

Assessing the Silent Frontlines: Exploring the Impact of DDoS Hacktivism in the Russo-Ukrainian War

Yagiz Yilmaz*, Orcun Cetin*, Omer Said Ozturk*, Emre Ekmekcioglu*, Budi Arief[†], Julio Hernandez-Castro[‡]

**Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul, Turkey*

Email: {yagizyilmaz, orcun.cetin, omersaid, eekmekcioglu}@sabanciuniv.edu

[†]School of Computing, University of Kent, Canterbury, UK

Email: b.arief@kent.ac.uk

[‡]Computer Systems Department, Universidad Politécnica de Madrid, Madrid, Spain

Email: jc.hernandez.castro@upm.es

Abstract—This study assessed the impact and effectiveness of Distributed Denial of Service (DDoS) attacks during a period of about four months of the Russo-Ukrainian war, by observing the exchanges between the opposing sides. The data collection phase took place between the 28th of November 2022 and the 15th of April 2023. In total, we monitored 1,257 websites and web applications targeted in the conflict, with 633 targeted by pro-Russian and 624 by pro-Ukrainian entities. Only a small fraction (1.27%) of the targets remained unaffected, whereas 30.63% faced complete shutdowns. When considering the extent of the attacks conducted by the belligerents in the war, the attacks by pro-Russian entities showed a slightly more successful overall impact, with 36.18% of their targets were taken down, compared to 25.00% on the opposite side. Businesses demonstrated greater resilience against DDoS attacks compared to governmental and educational institutions. An in-depth analysis revealed significant differences in target categories, despite both sides primarily targeting businesses. Our findings regarding the usage of DDoS protection services among the 1,257 analysed targets showed that only 13.37% used such services. Among these minority of users, 70.24% had protection from the beginning of our analysis, while 29.76% adopted it only after experiencing attacks. We also looked into the use of geolocation-based access policies on websites targeted by pro-Ukrainian entities. Our findings indicated that most of these websites do not implement geolocation-based access restrictions. To an extent, such restrictions could have been useful for preventing some unsophisticated attacks. Surprisingly, only a small percentage (4.50%) restricted access to solely Russian addresses, while a fraction (12.56%) seemed to implement adaptive access policies in response to cyberattacks. Lastly, and quite surprisingly for us, we discovered that a significant number of targets on the Russian side were using anti-DDoS services and technology provided by countries that have for a long time imposed economic and commercial sanctions on Russia. This may or may not be strictly illegal, but it is without question against the spirit of these sanctions.

Index Terms—DDoS, Russia, Ukraine, Cyberwar, Hacktivism

1. Introduction

Cyberspace has proven to be an effective platform for a wide range of social interactions, education, commerce, and even governance. However, cyberspace has been used not only for positive purposes, but also for malicious activities such as cybercrime, subversion, and sabotage. Of particular note, Distributed Denial of Service (DDoS) attacks are often utilised in modern armed conflicts and political tensions between nations. The growing occurrence of – and reliance on – DDoS attacks signals a significant shift in the landscape of modern conflict, underscoring the increasing importance of the cyber domain as part of the battlefield. This transition not only amplifies the complexity and reach of conflicts, but also introduces many novel cybersecurity challenges. In this context, DDoS attacks and other forms of cyberattack act as extensions of state power and serve as asymmetric tools for non-state actors, providing a method to accomplish strategic objectives without resorting to open warfare.

While the Russo-Ukrainian war has shown the extent of both sides' cyber capabilities, cyberattacks between these warring parties did not start with the invasion of Ukraine. Russia has been conducting cyber espionage and disruption campaigns against Ukraine since at least 2014, before their Crimean illegal occupation. For example, pro-Russian groups carried out an operation in December 2015 that resulted in a significant power outage in the Ivano-Frankivsk region of Ukraine [1]. Approximately half of the households in this area experienced a loss of electricity for several hours. This action illustrates a deliberate attempt to disrupt essential services and demonstrates the tangible impact of cyber warfare on critical infrastructure. However, there has been an apparent change in strategy in the most recent conflict. While attacks against critical systems is still a very relevant threat, there is now a broader range of targets, including government agencies, institutions, businesses, banks, media, NGOs and many others.

One of the pro-Russian entities' main approaches includes deploying destructive malware to dismantle critical infrastructure, launching DDoS attacks to paralyse govern-

mental and essential services, and engaging in espionage for gaining intelligence superiority [2]. In response, Ukraine has fortified its cyber defences through international support and the mobilisation of volunteer groups, such as the “IT Army of Ukraine” which not only defends against Russian cyberthreats but also takes the fight to Russia’s digital doorsteps using DDoS attacks [3]. Even though it is known that DDoS attacks are widely deployed in the Russia-Ukraine cyber conflict [4], their extent and effectiveness are still overlooked and unreported. This presents a research gap that we aim to address through our research.

This paper presents the first empirical study evaluating the effectiveness of DDoS attacks between the parties at war after the latest Russian invasion. To achieve this, we first identified and tracked the channels and groups where the two main sides of the conflict register their targets. We compiled 1,257 targets, where 633 were targeted by pro-Russian and 624 by pro-Ukrainian entities. Then, we monitored those targets for at least 14 days to record their availability (i.e. whether they were down or not), which helped us assess how the attacks affected them. During the monitoring, we utilised several servers deployed across the globe to ensure data quality and observe geolocation-based access policies.

In summary, our key contributions are as follows:

- We carried out the first empirical study of DDoS attacks during the Russo-Ukrainian conflict.
- We observed that the pro-Russian side had more impactful attacks than the pro-Ukrainian side, overall. Over one-third of those targeted by pro-Russians (36.18%) went completely down, whereas this ratio was only around one-quarter for those targeted by pro-Ukrainians.
- We found that the rates of DDoS protection usage were remarkably low: overall, below 14% showed any indications of using these services. Additionally, we observed that some targets in Russia were able to obtain DDoS protection services from companies based in countries (notably the US) that are supposed to be applying an embargo on Russia.
- Most of the pro-Russian websites did not apply geolocation-based access policies. Only about 5% were available to only Russian IPs, and about one-eighth applied a dynamic access restriction policy. Compared to the ones that employed any kind of access restriction policy, an overwhelming majority (82.94%) did not utilise such simple measures as safeguards.

We believe that some general but quite useful and actionable recommendations on how to improve the current *status quo* of the Russo-Ukrainian cyberwar can be extracted from this work. We hope that it will also lead to a number of more general strategies to put in place in future cyber conflicts. In any case, we present our analysis in Sections 5 and 6.

The rest of this paper is structured as follows. Section 2 provides an overview of the literature highlighting relevant subjects such as DDoS attacks, cyberattacks in the Russo-Ukrainian conflict, and the societal perspective of cyberwar.

Section 3 outlines the methodology we used in this research, including our target selection, data collection, and analysis methods. Section 4 presents the key results of our study. Section 5 discusses the implications of our results, as well as the limitations of our study. Lastly, Section 6 concludes our paper, and explores further research directions.

2. Related Work

From a technical point of view, DoS/DDoS attacks are widely acknowledged as one of the most difficult threats to parry. According to the Ponemon Institute, (D)DoS attacks are the most popular and costly type of cyberattack [5]. Over time, significant efforts have been made to tackle and reduce them; however, the pace of technological innovation constantly introduces fresh attack vectors.

Welzel et al. monitored the Command and Control (C&C) servers of two botnets, DirtJumper and Yoddos, to find out about their targets and to assess their effectiveness [6]. Their evaluation criteria were based on availability, and their approach included DNS monitoring, observation of TCP connection timings, and examination of HTTP responses and contents. They concluded that over 65% of the recorded victims were significantly affected by the attacks.

Wang et al. analysed more than 50,000 separate DDoS attacks over a span of seven months [7]. Their analysis involved the observation of 674 botnets belonging to 23 distinct families. Moreover, their data contained over 9,000 victim IP addresses in relation to 1,074 organisations across 186 countries. They discovered that most of the attacks ran over HTTP. They uncovered some routines linked to certain attacking sources, concretely regarding their geospatial distribution, which could – to a certain extent – help with the prediction of future attacks.

