European Journal of Information Technologies and Computer Science
Vol 5 | Issue 1 | January 2025
ISSN 2736-5492

RESEARCH ARTICLE

Computer Power Consumption while using Ad-Blocker on a System with AI Accelerators

Khan Awais Khan*, Mohammad Tariq Iqbal, and Mohsin Jamil

ABSTRACT

This study investigates the impact of ad-blockers on system power consumption in a computing environment equipped with an AI accelerator. The increasing prevalence of online advertisements has raised concerns about system performance and energy efficiency, prompting many users to turn to ad-blockers. However, the effectiveness of ad-blockers on power consumption, especially in systems equipped with specialized AI accelerators, remains underexplored. In this research, we evaluate the power usage, GPU utilization, and memory consumption of computers running ad-blockers on both Windows and Ubuntu operating systems. The study compared traditional CPU/GPU methods with AI-accelerated scenarios, using popular ad-blockers such as AdBlock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite. Results indicate that uBlock Origin and uBlock Origin Lite were the most efficient, significantly reducing power consumption and memory usage compared to other ad-blockers. However, multimedia-heavy websites presented challenges, with increased resource usage observed. The findings emphasize the importance of choosing appropriate ad-blockers to enhance energy efficiency, optimize system resources, and contribute to sustainable computing.

Keywords: Ad-blockers, AI accelerator, energy efficiency, memory usage.

Submitted: October 22, 2024
Published: January 21, 2025

10.24018/ejcompute.2025.5.1.144

Electrical and Computer Engineering Department, Memorial University of Newfoundland, Canada.

****Corresponding Author:***
e-mail: kakhan@mun.ca

1. Introduction

In the digital age, online advertising has become an integral part of the Internet ecosystem, generating significant revenue for content creators and service providers. However, the proliferation of online advertisements also presents various challenges for end-users, including privacy concerns, security risks, distraction, and diminished browsing experiences. Moreover, the increasing volume of advertisements has a direct impact on computer performance and energy consumption, as ads require additional computing power and data to render. Consequently, users have turned to ad-blockers as a solution to mitigate these issues and enhance their browsing experience by eliminating intrusive ads.

Ad-blockers are software tools designed to detect and prevent advertisements from loading on web pages, thereby reducing network bandwidth usage, enhancing browsing speed, and addressing privacy concerns. Despite these benefits, the operation of ad-blockers itself consumes system resources, such as CPU, GPU, and memory, contributing to increased power consumption. Understanding the implications of ad-blockers on system energy consumption is essential, especially in a world increasingly focused on environmental sustainability and reducing carbon footprints. As laptops, desktops, and mobile devices are widely used, optimizing their energy efficiency not only extends battery life but also minimizes the environmental impact of technology use.

Recent advancements in computer hardware have introduced specialized components, such as AI accelerators, designed to optimize specific computational tasks more efficiently than traditional CPUs and GPUs. AI accelerators, such as Tensor Processing Units (TPUs) and dedicated neural processing units (NPUs), leverage machine learning capabilities to handle complex and repetitive tasks, providing higher performance while consuming less power. These accelerators have shown significant potential in enhancing the efficiency of various workloads, including image recognition, natural language processing, and even content filtering. This raises an intriguing question: Can AI accelerators be leveraged to make ad-blocking more energy efficient?

This study aims to explore the power consumption of computers running ad-blockers on systems equipped with AI accelerators. Specifically, it investigates whether the use of AI acceleration can effectively reduce the energy footprint associated with ad blocking compared to conventional methods that rely solely on CPU or GPU resources. By examining power consumption under different scenarios—such as using traditional ad-blockers versus AI-optimized ad-blocking mechanisms—this research aims to provide valuable insights into optimizing power usage for enhanced sustainability.

Furthermore, this research considers the broader implications of power consumption in modern computing environments. As consumers increasingly rely on mobile devices and laptops, battery life has become a critical factor in determining user satisfaction. Ad-blockers, while improving browsing speed and privacy, might paradoxically drain more power if not implemented efficiently. Therefore, understanding how AI accelerators impact the energy usage of ad-blocking applications can lead to the development of more energy-efficient solutions, benefiting both individual users and contributing to a more sustainable technology landscape.

The results of this study will be useful for multiple stakeholders, including end-users seeking to maximize device battery life, developers aiming to create efficient ad-blocking solutions, and researchers focusing on sustainable computing practices. The research also addresses the importance of energy efficiency in the design and implementation of privacy-enhancing technologies, highlighting how hardware advancements can be utilized to balance privacy, performance, and sustainability in an increasingly connected world.

2. Literature Review

The rapid advancement of artificial intelligence (AI) has led to an increasing demand for efficient hardware accelerators capable of handling complex computations. Traditional processors struggle to meet the performance and energy efficiency requirements of modern AI applications. As a result, specialized AI hardware accelerators have emerged as critical components in the deployment of AI systems. A comprehensive overview of AI hardware accelerators, including their architectural designs, performance metrics, and integration challenges, is provided by Mishra *et al*. [1]. This foundational knowledge sets the stage for exploring specific advancements in energy-efficient accelerator designs and their applications across various domains.

Energy efficiency is a paramount concern in the design of deep learning accelerators, especially given the growing complexity of neural network models. The Eyeriss accelerator introduced by Chen *et al*. [2] is an energy-efficient reconfigurable hardware designed for deep convolutional neural networks (CNNs). Eyeriss employs a novel dataflow called Row Stationary, which optimizes data reuse and minimizes data movement, leading to significant energy savings.

