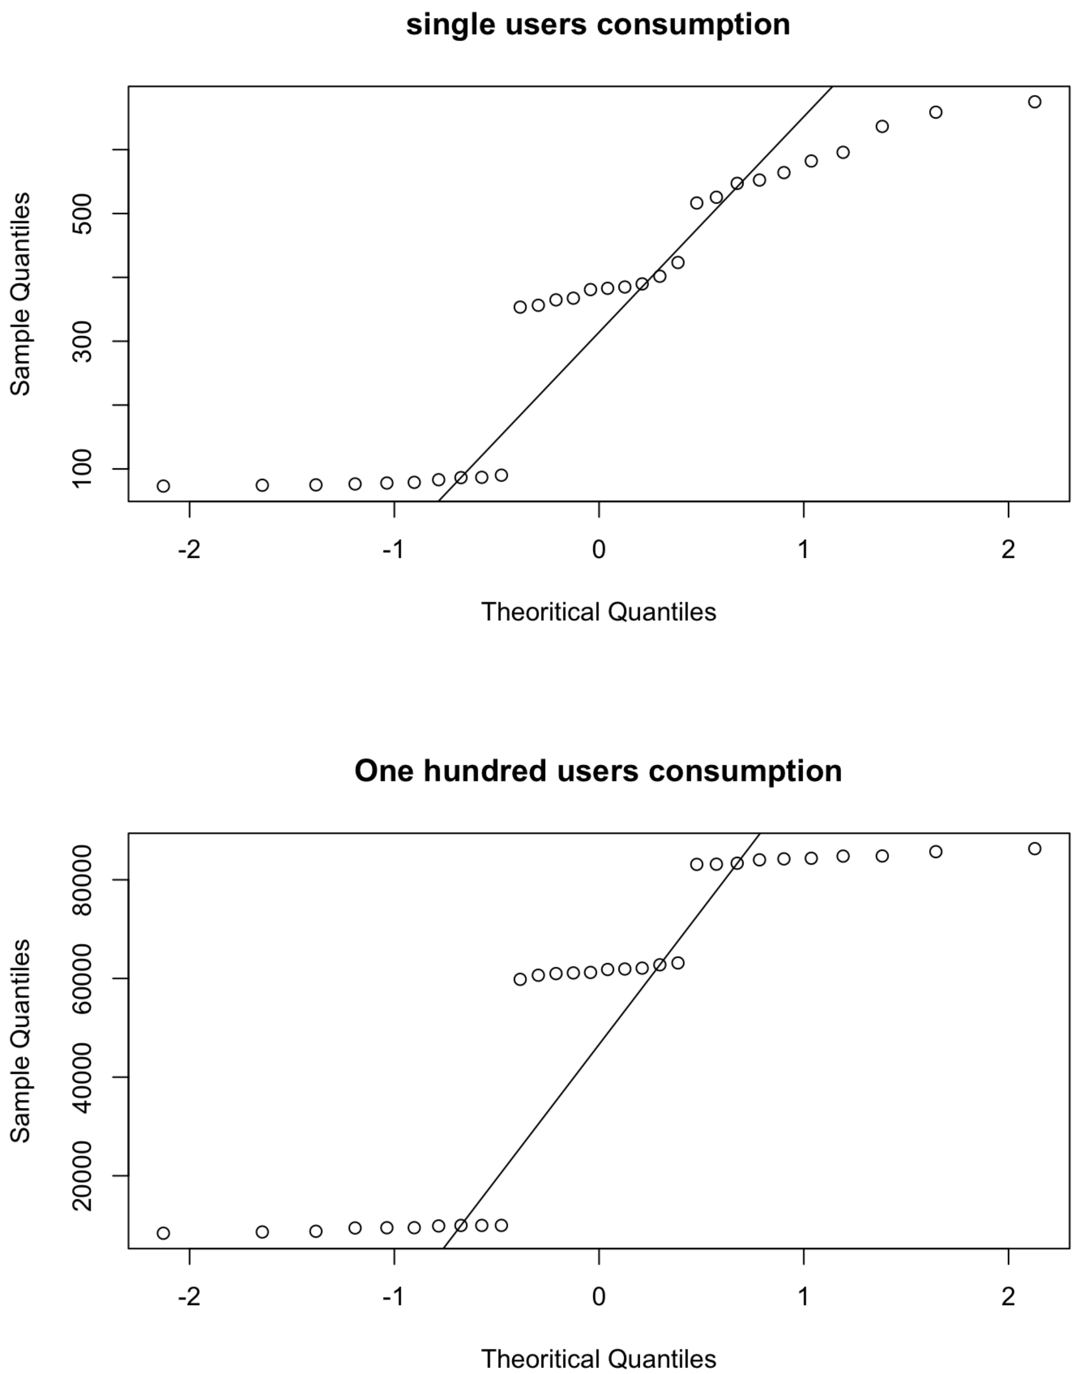

Figure 5.4: Normality check using qqplots.

6

Observation, Insights and Discussion

This chapter provides insights from our analysis of the energy impact on both the client and server sides. We provide answers to the set research questions and proposed strategies to reduce this energy impact.

From our experimental analysis, we found that different user behaviors and interactions influenced the energy impact on the client-side. Our systematic research reveals that the type of device and the energy efficiency of the browser engine, via which a shopping application is accessed influences the energy usage on the client side. On the server-side, the energy impact is heavily influenced by the number of concurrent users, the complexity of user interaction, the server's efficiency, and the server's ability to balance load effectively.

Our research shows that user shopping behavior has an impact on both server and client-side energy usage. We also found that obtaining single metrics for end-to-end energy usage in the computing continuum is not visible from the insights obtained from our research; existing research in 2 supports this due to the inter-connectivity of devices and their heterogeneity at each level of the computing continuum.

On the analysis of how different users' behaviors impact energy usage on the computing continuum, the answers to the proposed research question provide insight into the influences of these behaviors on energy usage on the computing continuum.

6.1 Observation, Findings, and Implications

We approach our research questions through a systematic study, where we set our research objectives which were to access and measure the energy impact of digital service, specifically, in the online shopping domain. This was to understand how users' shopping interaction and how user behavior impacts energy consumption. We performed holistic literature research where we identified gaps in the existing body of knowledge. The insight gained from the systematic study guided us in the development of our research question and subsequently the development of the research hypothesis. Based on the outcome of our systematic study and our experiment we attempt to answer the formulated research questions.

RQ1. What are the methodologies for measuring the end-to-end energy consumption of digital services across the computing continuum?