Kalkan et al. highlighted that Software-Defined Networking (SDN) was a relatively recent communication paradigm liable to various security threats, including but not limited to DDoS attacks [8]. They presented a model which increased the SDNs’ resilience against DDoS. Their model employed a joint entropy metric to detect and mitigate them. The central idea is to track the randomness of network packets to discern attacks, since DDoS typically involves great similarities in packet structure over very large traffic volumes. Their model consisted of three phases: nominal, preparatory, and active mitigation. The model profiled the period where no attack was present, made arrangements based on bandwidth properties, followed the detection of suspicious traffic, evaluated it, and chose whether to let it pass or block it, while ensuring uninterrupted benign traffic.

Awan et al. presented an interesting work on detecting DDoS attacks through big data techniques [9]. Their technique involved the use of a random forest and a multi-layer perceptron to predict whether there was an attack at the application layer. Moreover, the approach was able to detect attacks in real time, allowing for an early intervention.

Another approach aimed at detecting DDoS attacks was reported by Doriguzzi-Corin et al., who developed a practical and lightweight deep learning system [10]. They utilised

convolutional neural networks (CNNs) for the classification of network traffic, labelling the traffic as benign or malicious. They claimed that their system brought up a very low overhead in DDoS detection. Additionally, they also affirmed that their solution offered a 40x reduced processing time, compared to the state-of-the-art, while maintaining a matching detection accuracy. It is interesting to note that there are some open questions about whether accuracy is the most meaningful metric in this context of anomaly detection with unbalanced datasets.

Moving on from the technical aspects to a more societal standpoint, it is important to draw attention to the fact that cyberspace has also been a focal point in the scientific community, especially after its obvious eruption as a domain of war. It has been highlighted that modern conflicts quite frequently extend beyond the conventional physical/kinetic (real-world) domain, and spill into cyberspace too [11], [12].

Madnick discussed the potential implications on the future of warfare [13] of the current cyberattacks by pro-Russian parties targeting Ukraine. The author claimed that the age of cyber warfare is just emerging. Additionally, he described the conflict between Russia and Ukraine as a “live testing ground” used by Russians to build their next generation of cyber weapons.

An investigation conducted by Vu et al. on low-level cybercriminals who took sides with the belligerents in the Russo-Ukrainian conflict reported more than 300,000 defacement cases as well as nearly two million reflected DDoS attacks [14]. On top of these acts, the authors also considered data from hack forum posts and announcements. The authors believed that low-level cybercriminals briefly intensified their activities because of the conflict. Furthermore, their observations implied a limited amount of activity by high-profile actors (e.g. state-sponsored events).

Ashraf highlighted that academics have struggled to establish a singular definition for the “cyberwar” concept and underscored that there have been ambiguities due to this situation [15]. The lack of definitional clarity across disciplines makes interdisciplinary research and policymaking on complex issues challenging, as divergent definitions can lead to misaligned outcomes that do not reflect current realities. The author expressed that, in this case, this lack of clarity induced challenges for interdisciplinary research since the established definitions might differ by field. As a result, the author suggested utilising a framework rather than a singular definition, which would provide a baseline for evaluating cyberwar definitions across the literature.

3. Methodology

This section provides a thorough explanation of the methodology used in this study. First, we explain our approach for identifying DDoS targets in the context of the current Russo-Ukrainian conflict. Then, we introduce the data collection and analysis phases, and lastly, we describe the metrics used to evaluate the effectiveness of the DDoS attacks on the observed targets.



Figure 1. Disbalancer.com home page promoting their Liberator tool (captured on 27 December 2023)

3.1. Target Selection

Compiling a large and representative list of targets was a critical part in the assessment of the tug of the DDoS cyberwar. To accomplish that, we first started looking to determine which were the primary channels used by both sides to inform their supporters of the targets to attack. Recording the targets of pro-Ukrainian entities proved to be relatively straightforward, since one of the media they used was the Liberator tool [3], which will be introduced shortly. Nevertheless, identifying the targets registered by the pro-Russian parties was slightly more challenging, since they did not use a well-publicised infrastructure to distribute targets for their attacks.

Our preliminary investigations indicated that the media we needed to track in order to obtain information regarding targets could be narrowed down to a bunch of Telegram channels/groups as well as the Liberator tool mentioned earlier. Figure 1 shows the Liberator tool’s home page. Liberator is a tool provided by a cybersecurity initiative called *Disbalancer* serving as a “collaborative DDoS platform” backed up by volunteering participants on the pro-Ukrainian side [16]. The tool supports all major operating systems, including Linux, Windows and MacOS. It allows users to choose the exact number of threads their devices will allocate to the attack, letting them control the load on their systems. The application runs until the users kill the process. Liberator does not have any persistence mechanism, by default. Voluntary DDoS contributors can start and stop whenever they want. Considering the tool’s public visibility in the media [17]–[19], Liberator was one of the main channels for tracking and monitoring targets.

On top of a dedicated tool like Liberator, Telegram channels and groups also proved quite useful in tracking targets. It is known that some hacktivist groups use Telegram to coordinate and disseminate content and other information [3]. Our procedure involved looking for groups that had been recently advertising targets, which resulted in seven Telegram groups in total: one pro-Russian and six pro-Ukrainian. The main pro-Russian group we identified was all-encompassing in its targeting of victims. Other, smaller, pro-Russian groups showed a tendency to simply copy the targets proposed by this primary group. As a result, we

TABLE 1. TELEGRAM GROUPS CONSIDERED FOR TARGET COMPIATION

	Monitored Telegram Groups
Pro-Russian Entities	Cyber Army of Russia Reborn
Pro-Ukrainian Entities	Cyber Cerber, Cyber Palyanitsa, Haydamaki, Incourse 911, IT Army of Ukraine, Student Cyber Army

monitored this single spearhead group on the pro-Russian side. This is further discussed in Section 5.

Table 1 lists the Telegram groups monitored in this study. It is important to highlight that the sources considered for building the target list are not necessarily exhaustive. At the beginning of our research, we got in touch with “Cyberknow”, a hacktivist observer of the Russo-Ukrainian war (<https://x.com/Cyberknow20>), who provided us with some of the pro-Ukrainian Telegram groups shown in Table 1. We used this information as the base for creating the pro-Ukrainian target set. For the pro-Russian target list, we picked the most popular Telegram group available. Following the registration of targets, their availability and web contents were collected periodically for further evaluation.

For the rest of the paper, we will use specific wording to describe the attacks and the sides in the conflict, in order to improve clarity and avoid any confusion.

- Targets which were aimed to be shut down by Ukraine supporters will be referred to as “targeted by pro-Ukrainian entities”.
- Targets marked by Russian supporters will be indicated as “targeted by pro-Russian entities”.

There are two reasons why we pick such strict terminology. First, the supporters of Ukraine or Russia are not necessarily solely Ukrainian or Russian (e.g. Belarus is supporting Russia [20] and the US and most EU countries support Ukraine [21]). Second, the targets specified by either side have not necessarily declared support to the opposite side – the only constant here is that both sides are registering new targets daily.

3.2. Data Collection

The data collection procedure in this study consisted of two main phases. The first was compiling a list (as exhaustive as possible) of the targets distributed by either side; this was carried out between 28 November 2022 and 2 April 2023. The second phase involved collecting responses to our recurring HTTP requests, to document the state of the targets and evaluate the effectiveness of the attacks. This was conducted between 28 November 2022 and 15 April 2023, including an additional 13 days to monitor existing targets. The technical architecture employed in the data collection phase included five main components: central data repository, restricted Liberator unit, fake DNS server, Telegram channel observer, and monitoring unit.

Figure 2 illustrates the data collection architecture. In this figure, the orange arrows represent the gathering of




Figure 2. Representation of the technical architecture of our work

targets (first phase), while the green arrows represent the collection of HTTP responses from targets (second phase). The two units presented in Figure 2 will be explained in further detail in the following paragraphs.

We used MongoDB as our central data repository, to store all the metadata and HTML content obtained from the targets. The main reason we chose MongoDB was its built-in compression capabilities, which were required by the large volume of data we collected¹. MongoDB’s WiredTiger storage engine (<http://source.wiredtiger.com>) provided efficient Zlib compression, crucial during data collection [22].