Building on the need for energy efficiency in personalized applications, an accelerator facilitating in situ personalization on smart devices is proposed by Choi *et al*. [3]. This design allows for on-device training of deep CNNs, reducing dependency on cloud services and enhancing user privacy. The accelerator optimizes energy consumption by leveraging low-precision computations and efficient memory hierarchies.

In the quest for adaptability, the UNPU accelerator presented by Lee *et al*. [4] supports fully variable weight bit precision. UNPU dynamically adjusts the precision of weights during neural network operations, balancing the trade-off between energy efficiency and computational accuracy. This flexibility enables the accelerator to cater to a wide range of applications with varying precision requirements.

The separation of memory and processing units in traditional architectures leads to significant energy consumption due to data movement. In-memory computing addresses this challenge by performing computations within the memory elements themselves. The Pipelayer accelerator introduced in Song *et al*. [5] leverages resistive random-access memory (ReRAM) for deep learning applications. By minimizing data transfer between memory and processing units and supporting pipeline parallelism, Pipelayer reduces energy consumption and enhances throughput.

The challenges and solutions associated with in-memory computing for AI accelerators are extensively discussed in Cherupally *et al*. [6]. Limitations of existing memory technologies are examined, and architectural innovations are proposed to overcome them. This work emphasizes the importance of co-designing hardware and algorithms to fully exploit the benefits of in-memory computing, particularly in terms of energy efficiency and performance scalability.

To accommodate the diverse computational patterns of deep learning models, flexibility in dataflow mapping is essential. The MAERI architecture presented in Kwon *et al*. [7] enables flexible dataflow mapping over deep neural network (DNN) accelerators through reconfigurable interconnects. MAERI's design allows for the efficient execution of various neural network layers by adapting the dataflow to the specific requirements of each layer. This adaptability enhances both performance and energy efficiency, as the accelerator can optimize resource utilization dynamically.

Understanding the landscape of AI accelerators requires comprehensive benchmarking and analysis. A survey of machine learning accelerators is conducted in [8], examining architectural features, performance metrics, and application domains. This work provides valuable insights into the strengths and weaknesses of various accelerator designs, highlighting trends in the field.

A systematic study on benchmarking AI inference accelerators is presented by Jiang *et al*. [9], proposing methodologies for evaluating accelerator performance across different workloads and models. The importance of standardized benchmarks for fair comparisons is emphasized.

Significant surveys contributed by Capra *et al*. [10], [11] update the state of efficient hardware architectures for accelerating deep CNNs and discuss hardware and software optimizations for accelerating deep neural networks.

These works collectively offer a thorough understanding of the current state and evolution of AI hardware accelerators.

The deployment of AI accelerators extends beyond traditional computing applications into specialized fields like healthcare and biomedical engineering. Hardware implementations of deep network accelerators tailored for these applications are explored in Azghadi *et al*. [12]. This research highlights the unique challenges in these domains, such as the need for high precision and reliability and how customized hardware solutions can address these requirements while maintaining energy efficiency.

AI technologies play a significant role in optimizing energy consumption in smart devices and systems. A review of AI-empowered methods for smart energy consumption is provided in Himeur *et al*. [13], focusing on load forecasting, anomaly detection, and demand response. The work discusses how AI algorithms, when implemented efficiently on hardware accelerators, can lead to substantial energy savings in smart grids and buildings.

The complexity of modern chip design necessitates innovative methodologies to accelerate development cycles. A graph placement methodology leveraging machine learning to expedite chip design processes is introduced in Mirhoseini *et al*. [14]. By formulating chip design as a reinforcement learning problem, this approach achieves superior placement quality in a fraction of the time required by traditional methods. This advancement not only speeds up the development of AI accelerators but also contributes to their performance and energy efficiency.

Personalized recommendation systems are computationally intensive and require efficient processing of large datasets. The RecNMP architecture proposed by Ke *et al*. [15] accelerates personalized recommendation tasks using near-memory processing. By bringing computation closer to memory, RecNMP reduces data movement overhead and improves energy efficiency. This approach demonstrates the applicability of AI accelerators in enhancing user-centric services while managing energy consumption effectively.

Ad-blockers have become prevalent tools for users seeking to enhance their browsing experience by eliminating unwanted advertisements. The economic implications of ad-blocker platforms on advertisers and the internet ecosystem are analyzed by Ray *et al*. [16]. The study highlights the tension between user preferences and the revenue models of content providers.

The effects of ad-blocker adoption on digital piracy are explored by Datta and Madio [17], discussing whether ad-blockers serve as a deterrent to piracy by improving user experience on legitimate platforms or inadvertently encourage piracy by disrupting revenue streams.

A lab experiment assessing the impact of ad-blockers on consumer behavior is conducted by Frik *et al*. [18]. The findings suggest that while ad-blockers improve user experience, they also alter consumer engagement with content and advertisements.

The nuanced view of ad-blockers being beneficial or detrimental to the digital economy is provided by Aseri *et al*. [19], exploring the complex interplay between user experience, content monetization, and the sustainability of online services.

Beyond economic implications, ad-blockers have a significant impact on energy consumption and device performance. The contribution of open-source ad-blockers to energy conservation is investigated in Pearce [20]. The study demonstrates that blocking advertisements reduces data usage and processing demands, leading to lower energy consumption on user devices.

Building better mobile web browsers for ad blocking from an energy perspective is the focus of Heitmann *et al*. [21]. The analysis of different ad-blocking strategies and their effects on the energy efficiency of mobile browsers provides insights into optimizing both user experience and device battery life.