Our research as evidenced in section 2, previous studies have attempted different methodologies for tracking energy use in the computing continuum. Different methodologies have been proposed and implemented at different layers of the continuum. From the systematic studies, we identified the following as the methodologies for monitoring and measuring digital service energy impacts:

Energy profiling: These are mostly software tools used to track the energy consumption in the continuum. Here, energy impact data can be obtained at different levels of granularity.

Simulation and modeling method: These are mostly created as virtual prototypes to imitate the computing environment, where realistic workloads can be tested. These methodologies are used to estimate energy consumption in the continuum, based on a set of defined parameters and behavior tested.

Analytical method: This are empirical model that relies on real data to estimate energy consumption. The estimates from this method are based on energy data reported by system parameters. This method provides data-driven insights and has predictive capabilities.

Life Cycle Assessment (LCA): is a well-known methodology that can be used to measure the end-to-end energy consumption of digital services. This methodology gives a holistic overview of the energy impact associated with digital services from manufacturing to disposal. However, obtaining the correct data for the implementation of this model is challenging.

Ideally, different methodologies can be deployed at different layers of the continuum to measure the energy consumption of digital services, but each of these methodologies has its advantages and disadvantages, as such trade-offs should be considered before a methodology is deployed. [1]

[1] When considering the end-to-end measurement in the computing continuum, a single methodology is not sufficient to effectively measure the energy impact of digital services. However, the simulation technique which is the closest to offering an end-to-end overview of energy measurement only gives an estimate.

To gain a clear understanding of the end-to-end energy consumption of digital services across a computing continuum different methodologies and tools should be deployed at each layer of the continuum. Deploying different methodologies at each layer of the continuum helps in accounting for the challenges faced due to the heterogeneity of devices.

