

Getting Bored of Cyberwar: Exploring the Role of Low-level Cybercrime Actors in the Russia-Ukraine Conflict

Anh V. Vu
University of Cambridge
Cambridge, United Kingdom
anh.vu@cl.cam.ac.uk

Daniel R. Thomas
University of Strathclyde
Glasgow, United Kingdom
d.thomas@strath.ac.uk

Ben Collier
University of Edinburgh
Edinburgh, United Kingdom
ben.collier@ed.ac.uk

Alice Hutchings
University of Cambridge
Cambridge, United Kingdom
alice.hutchings@cl.cam.ac.uk

Richard Clayton
University of Cambridge
Cambridge, United Kingdom
richard.clayton@cl.cam.ac.uk

Ross Anderson
University of Cambridge & Edinburgh
Cambridge, United Kingdom
ross.anderson@cl.cam.ac.uk

ABSTRACT

There has been substantial commentary on the role of cyberattacks carried out by low-level cybercrime actors in the Russia-Ukraine conflict. We analyse 358k website defacement attacks, 1.7M UDP amplification DDoS attacks, 1764 posts made by 372 users on HACK FORUMS mentioning the two countries, and 441 Telegram announcements (with 58k replies) of a volunteer hacking group for two months before and four months after the invasion. We find the conflict briefly but notably caught the attention of low-level cybercrime actors, with significant increases in online discussion and both types of attacks targeting Russia and Ukraine. However, there was little evidence of high-profile actions; the role of these players in the ongoing hybrid warfare is minor, and they should be separated from persistent and motivated ‘hacktivists’ in state-sponsored operations. Their involvement in the conflict appears to have been short-lived and fleeting, with a clear loss of interest in discussing the situation and carrying out both website defacement and DDoS attacks against either Russia or Ukraine after just a few weeks.

CCS CONCEPTS

• **Social and professional topics** → **Computer crime**; • **Applied computing** → **Cyberwarfare**; • **Security and privacy** → Social aspects of security and privacy; • **Networks** → Denial-of-service attacks; • **Mathematics of computing** → Time series analysis.

KEYWORDS

Russia-Ukraine conflict; DDoS attacks; website defacement attacks; cybercrime; cyberwar; volunteer hacktivists; IT Army of Ukraine.

ACM Reference Format:

Anh V. Vu, Daniel R. Thomas, Ben Collier, Alice Hutchings, Richard Clayton, and Ross Anderson. 2024. Getting Bored of Cyberwar: Exploring the Role of Low-level Cybercrime Actors in the Russia-Ukraine Conflict. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3589334.3645401>



This work is licensed under a Creative Commons Attribution International 4.0 License.

WWW '24, May 13–17, 2024, Singapore, Singapore
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0171-9/24/05
<https://doi.org/10.1145/3589334.3645401>

1 INTRODUCTION

Researchers, politicians, and journalists have long been fascinated by ‘cyberwar’ – the spectre of armed conflict between nations spilling over into attacks conducted over the Internet [62]. ‘Colder’ forms of inter-state conflict are characterised by espionage and intelligence gathering, which may facilitate the degradation of online systems once hostilities commence [18]. Alongside this, there has been a thirty-year history of speculation around how the tools and techniques of the cybercrime underground – Distributed Denial of Service (DDoS) attacks, disruption and compromise of services, web defacements, and similar techniques – might allow civilians to play a role in a ‘hot’ war between developed nations [5]. Much of this speculation, drawing from criminological models of low-level cybercrime groups and on links between this underground and well-organised ‘hactivist’ movements, has argued these groups would play a crucial role, making the future of war hybrid, chaotic, and unpredictable [78]. The 2022 Russian invasion of Ukraine provides an opportunity to assess what has happened in practice.

Russia and Ukraine have a long history of electronic information warfare [45] and are among the most active cybercrime hubs [43]. When Russia invaded Ukraine on 24 February 2022, war-related attacks on the two countries were regularly reported [55]. A popular narrative is that the engagement of low-level cybercrime actors and volunteers could be a game changer and could undermine Russia’s war [22]. Some commentators predicted it will be the first full-scale cyberwar [6], its effects will last for decades [64], and youngsters would be drawn into a ‘cyberwar’ by joining IT Army of Ukraine – a group backed by the Ukrainian state to co-ordinate volunteers and civilians to help disrupt Russian assets [52, 65]. Some have suggested a ‘real cyberwar’, predicting hacktivist attacks on Russia would escalate further through 2022 [69]. These narratives regularly appear in the press and play a role in shaping domestic policy responses to cybercrime. Although less likely to grab headlines, a contrary narrative around ‘overhyped cyberwar’ suggests cyber operations in the conflict have been slow [46] and insignificant [37], while claims of an unprecedented level of cyberattacks and Russia’s much-vaunted cyber capabilities are questionable [24, 53, 81]. GCHQ commented the cyber conflict had not yet materialised [19] and pointed to the resilience of Ukraine’s defences [71].

Government-backed cyber operations [26, 49] and destructive attacks have continued [48, 83]. However, data about nation-state attacks is hard for academics to access, and actors behind significant

real-world attacks tend to take steps to avoid scrutiny. We are particularly interested in non-governmental activity contributed by many low-level but high-volume actors, focusing on the hypothetical ‘volunteer army’, where participants are mostly unskilled and their activity highly relies on off-the-shelf tools. We explore their role in the ‘cyberwar’ between Russia and Ukraine, in which both sides have substantial IT infrastructure, a thriving digital underground crime ecosystem, and significant access to offensive capacities.

We longitudinally and statistically measure activities linked with low-level cybercrime actors, including web defacements (§4) and DDoS attacks (§5). The findings are incorporated with analyses of hacking community reactions, including a pro-Ukraine volunteer group (§6). The role of these actors in the conflict is discussed in §7. This study was ethically approved (see Appendix §A). All data and scripts are available to researchers on request (see Appendix §B).

2 BACKGROUND AND RELATED WORK

Information warfare has long been part of ‘hybrid’ modern conflicts, especially around the control of communications [30, 41]. The enemy’s ability to spread news and propaganda can be degraded by targeting crucial sites, public services, broadcast and telecom infrastructure. Censorship is often used during wartime [59]; governments block access to global services, especially social networks and media platforms, to suppress unwanted narratives. Russia blocked news and anti-war domains when the conflict started [60, 68], and lost access to foreign sites [60] as well as service providers [35]. Ukrainian users experienced degraded network performance [34], while Ukrainian supporters tried unconventional channels such as online reviews to bypass censorship [51]. Attacks are not just online; Russian missiles hit TV towers in Kyiv in early March 2022 [73].

Some associations between kinetic warfare and ‘nationalistic’ cyberattacks have been reported. Ukrainian firms were hit by data wipers such as CaddyWiper and NotPetya [9, 86], DDoS attacks [21, 66] and phishing campaigns [75]; Ukraine supporters have used spam senders to distribute propaganda in Russia [8] and have stolen cryptocurrency from Russian wallets [23]. Ukrainian universities were hacked [85], the Ukrainian electricity grid was hit by Indus-troyer2 [29], and the Ukrainian satellite Internet was downed [61]. Attackers self-identifying with the Anonymous movement declared a ‘cyberwar’ on Russia [72] with attacks against Russian Ministry of Defence databases [14] and state TV channels [15], while Killnet struck back [47]. Russia intermittently received attacks instigated by volunteer hacktivists of the IT Army of Ukraine [22, 82].

While the security industry has reported some insights [1, 2, 26, 49], empirical quantitative academic work analysing the link between armed conflicts and cybercrime has been limited. A notable report is by a Czech university’s incident response team, showing negligible impact on their network after hundreds of users launched DDoS attacks against Russia for a week after the invasion [32].

One type of attack linked with the low-level cybercrime actors is website defacement [63], which accounted for around 20% of online attacks in 2014 [28] and is often organised into discrete campaigns [44]. Attackers (or defacers) gain unauthorised access using off-the-shelf tools and simple exploits, then alter sites’ appearance to demonstrate success [44]. Defacers have heterogeneous developmental trajectories [77]; they are often organised in groups [58]

and have been using online archives [40] as a ‘hall of fame’ to show off their achievements to gain reputation. Defacements are mostly hobbies or pranks with greetings to peers [84], but some advertise hacking tools or services to make money or display other motives such as a wish for community recognition, patriotic, religious and political views [7, 63]. Web defacement may cause economic harm [4, 16] and has occasionally been used as a proxy for terrorist and other serious activities [31]. Another simple type of large-scale attack linked to low-level cybercrime actors is amplified DDoS, with many attributed to DDoS-for-hire or ‘booter’ services [39]. Such services abound [36], and off-the-shelf DDoS tools are widely available; they were tailored and provided to pro-Ukrainian volunteers early in the conflict to attack Russian infrastructure and assets.

Unlike state-sponsored activities, defacement and DDoS attacks can be systematically collected and measured with reasonable coverage. Defacements are available on online archives [40], while DDoS attacks can be collected through honeypots [38, 54, 76]. Launching these attacks with ready-made tools is straightforward for those without much technical expertise. They can be executed quickly, at scale, and have instant, noticeable effects such as altering targets’ appearance, making them inaccessible, or taunting opponents with compromised sites. During wartime, the need to rapidly disseminate political messages and propaganda makes them attractive.

3 METHODS AND DATASETS

We use several quantitative datasets collected regularly and separately, spanning 1 January to 30 June 2022;¹ timestamps are normalised to UTC. To determine if the conflict has impacts resulting in different means (or mean ranks) of daily cyberattacks and hacking discussions, we separate the period into three eras; E_1 : before the invasion, from 1 January to 24 February 2022; E_2 : around one month immediately after the invasion, from 24 February to 31 March 2022; and E_3 : from 1 April to 30 June 2022. We then apply unpaired statistical tests, using One-way ANOVA or Kruskal-Wallis depending on the data distribution; the null hypothesis H_0 is there is no significant difference between the three eras. We use post-hoc tests: Tukey-Kramer for ANOVA or Dunn’s for Kruskal-Wallis to identify pairs causing changes if any. Effect sizes are measured by η^2 , ranging [0, 1]; $0 \leq \eta^2 < 0.01$: no effect; $0.01 \leq \eta^2 < 0.06$: small effect; $0.06 \leq \eta^2 < 0.14$: medium effect; $0.14 \leq \eta^2 \leq 1$: large effect [50].