An experimental study on the energy and bandwidth costs of web advertisements on smartphones is conducted in Albasir *et al*. [22]. The findings reveal that advertisements significantly increase energy consumption and data usage, underscoring the potential benefits of ad-blocking technologies in prolonging battery life and reducing costs for users.

3. Experimental Design Overview

The primary objective of this experiment was to evaluate the impact of using different ad-blockers on system power consumption, specifically focusing on systems equipped with AI accelerators. The study compared power usage between scenarios where ads were blocked versus scenarios without ad-blockers on both Windows and Ubuntu systems. Additionally, the experiment included comparisons between different ad-blockers across various websites, as well as between the Windows and Ubuntu operating systems. The hypothesis was that using an ad-blocker would reduce overall power consumption by minimizing the processing required for advertisements, particularly when using an AI-accelerated setup.

3.1. Hardware and Software Specifications

1. *Processor Type:* AMD Ryzen 9 4900H with Radeon Graphics, 3.30 GHz
2. *AI Accelerator Model:* Nvidia GeForce RTX 2060
3. *AI Accelerator and Specifications:* Built on the 12 nm process and based on the TU106 graphics processor In its TU106-200-KA-A1 variant, the card supports DirectX 12 Ultimate. The Second generation Tenson Cores (succeeding Volta's) work in cooperation with the RT cores and their AI features
4. *System Architecture:* 64-bit operating system, x64-based processor
5. *Memory Specifications:* Installed RAM: 24.0 GB (23.4 GB usable)
6. *Network Configuration:* Wired network connection, 1.5 Gbps speed
7. *Operating System and Version:*
   - Windows 11 Home, Version 23H2
   - Ubuntu (Version 24.04.1)
8. *Ad-Blockers Used:*

- Tested without an ad-blocker and with the following: AdBlock, AdBlock Plus, uBlock, uBlock Origin, and uBlock Origin Lite.

9. *Additional Software/Dependencies:*
    - Web Browser: Latest version of Google Chrome for both Windows and Ubuntu systems.

3.2. *Measurement Tools*

1. *Power Measurement Tool Method:*
    - For Windows, HWinfo was used to monitor both CPU and GPU power consumption.
    - For Ubuntu, a custom bash script was utilized to log CPU and GPU power, as detailed in the pseudocode below.
2. *Additional Metrics Measured:*
    - CPU power consumption
    - GPU power consumption
    - Memory usage

3.3. *Test Scenarios*

1. *Scenarios Tested:*
    - Power consumption without an ad-blocker (baseline scenario).
    - Power consumption with different ad-blockers: AdBlock, AdBlock Plus, uBlock, uBlock Origin, and uBlock Origin Lite.
    - Each scenario was tested on both Windows and Ubuntu systems.
2. *Number of Repetitions for Each Scenario:*
    - The same video was played from the following websites for each ad-blocker, and each test was repeated three times for statistical significance:
      - 9gag
      - ARYzap
      - Dailymotion
      - Kisscartoon
      - YouTube
      - Cricbuzz
      - Espncricinfo
      - TheNews
3. *Comparisons Performed:*
    - Ad Blockers on Different Websites: Power consumption data was collected for each ad-blocker across all the aforementioned websites. This allowed for a comparison of energy efficiency across different ad-blockers when exposed to varying content types.
    - Windows vs. Ubuntu: The experiment also aimed to compare system power consumption between Windows and Ubuntu operating systems for each scenario, highlighting any differences in how the two systems handle advertisements and ad-blocking software.

3.4. *Data Collection Procedure*

1. *Duration of Each Test Run:* The duration of each test run depended on the length of the video, but the same video length was used for each ad-blocker scenario to ensure consistency.
2. *Logging Method:* The data collection procedure for Ubuntu involved a custom script to log CPU and GPU power consumption along with memory usage. The pseudocode in Fig. 1 details the methodology used for logging the data.

4. Results

The results of this study are presented to evaluate the impact of different ad-blockers on system power consumption for systems with an AI accelerator, comparing the performance on both Windows and Ubuntu operating systems. Specifically, comparisons were made between power consumption without an ad-blocker and with different ad-blockers (AdBlock, AdBlock Plus, uBlock, uBlock Origin, and uBlock Origin Lite) across various websites.

4.1. *CPU Power Consumption by Ad-Blockers on Windows*

Table I shows CPU power consumption (in watts) for different websites using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite. The goal is to compare the effectiveness of these ad-blockers in reducing CPU power usage on Windows.

uBlock Origin and uBlock Origin Lite consistently resulted in the lowest CPU power consumption across most websites, showing their effectiveness in reducing power usage. For example, 9gag showed a reduction from 11.2 W (without an ad-blocker) to 8.2 W with uBlock Origin. In contrast, Adblock Plus demonstrated moderate reductions, while Adblock had mixed performance, especially with ARYzap, where it reduced consumption to 2.5 W.

A stacked bar chart is proposed to illustrate the comparative effectiveness of different ad-blockers in reducing CPU power consumption across multiple websites.

This visualization highlights the overall reduction in power usage achieved by each ad-blocker and provides insights into which websites contribute most to the total power savings. The stacked bar chart effectively conveys both the aggregate impact and individual website-specific efficiency of each ad-blocker, making it a comprehensive tool for comparative analysis.

4.2. *GPU Power Consumption by Ad-Blockers on Windows*

This table presents GPU power consumption (in watts) across different websites using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite. The objective is to evaluate the effectiveness of these ad-blockers in reducing GPU power consumption on the Windows operating system.