The restricted Liberator component was utilised to obtain registered target information and push this information to our central data repository along with the relay information indicating which Liberator server was disseminating the target information. For this, we ran the official distribution of the Liberator application in a Docker container while monitoring and parsing the network traffic in the container in which it resided. To avoid any kind of contribution to the DDoS attacks which were anticipated to be launched by the app on starting up, the Docker container’s default DNS server always returns the loopback IP address (127.0.0.1) as the A record whenever there is a demand for target resolution. Hence, the flood attack originating from the Liberator application never leaves our servers and, in particular, never reaches the intended target. We deployed five of these units in five different countries: China, Poland, Sweden, the United Kingdom and the United States. During our study, it was observed that our restricted Liberator units gathered targets from four main C&Cs, which were hosted on Amazon Web Services, Namecheap, and Hosting Ukraine (<https://www.ukraine.com.ua/uk/>).

The fourth component of our system was the Telegram channel observer. This was used to monitor Telegram channels from both the pro-Ukrainian and pro-Russian sides. To place Telegram chat hooks and fetch the targets when a message arrives, we used the Telethon library [23] for Python. A different script was run in parallel to monitor each side. During the data collection period, these scripts

1. Initially we used MongoDB mostly because we were keeping a large amount of data and it provides built-in compression capabilities. However, when the data collection was finished, it was somewhat inefficient to query it. This was why we later decided to shift to MySQL. In the future, it may be better to store the metadata in SQL-based solutions from the beginning, restricting MongoDB only to store bulky contents such as HTML responses.

parsed all URLs in the groups using regular expressions. Following the data collection, we removed false-positive URLs (such as global news and social media websites) from the target list, stopping them from making it into our central data repository for further consideration.

The monitoring unit was responsible for fetching the target list from the central data repository, periodically crawling those targets, and pushing the collected data. The HTTP crawler element in the monitoring unit was implemented in Go, with browser automation. The collected HTTP data had the purpose of getting the webpage contents as shown to a legitimate user. To avoid any restrictions during crawling (e.g. our requests being blocked), our crawler tries to mimic a typical user browsing the web, instead of a bot crawling the content. To replicate this behaviour, we tried two mainstream browser automation libraries in Go.

In our initial attempts, we tested *Playwright* due to its multi-browser support (Chromium, Firefox, Webkit) and also its community presence [24]. Nevertheless, due to the myriad of targets that needed to be crawled, we required an extensive amount of parallelism. Additionally, independent browser instances were needed while crawling to isolate previous activities for each of the targets, aiming to prevent side effects by cookies or caches. Providing this functionality using *Playwright* was infeasible with our limited infrastructure, primarily due to high memory and storage needs. As a result, in our final implementation, we decided to utilise the *chromedp* library [25] (using simple GET requests), which uses Chrome debug protocol to automate the Chrome executable installed in the current system. The HTTP crawlers had a 30-second timeout after a request, waiting for the response and then recording it, if any. The monitoring units' target was to maintain an approximate 15-minute interval between each data point for every target and server for 14 days. Yet, there were negligible deviations in some cases, such as when simultaneously checking several targets, which required parallelism and caused an overwhelming load on our systems. In total, we deployed six of these monitoring units across four countries.

We employed servers based in Germany, Netherlands, Turkey, and Russia for the collection of HTML contents from the targets. Initially, we also intended to use servers located in Ukraine for our study, and several tests were conducted on servers obtained from a Ukrainian service provider to explore this possibility. However, due to availability problems, we could not proceed with the servers in Ukraine. It is important to spread the locations of our data collection servers. Four of our servers were in the EU area: one in the Netherlands and three in Germany. The reason we picked two different EU countries was due to the possibility of filtering or censorship against any of those countries. We decided to use multiple replicas of the German servers to introduce some degree of redundancy in our system, assuring a more reliable uptime. Besides the EU, we chose Turkey because of its more neutral standing in the conflict. We had one server there. We chose Russia and deployed one server there to monitor if any kind of geolocation access rules were applied.

Considering the potential for measurement errors, and the need to avoid being blocked, we selected servers in different countries to host the *Liberator* and monitoring units, to mitigate those risks. Additionally, having multiple servers deployed globally improved the robustness of our measurement infrastructure. In the (infrequent) case of targeted websites being online but one or two of our servers being down, other servers would still collect valid data, and we could still use these data points in our measurements.

3.3. Data Analysis

The data collected in this study required multiple filtering and processing stages before it could be properly analysed. These procedures include the creation and utilisation of metadata, data cleaning, clustering, and the detection and handling of false positives.

The large amount of content stored in our MongoDB database required an efficient way to query and evaluate the data. The indexes built in MongoDB were considered first because no more write operations would be made on the database – index data structures slow down write operations. However, this initial attempt failed because query times became too slow and needed to improve significantly. As a result, we created a MySQL metadata table containing all the information in the MongoDB collection, except for HTML contents. This metadata allowed us to query and filter quickly and efficiently.

Regarding data cleaning, we observed that some targets obtained from the *Liberator* or Telegram groups were duplicated, invalid, or not at all being related to the war, hence requiring elimination. We employed the following process to deal with data cleaning. First, we stripped out any whitespaces and we converted the remaining characters to lowercase, then we discarded any duplication. A manual inspection followed this operation; we observed occurrences of arbitrary numbers, a whole IP range, a date, an email address, a lexically correct case with an invalid TLD, and some websites that, in principle, seemed not related at all to the war. After this, the remaining targets were deemed valid and composed of IP addresses and fully qualified domain names (FQDN), which is the minimal form of a valid hostname.

Figure 3 depicts the data cleaning process. Of particular note, we manually inspected the lists and removed 37 sites that were not, in principle, related to the war such as Youtube, Github, Tiktok, Paypal, Twitter, etc. but which were initially included in these lists – see step (c) in Figure 3. Additionally, we removed 86 redundant websites (i.e. websites with the same FQDNs or IP addresses, but with different directories) – see step (d) in Figure 3.

As explained earlier, we made multiple HTTP requests to the remaining valid targets to collect their responses for further evaluation. In this study, we defined the case of a target being up and working properly as a visitor being able to get the intended content. Our definition of a target being up necessitated checking the collected HTML contents served in HTTP response bodies. Even though we




Figure 3. Data cleaning process

frequently received a response body, it did not necessarily mean the target was working properly. Therefore, there might have been “false positives” in our data that must be identified and re-labelled correctly, prior to further investigation. Some cases considered false positives involved resource-unavailable messages by hosting providers, errors in the tech stack used by the website, empty HTML contents or bodies in the response, webpages not loading correctly, and access-blocked pages. Before attempting to detect and eliminate false positives, we decided to cluster the gathered HTML contents for each target, to make this operation feasible. The large size of the collected data, over 1 TB, required an efficient method to cluster them for each target. To cluster data, we developed and used a custom Golang implementation of the python package *html-similarity* [26]. We opted for Golang due to performance reasons. During the package development in Golang, some logic enhancements were also implemented to reduce the number of redundant operations, for further optimisation. The similarity score we used in clustering was 85%. We also tested a threshold of 90% initially. However, this score resulted in the separate clustering of similar contents that should have been clustered together. Hence, we lowered the threshold and manually inspected the results to check they were in alignment with those an expert human could produce.

Following the clustering operation for each target, the results needed to be examined and labelled. For this, three parties were involved: two of the investigators and GPT. The manual labelling was initially done by one of the investigators on a Flask-based web application, which displayed raw HTML content alongside its rendered version for each cluster and target. Some cases required translation; therefore, we included a utility in our tool powered by the Google Cloud Translation API. Error pages provided by DDoS protection services (e.g. Cloudflare), custom HTTP 4XX and 5XX pages, errors due to tech stack (pages indicating errors related to databases, frameworks, coding issues), pages with no content (empty HTML body), pages displaying “blocked” messages and defacements were marked as false positives during re-labelling. Pages indicating bot checks by services like Cloudflare and DDoS-Guard were not labelled as false positives because these services provide explicit messages

TABLE 2. CATEGORISATION OF TARGETS

Category	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
Business	110 (29.26%)	209 (42.31%)	319 (36.67%)
Government	46 (12.23%)	90 (18.22%)	136 (15.63%)
Education	53 (14.10%)	17 (3.44%)	70 (8.05%)
Travel	5 (1.33%)	64 (12.96%)	69 (7.93%)
IT	18 (4.79%)	32 (6.48%)	50 (5.75%)
News & Media	22 (5.85%)	16 (3.24%)	38 (4.37%)
Shopping	16 (4.26%)	16 (3.24%)	32 (3.68%)
Finance	10 (2.66%)	16 (3.24%)	26 (2.99%)
Sports	24 (6.38%)	0 (0.00%)	24 (2.76%)
Organisation	17 (4.52%)	6 (1.21%)	23 (2.64%)
Entertainment	8 (2.13%)	2 (0.40%)	10 (1.15%)
Other	47 (12.50%)	26 (5.26%)	73 (8.39%)
Total	376	494	870

when the host is having problems. Lastly, in the cases where the data fetched from the targets indicated no service or content initialised (e.g. default NGINX or Apache pages), the decision was based on the current status of the target when the data analysis was done. If the target was persistently in the same uninitialised state, its label remained the same. Otherwise, it was re-labelled as a false positive.