Web Defacement Attacks. We fully scrape the most popular active defacement archives during the period; see Table 1. We started with ZONE-H, the largest and most popular one (since March 2002) providing cybersecurity news and self-reported defacements along with hacking content [40]. We then took out the most active defacers from ZONE-H and investigated their online presence, including the archives on which they report attacks. Some defacers promote their attacks on Twitter and Telegram; we also looked there. We then shortlisted the five largest defacement archives by attack volume, including OWNZYOU (since January 2021), ZONE-XSEC (since May 2020), HAXOR-ID (since November 2019), and DEFACER-PRO (since June 2021). Smaller archives were historically active [44], but either vanished (HACK MIRROR and MIRROR ZONE) or have hosted different content (HACK-CN and MYDEFACE). While not all compromised sites

¹ No further substantial changes have been observed beyond the six-month mark. We thus decided to maintain that timeframe, which is sufficient to deliver our narratives.

Table 1: The complete collection of the 5 most popular defacement archives for 6 months from 1 January 2022 to 30 June 2022.

Archive URL	ZONE-H zone-h.org	OWNZYOU ownzyou.com	ZONE-XSEC zone-xsec.com	HAXOR-ID hax.or.id	DEFACER-PRO defacer.pro	Total 5 archives
Manual staff verification	✓	✓
Automatic validity sanitisation	.	.	✓	✓	✓	✓
Team information	.	.	✓	✓	✓	✓
Country of targeted victims	✓	✓	✓	✓	✓	✓
Originating country of defacers	.	.	.	✓	.	✓
Reasons and/or motivations	.	.	.	✓	.	✓
Types of vulnerability	.	.	.	✓	.	✓
Snapshots of defaced websites	✓	✓	✓	✓	✓	✓
Defacements (raw)	164 312	76 608	53 852	34 482	28 594	357 848
Defacements †	164 312	67 510	53 814	34 465	27 662	317 049
Valid defacements †	143 485 (87.32%)	47 657 (70.59%)	53 705 (99.80%)	34 439 (99.92%)	26 379 (95.36%)	274 963 (86.73%)
Invalid defacement †	20 827 (12.68%)	19 853 (29.41%)	109 (0.20%)	26 (0.08%)	1 283 (4.64%)	42 086 (13.27%)
Defacers (raw)	2 173	1 214	561	484	540	4 347
Defacers †	1 790	689	560	482	526	3 454
Defacers with valid reports †	1 655 (82.01%)	553 (54.00%)	541 (99.82%)	443 (99.55%)	486 (97.79%)	2 781 (77.44%)
Defacers with invalid reports †	843 (41.77%)	722 (70.51%)	24 (4.43%)	15 (3.37%)	147 (29.58%)	1 656 (46.12%)

✓ fully available; ✓ partly available; . not available; † duplicated defacements and defacer handles within and across different archives were unified.

get reported, measuring trends from the most reputed archives is likely indicative. The country of defaced sites is identified based on ccTLD, IP geolocation, and geolocation of the AS hosting the sites, excluding CDNs (Appendix §C). The defacement submission process is detailed in Appendix §D. We ensure data completeness and bypass challenges e.g., CAPTCHA and IP blocking (Appendix §E).

Further steps are performed to enhance data reliability. First, many on-hold submissions are valid but were never verified; we perform a semi-automatic validation using the messages left on defaced pages (see Appendix §F). Second, submissions may be reported to multiple archives to broaden their visibility. We de-duplicate across and within archives by hashing their content (see Appendix §G). Third, as ‘notifier’ can be arbitrary, typos can give a single attacker multiple identities; we correct typos across all archives using handles’ similarity and messages left on defaced pages (see Appendix §G). In total, 137 339 reports were verified by the archives, 97 652 were automatically validated by us and a further 39 972 were validated semi-automatically; 40 799 (11.00%) duplicate reports are merged across all archives. Of the remaining 317 049 reports, we analyse 274 963 validated submissions (86.73%, around 1 500 per day). Of these, 4 347 defacer handles are also unified to 3 454.

UDP Amplification DDoS Attacks. We use 1.7M DDoS attack records gathered by a honeypot network emulating protocols vulnerable to reflected UDP attacks [76]. A flow of packets is considered to be an attack if any sensor observes at least five packets for the same victim IP or IP prefix, and the attack is deemed to last from the first packet until the last packet preceding 15 minutes without further packets. In 2022, the median number of honeypots contributing data was 50, 95% CI [34, 51]; the median number of observed attacks per week was 35 000, 95% CI [11 900, 271 000] and on IP prefixes of 438, 95% CI [0, 3 480]; the median attack duration was 1.53 minutes, while the maximum was 11 300 minutes. The country of victims is identified by geolocation of the IP address and AS hosting it, excluding CDNs (see Appendix §C).

Underground Forum Discussions. Online forums are structured around subforums containing threads with multiple posts. To assess changes in discussion topics within the hacking community, we use a snapshot of the most popular hacking forum, HACK FORUMS from the CRIMEBB dataset [57]. The forum is a place for users to learn about attacks and trade cybercrime tools and services. Many are low-level actors; however, some have been prosecuted for cybercrime-related activities [56]. We extract all 123 threads within the six months consisting of at least one post with the keywords ‘Russia’ and/or ‘Ukraine’ (case-insensitive): 115 related to Russia, 108 related to Ukraine, in which 100 related to both. We then use all 1 279 posts from 84 highly relevant threads – those with titles directly having the keywords. For the rest 39 less-relevant threads, we count 485 posts directly consisting of the keywords. In total, 1 764 relevant posts made by 372 active forum users are analysed.

Volunteer Hacking Discussions. Two days after the invasion, the Ukrainian government called on pro-Ukraine ‘hacktivists’ to join the IT Army of Ukraine, which was stood up in an ad-hoc manner [65] to support the war effort [52, 74]. The most tangible outcome is a public Telegram channel used mainly to recruit and encourage volunteers to spread political propaganda and co-ordinate disruptive efforts against Russia. We confirmed with a Ukrainian government source that it is the official channel used for communication amongst Ukrainian civilians, with messages being forwarded to other unofficial satellite groups with far fewer subscribers.

The group, with over 200k subscribers, promotes lists of Russian targets (both in Ukrainian and English) most mornings with URLs and IP addresses posted on their channels. They encourage using various attack vectors to disrupt communication and financial systems by hitting banks, businesses, government, and logistics [82]. They provide guides and tools to launch attacks e.g., quickly fetching daily targets and granting access to individuals’ cloud resources for later coordinated attacks, claiming ordinary Russians have seen effects when they hit banks, exchanges [20], and cinemas [13].

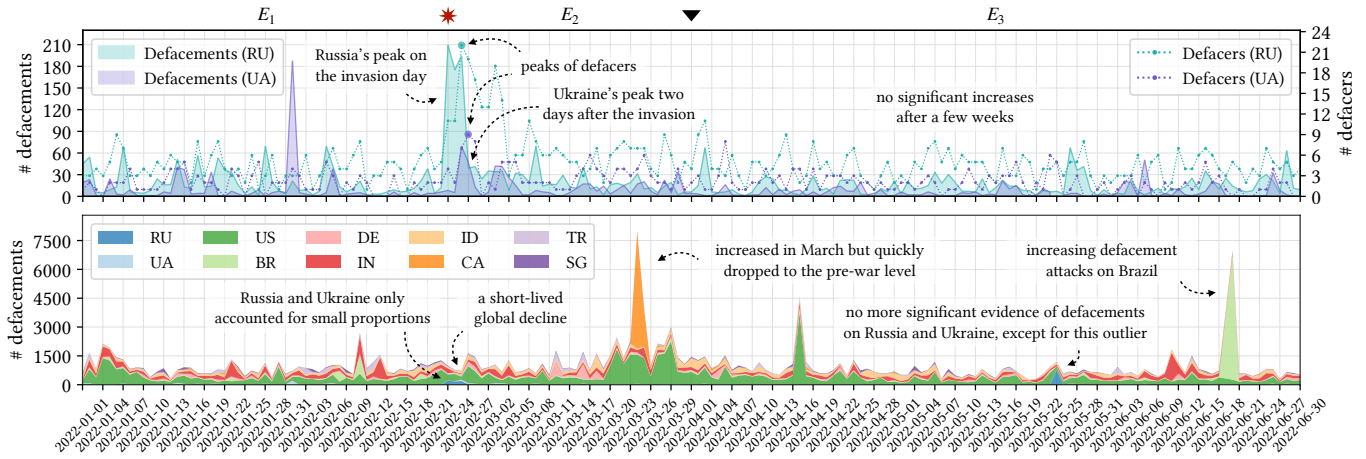


Figure 1: Number of defacements and defacers per day in the Russia-Ukraine scale (top) and the global scale (stacked, bottom).

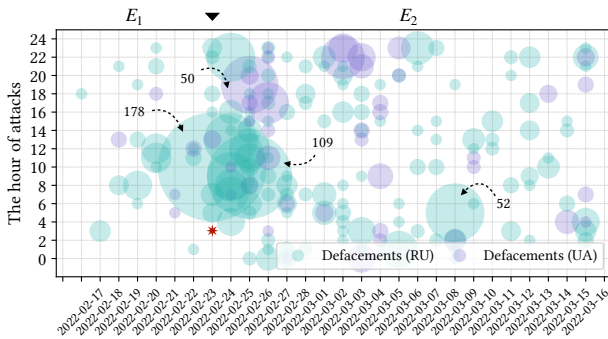


Figure 2: Number of defacements hitting Russia and Ukraine by hour around the invasion day (marked with the red star).