From Table II, uBlock Origin demonstrated the most consistent reduction in GPU power consumption across most websites, such as 9gag and Dailymotion, with notable reductions from 25.3 W (without an ad-blocker) to 21.5 W and 24.1 W, respectively. Adblock and Adblock Plus generally showed moderate reductions, while uBlock had mixed results, including increased GPU consumption

```
#!/bin/bash
filename="resource_log_$(date).csv"
interations=12
# Check if radeontop is installed
if ! command -v radeontop &> /dev/null; then
    echo "radeontop could not be found. Please install it with: sudo apt install
radeontop"
    exit 1
fi
# Function to get CPU package power
get_cpu_power() {
    sensors | grep 'PPT' | awk '{print $2}'
}
get_nvidia-gpu_power() {
    nvidia-smi --query-gpu=power.draw --format=csv,noheader
}
# Function to get memory usage
get_memory_usage() {
    free -m | grep "Mem" | awk '{print $3}'
}
log_entries=()
# Logging for 12 seconds, checking every second
for (( i=1; i<=interations; i++ )); do
    cpu_power=$(get_cpu_power)
    gpu_power=$(get_nvidia-gpu_power)
    memory_usage=$(get_memory_usage)
    # Create a CSV line
    log_entry="$i,$(date),$cpu_power, $gpu_power,$memory_usage"
    echo "----------------------------------"
    echo "Log Entry: $i, Date: $(date) ,CPU Power: $cpu_power, GPU Power: $gpu_power,
Memory Usage: $memory_usage"
    # Add the entry to the array
    log_entries+=("$log_entry")
    sleep 1  # Wait for 1 second before the next log
done
# Write the results to a CSV file
{
    echo "Log Entry, Date, CPU Power, GPU Power, Memory Usage"  # CSV header
    for entry in "${log_entries[@]}"; do
        echo "$entry"
    done
} > "$filename"
echo "Logging complete. Results saved to $filename."
```

Fig. 1. Details of the methodology used for logging the data.

TABLE I: CPU Power Consumption by Ad-Blockers on Windows

Websitws	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	11.2	9.1	10.4	8.5	8.2	8.2
ARYzap	10.3	2.5	10.0	10.0	7.5	9.5
Dailymotion	17.9	10.4	10.6	9.6	10.5	10.0
Kisscartoon	12.4	9.6	10.7	9.2	9.1	9.2
Youtube	4.0	3.8	4.1	5.7	4.2	4.1
Cricbuzz	11.7	11.1	9.6	10.4	9.8	9.6
Espncricinfo	18.6	18.6	17.5	17.4	17.6	17.5
TheNews	11.0	11.2	12.0	10.7	10.8	10.6

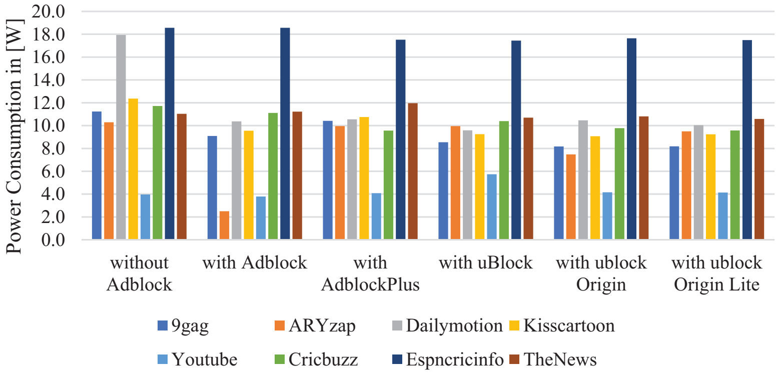

Fig. 2. Comparison of CPU power consumption by ad-blockers on windows.

TABLE II: GPU Power Consumption by Ad-Blockers on Windows

Websites	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	25.3	25.0	24.6	23.8	21.5	23.0
Aryzap	26.4	24.2	25.7	22.2	22.2	25.1
Dailymotion	34.4	28.0	27.6	27.2	24.1	26.0
Kisscartoon	28.6	23.2	25.5	25.1	24.2	23.6
Youtube	11.8	9.9	12.8	14.3	13.0	13.5
Cricbuzz	29.2	28.3	25.1	27.8	24.2	24.2
Espncricinfo	25.0	29.0	29.0	41.3	36.4	37.7
Thenews	35.0	27.3	26.9	32.5	25.3	26.3

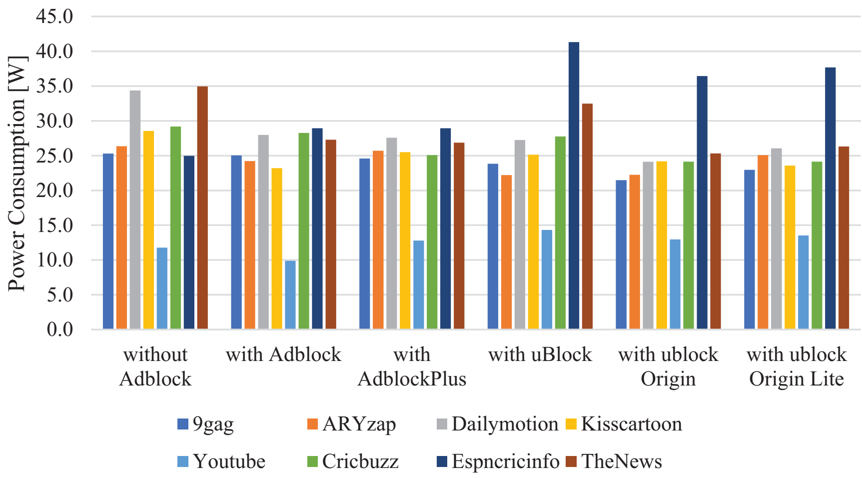

Fig. 3. Comparision of GPU power consumption by ad-blockers on windows.