On top of manual labelling, we chose to utilise OpenAI’s *gpt-3.5-turbo-16k-0613* [27] model since large language models (LLM) have flourished recently and have proven very helpful in automatising many tasks [28]. Additionally, the exceptional language processing abilities of GPT have already become a point of interest in the scientific community [29], [30]; we wanted to investigate and provide some insights into whether HTML document analysis is a good use case for the current state of GPT. Due to the context length of the model we use (it supports up to 4,096 tokens [27]), most of the HTML contents needed to be split into chunks before being fed into the model. Even though we used a model (i.e. *gpt-3.5-turbo-16k-0613*) with a larger context window, this issue was unavoidable due to the length of HTML bodies. Excluding the prompt, which was prepended to each chunk, we chose a 2,000-token window length. The token calculations were made using the Tiktoken utility with *cl100k_base* encoding [31]. The prompt we fed to GPT included the directive of “act like a cybersecurity analyst trying to understand DDoS attacks” and clarified what cases we define as “false positives”. The “majority vote” approach was applied to determine the final decision for the contents.

Inevitably, there were cases in which the investigator and GPT yielded different opinions. Overall, the similarity score between the investigators’ and GPT’s labels was 87.45%. To resolve those disagreements, a third party (another investigator), was involved in the process of establishing a consensus. The differences were observed in 12.54% of the data (n=504), and out of these, the human third party sided with GPT in 37 cases (7.34%), and with the first investigator in 467 cases (92.66%). No specific patterns (e.g., certain URLs or domains) were observed with regard to these differences.

TABLE 3. RECURRENCE AMONG TARGETS

Recurrence (days)	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
1	606 (95.73%)	325 (52.08%)	931 (74.07%)
2	23 (3.63%)	101 (16.19%)	124 (9.86%)
3	3 (0.47%)	45 (7.21%)	48 (3.82%)
4	0 (0.00%)	50 (8.01%)	50 (3.98%)
5	0 (0.00%)	28 (4.49%)	28 (2.33%)
6	0 (0.00%)	11 (1.76%)	11 (0.88%)
7	0 (0.00%)	18 (2.88%)	18 (1.43%)
8	0 (0.00%)	7 (1.12%)	7 (0.56%)
9	0 (0.00%)	2 (0.32%)	2 (0.16%)
10+	1 (0.16%)	37 (5.93%)	38 (3.02%)

Another process that needed to be done on the targets was categorisation to understand any temporal trends in target selection. For this purpose, we initially used the VirusTotal API [32], which compiles information from various vendors’ databases regarding website categories. The information collected through the API had some missing points and inconsistencies. Hence, we also ran a manual check on this data, to improve target categorisation accuracy.

Since this study aimed to evaluate the effectiveness of DDoS attacks, we needed to devise a metric. Hence, we introduced the metric considered in this study: Quality of Service (QoS). It was calculated as follows: Data points representing a binary state (1 for up and 0 for down) for each target collected by our servers (five worldwide) were aggregated. Then, using the aggregated data, hourly binary states were averaged. Finally, each target’s 24-hourly QoS scores were averaged into its daily QoS score.

To prevent issues related to packet loss (e.g., slow network connections) on our end that could adversely affect our results, we have deployed multiple servers worldwide. This approach improves the reliability of our measurement infrastructure, which is introduced in detail in Section 3.2. If targeted websites were online but one or two of our servers were down, other servers would still collect data, and we used these data points in our measurement.

4. Results

In this section, our main findings are presented. First, the targets are categorised based on the sectors they belong to, along with their recurrence (see Section 4.1). Second, the effectiveness of the DDoS attacks on recorded targets is shown. Third, DDoS protection service utilisation is presented, and lastly, whether connecting to the targets from Russia changes access is also explored.

4.1. Targets

Between 28 November 2022 and 2 April 2023, our trackers recorded 1,388 candidate targets in the channels we observed. After the sanitisation process (mainly removing near duplicates, i.e. the same hostname but different directory), 1,257 targets were left. The remaining targets were composed of 870 FQDNs and 387 IP addresses.

TABLE 4. EFFECTIVENESS OF DDoS ATTACKS

	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
Completely Shut-Down	36.18% (n = 229)	25.00% (n = 156)	30.63% (n = 385)
Affected	62.56% (n = 396)	73.72% (n = 460)	68.10% (n = 856)
Completely Unaffected	1.26% (n = 8)	1.28% (n = 8)	1.27% (n = 16)
Total	100.00% (n = 633)	100.00% (n = 624)	1,257

To see whether there were any trends regarding victimisation, we checked 870 FQDNs in our data in terms of “categories”. When the categories that the targets reside in are considered, we had 12 in total. We observed that both sides primarily targeted businesses. Except for businesses, only the “Entertainment” category had the same ranking of importance for both sides (11 out of 12). Overall, we observed a statistically significant difference in the distribution of categories when picking targets ($\chi^2(11, n = 870) = 145.75$, $p < .001$). This implies that the two sides had noticeably different priorities regarding targets. Table 2 summarises the statistics regarding the sides and targets’ categories.

When we studied target recurrence, where recurrence is defined as being registered multiple times as a target, we observed rather interesting outcomes. The pro-Ukrainian entities registered targets more repetitively, compared to the pro-Russian entities. We believe one possible reason is that the number of channels we tracked was possibly more extensive on the pro-Ukrainian side. The repetitive target registrations might have also impacted the attack duration: on average, targets of the pro-Ukrainian entities were monitored for 25.25 days, whereas this number was 14.47 for the pro-Russian side.

In total, we recorded 3,235 valid target registrations. Of those, 2,564 (79.26%) belonged to the pro-Ukrainian side, and 671 (20.74%) belonged to the pro-Russian side. About a quarter of targets registered by the pro-Ukrainian side were unique, while this rate was about 95% for the other side.

The descriptive statistics and the daily distribution of target registration are shown in Table 3 and Figure 4. It appeared that the pro-Russian entities were more organised (having a central team to direct the effort) than the pro-Ukrainian entities (which would follow a more ad hoc and grassroots movement approach, even though they have the Liberator tool, which is meant to help non-technical people get involved in the DDoS volunteering effort).

4.2. Effectiveness of the DDoS Attacks

Table 4 presents a summary of the effectiveness of the DDoS attacks observed in our study, showing that the attacks affected an overwhelming majority of targets, regardless of their side on the conflict. Overall, only a tiny portion, a mere 1.27% of all targets, had a QoS score of 100%, indicating they were not affected to any extent.




Figure 4. Daily registered targets by side

However, when we focus on the exact opposite outcome, hence on targets that went completely down, the ratio was quite significant: almost one-third. When the ratios of the targets entirely shut down by each side were compared, we observed that the ones targeted by pro-Russians were affected more severely. Our findings implied that there

is statistically significant evidence supporting that getting completely shut down depends on the side you are in the conflict ($\chi^2(1, n = 1,257) = 17.95, p < .001$). In contrast, we found no evidence for remaining totally unaffected by the attack depending on the sides ($\chi^2(1, n = 1,257) = 1.3332e-28, p > .05$).