RQ2. Does the energy consumption of digital services vary across the different components of the computing continuum?

In our research on how energy usage varies across different components of the continuum, the end-user device and server energy usage were the components of interest. Energy usage at these layers differs significantly because they are influenced by different factors. Although we did not isolate different user behaviors on the client-side, the energy impact associated with 5 minutes of typical interaction with a shopping platform earned an energy impact ranging from 700 to 900. While, on the server-side for a single user a typical 5 minutes of product browsing, product filtering and cart updates the energy impact scores are 600, 300, and 100 respectively. These numbers are composite score that takes into account the energy consumption of different components of the experimental device. The lower the composite score the better.

Since behaviors on the client-side were not isolated we can not directly compare the degree variation. However, We identified that different factors at each layer contribute to variations in energy impact. As shown in 4 client-side energy usage is mostly influenced by the energy efficiency of the browser through which the service is accessed; the device type and the nature of the interaction performed by the user impact client-side energy usage. On the server side, factors such as the workload, the number of concurrent users, the nature of user interaction, and server efficiency significantly influence energy usage.

RQ3. To what extent does users behavioral pattern impact energy usage?

Research question three aims to provide insight into which specific user behavior has the most significant impact on energy usage across the computing continuum. We found that varying shopping behaviors utilize resources differently, with the duration and device type of service access also influencing this usage. However, on the client-side, browsing product pages may have a significant impact on energy usage due to page rendering.

On the server-side 5 Product browsing had the highest energy impact, highlighting the cumulative effect of frequent database access and content delivery. This is followed by the energy impact of product filtering, as product search requires an algorithm and database matches. The cart update's energy impact though visible, this behavior appears to be more energy-efficient than product browsing and product filtering on the server-side.

We could see that each user interaction behavior had an impact on energy usage on the computing continuum, though the level severity differs. In contrast to the belief that cart updates would have the greatest impact on energy usage due to database updates and computations, we found that product browsing and filtering took the lead in this analysis. Product browsing leading in this energy impact can be attributed to the database fetches, rendering requirements of each user interaction, and the demand for rendering complex multi-media content.

RQ4. What are the environmental impacts of online shopping?
As part of our research, we looked at the environmental impact of online shopping on the computing continuum. Aside from the carbon impact associated with the last mileage delivery. Carbon footprints are left by the end-user devices and the data center that hosts the shopping applications. The quest for high-quality service, rapid response, service availability, and scalability compounds these impacts. On the client-side, the convenience of online shopping tends to change users' shopping habits, leading to smaller and more frequent purchases which increases the environmental impact left behind.

6.2 Strategies to mitigate Online Shopping Energy Impact

Based on our research insights, we propose educating digital service users on the impact of their shopping behavior on the continuum. Frequent shoppers should prioritize the use of energy-efficient devices to access online shopping applications; they should also make use of energy-saving features during online shopping sessions. Online shoppers should also be encouraged to bulk purchases to reduce frequent and random visits to online stores.

We recommend light web page development for service designers and developers to minimize energy consumption during page loading and rendering time. Also, the capability of adaptive web applications that adapt effectively to the device be explored to reduce the issue of resource contention. We should optimize algorithms that frequently require database access and updates to reduce energy costs due to computation overhead.

7

Conclusion

7.1 Summary

The objective of this research was to investigate how user behavior, impacts energy consumption in the computing continuum and to propose strategies to mitigate the negative effects of these behaviors. This research is relevant due to the growth in the number of digital services. Understanding how user interaction with these services affects energy usage will help reduce the carbon footprint left by these services due to user activities. Through awareness creation on issues related to sustainability due to digitalization, this research aims to promote friendly environmental uses of digital solutions.

We presented insights on how different online shopping behaviors and interactions impact the energy usage of the computing continuum. We started by theoretically reviewing existing literature on digital services and continuum energy usage. From preliminary research, we streamlined our work to understand how user behavior in online shopping impacts continuum energy usage.