TABLE III: CPU Power Consumption by Ad-Blockers on Ubuntu

Websites	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	20.36585	20.36585	13.60976	13.36585	12.175	10.15
ARYzap	10.83077	12.00833	12.10833	12.26667	12.39167	10.325
Dailymotion	21.84566	12.91489	14.19149	12.55239	11.38298	13
Kisscartoon	12.62069	16.43704	15.68302	15.87925	16.77358	16.21509
Youtube	12.62069	16.43704	15.68302	15.87925	16.77358	16.21509
Cricbuzz	19.55833	13.40833	15.725	13.925	14.68333	13.55833
Espncricinfo	23.53623	22.15942	22.62319	23.5942	21.85507	21.2029
TheNews	12.68333	14.7	13.40833	14.45833	13.44167	14.36667

on Espncricinfo (41.3 W). YouTube showed a reduction in GPU power from 11.8 W without an ad-blocker to 9.9 W with Adblock, indicating its efficiency for lighter media content.

Fig. 2 is proposed to illustrate the comparative effectiveness of different ad-blockers in reducing GPU power consumption across multiple websites.

4.3. *CPU Power Consumption by Ad-Blockers on Ubuntu*

Table III presents CPU power consumption (in watts) across different websites when using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite, on the Ubuntu operating system. Fig. 4 compares the effectiveness of these ad-blockers in reducing CPU power usage on Ubuntu.

uBlock Origin Lite and uBlock Origin achieved the lowest CPU power consumption across several websites, notably on 9gag and Dailymotion, reducing the power from 20.37 W (without an ad-blocker) to 10.15 W and 11.38 W, respectively. In contrast, Adblock Plus and uBlock showed mixed performance, with some websites (e.g., Espncricinfo) experiencing slightly increased power consumption compared to no ad-blocker. This indicates that uBlock Origin Lite was the most energy-efficient option overall on Ubuntu.

Fig. 3 is used to illustrate the comparative effectiveness of different ad-blockers in reducing CPU power consumption across multiple websites.

4.4. *GPU Power Consumption by Ad-Blockers on Ubuntu*

Table IV presents GPU power consumption (in watts) across different websites when using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite, on the Ubuntu operating system. Fig. 4 compares the effectiveness of these ad-blockers in reducing GPU power usage on Ubuntu.

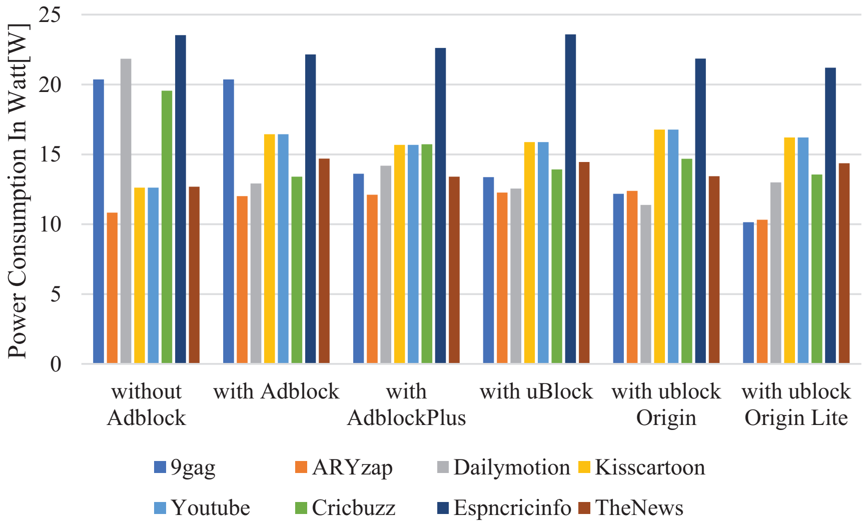

Fig. 4. Comparison of CPU power consumption by ad-blockers on Ubuntu.

TABLE IV: GPU Power Consumption by Ad-Blockers on Ubuntu

Websites	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	28.24697	15.54487	12.68321	13.96951	13.11368	11.95289
ARYzap	11.46285	17.0395	11.42833	9.560917	9.643083	8.68975
Dailymotion	19.6329	13.51326	8.478333	11.70014	11.39044	14.532
Kisscartoon	12.48525	11.95083	8.859167	11.3735	9.77025	9.95225
Youtube	15.2075	24.37908	23.88533	24.1685	24.80083	24.63283
Cricbuzz	28.932	15.12467	18.04992	13.96042	14.82758	13.6695
Espncricinfo	35.55667	32.83754	33.2113	35.75812	33.97362	33.5758
TheNews	12.3875	16.39342	16.85258	20.1735	16.59958	17.98608