TABLE 5. QoS SCORES BY CATEGORIES

Category	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
Business	82.43%	71.92%	75.54%
Government	72.33%	53.89%	60.12%
Education	26.21%	77.28%	38.61%
Travel	94.51%	52.46%	55.51%
IT	81.04%	78.52%	79.42%
News & Media	84.41%	69.27%	78.03%
Shopping	70.17%	71.45%	70.81%
Finance	79.72%	66.58%	71.64%
Sports	87.23%	N/A	87.23%
Organisation	76.21%	88.68%	79.46%
Entertainment	80.58%	49.61%	74.38%
Other	72.40%	55.23%	66.28%

TABLE 6. DESCRIPTIVE STATISTICS REGARDING QoS

	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
Count	633	624	1,257
Mean	52.720	56.089	54.392
Mode	0.000	0.000	0.000
SD	46.041	44.632	45.361
Min	0.000	0.000	0.000
25%	0.000	0.002	0.000
50%	78.843	80.377	80.029
75%	98.065	98.716	98.369
Max	100.000	100.000	100.000

Moving away from these extremes, namely being completely shut down or unaffected, around two-thirds of all targets (68.10%) indicated varying degrees of QoS issues.

The QoS metric which we used to evaluate the effectiveness of the attacks had a slight negative skewness for both sides, as well as the targets combined (-0.18 for the ones targeted by pro-Russian entities, -0.30 for pro-Ukrainian side targets, and lastly, -0.24 combined).

Table 5 shows the QoS scores across multiple categories. An explanation regarding how the QoS scores were calculated is provided in the final paragraph of Section 3.3. As can be seen in this table, various degrees of damage were inflicted on the top categories of targets. A significant point of interest for both sides is that businesses did relatively well against the attacks, with a combined QoS score of 75.54%. Governmental websites were not as resilient as business websites, especially the ones targeted by pro-Ukrainian entities. However, the most critical point to highlight regarding the top three categories was the education category. Specifically, targets of pro-Russian entities suffered severely, indicated by almost one-quarter of the QoS score.

Table 6 provides descriptive statistics of the QoS metric. The “Count” values in Table 6 were obtained simply by distinguishing the targets, whether they were targeted by pro-Russian or pro-Ukrainian entities. The “Mean” values were calculated by averaging the daily QoS scores for all of the respective target groups (pro-Russian or pro-Ukrainian), for the whole observation period. The rest of Table 6 was computed in a similar way to how the “Mean” values were.




Figure 5. Weekly Quality of Service scores

When the distribution of the QoS scores of both sides was compared, a statistically significant difference was observed, implying a lower score for the targets of pro-Russian entities (Mann–Whitney $U = 181473$, $n_1 = 633$, $n_2 = 624$, $p = .006$ one-sided). Hence, we can conclude that the pro-Russian attacks were more successful, in general, than the pro-Ukrainian ones.

Figure 5 displays our success rate in accessing the content provided by the targets on a weekly basis. Although there were fluctuations, the targets attacked by pro-Russian entities depicted a more “stable” QoS score, as the trend-line also supports. When we examined the targets of pro-Ukrainian entities, starting from the ninth week of 2023, their QoS score made a jump, implicating that their attacks were becoming less effective. Lastly, it is worth emphasising that when the weekly scores are considered alone, it could be easier for one to interpret that the targeted by pro-Ukrainian entities had a slightly better QoS score that might indicate their attacks were more effective; however, a critical remark here is that each week had a different weight, which affects the overall result significantly.

When the post-attack states of the affected targets were investigated, we observed a great degree of variance. The most prevalent sign of failure we encountered was timeouts. Overall, they comprised slightly less than two-fifths of the data points where our attempts to access the content hosted on the targets failed. Even though timeouts were the most frequent case when both sides were combined, they were particularly relevant on the ones targeted by pro-Ukrainian entities, with more than half of the failure data points caused by them. Timeouts were the second most prevalent issue for the ones targeted by pro-Russian entities, which came after HTTP 4XX errors that implied client-side errors such as content not being found, unauthorised or restricted access, and too many requests to a server. These two causes of failure composed more than half of the data points where access attempts were unsuccessful overall.

After timeouts and HTTP 4XX errors, other observed types of errors were less prominent; each had a frequency of less than 10%. Overall, “ERR_NAME_NOT_RESOLVED” took third place, indicating a DNS error where the IP address corresponding to a hostname could not be determined. This

TABLE 7. OBSERVED REASONS FOR FAILURE AFTER DDoS ATTACKS

	Targeted By		
	Pro-Russian Entities	Pro-Ukrainian Entities	Combined
Timeout	16.35%	50.92%	38.79%
HTTP 4XX	29.02%	11.44%	17.61%
ERR_NAME_NOT_RESOLVED	7.57%	3.95%	5.22%
ERR_CONNECTION_REFUSED	4.10%	4.32%	4.24%
Unavailable Message by Hosting Provider	6.85%	2.54%	4.05%
Cloudflare 1003	4.38%	2.98%	3.47%
HTTP 5XX	4.08%	2.04%	2.76%
Not Configured (Nonpersistent)	1.75%	3.08%	2.61%
Tech Stack Related Error	6.74%	0.30%	2.56%
ERR_ADDRESS_UNREACHABLE	2.68%	1.85%	2.14%
Empty HTML Body	3.63%	1.29%	2.11%
DDoS-Guard: Service Not Identified	0.00%	3.00%	1.95%
Unable to Load	1.98%	1.81%	1.87%
Cloudflare 1020	2.35%	1.46%	1.77%
ERR_EMPTY_RESPONSE	0.91%	1.95%	1.59%
Moved	0.00%	1.57%	1.02%
StormWall	0.00%	1.20%	0.78%
ERR_CONNECTION_RESET	1.18%	0.40%	0.68%
Locked by Hosting Provider	0.66%	0.46%	0.53%
Other	5.78%	3.42%	4.25%

error was more frequent in the ones targeted by the pro-Russian entities, almost doubling the other side. The other types of errors included browser network errors [33], tech stack-related errors, DDoS protection service errors, hosting provider-related errors and others, including unspecified errors, redirections, and getting blocked due to attacks.

Table 7 shows a breakdown of the types of errors we encountered. One potential speculation here is that the targets compiled by pro-Ukrainian entities tend to be repeated (i.e., multiple pro-Ukrainian volunteers might register the same Russian targets). This would then cause the high “Timeout” percentage value of those targets identified by pro-Ukrainian entities since some of the Russian websites were already down by the time further effort was made to reach them. Another speculation is that different hosting suites, services, and infrastructures might have different failure messages or even custom error messages. Encountering DDoS protection service messages (e.g. Cloudflare, DDoS-Guard) may indicate that a target has recently adopted such services, but has not fully configured their system during our monitoring.

4.3. Utilisation of DDoS Protection Services

On top of attack effectiveness on targets, we also investigated whether they used any DDoS protection services, and if they did before the attack or started using those services afterwards. We observed indicators of utilisation of these services. The “indicators” we considered were based on the keywords/patterns present in the HTML contents we collected, which were also used to evaluate the effectiveness of the attacks. Although this may not be an exhaustive list of indicators, when combined with our later manual inspection (in which we did cluster HTML contents), we are confident

```
center">Performance & security by <a rel="noopener noreferrer"
href="https://www.cloudflare.com?
utm_source=challenge&utm_campaign=m"
target="_blank">Cloudflare</a></div> </div> <span
id="trk_jschal_js"></span></body></html>
alt=""><div id="description">We have detected suspicious network
activity from your device.<br>Please verify to continue.</div></div>
<div id="link-ddg"><div class="logo-foot"></div><span>Protection and Acceleration by</span> <a
href="https://ddos-guard.net" target="_blank" id="link">DDoS-
Guard</a></div></div><script>var
class="attribution">Performance, security and DDoS protection by <a
href="https://cloud-shield.ru/?from=iua-captcha-en"
target="_blank">Cloud-Shield.ru</a></div></td></tr></tbody></table>
<div class="gorizontal-vertikal"></div> <script> var
utm_set = null; function setup_utm(){ if (utm_set == null) return null;
```

Figure 6. Four examples of HTTP response to our query for determining the DDoS protection service utilised by a target

TABLE 8. STATUS OF DDoS PROTECTION AMONG TARGETS

	Provider	Targeted By		
		Pro-Russian entities	Pro-Ukrainian entities	Combined
Had from day one of the attack	Cloudflare	61	28	89
	DDoS-Guard	2	19	21
	Cloud-Shield	0	2	2
	StormWall	0	6	6
Obtained after the attack	Cloudflare	22	11	33
	DDoS-Guard	0	16	16
	Cloud-Shield	0	0	0
	StormWall	0	1	1
All time	Total	85	83	168

that we have included and classified with high accuracy the vast majority of the DDoS protection services in use.