We believe the involved ‘volunteer hacktivists’ are mostly low-level actors, as much of their activity depends on tools provided by the group. We collect 441 announcements with 57 757 replies and 900k emoji reactions posted in the channel from its inception until 30 June 2022 using Telethon, which interacts with official Telegram APIs to fully capture messages and metadata. We then used regular expressions to extract promoted IP addresses and domains; subdomains such as `www.xyz.ru` and `smtp.xyz.ru` are combined. Besides Russian and Belarusian domains (`.ru`, `.su`, `.by`), top-level domains (e.g., `.tv`, `.com`) are also targeted. URL shorteners (e.g., `goo.gl`) and online services (e.g., `youtube.com`) are excluded, resulting in 3 845 targets: 2 291 IP addresses (59.58%) and 1 554 domains (40.42%).

4 THE EVIDENCE FROM WEB DEFACTEMENTS

We measure the dynamic of defacement, both on the global and Russia-Ukraine scales. Figure 1 shows the number of defacements per day as the conflict progressed. Figure 2 breaks down changes by hour for the most active four-week period from 17 February 2022.

The Russia-Ukraine Scale. Website defacements targeting Russia immediately peaked on the invasion day at 209 (14.48% of all defacements on that day, while it was 0.60% the day before). The first big wave was at around 10 AM (7 hours after the invasion) with 178

attacks caused by a single defacer, followed by smaller waves on the same day. Two follow-up waves occurred at 1 PM on 25 February and 9 AM on 26 February with 43 and 109 attacks, respectively. The number of defacers targeting Russia peaked 2 days later: while only 11 defacers accounted for the peak on 24 February, there were 22 on 26 February. No notable change in defacements hitting Ukraine was seen on the invasion day, but a peak of 69 attacks occurred 2 days later (6.30% of all defacements on that day, while it was 0.47% the day before). The largest wave was at around 7 PM on 26 February (26 attacks), followed by medium waves at 5 PM on 27 February (26 attacks) and 10 PM on 3 March (29 attacks). Defacers targeting Ukraine peaked with 9 on 27 February, 1 day after the largest wave.

There was a spike of 771 web defacements by five defacers targeting Russia on 25 May. Of these, 764 were claimed by a single defacer compromising a server hosting 760 sites. This outlier appears to be unique; it was removed from the graph for better visualisation. The peak of 187 web defacements hitting Ukraine on 1 February 2022 by four defacers did not have a single cause and did not lead to a sharp increase in the number of defacers in the following days.

For both Russia and Ukraine, Kruskal-Wallis tests suggest a statistically significant difference in the number of web defacements and defacers per day through E_1 , E_2 , and E_3 , see Table 2 (Appendix §H). Post-hoc analysis indicates a significant difference between the pre-invasion (E_1) and one-month-post-invasion (E_2) periods. $p(E_2, E_3)$ is also significant, but *not* $p(E_1, E_3)$, suggesting the situation returned to the pre-war levels after the second era. The effect sizes η^2 are all between medium and large, ranging [0.06, 0.14].

The Global Scale. The number of web defacement attacks targeting Russia and Ukraine appears to be trivial when set against the global scale. Among 274 963 analysed web defacements, only 5 899 (2.15%) targeted the two countries (4 340 for Russia and 1 559 for Ukraine). The top 10 countries account for 69.85% of all web defacements; sites hosted in the US by global-scale vendors (e.g., Cloudflare, Amazon) are excluded (see Appendix §C), but the US still consistently suffers the majority of web defacements. Since January 2022, the US accounts for 26.95% of defacement attacks, followed by India with 11.47% and Indonesia with 8.41%, while Russia and Ukraine only account for small proportions: 1.58% and 0.57%, respectively.

There was a short-lived decline in defacement attacks worldwide on the invasion day (from around 1 400 to 1 000), while it peaked for Russia from nearly zero to 209 (14.48% of all defacements). This suggests a genuine change in the way defacers chose their targets, precipitated by the war. The US is consistently the largest target but only accounts for 21.97% on that day. During the last two weeks of March 2022, the number of defacements significantly increased at the global scale, with many defacements targeting the US. However, much like the Russia-Ukraine scale, the effect lasted for only a few weeks. The unusual peaks against Brazil happened in late June (also for DDoS attacks, see Section §5), without a clear explanation.

The patterns seen from the Kruskal-Wallis and post-hoc tests in the Russia-Ukraine scale do not apply for most top countries, see Table 2 (Appendix §H). ANOVA/Kruskal-Wallis tests on the defacement and defacer counts are not *both* significant for Brazil, Germany, India, and Singapore; no significant changes are seen between the pre-invasion and one-month-post-invasion eras $\langle E_1, E_2 \rangle$ for Canada and Turkey. Indonesia has a similar pattern of post-hoc tests, but one of the effect sizes is small. The only country following a close pattern is the US with medium effect sizes, yet Section §5 will show this did not hold for DDoS attacks on the US.

The evidence above suggests a genuine increase of website defacements against the two countries shortly after the invasion, standing out significantly from other top countries. Russia was the first to be hit at scale, followed by Ukraine a few days later. However, this effect was fairly short-lived for both countries, lasting for only a few weeks before returning to pre-war levels, presumably as defacers ran out of targets or just lost interest in carrying out attacks. The number of defacers involved was small, but for a while, they turned from indiscriminate to more targeted attacks.

Defacement Motives. The conflict caught the attention of existing defacers, who performed many attacks against other countries but not Russia and Ukraine until just after the invasion, suggesting their choice of targets was influenced. We also found some ‘new faces’ e.g., the second most active defacer targeting Russia after the war began first appeared in mid-February, peaked on the invasion day, stayed significant for three days then declined quickly. While some minor players at the global scale made a significant contribution to the rise in attacks on Russia and Ukraine, the three most active defacers globally made a trivial contribution (less than 10) against either country. We do not verify findings on their general motives (see §2), but to gain conflict-related insights we analyse the contents of 4 340 defacements targeting Russia and 1 559 hitting Ukraine.

We annotate motives based on 1 341 unique messages left on the defaced pages. We consider a political sentiment and mark it as supporting Russia/Ukraine if a support/objection is expressed e.g., ‘*We stand with Ukraine!*’. We mark messages consisting of defacer signatures e.g., ‘*Hacked by ABC*’ without clear motives, or just greetings to peers as being for self-aggrandisement. Messages advertising hacking tools and services or asking for ransom are marked financially motivated e.g., ‘*Contact me for shells*’. We label messages expressing favourite mottos or moods as self-expression e.g., ‘*Not much I want, hope my life will be better*’, and exclude 1 278 messages (21.66%) containing empty or random messages.

We find diverse motives, but despite targeting Russia and Ukraine, most messages do not refer to the conflict. 2 723 (46.16%) were for

self-aggrandisement, 1 219 (20.66%) self-expression, 143 (2.42%) related to other conflicts (such as Israel-Palestine), 58 (0.98%) related to patriotism, and 89 (1.51%) were financially motivated (mainly from the two most active defacers globally who did not change their targeting due to the conflict). Some defacers did leave conflict-related messages: 286 (4.85%) supporting Ukraine, roughly 2.8 times higher than those supporting Russia at 103 (1.75%). Notably, some defacers supported Russia, yet also defaced Russian sites, saying they wished to alert and help secure the systems (22 attacks) – ‘*I have secured this domain, I love Russia*’, was a message the third most active pro-Russia defacer left on a Russian website. Likewise, other defacers supported Ukraine yet defaced Ukrainian sites (12 attacks) e.g., ‘*Hello Volodymyr Zelensky, I’m sorry to hack your site. I just wanted to tell you that people need a president like you. We support Ukraine*’. Such signatures are likely intentionally war-related as Russia and Ukraine were not frequently targeted before.

5 THE EVIDENCE FROM DDOS ATTACKS

We now examine if there were also significant changes in DDoS attack volumes targeting Russia and Ukraine after the conflict. Figure 3 shows the number of DDoS attacks in both Russia-Ukraine and global scales over the three eras, while Figure 4 shows their changes by hour during the most active four weeks from 24 February.

The Russia-Ukraine Scale. DDoS attacks lagged defacement by about a week, but occurred in higher volumes and lasted longer; most happened after 7 AM. The number of both DDoS attacks and victims targeting Russia first increased on 2 March (six days after the invasion) with 851 victims, 511 of them at around 6 PM. The attacks peaked four days later with 1 137 victims. High activity levels continued through 23 March, with the biggest wave occurring at around 2 PM on 8 March with 755 victims. Smaller waves continued regularly during the next few weeks. Regarding DDoS attacks hitting Ukraine, significant waves started around a week after Russia’s first big wave (some small spikes targeting Ukraine before Russia were insignificant) with the first notable spike on 10 March having 526 victims, then became prevalent during two weeks from 18 to 31 March: big waves were on 18 March at around 12 PM, 1 PM and 4 PM with 257, 476, and 700 victims, respectively. Other big and medium waves lasted until the end of March, with the biggest peak on 31 March when 1 296 victims were hit. The increased volume only continued for about a month before declining sharply.

Kruskal-Wallis tests suggest statistically significant changes between the daily number of DDoS attacks and victims through the three eras for both Russia and Ukraine, much like with defacements; see Table 3 (Appendix §H). Post-hoc analysis shows high significance levels of $\langle E_1, E_2 \rangle$ and $\langle E_2, E_3 \rangle$, suggesting notable changes between the pre-invasion vs. one-month-post-invasion periods, and the one-month-post-invasion periods vs. the period after that. The main difference between Russia and Ukraine is that the situation for Ukraine returned to pre-invasion levels after one month i.e. $\langle E_1, E_3 \rangle$ is not significantly different, while we still see some difference with Russia. Indeed, the number of DDoS attacks hitting Russia was still slightly higher than before the invasion (see Figure 3). The effect size is large for Russia, while it is medium for Ukraine.