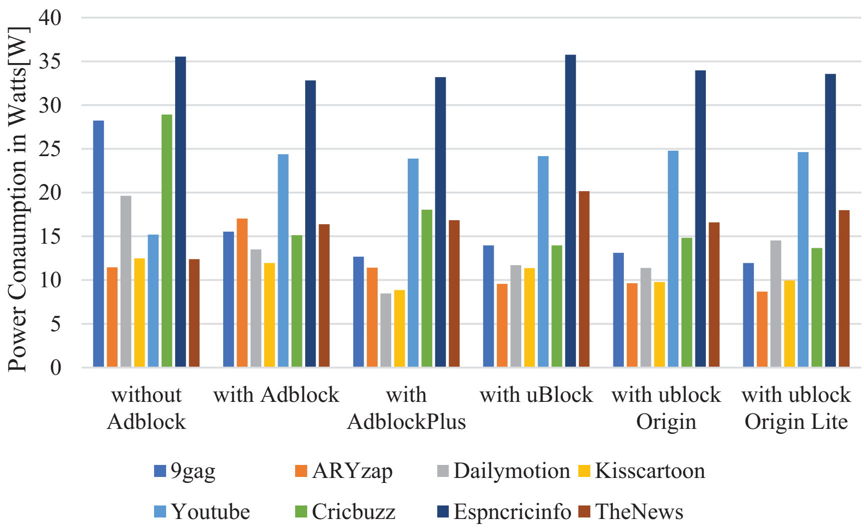

Fig. 5. Comparison of GPU power consumption by ad-blockers on Ubuntu.

In Table IV, uBlock Origin Lite and Adblock Plus demonstrated the most consistent reduction in GPU power consumption for most websites, particularly for 9gag and Kisscartoon, with reductions from 28.25 W (without an ad-blocker) to 11.95 W and 8.86 W, respectively. ARYzap also showed the lowest GPU power consumption with uBlock Origin Lite at 8.69 W. Conversely, YouTube did not show any significant reduction, and GPU power consumption remained similar across all ad-blockers.

Fig. 5 provides a visual comparison of the power consumption data across all ad-blockers and websites.

4.5. Memory Usage by Ad-Blockers on Windows (in MBs)

Table V presents the memory usage (in megabytes) for different websites while using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite, on the Windows operating system. The aim is to compare the effectiveness of these ad-blockers in

TABLE V: Memory Usage by Ad-Blockers on Windows (in MBs)

Websites	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	8106.538	8171.436	8156.769	7886.61	7799.132	7786.368
ARYzap	10592.24	10522.94	10578.93	10497.67	10329.94	10299.97
Dailymotion	11378.77	9039.435	8975.851	8935.417	10452.69	10382.44
Kisscartoon	8014.356	8158.286	8099.026	7699.921	7759.095	7776.5
Youtube	5535.633	5796.438	5879.622	5821.586	10375.21	10342.7
Cricbuzz	8065.493	7958.654	7965.456	7862.528	7902.571	7957.448
Espncricinfo	8372.398	8372.398	8059	7892.362	7858.898	7693.081
TheNews	7975.206	8033.767	8139.765	7880.556	7793.475	7759.271

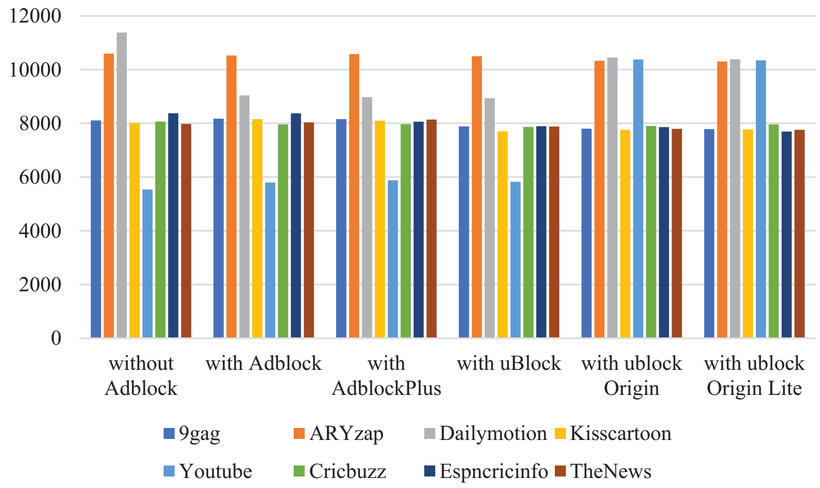

Fig. 6. Comparison of memory usage by ad-blockers on Windows [MBs].

reducing memory usage, which indicates the impact of each ad-blocker on system resource efficiency.

uBlock Origin and uBlock Origin Lite resulted in the most consistent reduction in memory usage across most websites. For instance, 9gag saw a reduction in memory usage from 8106.54 MB (without an ad-blocker) to 7786.37 MB with uBlock Origin Lite. Kisscartoon also saw a decrease, with uBlock resulting in the lowest memory usage at 7699.92 MB. Conversely, YouTube exhibited increased memory usage for most ad-blockers, with uBlock Origin reaching 10375.21 MB, highlighting potential inefficiencies for multimedia-rich content.

Bar chart has been used to visually compare memory usage for each ad-blocker across different websites (Fig. 6).

4.6. Memory Usage by Ad-Blockers on Ubuntu (in MB)

This table presents the memory usage (in megabytes) for different websites while using various ad-blockers, including Adblock, Adblock Plus, uBlock, uBlock Origin, and uBlock Origin Lite, on the Ubuntu operating system. The goal is to compare the effectiveness of these ad-blockers in reducing memory usage, which provides insight into the system resource efficiency of each ad-blocker.

Table VI, uBlock Origin and uBlock Origin Lite demonstrated the lowest memory usage across multiple websites, particularly for 9gag and Kisscartoon, with reductions from 3248.28 MB (without an ad-blocker) to 2876.96 MB and 2960.17 MB, respectively. Dailymotion also showed significant reductions with uBlock at 3187.80 MB. In contrast, YouTube and Cricbuzz displayed relatively stable or increased memory usage across all ad-blockers, with no major savings observed.

A stacked bar chart has been employed to illustrate the comparative memory usage of different ad-blockers across multiple websites (Fig. 7).