Figure 6 displays four examples of these indicators. Except for Stormwall, DDoS protection companies would typically provide a text highlighting that the targeted websites were protected by their services. However, Stormwall only displayed a loading animation hosted on their end during the security checks performed.

Based on the data we collected, we found that 168 unique targets out of 1,257 (13.37%) had clear signs of employing DDoS protection services. Our investigation, somewhat surprisingly, showed that the targets had been getting services from only four providers: Cloudflare (US-based), DDoS-Guard (Russia-based), Cloud-Shield (Saudi Arabia-based), and lastly StormWall (Slovakia-based).

Beyond considering what percentage of the targets used DDoS protection, we also examined the adoption evolution of these anti DDoS countermeasures. We found that 118 (70.24%) of the targets had been using DDoS protection services since they were registered as targets, and only 50 (29.76%) opted to obtain protection after the attacks. The ones targeted by pro-Russian had a relatively higher rate (53.39%) of DDoS protection service usage on the first day of the attacks than those targeted by pro-Ukrainian (46.61%). In contrast, the ones targeted by pro-Ukrainian entities had a higher later adoption rate (56.00%) of protection services compared to pro-Russian entities’ targets (44.00%). Cloudflare, an American company, was the most frequently used service provider for both sides of the conflict.

TABLE 9. A SUMMARY OF GEOLOCATION-BASED ACCESS RESTRICTIONS FOR PRO-RUSSIAN ENTITIES

Access Restriction Type	Count	Percentage
Allowed Only to Russia	19	4.50%
Dynamic Access Restrictions	53	12.56%
No Access Restrictions	350	82.94%
Total	422	100%

Table 8 provides a summary of the DDoS protection usage, including the breakdown of which DDoS protection providers being used (and in how many instances) by each side of the conflict, before and after the attack.

Lastly, we observed only two targets that changed anti-DDoS providers during our study. Both of these shifted to Cloudflare from DDoS-Guard. One of these cases occurred on the first day being registered as a target. The other case was more interesting: on day two, after receiving the first wave of attacks, one of the targets obtained service from DDoS-Guard and two days later switched to Cloudflare.

4.4. Geolocation-Based Access Restrictions for Pro-Russian Entities

We observed that some pro-Russian websites implemented geolocation-based access control measures. According to a Nikkei Asia article [34], websites associated with the Russian government strategically restrict access from foreign locations as a safeguard against cyberthreats, including DDoS attacks. This, of course, can be easily circumvented by using proxies or VPNs, but could deter some unsophisticated attackers. Out of the 1,257 targets that we introduced in earlier sections, we were able to collect data for 883 in our Russia server. Of these, 422 were identified as targets by pro-Ukrainian groups and remained accessible from our Russian server on at least one occasion.

Table 9 shows descriptive statistics of pro-Russian websites regarding their access policies by geolocation. An overwhelming majority of the targets (82.94%) indicated no sign of access restrictions by the location we attempted to access. In addition, we found out that only 4.50% of these websites were exclusively accessible from within Russia. These were predominantly associated with governmental functions or related to the travel industry.

The remaining websites (12.50%) employed adaptive restriction policies, periodically allowing access solely to Russian users before subsequently becoming accessible to the broader internet audience. Typically, these websites tend to tighten their access restrictions in response to DDoS attacks and relax them once the attack has subsided. The majority of these websites spanned categories such as business, government, and finance. A smaller portion comprised news organisations, travel agencies, shopping platforms, and educational institutions.

5. Discussions & Limitations

As already discussed in the results section, the pro-Russian side displayed less recurrence during target regis-

tration. Although their target registrations were about one-fourth of the pro-Ukrainian side, the count of unique targets was almost identical. Consequentially, this might have caused a longer duration of attacks against pro-Russian websites. However, that does not necessarily mean those targets suffered more, as shown by the effectiveness of the attacks presented in the results. One possible reason for recurrence might be that a target previously registered was not successfully shut down. As a result, it has been re-targeted to increase the likelihood of completely shutting down the website. Such cases can be seen in the timelines provided in Appendix A.

The sides’ interests – in terms of target categories – did not show much similarity. In the category distribution, “News & Media” had a relatively low significance for both sides. We believe this is an interesting piece of information, since these tools can easily be utilised for mass propaganda [35], [36], such as promoting an ideology and spreading misinformation [37]. Moving on from that, the prevalence of businesses and governmental targets made sense because it could serve multiple purposes, such as disrupting the services provided or demonstrating power. Regarding businesses, we observed some cases where online arms dealers were being particularly targeted. The reason was probably slowing the troops’ recruitment process down because it was reported that frequently under-equipped Russian soldiers had to buy their own gear [38], [39]. Another interesting observation was the frequency of the education category. The pro-Russian entities targeted them four times more frequently compared to the pro-Ukrainians. Although the motives are unclear, the reason might be the dissemination of propaganda or undermining public confidence.

During the data analysis, we observed that some targets were also defaced in addition to becoming DDoS targets. These defacement attacks were claimed by a hacking group, which we tracked to learn from their other targets. Those incidents show that the fight in a cyberwar is by no means limited to DDoS attacks.

The global sanctions and embargoes on Russia meant economic isolation and trade restrictions. However, we observed that some of the pro-Russian targets were able to obtain DDoS protection services from foreign countries, including the United States and the EU. This may be in flagrant violation of the sanctions imposed on Russia. If anything, this observation shows the many challenges regarding the applicability of sanctions in cyberspace.

GPT’s use to detect false positives by checking the response bodies obtained from the targets could be considered a promising approach. Even though it performed relatively well and displayed the potential of LLMs, we could not recommend it for labelling HTML content, in its current form. The main issue, we believe, was the limited context window length. Lengthy HTML documents caused the GPT to quickly get out of context, which resulted in wrong evaluations. Considering the complexity of modern websites, it would not make sense to expect them to get any shorter. In contrast, expecting them to get longer would be logical due to the modern one-page application structure and

their increasing complexity. Hence, this issue might remain to be a challenge in the future.

Nearly one-third of the targets (30.63%) were entirely shut down, meaning they were not functional when we started to monitor them. We believe this implies multiple different scenarios for both sides of the conflict. First, the targets might have been registered after the attacks upon them had already been launched. This case is purely about re-victimisation: it is possible that a downed website might be resurrected in the future, therefore it is being added again in anticipation that it might need to be taken down again. Second, the Liberator tool has several servers across the globe, and there might be some delay before all of them are updated with the new targets. Furthermore, DDoS attacks might take some time to succeed (or be propagated among the Liberator servers). In the case here, some volunteers might have added duplicate entries to the target list, which are then shown as already down when we observed them. Third, similar to the Liberator case, some attackers in a Telegram group might have started an attack on a target, but they felt like they needed more help to take the target down, so they added this target to the list again to encourage more people to attack it.

The most significant limitation of this study was the absence of data collected from a Ukrainian server. Unfortunately, we could not deploy our infrastructure there due to technical difficulties directly linked with the war, such as general availability problems and payment issues. However, considering the political outlook, we believe that even if we had the data collected from the Ukraine, it would not yield unique cases like the data collected from Russia. This is because the world is mostly united around Ukraine, and Russia is left almost alone and isolated, yielding a more distinctive position in the conflict. Furthermore, another imbalance was the number of the Telegram groups tracked to identify targets by both sides. However, despite this imbalance, the final numbers of targets per side were almost identical. Additionally, tracking other smaller groups in the pro-Russian side would yield little to no value since they tend to disseminate the same targets. Another limitation was utilising GPT-3.5 despite GPT-4 being available, due to limited resources. In a similar use case, GPT-4 would likely yield a slightly better performance. Our study did not include a “control set” to compare the availability of non-targets with those that had been targeted. Having a large control set could be beneficial in future work.