The Global Scale. We again see concentrations in DDoS attacks, with the top 10 countries accounting for 70.49% of all victims.




Figure 3: Number of DDoS attacks and victims per day in the Russia-Ukraine scale (top) and global scale (stacked, bottom).




Figure 4: Number of DDoS attacks on Russia and Ukraine by hour around the invasion day (marked with the red star).

The US still dominates (24.68%), followed by Brazil (11.99%) and Bangladesh (8.10%). Ukraine took 1.57%, while Russia lies 8th at 3.61%. Our DDoS and defacement datasets show some correlations. Three of the top 10 countries for defacements are also in the top 10 for DDoS targets: the US, Germany, and Brazil. As with defacements, DDoS attacks rose globally during the last two weeks of March 2022 with large numbers targeting Bangladesh, but this effect is insignificant; see Table 3 (Appendix §H). The unusual peaks of both defacements and DDoS against Brazil in late June are notable; Brazil is often ranked among the top cybercrime hubs worldwide [43], yet we lack a convincing causality. A similar peak observed on the Russia-Ukraine scale can also be seen at the global scale following the invasion. As with defacements, DDoS attacks thrived on a global scale in March, yet quickly returned to the previous levels after a few weeks with no lasting global changes and no significant evidence of further waves, whether targeting Russia or Ukraine.

The patterns seen from the Kruskal-Wallis and post-hoc tests in the Russia-Ukraine scale do not apply for most top countries, see Table 3 (Appendix §H). Kruskal-Wallis tests are not *all* significant for the US, Bangladesh, and Poland, despite the US accounting for the largest number of attacks and there was a visual increase for Bangladesh (as the tests compare mean ranks instead of means).

No significant changes are seen between the pre-invasion and one-month-post-invasion periods (E_1, E_2) for France and the UK. Brazil and Germany have a similar phenomenon of post-hoc tests in the Russia-Ukraine scale, yet one effect size is small. China is the only country following that phenomenon with large effect sizes; the main difference is that the changes in $p(E_2, E_3)$ are not significant.

As with defacements, the evidence above suggests a genuine increase in DDoS attacks targeting Russia and Ukraine as the conflict began, standing out significantly from most top countries. Russia was still the first to be hit at scale, followed by Ukraine shortly after. The outbreak of both defacement and DDoS attacks on Russia and Ukraine was significant and timely, but fairly short-lived. While defacements returned to previous levels after a few weeks, DDoS attacks on Russia remained marginally higher than pre-war levels.

6 HACKING COMMUNITY REACTIONS

HACK FORUMS Discussions. There was an immediate surge of posts on HACK FORUMS mentioning the two countries after the invasion, from near zero to over 120 per day; see Figure 5. Kruskal-Wallis tests confirm the significance $H(2) = 72.98, p < .0001$, with a large effect size $\eta^2 = 0.40$; pairwise post-hoc tests for $\langle E_1, E_2 \rangle$ and $\langle E_2, E_3 \rangle$ are both significant ($p < .0001$), but not $\langle E_1, E_3 \rangle$ ($p = .8501$). The number of posting users shows a similar story: Kruskal-Wallis test reports $H(2) = 77.54, p < .0001$ with a large effect size $\eta^2 = 0.42$; pairwise post-hoc tests for $\langle E_1, E_2 \rangle$ and $\langle E_2, E_3 \rangle$ are both significant ($p < .0001$), but not $\langle E_1, E_3 \rangle$ ($p = .6657$). This fits the evidence seen with defacement and DDoS attacks: both posting activity and users returned to the pre-war level after a few weeks, presumably as users quickly lost interest and moved on to other discussion topics.

This posting volume is tiny when set against the 62M-post size of HACK FORUMS, showing trivial contributions of the Russia-Ukraine discussions to the overall landscape (as with the previous evidence seen from defacement and DDoS attacks). These posts are centralised: 97.22% belongs to the top 5 popular subforums. Ranked 1st is ‘science, religion, philosophy, and politics’, accounting for 53.40%; ranked 2nd is ‘news and happenings’ with 33.28%; ‘website hacking’ ranked 3rd, followed by ‘crypto currency’, then general chats. We

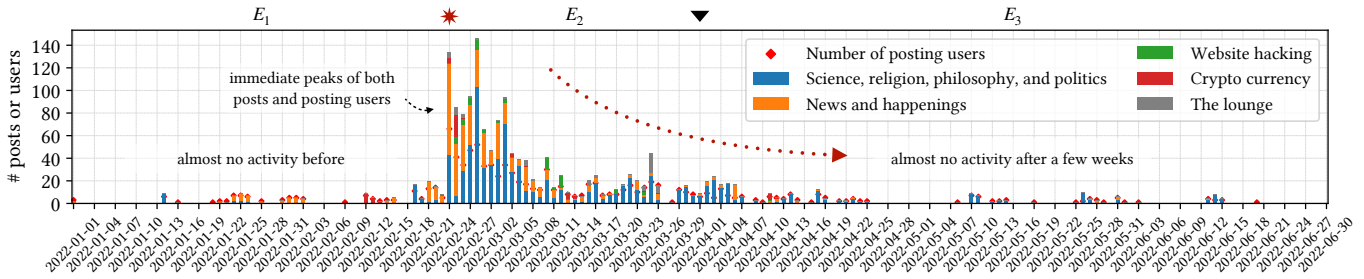


Figure 5: Number of daily posts and posting users on Hack Forums mentioning Russia and/or Ukraine (the top 5 subforums).

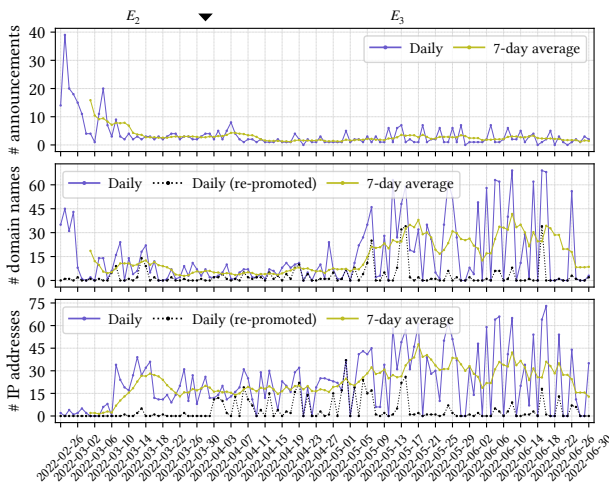


Figure 6: The number of announcements and (re-promoted) targets in the IT Army of Ukraine Telegram channel by day.




Figure 7: Levels of user engagement daily in the IT Army of Ukraine Telegram channel. Values are min-max normalised.

see some ‘news and happenings’ posts in the past, but mostly no ‘science, religion, philosophy, and politics’ posts until the invasion.

Targets Promoted by the IT Army of Ukraine. Many announcements and targeted domains were posted in the first two weeks after the invasion, beginning on 26 February, peaking on 27 February with 40 announcements and 45 domains promoted (IP addresses were not regularly included until later), see Figure 6. Yet they quickly declined to consistently less than 10 per day after two weeks with some days (e.g., 24 and 26 April) having no posts. The number of subscribers also dropped from 300k to around 160k in October 2023.

While the number of announcements dropped, the number of targets steadily increased, particularly in May and June 2022 with

multiple-target posting. Activities were unstable at that time; targets got promoted less frequently and occasional days had no targets. Targets were mostly fresh in the first two weeks, but then a considerable proportion got re-promoted on multiple days e.g., all advertised IP addresses and most domains were re-posted during 4–6 May. Along with frequent zero-target days, this suggests the group might run out of new targets or get bored with finding them.

Community reactions and engagement tell much the same story as with DDoS and defacement attacks (see Figure 7). While more targets were promoted in May and June, volunteers appeared to have largely lost interest, despite their intense activity in the first few weeks. The decline in reaction was consistent across all engagement types: views, emojis, forwards, and replies. Older announcements may have more time to accrue views as people scroll up the channel, but the emojis, forwards, and replies require user intent. We believe the figures reflect a genuine decline in user engagement over time.

The group first provided instructions about tools and guidance to carry on attacks against Russian payment systems on 9 March (two weeks after the invasion), attracting high levels of engagement: 240k views, 2.6k emojis, 1.2k forwards, and 421 replies from 197 users. The next was on 1 April: while the number of replies and forwards was similar to the first, other kinds roughly halved. From mid-May to late June, instructions were posted four more times, yet users were around four times less engaged than the first in March, indicating a loss of interest despite the operator’s intensive efforts.

Target Selection. Targets were often themed, patterning around particular weekdays e.g., online news and propaganda, food delivery, and entertainment are often attacked at weekends to maximise impact as people spend more time online. Themes were also occasionally set with re-promoted old targets, leading to wide variations in the number of new targets, particularly from May onwards (zero on some days, see Figure 6). Subscribers can suggest new targets, but group organisers post most of them. During the last two months, they often re-promoted targets linked to old posts, which could be as simple as ‘we continue to work with yesterday’s targets’, suggesting that they might have some difficulties in hunting new targets.

We use categories linked with targets by default; when unavailable, we rely on root domains e.g., .tv and .gov are likely news and government sites. Categories of generic domains (e.g., .net, .com) are identified by direct visits (via Russian IP relays) or querying Internet archives if they are down. Some targets were indeed down while previously active, suggesting attacks might have succeeded e.g., ksrif.ru (the Constitutional Court of the Russian Federation) was down for a while, and data.gov.ru was both defaced and DDoSed.