Data collected on user behavior and its associated energy impact were analyzed, and results show that in the context of online shopping product browsing has the highest impact on energy on the server-side. Based on the insight from the data collected and analyzed we develop a statistical model that predicts server-side energy usage based on different user behaviors. This model is usable for predictive analysis and informing developers about the energy efficiency/need of their digital services. This research provides insights that address a critical aspect of sustainable IT by focusing on user interactions, behavioral patterns, and their environmental implications.

Besides the behavioral influence of users on energy usage, we identified that browser energy efficiency impacts client-side energy. On the server side, resource contention and saturation, as well as the number of concurrent users and frequency of access, impacts energy usage.

We addressed challenges associated with empirical studies in computer science in terms of energy monitoring and measurement tools, encouraging the design of tools that are cross-platform compatible and suited to new hardware. We summarized the implications of our findings for service users, developers, and policymakers.

The insight from the theoretical and empirical studies helped us understand shopping behaviors with the highest energy impact and the factors that influence this consumption.

7.2 Contributions

This research provides insights into how user interactions during online shopping sessions impact energy usage in the continuum, which can aid developers in code optimization to reduce resource usage, reduce latency, and improve load balancing. Quantifying the energy impact of different user behaviors highlights behaviors with high energy spikes in online shopping while pointing to where consumption can be reduced to lower carbon emissions. With insights from this research, energy-efficient user behaviors are encouraged, such as a reduction in incessant small shopping sessions to promote sustainable use of digital services. Educating users about the energy impact of their digital behaviors to foster more conscientious usage patterns. For digital service providers, understanding user behavior can lead to optimization strategies that can be leveraged to reduce energy impact due to users' behavior. Policymakers can leverage findings from this research to make informed decisions and regulations that promote energy efficiency in the digital service industry.

7.3 Threats to validity

During the research, data collection, and analysis phase, we identified some threats that can influence the validity of our results. These threats cut across, internal, external, and construct validity. User behaviors and technology evolve so these findings might not hold in the future. The experimental results depend highly on the data collected from *powermetrics*, which gives a composite score of energy impact per process. Results may vary with different tools. However, within the setup and technical specification described in 3.2, 3.3 this result can be replicated. The experiment was performed in a controlled environment, results might differ in real-world scenarios. The experimental platform poses a threat to the validity of these results as the result might be different given a different setup. The accuracy and reliability of tools and methods used to measure energy consumption and user behavior may affect the generality of our measurements.

Also, operating a client and server on the same device poses both advantages and implications to the experiment. The issue of resource contention will arise due to the computational resources being shared. This contention is capable of elevating the amount of energy used as well as impacting the battery life of the experimental device. To mitigate this threat, we ensured that only the setup needed for the experiment was running. Also, we allow for cooling time between the experiments to reduce the thermal effect.

On the brighter side consolidating the client and server on a single machine tends to reduce the energy footprint associated with network communication and idle states.

On the client side, accurately defining and measuring specific user behaviors was challenging, which could lead to misinterpretation of user actions. As such we did not isolate users behavior on the client side.

7.4 Future Work

This research serves as a foundation for future studies to explore the impact of user behavior on new technologies' energy usage.

We proposed developing efficient, cross-platform tools that are hardware and operating system-independent for monitoring and measuring energy consumption data.

We also proposed research into how specific user interface design can impact energy consumption in the context of online shopping.

References

[1] Unctad annual report 2021, 2021. URL https://unctad.org/annual-report-2021. 2

[2] Ehsan Ahvar, Anne-Cécile Orgerie, and Adrien Lebre. Estimating energy consumption of cloud, fog, and edge computing infrastructures. *IEEE Transactions on Sustainable Computing*, 7(2):277–288, 2019. 1, 9

[3] AIOTI. Aioti computing continuum. White paper, 2022. URL https://ioti.eu/wp-content/uploads/2022/04/AIOTI-Computing-Continuum-Final.pdf. Accessed: April 4, 2024. 1

[4] Mohammad Aldossary. A review of dynamic resource management in cloud computing environments. *Computer Systems Science & Engineering*, 36(3), 2021. 6

[5] Autorité de régulation des communications électroniques et des postes. Achieving digital sustainability. https://en.arcep.fr/uploads/tx_gspublication/achieving-digital-sustainability-report-dec2020.pdf, 2020. 1