Finally, there are some inherent limitations when researching in the middle of a war, because information is scarce, misinformation abounds, and basic infrastructure turns unreliable, so the data collection phase presented more challenges and limitations than those of similar studies between non-warring countries.

6. Conclusions

This study investigated the impact of DDoS attacks that emerged from the latest Russo-Ukrainian conflict. The initial question we tried to grasp was how effective and

numerous the attacks were, and whether they triggered geolocation-based access policies or any other defensive countermeasures and which ones were most successful. The effectiveness evaluation was based on the QoS metric, which represented the extent to which our attempts to access content provided by a target were successful. Between 28 November 2022 and 2 April 2023, we recorded 1,257 valid targets and monitored their status for 14 days to evaluate the attacks’ effectiveness. The distribution of unique targets among the sides was fairly balanced despite the pro-Ukrainian side being four times more prolific in target registration: 624 were targeted by pro-Ukrainians, and pro-Russians targeted 633. Upon those targets, we observed that the pro-Russians conducted more effective attacks.

Both sides mainly targeted business websites. Governmental websites followed them. Out of these, we observed that businesses were more resilient against attacks than government sites. Among all, pro-Ukrainian educational targets displayed the least resilience against attacks with only a 26.21% QoS score.

Another point we observed was low rates of DDoS protection utilisation by the targets, including both from the beginning of the attacks against them and later adoption.

Lastly, we observed an indication of geolocation-based access policies in nearly one-fifth of pro-Russians. Most of these cases involved dynamically changing policies, whereas a small portion involved permanently restricting access from outside Russia.

For future work, it will be beneficial to do the tracking – and to collect pertinent data – from one or more Ukrainian servers. This can reveal valuable insights regarding how a country’s internet infrastructure might perform while under an intense pressure caused by war. In addition to this, it will also be useful to compile more data sources for targeting the attacks (i.e. servers being used for coordinating the DDoS attacks) and look into detail how they identify and prioritise their targets.

With the recent advances in LLMs (such as ChatGPT), it is also worthwhile to explore the possibility of leveraging LLMs to automate the whole process, from target selection, to crafting the most effective DDoS attack based on what is known about the target.

As a final note, it is important to remind the reader of the increasingly relevant role that DDoS attacks play in cyber-war. For this reason, we believe it is important that countries develop their own anti-DDoS capabilities and infrastructure, so that they do not have to rely, particularly at a time of crisis, on any external provider help that can be unreliable due to changing political alliances or tensions. This should be a matter of national security for any developed country, and it is hence surprising that only four anti-DDoS services were used by both sides, and even more so that some of them were based on countries that had imposed sanctions on Russia. Whether this is in flagrant violation of the sanctions or not is a different matter, but no doubt this goes against their spirit. Anti-DDoS technology should be considered an important cyberweapon, so aiding an adversary with access to yours should not be taken lightly.

Acknowledgement

This work was partly supported by the funding received from the European Commission under the Digital Europe Programme through the Cybersecurity for Industry 4.0 Technologies in Operational Technology (CyberSec4OT) project, Grant Agreement no. 101190037. The authors would also like to thank the anonymous reviewers and the shepherd for their constructive feedback.

References

- [1] D. Serpanos and T. Komninos, "The Cyberwarfare in Ukraine," *Computer*, vol. 55, no. 7, pp. 88–91, 2022.
- [2] P. Forsström, D. Adamsky, D. Minzarari, C. Reach, R. Thornton, M. Miron *et al.*, "Russia's War on Ukraine: Strategic and Operational Designs and Implementation," *Finnish National Defence University*, 2023. [Online]. Available: <https://urn.fi/URN:ISBN:978-951-25-3400-5>
- [3] S. Soesanto, "The IT Army of Ukraine: Structure, Tasking, and Eco-System," *CSS Cyberdefense Reports*, pp. 1–31, June 2022.
- [4] B. Van Niekerk, "The Evolution of Information Warfare in Ukraine: 2014 to 2022," *Journal of Information Warfare*, vol. 22, no. 1, 2023.
- [5] Ponemon Institute, "The Cost of Denial-of-Services Attacks," https://cdn2.hubspot.net/hubfs/5104295/ArxNimbus_April2019/Pdf/the-cost-of-denial-of-services-attacks.pdf, Mar. 2015.
- [6] A. Welzel, C. Rossow, and H. Bos, "On Measuring the Impact of DDoS Botnets," in *Proc. of the Seventh European Workshop on System Security*, 2014, pp. 1–6.
- [7] A. Wang, W. Chang, S. Chen, and A. Mohaisen, "Delving Into Internet DDoS Attacks by Botnets: Characterization and Analysis," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2843–2855, 2018.
- [8] K. Kalkan, L. Altay, G. Gür, and F. Alagöz, "JESS: Joint Entropy-Based DDoS Defense Scheme in SDN," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 10, pp. 2358–2372, 2018.
- [9] M. J. Awan, U. Farooq, H. M. A. Babar, A. Yasin, H. Nobanee, M. Hussain *et al.*, "Real-Time DDoS Attack Detection System Using Big Data Approach," *Sustainability*, vol. 13, no. 19, p. 10743, 2021.
- [10] R. Doriguzzi-Corin, S. Millar, S. Scott-Hayward, J. Martinez-del Rincon, and D. Siracusa, "LUCID: A Practical, Lightweight Deep Learning Solution for DDoS Attack Detection," *IEEE Trans. Netw. Service Manag.*, vol. 17, no. 2, pp. 876–889, 2020.
- [11] North Atlantic Treaty Organization, "Warsaw Summit Communiqué" press release — nato.int," https://www.nato.int/cps/en/natohq/official_texts_133169.htm, July 2016.
- [12] M. Robinson, K. Jones, and H. Janicke, "Cyber warfare: Issues and challenges," *Computers & Security*, vol. 49, pp. 70–94, 2015.
- [13] S. Madnick, "What Russia's Ongoing Cyberattacks in Ukraine Suggest About the Future of Cyber Warfare," *Harvard Business Review*, vol. 7, 2022.
- [14] A. V. Vu, D. Thomas, B. Collier, A. Hutchings, R. Clayton, and R. Anderson, "Getting Bored of Cyberwar: Exploring the Role of the Cybercrime Underground in the Russia-Ukraine Conflict," *arXiv preprint arXiv:2208.10629*, 2022.
- [15] C. Ashraf, "Defining cyberwar: towards a definitional framework," *Defense & Security Analysis*, vol. 37, no. 3, pp. 274–294, 2021.
- [16] D. Black, "The cyber army needs you, urges new app — cybernews.com," <https://cybernews.com/news/the-cyber-army-needs-you-urges-new-app/>, Mar. 2022.
- [17] M. Burgess, "Russia Is Being Hacked at an Unprecedented Scale — wired.co.uk," <https://www.wired.co.uk/article/russia-hacked-attacks>, Apr. 2022.
- [18] A. Baydakova, "Inside the Ukrainian Crypto Startup Waging Cyberwar on Russia — coindesk.com," <https://www.coindesk.com/layer2/2022/03/30/inside-the-ukrainian-crypto-startup-waging-cyberwar-on-russia/>, Mar. 2022.
- [19] A. Sharma, "A YouTuber is encouraging you to DDoS Russia—how risky is this? — bleepingcomputer.com," <https://www.bleepingcomputer.com/news/security/a-youtuber-is-encouraging-you-to-ddos-russia-how-risky-is-this/>, May 2022.
- [20] A. Kudrytski, "Why Belarus Is Backing Russia in Its War in Ukraine — washingtonpost.com," https://www.washingtonpost.com/business/energy/2023/03/27/why-is-belarus-supporting-russia-in-its-war-in-ukraine/291b4074-cc9b-11ed-8907-156f0390d081_story.html, Mar. 2023.
- [21] "EU support for Ukraine — european-union.europa.eu," https://european-union.europa.eu/priorities-and-actions/eu-support-ukraine_en.
- [22] A. Kamsky, "New Compression Options in MongoDB 3.0 — mongodb.com," <https://www.mongodb.com/blog/post/new-compression-options-mongodb-30>, Nov. 2021.
- [23] LonamiWebs, "LonamiWebs/Telethon: Pure Python 3 MTProto API Telegram client library, for bots too! — github.com," <https://github.com/LonamiWebs/Telethon>, 2016.
- [24] playwright community, "playwright-community/playwright-go: Playwright for Go a browser automation library to control Chromium, Firefox and WebKit with a single API. — github.com," <https://github.com/playwright-community/playwright-go>, 2020.
- [25] chromedp, "chromedp/chromedp: A faster, simpler way to drive browsers supporting the Chrome DevTools Protocol. — github.com," <https://github.com/chromedp/chromedp>, 2017.
- [26] E. Marca, "matiskay/html-similarity: Compare html similarity using structural and style metrics — github.com," <https://github.com/matiskay/html-similarity/>, 2020.
- [27] OpenAI, "Models — platform.openai.com," <https://platform.openai.com/docs/models/gpt-3-5>, 2023.
- [28] Z. Xi, W. Chen, X. Guo, W. He, Y. Ding, B. Hong *et al.*, "The Rise and Potential of Large Language Model Based Agents: A Survey," *arXiv preprint arXiv:2309.07864*, 2023.
- [29] J. Ye, X. Chen, N. Xu, C. Zu, Z. Shao, S. Liu *et al.*, "A Comprehensive Capability Analysis of GPT-3 and GPT-3.5 Series Models," *arXiv preprint arXiv:2303.10420*, 2023.
- [30] G. Suri, L. R. Slater, A. Ziaee, and M. Nguyen, "Do Large Language Models Show Decision Heuristics Similar to Humans? A Case Study Using GPT-3.5," *Journal of Experimental Psychology: General*, 2024.
- [31] OpenAI, "How to count tokens with tiktoken — cookbook.openai.com," https://cookbook.openai.com/examples/how_to_count_tokens_with_tiktoken, 2023.
- [32] VirusTotal, "VirusTotal API v3 Overview — docs.virustotal.com," <https://docs.virustotal.com/reference/overview>, 2022.
- [33] Chromium, "net_error_list.h — source.chromium.org," https://source.chromium.org/chromium/chromium/src/+main:net/base/net_error_list.h.
- [34] A. Teraoka and R. Namiki, "Russia walls off government websites from nonfriendly countries — asia.nikkei.com," <https://asia.nikkei.com/Politics/Ukraine-war/Russia-walls-off-government-websites-from-nonfriendly-countries>, Mar. 2022.
- [35] E. S. Herman and N. Chomsky, "Manufacturing Consent," in *Power and Inequality*. Routledge, 2021, pp. 198–206.
- [36] A. Mullen and J. Klaehn, "The Herman–Chomsky Propaganda Model: A Critical Approach to Analysing Mass Media Behaviour," *Sociology Compass*, vol. 4, no. 4, pp. 215–229, 2010.
- [37] C. Fuchs, "Propaganda 2.0: Herman and Chomsky's Propaganda Model in the Age of the Internet, Big Data and Social Media," in *Propaganda Model Today: Filtering Perception and Awareness: Filtering Perception and Awareness*. University of Westminster Press London, 2018, pp. 71–91.