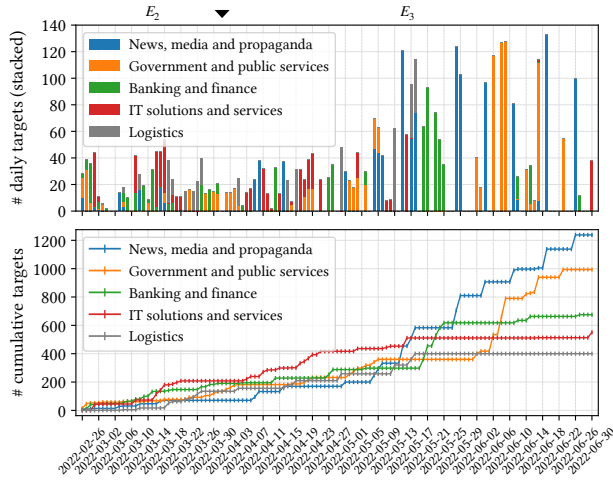


Figure 8: Number of daily stacked targets (top) and cumulative targets (bottom) being promoted in the top 5 categories.

Categories vary, yet five dominate 80.21% of all targets, see Figure 8. ‘*News, media and propaganda*’, including TV broadcasting, has been consistently promoted since the war began but only became the most common one in May when it overtook ‘*IT solutions and services*’. ‘*Government and public services*’, which includes military, state-owned websites, and public services for civilians such as parking and lighting (including governments imposed on occupied territories) has also been regularly targeted throughout, but they only grew rapidly towards the end of the period, making this the second most common category overall. ‘*Banking and finance*’ ranked 3rd, including banks, stock exchanges, electronic payment, accounting, credit services, credit card providers, trading, bidding, investment platforms, funding agencies, and the micro-finance industry. ‘*IT solutions and services*’ ranked 4th, including software solutions supporting governments, digital signature and information security services such as DDoS-Guard. This sector was actively promoted early on but was targeted far less thereafter. ‘*Logistics*’ ranks 5th, including airlines and aviation, travel, shipping, and food delivery. Other categories by popularity include (online) marketplaces and stores (e.g., job markets, real estate, e-commerce, drug stores, app stores), manufacturers and trading (e.g., military footwear, shoes, wood and roofing materials), education, insurance, telecommunications (e.g., Internet providers), businesses and state companies (e.g., energy and steel manufacturers), online forums, entertainment (e.g., cinemas), and non-governmental organisations.

Crossover with Observed Attacks. The IT Army of Ukraine maintains a dashboard of targets’ status, claiming many are down due to their actions. To find whether the attacks involved reflected DDoS or defacement, we correlate our attack records with promoted targets since the Telegram group started. We consider a defacement overlap when either its URL or IP address matches promoted targets, while for DDoS attacks, only IP addresses are used. There was very little overlap with website defacement attacks: among 3 845 targets promoted by the group, there are only 59 valid matches (1.53%), including 7 domain matches (0.18%) and 52 IP matches (1.35%). Notably, no overlaps occur on the day targets are

promoted, suggesting that defacers chose their targets themselves independently (by scanning sites based on ccTLDs); these targets are largely unimportant and irrelevant to the conflict. For DDoS attacks, we observe 707 (30.86%) total overlaps among 2 291 promoted IP addresses, which is considerable. Unlike defacements, some are executed the same day they are promoted; we find many same-day overlaps in late March, early April, and during May, peaking on 19 March 2022 with 22 victims overlapping. However, the crossover dropped quickly, becoming less frequent from late May while many new targets were still actively advertised. This suggests a loss of interest by volunteers in attacking targets promoted by the group.

7 CONCLUDING REMARKS

The role of the low-level cybercrime actors studied in this paper amounts to essentially trivial acts of solidarity and opportunistic competition. Their primary impact is probably to disseminate political propaganda, with little measurable evidence to suggest these actors are making any persistent contribution to the conflict, even in a major war between two nations with a long history of cyber warfare. Their role and capacity in future nation-state conflicts should not be confounded with state hacking or political ‘hacktivism’. Our diverse, separately collected datasets all point to a narrative that notable attention was temporarily drawn to Russia and Ukraine rather than other countries. Neither the engagement on HACK FORUMS nor Telegram, or the outbreak of defacements and DDoS attacks, was long-lasting – presumably as participants just lost interest, despite their choice of targets being influenced by the war for a while. This is in line with other work suggesting that boredom is an important factor in people leaving cybercrime [11].

We do not dispute claims about the prevalence of state-sponsored attacks such as malware and phishing [26, 49], but rather provide additional perspectives on the role of low-level cybercrime actors. Some cybercrime-related activities may indeed contribute to the war effort. Leaks, especially of high-profile datasets gathered from Russian public services, have consistently made headlines. They may or may not be connected to civilians, hacktivists, state actors, or other groups. Much as with ransomware, their low numbers and vast disparities in impact make them far less cross-comparable. Our findings fit a more general pattern in the cybercrime ecosystem increasingly characterised by an entrepreneurial, service-based economy which is becoming alienated from traditional hacker culture’s concerns with technical learning and dissent [3]. Committed, persistent hacktivists appear to be separate from low-level crime communities whose interest seems to have been fleeting and easily diverted by trending news. They did indeed briefly get involved in the war effort using off-the-shelf exploit tools, but their role on the ‘hard’ digital frontline remains rather limited. These are best seen as actions in the theatre of protest, ‘soft power’ and solidarity. In a ‘cyberwar’, efforts should be prioritised against high-profile and state-sponsored threats instead of such low-level cybercrime actors, even though they may cause immediately noticeable effects.

ACKNOWLEDGMENTS

We thank the anonymous reviewers. This work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant No 949127.