The results of this study show that uBlock Origin and uBlock Origin Lite are the most effective ad-blockers for reducing CPU and GPU power consumption as well as memory usage across both Windows and Ubuntu. On Windows, uBlock Origin and uBlock Origin Lite reduced CPU power consumption for 9gag by 26.8%, from 11.2 W to 8.2 W, and for Kisscartoon by 26.6%, from 12.4 W to 9.1 W. On Ubuntu, these ad-blockers performed even better, reducing CPU power consumption for 9gag by 50.2%, from 20.37 W to 10.15 W, and for Dailymotion by 47.9%, from 21.85 W to 11.38 W.

In terms of GPU power consumption, uBlock Origin and uBlock Origin Lite also achieved significant reductions. On Windows, Dailymotion's GPU consumption decreased by 30%, from 34.4 W to 24.1 W. On Ubuntu, uBlock Origin Lite reduced 9gag's GPU consumption by 57.7%, from 28.25 W to 11.95 W. However, for multimedia-rich websites like YouTube, GPU power consumption increased with most ad-blockers, indicating that blocking dynamic video ads can add processing overhead.

For memory usage, uBlock Origin and uBlock Origin Lite consistently reduced consumption. On Windows, Espncricinfo's memory usage decreased by 8.1%, from 8372.40 MB to 7693.08 MB with uBlock Origin Lite. On Ubuntu, Kisscartoon's memory usage dropped by 12.9%, from 3400.98 MB to 2960.17 MB with uBlock Origin

TABLE VI: Memory Usage by Ad-Blockers on Ubuntu (in MB)

Websites	Without adblock	With adblock	With adblockplus	With ublock	With ublock origin	With ublock origin lite
9gag	3248.28	3248.28	3330	3348.32	2876.96	2942.08
ARYzap	3347.077	3341	3516.033	3417.592	3333.267	3277.467
Dailymotion	3552.713	3532.893	3407.667	3187.8	3350.987	3472.947
Kisscartoon	3400.978	3106.45	3129.908	3086.767	3029.767	2960.167
Youtube	3353.459	3641.898	3551.536	3561.985	3561.691	3478.563
Cricbuzz	3611.075	3675.842	3745.183	3675.492	3656.85	3669.558
Espncricinfo	3313.812	3444.348	3864.536	3692.739	3612.638	3128.754
TheNews	3506.458	3658.633	3672.583	3698.392	3714.283	3534.683

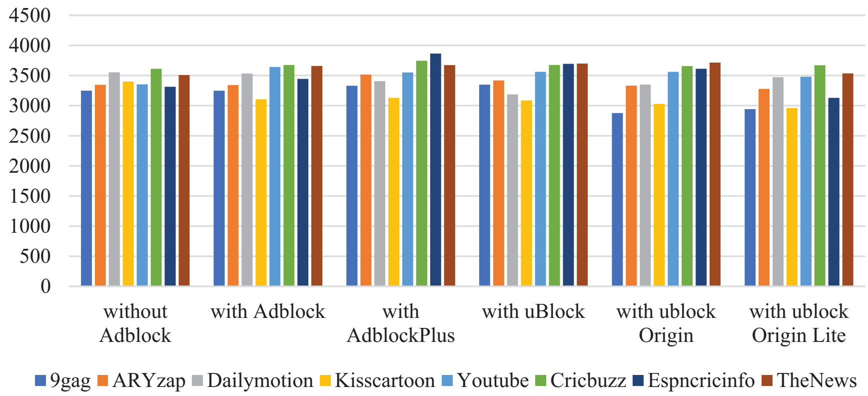

Fig. 7. Comparison of memory usage by ad-blockers on Ubuntu [MBs].

Lite. However, for YouTube, memory usage increased significantly, indicating potential inefficiencies in handling complex ad content.

Overall, uBlock Origin and uBlock Origin Lite were the most efficient in reducing power and memory usage, with Ubuntu generally showing better results than Windows. Multimedia-heavy websites, however, presented challenges, often resulting in increased resource usage due to the processing overhead involved in blocking dynamic ads. These results highlight the importance of choosing the right ad-blocker based on the type of content and operating system to achieve the best energy efficiency and resource optimization.

5. Conclusion and Future Work

The findings from this study provide valuable insights into the effects of ad-blockers on system power consumption in computing environments equipped with AI accelerators. uBlock Origin and uBlock Origin Lite consistently outperformed other ad-blockers, resulting in the most substantial reductions in both CPU and GPU power consumption as well as memory usage on both Windows and Ubuntu. On Windows, the CPU power consumption for websites like 9gag was reduced by 26.8% using uBlock Origin. Similarly, GPU power consumption on Ubuntu saw a 57.7% reduction for 9gag using uBlock Origin Lite. These results demonstrate that AI accelerators, when combined with the right ad-blockers, can effectively optimize power usage. However, the analysis also highlighted that multimedia-rich websites, such as YouTube, did not experience significant power reductions. Instead, resource usage increased due to the overhead of processing dynamic video content. This suggests that ad-blockers introduce additional processing requirements when handling multimedia advertisements, which can lead to inefficiencies.

Overall, the study concludes that uBlock Origin and uBlock Origin Lite are the most effective ad-blockers in terms of energy efficiency. Moreover, Ubuntu generally demonstrated better results than Windows, indicating that operating system choice plays a significant role in resource optimization. The results emphasize the need for careful selection of ad-blockers based on content type and system architecture to achieve the best energy efficiency. This research has implications for developers of ad-blocking technologies, system manufacturers, and end-users seeking to maximize energy efficiency and system performance.