[6] Jayant Baliga, Robert WA Ayre, Kerry Hinton, and Rodney S Tucker. Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1):149–167, 2010. 8

[7] Saeedeh Baneshi, Anuj Pathania, Benny Akesson, Andy Pimentel, and Ana-Lucia Varbanescu. Analyzing per-application energy consumption in a multi-application computing continuum. *fmec*, 2024. 7

[8] Anton Beloglazov, Jemal Abawajy, and Rajkumar Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future generation computer systems*, 28(5):755–768, 2012. 6

[9] Victor R Basili1 Gianluigi Caldiera and H Dieter Rombach. The goal question metric approach. *Encyclopedia of software engineering*, pages 528–532, 1994. 13

[10] Ana Cardoso. Reducing the energy use of video gaming: energy efficiency and gamification. 2020. 8

[11] Aaron Carroll and Gernot Heiser. An analysis of power consumption in a smartphone. In *2010 USENIX Annual Technical Conference (USENIX ATC 10)*, 2010. 7

[12] Eduardo Cuervo, Aruna Balasubramanian, Dae-ki Cho, Alec Wolman, Stefan Saroiu, Ranveer Chandra, and Paramvir Bahl. Maui: making smartphones last longer with code offload. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, pages 49–62, 2010. 8

[13] Sarah Darby. The effectiveness of feedback on energy consumption. 2006. 9

[14] João de Macedo, João Aloísio, Nelson Gonçalves, Rui Pereira, and João Saraiva. Energy wars-chrome vs. firefox: Which browser is more energy efficient? In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, pages 159–165, 2020. 23

[15] Florian Dost and Erik Maier. E-commerce effects on energy consumption: A multi-year ecosystem-level assessment. *Journal of Industrial Ecology*, 22(4):799–812, 2018. 9, 11

[16] European Commission. Eu commission, 2022. URL `https://commission.europa.eu/document/download/51991f3f-a49b-4f4d-811e-c854449169d8_en?filename=com_2022_548_3_en.pdf`. 3

[17] European Commission. Commission proposes new rules on cross-border public services. Press release, March 2024. URL `https://ec.europa.eu/commission/presscorner/detail/en/IP_24_585`. Accessed: April 4, 2024. 1

[18] European Union. Eurupa dataset, Year of Publication. URL `https://data.europa.eu/doi/10.2833/478689`. 1

[19] Muhammad Fahad, Arsalan Shahid, Ravi Reddy Manumachu, and Alexey Lastovetsky. A comparative study of methods for measurement of energy of computing. *Energies*, 12(11):2204, 2019. 22, 25

[20] Kirsten Gram-Hanssen. Efficient technologies or user behaviour, which is the more important when reducing households' energy consumption? *Energy Efficiency*, 6(3): 447–457, 2013. 10

[21] Loic Guegan and Anne-Cécile Orgerie. Estimating the end-to-end energy consumption of low-bandwidth iot applications for wifi devices. In *2019 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 287–294. IEEE, 2019. 1, 7

[22] Xiaoying Guo, Kwek Choon Ling, and Min Liu. Evaluating factors influencing consumer satisfaction towards online shopping in china. *Asian social science*, 8(13):40, 2012. 9

[23] Mikko VJ Heikkinen, Jukka K Nurminen, Timo Smura, and Heikki Hämmäinen. Energy efficiency of mobile handsets: Measuring user attitudes and behavior. *Telematics and Informatics*, 29(4):387–399, 2012. 9

[24] Kerry Hinton, Jayant Baliga, Michael Feng, Robert Ayre, and Rodney S Tucker. Power consumption and energy efficiency in the internet. *IEEE Network*, 25(2):6–12, 2011. 8

[25] David Huckebrink, Jonas Finke, and Valentin Bertsch. How user behaviour affects emissions and costs in residential energy systems—the impacts of clothing and thermal comfort. *Environmental Research Communications*, 5(11):115009, 2023. 4, 10