REFERENCES

- [1] NETSCOUT. 2022. DDoS Threat Landscape – Russia.
- [2] NETSCOUT. 2022. DDoS Threat Landscape – Ukraine.
- [3] Ross Anderson, Richard Clayton, Rainer Böhme, and Ben Collier. 2021. Silicon Den: Cybercrime Is Entrepreneurship. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Springer, New York, NY, USA, 28 pages.
- [4] Jason Andress and Steve Winterfeldt. 2011. *Cyber Warfare: Techniques, Tactics and Tools for Security Practitioners*. Elsevier, Amsterdam, The Netherlands.
- [5] John Arquilla and David Ronfeldt. 1993. Cyberwar Is Coming! *Comparative Strategy* 12 (1993), 141–165.
- [6] Atlantic Council. 2022. Vladimir Putin’s Ukraine Invasion Is the World’s First Full-Scale Cyberwar.
- [7] Sudipta Banerjee, Thomas Swearingen, Ruth Shillair, Johannes M. Bauer, Thomas Holt, and Arun Ross. 2022. Using Machine Learning to Examine Cyberattack Motivations on Web Defacement Data. *Social Science Computer Review* 40 (2022), 914–932.
- [8] BBC. 2022. Ukraine: Spam Website Set Up to Reach Millions of Russians.
- [9] Bleeping Computer. 2022. New Data-Wiping Malware Used in Destructive Attacks on Ukraine.
- [10] British Society of Criminology. 2015. Statement of Ethics.
- [11] Ben Collier, Richard Clayton, Alice Hutchings, and Daniel Thomas. 2021. Cybercrime Is (Often) Boring: Infrastructure and Alienation in a Deviant Subculture. *The British Journal of Criminology* 61 (2021), 1407–1423.
- [12] Miguel Cozar, David Rodriguez, Jose M. Del Alamo, and Danny Guaman. 2022. Reliability of IP Geolocation Services for Assessing the Compliance of International Data Transfers. In *Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS&P)*. IEEE, New York, NY, USA, 181–185.
- [13] Cybernews. 2022. Ukraine’s Cyber Army Hits Russian Cinemas.
- [14] Cybernews. 2022. Anonymous Leaks Database of the Russian Ministry of Defence.
- [15] Daily Mail. 2022. Anonymous Claims It Has Hacked Russian State TV Broadcasts to Show War Footage From Ukraine.
- [16] Dorothy E. Denning. 2011. Cyber Conflict as an Emergent Social Phenomenon. In *Corporate Hacking and Technology-Driven Crime: Social Dynamics and Implications*. IGI Global, Hershey, PA, USA, 170–186.
- [17] Periwinkle Doerfler, Andrea Forte, Emiliano De Cristofaro, Gianluca Stringhini, Jeremy Blackburn, and Damon McCoy. 2021. “I’m a Professor, Which isn’t Usually a Dangerous Job”: Internet-facilitated Harassment and Its Impact on Researchers. *Proceedings of the ACM on Human-Computer Interaction (HCI)* 5, CSCW2 (2021), 1–32.
- [18] Bernard Everett. 2013. Optically Transparent: The Rise of Industrial Espionage and State-Sponsored Hacking. *Computer Fraud & Security* 2013 (2013), 13–16.
- [19] Financial Times. 2022. Prospect of Russian Cyber War May Have Been ‘Overhyped’, Says UK Spy Chief.
- [20] Forbes. 2022. Moscow Exchange, Sberbank Websites Knocked Offline—Was Ukraine’s Cyber Army Responsible?
- [21] Forbes. 2022. ‘Most Severe’ Cyberattack Since Russian Invasion Crashes Ukraine Internet Provider.
- [22] Foreign Policy. 2022. Don’t Underestimate Ukraine’s Volunteer Hackers.
- [23] Forkast News. 2022. Ukraine Seizes Funds From Crypto Wallet Funding Russian Forces.
- [24] Austen D. Givens, Max Gorbachevsky, and Anita C. Biernat. 2023. How Putin’s Cyberwar Failed in Ukraine. *Journal of Strategic Security* 16 (2023), 96–121.
- [25] Oana Goga, Howard Lei, Sree Hari Krishnan Parthasarathi, Gerald Friedland, Robin Sommer, and Renata Teixeira. 2013. Exploiting Innocuous Activity for Correlating Users Across Sites. In *Proceedings of the ACM World Wide Web Conference (WWW)*. ACM, New York, NY, USA, 447–458.
- [26] Google. 2023. Fog of War: How the Ukraine Conflict Transformed the Cyber Threat Landscape.
- [27] Matthieu Gouel, Kevin Vermeulen, Olivier Fourmaux, Timur Friedman, and Robert Beverly. 2021. IP Geolocation Database Stability and Implications for Network Research. In *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA)*. International Federation for Information Processing (IFIP), Laxenburg, Austria, 10 pages.
- [28] Hackmageddon. 2015. Cyber Attacks Statistics, August 2014.
- [29] HeadMind. 2022. Industroyer2: The Russian Cyberattack on Ukraine Infrastructure.
- [30] Frank G. Hoffman. 2007. *Conflict in the 21st Century: The Rise of Hybrid Wars*. Potomac Institute for Policy Studies, Arlington, VA, USA.
- [31] Thomas J. Holt, Noah D. Turner, Joshua D. Freilich, and Steven M. Chermak. 2022. Examining the Characteristics That Differentiate Jihadi-Associated Cyberattacks Using Routine Activities Theory. *Social Science Computer Review* 40 (2022), 1614–1630.
- [32] Martin Husák, Martin Laštovička, and Tomáš Plesník. 2022. Handling Internet Activism During the Russian Invasion of Ukraine: A Campus Network Perspective. *Digital Threats: Research and Practice (DTRAP)* 3 (2022), 1–5.
- [33] Intricately Market Analysts. 2020. CDN Industry: Trends, Size, and Market Share.
- [34] Akshath Jain, Deepayan Patra, Peijing Xu, Justine Sherry, and Phillipa Gill. 2022. The Ukrainian Internet Under Attack: An NDT Perspective. In *Proceedings of the ACM Internet Measurement Conference (IMC)*. ACM, New York, NY, USA, 166–178.
- [35] Mattijs Jonker, Gautam Akiwate, Antonia Affinito, kc Claffy, Alessio Botta, Geoffrey M. Voelker, Roland van Rijswijk-Deij, and Stefan Savage. 2022. Where. Ru? Assessing the Impact of Conflict on Russian Domain Infrastructure. In *Proceedings of the ACM Internet Measurement Conference (IMC)*. ACM, New York, NY, USA, 159–165.
- [36] Mohammad Karami and Damon McCoy. 2013. Rent to Pwn: Analyzing Commodity Booter DDoS Services. *USENIX Login* 38 (2013), 20–23.
- [37] Nadiya Kostyuk and Erik Gartzke. 2022. Why Cyber Dogs Have Yet to Bark Loudly in Russia’s Invasion of Ukraine. *Texas National Security Review* 5 (2022), 113–126.
- [38] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID)*. Springer, New York, NY, USA, 615–636.
- [39] Johannes Krupp, Mohammad Karami, Christian Rossow, Damon McCoy, and Michael Backes. 2017. Linking Amplification DDoS Attacks to Booter Services. In *Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID)*. Springer, New York, NY, USA, 427–449.
- [40] Michael Kurzmeier. 2020. Towards a Concept for Archiving Hacked Websites. *NPPSH Reflections* 3 (2020), 35–58.
- [41] Chiara Libiseller. 2023. ‘Hybrid Warfare’ as an Academic Fashion. *Journal of Strategic Studies* 46 (2023), 858–880.
- [42] Jing Liu, Fan Zhang, Xinying Song, Young-In Song, Chin-Yew Lin, and Hsiao-Wuen Hon. 2013. What’s in a Name?: An Unsupervised Approach to Link Users Across Communities. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*. ACM, New York, NY, USA, 495–504.
- [43] Jonathan Lusthaus, Miranda Bruce, and Nigel Phair. 2020. Mapping the Geography of Cybercrime: A Review of Indices of Digital Offending by Country. In *Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS&P)*. IEEE, New York, NY, USA, 448–453.
- [44] Federico Maggi, Marco Balduzzi, Ryan Flores, Lion Gu, and Vincenzo Ciancaglini. 2018. Investigating Web Defacement Campaigns at Large. In *Proceedings of the Asia Conference on Computer and Communications Security (ASIACCS)*. ACM, New York, NY, USA, 443–456.
- [45] Margarita Jaitner. 2015. Russian Information Warfare: Lessons From Ukraine. In *Cyber War in Perspective: Russian Aggression against Ukraine*. NATO Cooperative Cyber Defence Centre of Excellence (CCDCOE), Tallinn, Estonia, 87–94.
- [46] Lennart Maschmeyer and Myriam Dunn Cavelt. 2022. *Goodbye Cyberwar: Ukraine as Reality Check*. Technical Report. Center for Security Studies (CSS), ETH Zürich.
- [47] Metro News. 2022. Anonymous Servers Down as Russian Killnet Hackers Strike Back.
- [48] Microsoft. 2022. Defending Ukraine: Early Lessons From the Cyber War.
- [49] Microsoft Threat Intelligence. 2023. Microsoft Digital Defense Report: Building and Improving Cyber Resilience.
- [50] Jeremy Miles and Mark Shevlin. 2000. *Applying Regression and Correlation: A Guide for Students and Researchers*. Sage, London, UK.
- [51] José Miguel Moreno, Sergio Pastrana, Jens Helge Reelfs, Pelayo Vallina, Andriy Panchenko, Georgios Smaragdakis, Oliver Hohlfeld, Narseo Vallina-Rodriguez, and Juan Tapiador. 2023. Reviewing War: Unconventional User Reviews as a Side Channel to Circumvent Information Controls. (2023). arXiv:2302.00598
- [52] Mykhailo Fedorov. 2022. We Are Creating an IT Army.
- [53] Nature News Explainer. 2022. Where Is Russia’s Cyberwar? Analysts Decipher Its Strategy.
- [54] Marcin Nawrocki, John Kristoff, Raphael Hiesgen, Chris Kanich, Thomas C. Schmidt, and Matthias Wählisch. 2023. SoK: A Data-Driven View on Methods to Detect Reflective Amplification DDoS Attacks Using Honey pots. In *Proceedings of the IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, New York, NY, USA, 576–591.
- [55] New York Times. 2022. Russia Uses Cyberattacks in Ukraine to Support Military Strikes, Report Finds.
- [56] Sergio Pastrana, Alice Hutchings, Andrew Caines, and Paula Buttery. 2018. Characterizing Eve: Analysing Cybercrime Actors in a Large Underground Forum. In *Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID)*. Springer, New York, NY, USA, 207–227.
- [57] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, and Richard Clayton. 2018. CrimeBB: Enabling Cybercrime Research on Underground Forums at Scale. In *Proceedings of the ACM World Wide Web Conference (WWW)*. ACM, New York, NY, USA, 1845–1854.
- [58] Robert C. Perkins, Marie Ouellet, Christian J. Howell, and David Maimon. 2022. The Illicit Ecosystem of Hacking: A Longitudinal Network Analysis of Website Defacement Groups. *Social Science Computer Review* 0 (2022), 390–409.

- [59] Byron Price. 1942. Governmental Censorship in War-Time. *American Political Science Review* 36 (1942), 837–849.
- [60] Reethika Ramesh, Ram Sundara Raman, Apurva Virkud, Alexandra Dirksen, Armin Huremagic, Dirk Rodenburg, David Fifield, Rod Hynes, Doug Madory, and Roya Ensafi. 2023. Network Responses to Russia’s Invasion of Ukraine in 2022: A Cautionary Tale for Internet Freedom. In *Proceedings of the USENIX Security Symposium (USENIX Security)*. USENIX Association, California, CA, USA, 2581–2598.
- [61] Reuters. 2022. Russia Downed Satellite Internet in Ukraine - Western Officials.
- [62] Thomas Rid. 2012. Cyber War Will Not Take Place. *Journal of Strategic Studies* 35 (2012), 5–32.
- [63] Marco Romagna and Niek Jan van den Hout. 2017. Hacktivism and Website Defacement: Motivations, Capabilities and Potential Threats. In *Virus Bulletin International Conference*. Virus Bulletin, Abingdon, Oxfordshire, UK, 1–10.
- [64] Dimitrios Serpanos and Theodoros Komninos. 2022. The Cyberwarfare in Ukraine. *Computer* 55 (2022), 88–91.
- [65] Stefan Soesanto. 2022. *The IT Army of Ukraine: Structure, Tasking, and Eco-System*. Technical Report. Center for Security Studies (CSS), ETH Zürich.
- [66] State Sites of Ukraine. 2022. Ukraine Has Suffered Over 3000 DDoS Attacks.
- [67] Statista. 2022. Number of Websites in Russia as of May 2022, by Domain Name.
- [68] Statista. 2022. Number of Web Domains Blocked in Russia Since the Start of the Ukraine Invasion on February 24, 2022 as of April 25, 2022, by Category.
- [69] TASS Russian News Agency. 2022. Number of DDoS Attacks on Russia in 2022 Will Grow 30-Fold.
- [70] TechCrunch. 2022. Web Scraping Is Legal, U.S. Appeals Court Reaffirms.
- [71] The Economist. 2022. The Head of GCHQ Says Vladimir Putin Is Losing the Information War in Ukraine.
- [72] The Guardian. 2022. Anonymous: The Hacker Collective That Has Declared Cyberwar on Russia.
- [73] The Guardian. 2022. Ukraine Says Russia Targeting Civilians as Missiles Hit Kyiv TV Tower.
- [74] The Guardian. 2022. ‘It’s the Right Thing to Do’: The 300,000 Volunteer Hackers Coming Together to Fight Russia.
- [75] The Hacker News. 2022. Google: Russian Hackers Target Ukrainians, European Allies via Phishing Attacks.
- [76] Daniel R. Thomas, Richard Clayton, and Alastair R. Beresford. 2017. 1000 Days of UDP Amplification DDoS Attacks. In *Proceedings of the APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, New York, NY, USA, 79–84.
- [77] Steve G.A. van de Weijer, Thomas J. Holt, and E. Rutger Leukfeldt. 2021. Heterogeneity in Trajectories of Cybercriminals: A Longitudinal Analyses of Web Defacements. *Computers in Human Behavior Reports* 4 (2021), 100113.
- [78] David Wall. 2007. *Cybercrime: The Transformation of Crime in the Information Age*. Wiley, Hoboken, NJ, USA.
- [79] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L. Mazurek, Manya Sleeper, and Kurt Thomas. 2022. SoK: A Framework for Unifying at-Risk User Research. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*. IEEE, New York, NY, USA, 2344–2360.
- [80] Zachary Weinberg, Shinyoung Cho, Nicolas Christin, Vyas Sekar, and Phillipa Gill. 2018. How to Catch When Proxies Lie: Verifying the Physical Locations of Network Proxies With Active Geolocation. In *Proceedings of the ACM Internet Measurement Conference (IMC)*. ACM, New York, NY, USA, 203–217.
- [81] Marcus Willett. 2022. The Cyber Dimension of the Russia–Ukraine War. *Survival* 64 (2022), 7–26.
- [82] Wired. 2022. Ukraine’s Volunteer ‘IT Army’ Is Hacking in Uncharted Territory.
- [83] Wired. 2023. Ukraine Suffered More Data-Wiping Malware Last Year Than Anywhere, Ever.
- [84] Hyung-jin Woo, Yeora Kim, and Joseph Dominick. 2004. Hackers: Militants or Merry Pranksters? A Content Analysis of Defaced Web Pages. *Media Psychology* 6 (2004), 63–82.
- [85] Wordfence. 2022. Ukraine Universities Hacked as Russian Invasion Started.
- [86] ZDNET. 2022. CaddyWiper: More Destructive Wiper Malware Strikes Ukraine.