In terms of future work, while this study provided substantial insights into the impact of ad-blockers on system power consumption, there are several areas worth exploring further. Firstly, the impact of ad-blockers on a broader variety of hardware configurations, including systems with more advanced AI accelerators, should be investigated to generalize the findings further. Additionally, examining the effect of ad-blockers on other operating systems, such as macOS or mobile platforms, could provide a more comprehensive understanding of their effectiveness across different environments. Future studies could also focus on optimizing ad-blocking algorithms to minimize processing overhead, particularly for multimedia content, to ensure that energy-saving benefits extend to video-rich websites. Finally, a more detailed analysis of network traffic and data usage implications in combination with power metrics could provide a holistic view of the benefits and trade-offs of using ad-blockers in AI-accelerated systems.

Conflict of Interest

The authors declare that they do not have any conflict of interest.

References

[1] Mishra, Cha J, Park H, Kim S. *Artificial Intelligence and Hardware Accelerators*. Cham, Switzerland: Springer; 2023.

[2] Chen Y-H, Krishna T, Emer JS, Sze V. Eyeriss: an energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE J Solid-State Circ*. 2017 Jan;52(1):127–38.

[3] Choi S, Sim J, Kang M, Choi Y, Kim H, Kim LS. An energy-efficient deep convolutional neural network training accelerator for in situ personalization on smart devices. *IEEE J Solid-State Circ*. 2020 Oct;55(10):2691–702.

[4] Lee J, Kim C, Kang S, Lee S, Won J, Lee H, *et al.* UNPU: an energy-efficient deep neural network accelerator with fully variable weight bit precision. *IEEE J Solid-State Circ*. 2019 Jan;54(1):173–85.

[5] Song L, Qian X, Li H, Chen Y. Pipelayer: a pipelined ReRAM-based accelerator for deep learning. *Proceedings of the IEEE International Symposium High Performance Computer Architecture (HPCA)*, pp. 541–52, 2017.

[6] Cherupally SK, Ramaswamy N, Juma J, Ye Y, Yadav AV, Govindan R, *et al.* In-memory computing for AI accelerators: challenges and solutions. *Semicond Sci Technol*. 2022;37(3):034001.

[7] Kwon H, Samajdar A, Krishna T. MAERI: enabling flexible dataflow mapping over DNN accelerators via reconfigurable interconnects. *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 461–75, 2018.

[8] Reuther A, Michaleas P, Jones M, Kepner J, Samsi S, Gadepally V. Survey of machine learning accelerators. *Proceedings of the IEEE High Performance Extreme Computing Conference (HPEC)*, pp. 1–12, 2020.

[9] Jiang, Li J, Liu F, Gao W, Wang L, Lan C, *et al.* A systematic study on benchmarking AI inference accelerators. *CCF Trans High Perform Comput*. 2022;4:87–103.

[10] Capra M, Bussolino B, Marchisio A, Shafique M, Masera G. An updated survey of efficient hardware architectures for accelerating deep convolutional neural networks. *Future Internet*. 2020 Jul;12(7):113.

[11] Capra M, Bussolino B, Marchisio A, Macii E, Martina M, Masera G, *et al.* Hardware and software optimizations for accelerating deep neural networks: survey of current trends, challenges, and the road ahead. *IEEE Trans Comput*. 2021 Aug;70(8):1197–212.

[12] Azghadi MR, Lammie C, Eshraghian JK, Wang L, Kavehei O. Hardware implementation of deep network accelerators towards healthcare and biomedical applications. *IEEE Trans Circ Syst*. 2020 Aug;67(8):2626–37.

[13] Himeur M, Bensaali B, Amira A, Bourguiba A, Elhoseny M, Karray F. AI-empowered methods for smart energy consumption: a review of load forecasting, anomaly detection, and demand response. *Int J Precis Eng Manuf-Green Technol*. 2023 Jan;10(1):25–45.

[14] Mirhoseini A, Goldie A, Yazgan M, Jiang JW, Songhori E, Wang S, *et al.* A graph placement methodology for fast chip design. *Nature*. 2021 Jun;594:207–12.

[15] Ke L, Gupta U, Cho BY, Brooks D, Wu C, Lee HS. RecNMP: accelerating personalized recommendation with near-memory processing. *Proceedings pf the ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 790–803, 2020.

[16] Ray A, Ghasemkhani H, Kannan KN. Ad-blockers, advertisers, and internet: the economic implications of ad-blocker platforms. *Proceedings of the International Conference on Information Systems (ICIS)*, 2017.

[17] Datta B, Madio L. Effects of ad-blockers adoption on digital piracy: a blessing or a curse? LCII Working Paper 2017-5, 2017.

[18] Frik A, Haviland A, Acquisti A. The impact of ad-blockers on consumer behavior: a lab experiment. *Proceedings of the 18th Workshop on the Economics of Information Security (WEIS)*, 2019.

[19] Aseri M, Dawande M, Janakiraman G. Ad-blockers: a blessing or a curse? *Inf Syst Res*. 2020 Sep;31(3):889–908.

[20] Pearce JM. Energy conservation with open source ad-blockers. *Technologies*. 2020 Jun;8(2):27.

[21] Heitmann N, Pirker B, Park S. Towards building better mobile web browsers for ad blocking: the energy perspective. *Proceedings of the ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, pp. 80–90, 2020.

[22] Albasir A, Naik K, Plourde B. Experimental study of energy and bandwidth costs of web advertisements on smartphones. *Proceedings of the 11th IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, pp. 693–8, 2014.