[26] International Energy Agency. Electricity 2024: Analysis and forecast to 2026, 2024. URL `https://iea.blob.core.windows.net/assets/6b2fd954-2017-408e-bf08-952fdd62118a/Electricity2024-Analysisandforecastto2026.pdf`. 1

[27] George KAMIYA, Paolo BERTOLDI, et al. Energy consumption in data centres and broadband communication networks in the eu. 2024. 1

[28] Yunbo Li, Anne-Cécile Orgerie, Ivan Rodero, Betsegaw Lemma Amersho, Manish Parashar, and Jean-Marc Menaud. End-to-end energy models for edge cloud-based iot platforms: Application to data stream analysis in iot. *Future Generation Computer Systems*, 87:667–678, 2018. 1, 7, 8

[Locust.io] Locust.io. An open source load testing tool. `https://locust.io/`. 18

[29] Susan Lund, Anu Madgavkar, Jan Mischke, and Jaana Remes. What's next for consumers, workers, and companies in the post-covid-19 recovery. *McKinsey & Company*, 2021. 2, 11

[30] Elham Maghsoudi Nia, Queena K Qian, and Henk J Visscher. Analysis of occupant behaviours in energy efficiency retrofitting projects. *Land*, 11(11):1944, 2022. 4, 9

[31] Mozilla Developer Network. Firefox source docs: Power metrics, 2024. URL `https://firefox-source-docs.mozilla.org/performance/powermetrics.html`. Accessed: 2024-08-01. 16

[32] .NET Foundation. eshoponcontainers, 2024. URL `https://github.com/dotnet/eShop`. Accessed: 2024-08-01. 16

[33] Sergio Pérez, Patricia Arroba, and José M Moya. Energy-conscious optimization of edge computing through deep reinforcement learning and two-phase immersion cooling. *Future Generation Computer Systems*, 125:891–907, 2021. 7

[34] Chris Preist, Dan Schien, Paul Shabajee, Stephen Wood, and Christopher Hodgson. Analyzing end-to-end energy consumption for digital services. *Computer*, 47(5):92–95, 2014. 1, 7, 8, 10

[35] Kurt W Roth and Kurtis McKenney. *Energy consumption by consumer electronics in US residences.* TIAX LLC, 2007. 7

[36] Robert Rusek, Joaquim Melendez Frigola, and Joan Colomer Llinas. Influence of user behavior on energy consumption and its relation with comfort. a case study based on sensor and crowd-sensed data. 2021. 9

[37] Jeffrey D Sachs, Christian Kroll, Guillame Lafortune, Grayson Fuller, and Finn Woelm. *Sustainable development report 2022.* Cambridge University Press, 2022. 2

[38] Bardia Safaei, Amir Mahdi Hosseini Monazzah, Milad Barzegar Bafroei, and Alireza Ejlali. Reliability side-effects in internet of things application layer protocols. In *2nd International Conference on System Reliability and Safety (ICSRS 2017)*, pages 207–212, 2017. doi: 10.1109/ICSRS.2017.8272822. 2

[39] Daniel Schien, Paul Shabajee, Mike Yearworth, and Chris Preist. Modeling and assessing variability in energy consumption during the use stage of online multimedia services. *Journal of Industrial Ecology*, 17(6):800–813, 2013. 9, 11

[40] Weisong Shi, Jie Cao, Quan Zhang, Youhuizi Li, and Lanyu Xu. Edge computing: Vision and challenges. *IEEE internet of things journal*, 3(5):637–646, 2016. 7

[41] Statista. Online shopping - statistics & facts. `https://www.statista.com/topics/871/online-shopping/`, Year of Access. 2

[42] Kamlesh Tiwari and Mohammad Shadab Khan. Sustainability accounting and reporting in the industry 4.0. *Journal of cleaner production*, 258:120783, 2020. 8

[43] Ihsan Ullah, Hyun-Kyo Lim, Yeong-Jun Seok, and Youn-Hee Han. Optimizing task offloading and resource allocation in edge-cloud networks: a drl approach. *Journal of Cloud Computing*, 12(1):112, 2023. 8

[44] Eric Williams. Environmental effects of information and communications technologies. *nature*, 479(7373):354–358, 2011. 11