A ETHICAL CONSIDERATIONS

This work is presented objectively to minimise risk to researchers. The collection, sharing, and analysis of web defacement, amplified DDoS attacks, and Telegram chats have been formally approved by the Department of Computer Science & Technology’s Research Ethics Committee. We do not attempt to gather private data; only publicly accessible data are collected. A 2022 US court ruling suggests scraping public data is legal [70]; our scraper does not overload websites. The amplified DDoS attack honeypots absorb attack packets without relaying them, thus reducing harm to victims.

Studying an ongoing conflict may harm individuals whose attacks are reported, while researchers might face retaliation from attackers due to leaking insights into their activities and community. To avoid potential harm, our experiments operate ethically and collectively, only presenting aggregated findings without identifying individuals. We did not ask for consent from Telegram users or web defacers when using scraped data, as sending thousands of messages would be impractical. We believe they are aware that content posted online will be publicly visible. This approach accords with the British Society of Criminology’s Statement on Ethics [10].

B DATA LICENSING

We have robust procedures and long experience in licensing our data in various jurisdictions. Our quantitative data, analysis scripts, and scrapers are available for academic researchers under a license agreement with the Cambridge Cybercrime Centre to prevent misuse and to ensure the data will be treated ethically, as access to sensitive data might risk researchers and the actors involved [17, 79].

C DETERMINING ATTACK GEOLOCATION

Accurately mapping IP addresses to countries is challenging as IP geolocation is not always stable and trustworthy [27]; providers prefer locating servers in countries with cheap hosting [80]. Attack geolocation can thus be determined differently e.g., ZONE-H says an IP is in Germany, while ZONE-XSEC detects Singapore and DEFACER-ID cannot tell. IP geolocation may be more reliable at the country level [12], but this is only part of the truth as websites are nowadays commonly hosted on content delivery networks (CDNs), where original IP addresses are hidden and the geolocation is of the CDNs.

For example, a ‘.ru’ website is supposed to be Russian, but it might be physically hosted in Vietnam, operated by a person living in Hong Kong, while proxied through Cloudflare with an IP address in the US. Relying on only one aspect might be risky, as both IP and domain can lie. We use data fusion to enhance the accuracy, prioritising: (1) top-level domain; (2) IP geolocation at collection time (MaxMind GeoIP2² for web defacements, and a database we maintain based on Regional Internet Registry data for DDoS attacks; and (3) geolocation of the AS hosting the IP address. If a website’s IP address belongs to a CDN, its geolocation is determined solely by ccTLD, as any geolocation of IP address or ASN will be unreliable.

The top three CDNs are Cloudflare, Amazon Web Services, and Akamai, serving around 89% of customers [33]. We ignore the trivial market shares of their competitors but count DDoS-Guard as it is Russian-based, which may affect the measurement. We expect the four to cover over 90% of customers. We found 4.87% of defacements are hosted on these CDNs by 14 262 prefixes as of the writing date: 1 698 of Cloudflare; 7 483 of Amazon Webservice; 5 056 of Akamai; and 25 of DDoS-Guard (these prefixes and AS number mappings are collected on Hurricane Electric Internet Services). For defacements, we prefer ccTLDs over IP geolocation as attackers often target sites by massively scanning domain ccTLDs, such as ‘.ru’ and ‘.ua’, rather than checking whether IP addresses are hosted in those countries.

² GeoIP2 is freely accessible at <https://maxmind.com/>. It offers both free and paid licenses, with the paid one being slightly more accurate and up-to-date. It claims to provide over 99.8% country-level and over 60% city-level accuracy, yet that varies from country to country e.g., 79% for Russia and 65% for Ukraine, within a 250km radius.

Accurate measurement of frequent ccTLDs used in Russia and Ukraine is complex; many Ukrainian firms use Russian services and vice versa. The domain most frequently used in Russia is '.ru' [67]. A similar report for Ukraine is unavailable, but we believe incorporating ccTLDs with IP and AS geolocation is reasonable as choosing targets based on ccTLDs is a common way used by low-lever actors.

D DEFACEMENT SUBMISSION PROCESS

The defacement submission is mostly automatic: users specify a 'notifier', team, defaced URL, vulnerability types, and hacking incentives. At that point, a record is made with details of the compromised system, its IP address and location, and a snapshot of the defaced page (often consisting of messages that may include political and ideological propaganda [7]). Messages can be hidden using identical font colours as the background, but are detectable through HTML. New reports are kept away from the dashboard until being verified by staff or bots. Although 'notifier' can be arbitrarily entered, defacers are incentivised to use consistent handles to cultivate fame and reputation; we thus consider 'notifier' to be reliable enough to differentiate them. The defacement snapshots, including messages left, are highly reliable as they are captured at reporting time.

E WEBSITE DEFACEMENT COLLECTION

Data completeness and reliability are critical for longitudinal measurement. Scraping *complete* defacement archives at scale, especially ZONE-H, is non-trivial, and was not guaranteed in prior work. Some purchased the data [44], but it is not sustainable and is ethically questionable. This is challenging as (1) ZONE-H uses CAPTCHA to prevent bots, (2) its dashboard sets a limit of 50 pages where older data is hidden, and (3) on-hold records may not appear promptly, leading to potential misses. The only way to get a complete scrape is by iterating all submission IDs, with IDs of valid and invalid reports often mixed. This issue, along with IP blacklisting and bot prevention, generates a non-trivial workload. We responded by (1) developing an efficient text CAPTCHA solver for ZONE-H using image-processing techniques, (2) routing our scraper through multiple proxies, and (3) iterating all submission IDs in turn. Raw data is stored in a database to avoid unnecessary future requests.

Five most trusted archives are included; an active one DEFACER-ID (since February 2016) is excluded as (1) the valid submission volume during the period is small (less than 27k); (2) unclear staff verification, no validity sanitisation on submission, no validity signal in defaced pages (in fact, over half of these have been deemed invalid by the archive); (3) defaced snapshots and defacers' messages are missing; and (4) victim geolocation is mostly lacking; determining it after the fact is problematic as sites could have been relocated.

F VALIDATING ON-HOLD DEFACEMENTS

How defacement reports are validated is not clearly stated. ZONE-H reports are kept on hold until being manually verified by staff; ZONE-XSEC, DEFACER-PRO, and HAXOR-ID use automatic validation, insisting messages left on defaced pages contain keywords linked to hacking activities (e.g., 'Hacked by ABC'); while it is unclear for OWNZYOU. Defacers may game the system by putting comments on blogs, or submitting search queries (e.g., ?search='Hacked by ABC'), which

Algorithm 1 Semi-automatic website defacement validation

```

1: procedure VALIDATE_DEFACEMENTS
2:   for each  $a \in \text{verifiedDefacements}()$  do           ▶ verified by archives
3:      $a.\text{status} \leftarrow 0$                              ▶ originally validated
4:   end for
5:   for each  $a \in \text{filteredDefacements}()$  do           ▶ filtered by terms
6:      $a.\text{status} \leftarrow 1$                              ▶ automatically validated
7:   end for
8:    $P \leftarrow \text{pendingGroups}()$                        ▶ groups of pending attacks
9:    $V \leftarrow \text{verifiedGroups}()$                      ▶ groups of verified attacks
10:  for each  $p \in \mathcal{P}$  do
11:     $T \leftarrow \{\}$ 
12:    for each  $v \in \mathcal{V}$  do
13:       $d \leftarrow \text{levenshtein}(p, v)$                  ▶ similarity with verified ones
14:       $T \leftarrow \text{topSimilar}(d, T)$                  ▶ extract top similar ones
15:    end for
16:     $\text{showSimilarDefacements}(T)$                        ▶ to assist the annotator
17:     $s \leftarrow \text{annotation}()$                          ▶ annotate the validity
18:    for each  $a \in p$  do                               ▶ each defacement in this group
19:       $a.\text{status} \leftarrow s$                            ▶ update validation status
20:    end for
21:    if  $\text{isValidated}(s)$  then                           ▶ if it is manually validated
22:       $V \leftarrow V \cup p$                              ▶ add to validated groups
23:    end if
24:  end for
25: end procedure

```

occasionally get through automatic sanitisation but our further validation excludes them. Manual staff review on ZONE-H may be slow, while automatic verification of the others is error-prone. Unverified records may be kept on hold forever, leading to incomplete data. Consequently, collecting only defacements shown in the dashboard is inadequate, making a complete dataset challenging to gather. To enhance data completeness and reliability, a semi-automatic validation is performed to check if on-hold reports are in fact valid.