- [38] T. Lister and K. Krebs, “Russians buy boots and body armor for the troops, as the Kremlin tries to fix the campaign’s problems — edition.cnn.com,” <https://edition.cnn.com/2022/12/22/europe/russians-crowdfund-soldiers-ukraine-cmd-intl/index.html>, Dec. 2022.
- [39] “‘We Have to Buy Everything’: Russian Soldiers Under-Equipped In Ukraine War — themoscowtimes.com,” <https://www.themoscowtimes.com/2022/05/20/we-have-to-buy-everything-ourselves-how-russian-soldiers-go-off-to-fight-a77751>, May 2022.

Appendix A.

Timelines displaying daily states of the targets

This appendix includes two representative samples from a series of 14 plots, depicting the timelines of the unique targets being affected by the DDoS attacks considered in this study. A set of seven plots relates to the websites targeted by pro-Russian entities, while another set of seven plots shows those targeted by pro-Ukrainian entities. Due to space limitations, only two plots are included in this paper.

- Figure 7 depicts a set of 100 sample timelines of the targets attacked by pro-Russian entities.
- Similarly, Figure 8 shows a set of 100 sample timelines of the targets hit by pro-Ukrainian entities.
- The x-axis represents the monitored time frame, with red dotted lines indicating the start and end dates of the whole study period.
- The y-axis shows the unique targets.
- Within each plot, the timelines depict the availability of the targeted websites. Red boxes signify that a target was completely down on a specific day (QoS score of 0%), while green boxes indicate that it was operational (QoS score of 100%). The other colours (yellow or orange) indicate a partial availability (i.e. QoS score between 0% and 100%, with an orange box representing a lower QoS score than that of a yellow box).
- Boxes marked with a cross denote that the target (specified in the y-axis) was registered on those particular days (specified in the x-axis).

The full set of these 14 plots is available on the following GitHub repository: <https://github.com/Assessing-the-Silent-Frontlines/Exploring-the-Impact-of-DDoS-Hackivism-in-the-Russo-Ukrainian-War>.

We observed the 1,257 targets (633 targeted by pro-Russian entities, and 624 targeted by pro-Ukrainian entities) for at least 14 days each. The starting point of our observation of each target depends on when that target was being registered. For instance, in Figure 7, TARGET-001 was registered on 18 December 2022, while TARGET-018 was registered on 28 November 2022. In some cases, the same target might be re-registered later on, which extended the observation period of that target (for example, see TARGET-95 in Figure 7, which was initially registered on 5 February 2023, but later on, it got registered again on 14 February 2023 and 18 March 2023). Each registration would trigger a 14-day observation period being added.

To provide readers with a better understanding of how to read and interpret the additional plots presented in Figures 7 and 8, we provide highlights of a selection of interesting cases. As presented in Section 4.1, we observed that some targets had been registered multiple times in a “recurring” fashion, especially those targeted by pro-Ukrainian entities. In those cases, different scenarios were observed:

- The attacks had no impact at all (e.g., see TARGET-079 in Figure 8, between 30 January and 21 February 2023, where the target stayed green all the time).
- The attacks had minimal impact on the target and were unable to shut it down, even though the target had been repeatedly registered for further attacks (e.g., see TARGET-015 in Figure 8, between 14 December 2022 and 11 January 2023, in which there were variations in the QoS of the target, but the target largely stayed green, with one yellow).
- The attacks had a considerable impact, and this state was preserved by additional waves of attacks (e.g., see the timeline of TARGET-097 in Figure 8).
- The attacks were able to shut down the target completely for an extended period (e.g., see TARGET-074 in Figure 8, between 24 February and 31 March 2023, where the target stayed red throughout).
- The initial attacks did not succeed; yet another wave with another registration of the target had a considerable effect (e.g., see TARGET-100 in Figure 8, between 4 December and 28 December 2022).
- The initial attacks took several days to impact the target significantly, and subsequent registrations maintained a low QoS score of the target (e.g., see the timeline of TARGET-035 in Figure 8).

Whether the attacks were successful or not, the cases involving repetitive registration of targets might imply that these targets were considered more *critical* than others.

In cases involving a single target registration (i.e. no repeated registrations), we observed that:

- The target was down for the whole observation period (e.g., TARGET-001 in Figure 7).
- The target was initially shut down, then showed some resilience for a few days, and was shut down again (e.g., TARGET-009 in Figure 7).
- The target had a fluctuating QoS score in the observation period (e.g., see TARGET-012 in Figure 7).
- The target initially suffered substantially, but it managed to recover (e.g., see TARGET-043 in Figure 7).
- The target resisted for a small period of time (less than a day) and got shut down pretty quickly (e.g., see TARGET-059 in Figure 7).
- The target had negligible drops in its QoS score in the first few days (see TARGET-060 in Figure 7).
- The target initially got affected by the attacks – albeit minimally (light orange QoS); however, it managed to keep high QoS scores after the initial disruption (e.g., see TARGET-073 in Figure 7).
- The target was not affected at all (e.g., TARGET-074 in Figure 7), and stayed green all the time.




Figure 7. Timelines displaying the states of the first 100 targets registered by pro-Russian entities (rotated 90° to the right)




Figure 8. Timelines displaying the states of the first 100 targets registered by pro-Ukrainian entities (rotated 90° to the right)