[45] Claes Wohlin, Martin Höst, and Kennet Henningsson. Empirical research methods in web and software engineering. *Web engineering*, pages 409–430, 2006. 8

[46] World Bank. With decisive actions, net zero energy is within reach in the emerging europe and central asia, says new world bank report. Press release, February 2024. URL `https://www.worldbank.org/en/news/press-release/2024/02/21/with-decisive-actions-net-zero-energy-is-within-reach-in-the-emerging-europe-and-cen`
Accessed: April 4, 2024. 1

[47] Jie Xu, Lixing Chen, and Shaolei Ren. Online learning for offloading and autoscaling in energy harvesting mobile edge computing. *IEEE Transactions on Cognitive Communications and Networking*, 3(3):361–373, 2017. 7

Appendix

We fit our energy impact data into our proposed model, the regression results show that at 99.96% our model can explain the variation in the Energy Impact values observed.

The results terminologies are explained below:

Residual value: shows the difference between observed and predicted value.

Coefficient: the estimated value of each predictor.

Intercept: The baseline measurement in this case a single user.

Estimate: explains how much impact each variable has on energy holding other parameters constant.

Residual Std.Error: the amount by which model prediction differs from observed values.

Multiple R-Square: percentage of variance explained by the model. Adjusted R-Squared: how much the model fits.

F-statistic: test the overall significance model. It compares the model fit against the intercept. Check if the model best explains variation in the dependent variable.

p-value: shows the significant level of the model.

All raw data files and an analysis script, including the replication package are available at: https://github.com/nsybee/User-behavior-impact-of-digital-services.git

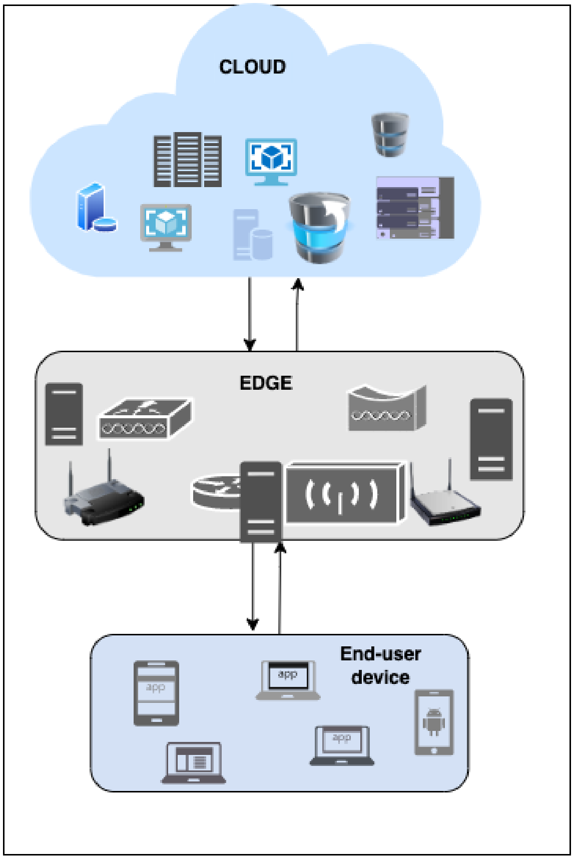

Figure 7.1: Computing continuum layers

```
Call:
lm(formula = EnergyImpact ~ task + Time * number_of_users, data = browsing_reg)

Residuals:
     Min       1Q   Median       3Q      Max
-1268.72  -111.73    -4.64   141.58  1920.76

Coefficients: (3 not defined because of singularities)
                         Estimate Std. Error t value Pr(>|t|)
(Intercept)                 585.3      178.1   3.286  0.00227 **
task10_product_browsing    5221.2      251.9  20.727  < 2e-16 ***
task100_product_browsing  83806.6      251.9 332.687  < 2e-16 ***
task5_product_browsing     2289.4      251.9   9.088 7.49e-11 ***
Time                           NA         NA      NA       NA
number_of_users                NA         NA      NA       NA
Time:number_of_users           NA         NA      NA       NA
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 563.3 on 36 degrees of freedom
Multiple R-squared:  0.9998,    Adjusted R-squared:  0.9998
F-statistic: 5.223e+04 on 3 and 36 DF,  p-value: < 2.2e-16
```