Our strategy is shown in Algorithm 1. First, reports verified by archives are considered valid. Second, messages on the defaced pages of on-hold submissions are used to decide the validity, as defacers often leave signatures for reputation e.g., 'Hacked by CoolHacker'. If messages include defacer handles and specifically contain common hacking terms: 'hacked by', 'h4ck3d by', 'h4cked by', 'p4wn3d by', 'pwn3d by', 'pwnd by', 'pwned by', 'pwndz by', 'owned by', 'own3d by', 'touched by', and 'kissed by', we consider them to be valid e.g., a message 'This website was hacked, contact me t.me/coolhacker' posted by a notifier 'CoolHacker' is considered valid. This method is looser than an exact comparison with 'Hacked by CoolHacker', but is still highly accurate; 100 randomly checked samples were all correct. Third, the remaining submissions are manually validated by looking for defacer signatures; some are obvious but some are complicated. Candidates are grouped by normalised handles and messages (redundant spaces removed), and then for each, the 10 most similar validated defacements are suggested to the annotator. Levenshtein distance is used to estimate the similarity between two messages, which is helpful as messages are often slightly modified from templates. If no message is found (instead images, iframes, or javascript), or leftover signatures cannot be spotted, a browser opens the defaced page to assist annotators.

Table 2: Significance levels of the impact on daily defacements and defacers targeting Russia, Ukraine, and top countries.

Country	Tests for the number of web defacements per day					Tests for the number of web defacers per day				
	ANOVA / Kruskal-Wallis	$\langle E_1, E_2 \rangle$	$\langle E_1, E_3 \rangle$	$\langle E_2, E_3 \rangle$	η^2	ANOVA / Kruskal-Wallis	$\langle E_1, E_2 \rangle$	$\langle E_1, E_3 \rangle$	$\langle E_2, E_3 \rangle$	η^2
Russia	$H(2) = 12.24, p < .01$	$p < .05$	$p = .3544$	$p < .001$	0.06	$H(2) = 26.57, p < .0001$	$p < .0001$	$p = .9083$	$p < .0001$	0.14
Ukraine	$H(2) = 17.86, p < .001$	$p < .001$	$p = .8377$	$p < .0001$	0.09	$H(2) = 13.64, p < .01$	$p < .01$	$p = .9286$	$p < .001$	0.07
The US	$H(2) = 17.84, p < .001$	$p < .01$	$p = .1435$	$p < .0001$	0.09	$H(2) = 24.30, p < .0001$	$p < .001$	$p = .5961$	$p < .0001$	0.13
Brazil	$H(2) = 3.60, p = .1656$	$p = .6481$	$p = .1725$	$p = .0912$	0.01	$H(2) = 11.68, p < .01$	$p = .3405$	$p < .05$	$p < .01$	0.05
Germany	$H(2) = 3.43, p = .1796$	$p = .2339$	$p = .5269$	$p = .0639$	0.01	$F(2, 178) = 3.24, p < .05$	$p = .7584$	$p = .1858$	$p = .0568$	0.04
India	$H(2) = 4.21, p = .1221$	$p = .9049$	$p = .0670$	$p = .1423$	0.01	$H(2) = 3.90, p = .1424$	$p = .0746$	$p = .8734$	$p = .0704$	0.01
Indonesia	$H(2) = 10.90, p < .01$	$p < .01$	$p = .1566$	$p < .05$	0.05	$H(2) = 17.93, p < .001$	$p < .001$	$p = .5517$	$p < .001$	0.09
Canada	$H(2) = 8.13, p < .05$	$p = .0944$	$p = .2458$	$p < .01$	0.03	$H(2) = 9.51, p < .01$	$p = .1020$	$p = .1492$	$p < .01$	0.04
Turkey	$H(2) = 20.07, p < .0001$	$p = .1171$	$p < .0001$	$p < .05$	0.10	$H(2) = 13.26, p < .01$	$p = .5501$	$p < .001$	$p < .05$	0.06
Singapore	$H(2) = 3.83, p = .1473$	$p = .7583$	$p = .0677$	$p = .2085$	0.01	$H(2) = 5.90, p = .0524$	$p = .3056$	$p < .05$	$p = .3218$	0.02

Table 3: Significance levels of the impact on daily DDoS attacks and victims targeting Russia, Ukraine, and top countries.

Country	Tests for the number of DDoS attacks per day					Tests for the number of DDoS victims per day				
	ANOVA / Kruskal-Wallis	$\langle E_1, E_2 \rangle$	$\langle E_1, E_3 \rangle$	$\langle E_2, E_3 \rangle$	η^2	ANOVA / Kruskal-Wallis	$\langle E_1, E_2 \rangle$	$\langle E_1, E_3 \rangle$	$\langle E_2, E_3 \rangle$	η^2
Russia	$H(2) = 60.67, p < .0001$	$p < .0001$	$p < .0001$	$p < .0001$	0.33	$H(2) = 57.13, p < .0001$	$p < .0001$	$p < .01$	$p < .0001$	0.31
Ukraine	$H(2) = 12.59, p < .01$	$p < .01$	$p = .4593$	$p < .001$	0.06	$H(2) = 15.16, p < .001$	$p < .001$	$p = .8765$	$p < .001$	0.07
The US	$H(2) = 6.98, p < .05$	$p = .5592$	$p < .05$	$p = .1182$	0.03	$H(2) = 4.43, p = .1093$	$p = .2527$	$p < .05$	$p = .5594$	0.01
Brazil	$H(2) = 9.81, p < .01$	$p < .01$	$p = .2006$	$p < .05$	0.04	$H(2) = 13.12, p < .01$	$p < .001$	$p < .05$	$p < .05$	0.06
Germany	$H(2) = 9.49, p < .01$	$p < .01$	$p < .01$	$p = .6609$	0.04	$H(2) = 17.24, p < .001$	$p < .001$	$p < .001$	$p = .2039$	0.09
Bangladesh	$H(2) = 3.96, p = .1379$	$p < .05$	$p = .2198$	$p = .2785$	0.01	$H(2) = 4.43, p = .1090$	$p < .05$	$p = .1353$	$p = .3585$	0.01
China	$H(2) = 80.16, p < .0001$	$p < .001$	$p < .0001$	$p < .0001$	0.44	$H(2) = 65.91, p < .0001$	$p < .0001$	$p < .0001$	$p = .0674$	0.36
France	$H(2) = 16.04, p < .001$	$p = .2586$	$p < .001$	$p < .05$	0.08	$H(2) = 9.96, p < .01$	$p = .1519$	$p < .01$	$p = .2366$	0.04
UK	$H(2) = 13.90, p < .01$	$p = .4892$	$p < .001$	$p < .05$	0.07	$H(2) = 7.94, p < .05$	$p = .5258$	$p < .01$	$p = .0976$	0.03
Poland	$H(2) = 0.34, p = .8423$	$p = .8841$	$p = .6759$	$p = .6002$	0.00	$H(2) = 0.04, p = .9809$	$p = .9081$	$p = .9365$	$p = .8449$	0.00

This assistance effectively reduces the annotation effort. One challenge is defaced pages redirecting to another page that can be modified dynamically: when they are down, the submission may point to a non-existent site, but a careful check could reveal evidence of defacers. We also consider a site is touched if its content is unchanged but the page title is modified to indicate hacking activities. We ignore cases that lack evidence to ensure those flagged valid are indeed valid. We do not use complex machine learning techniques as message texts contain lots of noise; given a small number of samples (around 10k), machine learning is not more effective than a rule-based approach. Sometimes, defacements appeared to be already verified at the collection time, but became invalid afterwards; we re-validated them months after the initial collection to make sure their status had been finalised by the archives.

G UNIFYING DEFACTMENTS AND DEFACERS

We hashed then unified defacements across all archives based on reporting dates, original defacer handles, root victim domains, and messages left on the defaced page. Hashing reporting dates may cause repeat counts if defacers resubmit to other archives after a few days, but excluding them may lead to missing defacements of the same URL on different days due to repeat victimisation. We also unified defacers across all archives, as users tend to pick similar pseudonyms on different platforms [25]. As the unification needs

to be accurate and the number of unique defacers does not exceed a few thousand, machine learning is not appropriate. We instead used a semi-automatic approach combining automated handle similarity analysis with manual review. As typos may occur (e.g., missing characters, character orders, case sensitivity), similar handle pairs are first extracted using Levenshtein distance, which is set to not exceed 25% of the length. Ten messages left by defacers in each pair are then sampled to assist the annotation. A pair of handles are unified under a single nickname if their messages are closed enough, assessed based on handle appearances, message semantics, stylometry, synonyms, typos, team, nationality, languages, and handle rarity (rare ones such as ‘cj2ks’ are more likely to be used by a single person, while common ones such as ‘glory’ are more likely to be shared by multiple individuals [42]). Messages left are diverse: many defacers leave identical messages on different archives, while some are relatively similar, some are distinct and contain the defacer names, and some consist of their phone numbers. We only confirm when having sufficient evidence, uncertain pairs are left unmatched.

H STATISTICAL SIGNIFICANCE OF IMPACT

Tables 2 and 3 detail the statistical significance levels of the impact on defacement and DDoS attack counts over the three eras, using One-way ANOVA or Kruskal-Wallis tests. Post-hoc analysis Tukey-Kramer is used for ANOVA, while Dunn’s is used for Kruskal-Wallis.