Figure 7.2: Regression result of product browsing

```
Call:
lm(formula = EnergyImpact ~ task + Time * number_of_users, data = reg_updates)

Residuals:
     Min       1Q   Median       3Q      Max
-1018.26   -11.73     0.07    28.01   572.69

Coefficients: (3 not defined because of singularities)
                      Estimate Std. Error t value Pr(>|t|)
(Intercept)              80.27      94.53   0.849   0.4014
task10_Updating_cart    604.96     133.69   4.525 6.35e-05 ***
task100_Updating_cart  9292.46     133.69  69.507  < 2e-16 ***
task5_Updating_cart     261.12     133.69   1.953   0.0586 .
Time                        NA         NA      NA       NA
number_of_users             NA         NA      NA       NA
Time:number_of_users        NA         NA      NA       NA
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 298.9 on 36 degrees of freedom
Multiple R-squared:  0.9948,	Adjusted R-squared:  0.9943
F-statistic:  2275 on 3 and 36 DF,  p-value: < 2.2e-16
```

Figure 7.3: Regression result of Cart Update

```
Call:
lm(formula = EnergyImpact ~ task + Time * number_of_users, data = reg_filtering)

Residuals:
     Min       1Q   Median       3Q      Max
-1736.20  -101.36    -6.38    66.85  1590.79

Coefficients: (3 not defined because of singularities)
                          Estimate Std. Error t value Pr(>|t|)
(Intercept)                  380.5      162.3   2.344   0.0247 *
task10_product_filtering    3160.2      229.6  13.767 6.49e-16 ***
task100_product_filtering  61174.2      229.6 266.487  < 2e-16 ***
task5_product_filtering     1335.3      229.6   5.817 1.22e-06 ***
Time                            NA         NA      NA       NA
number_of_users                 NA         NA      NA       NA
Time:number_of_users            NA         NA      NA       NA
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 513.3 on 36 degrees of freedom
Multiple R-squared:  0.9996,	Adjusted R-squared:  0.9996
F-statistic: 3.385e+04 on 3 and 36 DF,  p-value: < 2.2e-16
```

Figure 7.4: Caption

```
Call:
lm(formula = EnergyImpact ~ task + Time * number_of_users, data = reg_hundred_users)

Residuals:
     Min       1Q   Median       3Q      Max
-1736.19  -607.78    71.39   518.45  1920.76

Coefficients: (3 not defined because of singularities)
                          Estimate Std. Error t value Pr(>|t|)
(Intercept)                84392.0      287.0  294.02   <2e-16 ***
task100_product_filtering -22837.3      405.9  -56.26   <2e-16 ***
task100_Updating_cart     -75019.3      405.9 -184.81   <2e-16 ***
Time                            NA         NA      NA       NA
number_of_users                 NA         NA      NA       NA
Time:number_of_users            NA         NA      NA       NA
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 907.7 on 27 degrees of freedom
Multiple R-squared:  0.9992,	Adjusted R-squared:  0.9992
F-statistic: 1.795e+04 on 2 and 27 DF,  p-value: < 2.2e-16
```

Figure 7.5: Regression result of 100 users

```
Call:
lm(formula = EnergyImpact ~ task + Time * number_of_users, data = reg_five_user)

Residuals:
     Min       1Q   Median       3Q      Max
-129.849  -26.781   -3.686   22.816  151.271

Coefficients: (3 not defined because of singularities)
                       Estimate Std. Error t value Pr(>|t|)
(Intercept)             2874.76      23.12  124.36   <2e-16 ***
task5_product_filtering -1158.95      32.69  -35.45   <2e-16 ***
task5_Updating_cart     -2533.37      32.69  -77.49   <2e-16 ***
Time                         NA         NA      NA       NA
number_of_users              NA         NA      NA       NA
Time:number_of_users         NA         NA      NA       NA
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 73.1 on 27 degrees of freedom
Multiple R-squared:  0.9955,    Adjusted R-squared:  0.9952
```

Figure 7.6: Regression result of 5 users