

Helping hands: Measuring the impact of a large threat intelligence sharing community

Bouwman, X.B.; Le Pochat, Victor; Foremski, Pawel; Van Goethem, Tom ; Hernandez Ganan, C.; Moura, Giovane C.M.; Tajalizadehkhoob, Samaneh; Joosen, Wouter; van Eeten, M.J.G.

Publication date

2022

Document Version

Final published version

Published in

31st USENIX Security Symposium

Citation (APA)

Bouwman, X. B., Le Pochat, V., Foremski, P., Van Goethem, T., Hernandez Ganan, C., Moura, G. C. M., Tajalizadehkhoob, S., Joosen, W., & van Eeten, M. J. G. (2022). Helping hands: Measuring the impact of a large threat intelligence sharing community. In *31st USENIX Security Symposium* (pp. 1149-1165). USENIX Association. <https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Helping hands: Measuring the impact of a large threat intelligence sharing community

Xander Bouwman, *Delft University of Technology*; Victor Le Pochat, *imec-DistriNet, KU Leuven*; Pawel Foremski, *Farsight Security, Inc. / IITiS PAN*; Tom Van Goethem, *imec-DistriNet, KU Leuven*; Carlos H. Gañán, *Delft University of Technology and ICANN*; Giovane C. M. Moura, *SIDN Labs*; Samaneh Tajalizadehkhoob, *ICANN*; Wouter Joosen, *imec-DistriNet, KU Leuven*; Michel van Eeten, *Delft University of Technology*

<https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman>

This paper is included in the Proceedings of the
31st USENIX Security Symposium.

August 10–12, 2022 • Boston, MA, USA

978-1-939133-31-1

Open access to the Proceedings of the
31st USENIX Security Symposium is
sponsored by USENIX.

Helping hands: Measuring the impact of a large threat intelligence sharing community

Xander Bouwman¹, Victor Le Pochat², Pawel Foremski³, Tom Van Goethem², Carlos H. Gañán^{1,5},
Giovane C. M. Moura⁴, Samaneh Tajalizadehkhoob⁵, Wouter Joosen², and Michel van Eeten¹

¹Delft University of Technology

²imec-DistriNet, KU Leuven

³Farsight Security, Inc. / IITiS PAN

⁴SIDN Labs

⁵ICANN

Abstract

We tracked the largest volunteer security information sharing community known to date: the COVID-19 Cyber Threat Coalition, with over 4,000 members. This enabled us to address long-standing questions on threat information sharing. First, does collaboration at scale lead to better coverage? And second, does making threat data freely available improve the ability of defenders to act? We found that the CTC mostly aggregated existing industry sources of threat information. User-submitted domains often did not make it to the CTC's blocklist as a result of the high threshold posed by its automated quality assurance using VirusTotal. Although this ensured a low false positive rate, it also caused the focus of the blocklist to drift away from domains related to COVID-19 (1.4%-3.6%) to more generic abuse, such as phishing, for which established mitigation mechanisms already exist. However, in the slice of data that was related to COVID-19, we found promising evidence of the added value of a community like the CTC: just 25.1% of these domains were known to existing abuse detection infrastructures at time of listing, as compared to 58.4% of domains on the overall blocklist. From the unique experiment that the CTC represented, we draw three lessons for future threat data sharing initiatives.

1 INTRODUCTION

For years now, research has consistently found that threat data feeds each cover just a fraction of the landscape. Numerous comparisons have been made among different threat intelligence sources and they all find very little overlap: feeds are dominated by data points that appear only in a single source and in no other one. This holds across the spectrum of threat data, from freely available blocklists and abuse feeds [1, 2] to closed industry sources [3], all the way up to the most expensive feeds at the high end of the market [4]. It points to poor coverage of the threat landscape, a problem for the industry and – more importantly – for its customers. To illustrate: anti-phishing companies missed a large portion of

the phishing sites targeting their customers, while those sites were being discovered by their competitors [5].

While prior work on threat intelligence sources demonstrated the problem of low coverage, it does not discuss how it could be overcome. One obvious and often proposed solution is more data sharing. This typically takes place in informal trusted communities of specialists or in formalized sharing agreements among firms, such as the Cyber Threat Alliance [6] and the Anti-Phishing Working Group [7], which often demand reciprocity, or ‘quid pro quo’, in order to avoid free-riding behavior. This forms a high entry barrier for access to the shared data, as not everyone has enough to contribute in order to gain access, and so benefits are typically limited to corporate entities. Non-participants can only get access to the separate services of these firms. In economic terms, such sharing arrangements create club goods, not public goods.

A potentially more effective form of information sharing would go beyond these boundaries: open to any contributor, with free access for anyone to the pooled data. Under normal conditions, market incentives would prevent such a public good from emerging. Generating quality threat information costs money and firms have to recoup their investments. However, the 2020 global pandemic gave rise to a real-world experiment that temporarily suspended normal economics in data sharing: the COVID-19 Cyber Threat Coalition (CTC) [8]. The CTC was a volunteer-based response where firms and individuals shared threat intelligence on pandemic-related threats posed by cybercriminals as well as nation states [9]. The coalition's mission was “to operate the largest professional-quality threat lab in the history of cybersecurity” and reduce gaps in availability and coverage of existing defense mechanisms [8, 10]. After 3 months, over 4,000 individuals and organizations had signed up, with companies like Symantec, Microsoft and Cofense contributing data [11, 12]. Contrary to sharing based on quid pro quo, the resulting blocklist was freely available to anyone.

Does such large-scale open data sharing actually improve our defenses against threats? To the best of our knowledge, research on the effectiveness of large data-sharing arrangements

is extremely sparse. More than a decade ago, Moore and Clayton studied the inner workings of crowdsourcing at PhishTank, but not its impact on mitigation [13]. Thomas et al. found that a Google-based threat exchange could have more impact than existing standalone anti-abuse pipelines [14] but access to the pooled data in the exchange was restricted. A recent study on VirusTotal did not analyse the information-sharing aspects of the service nor its impact, but evaluated the aggregation of detection results of the participating vendors [15].

In this paper, we set out to extend the literature on data-sharing arrangements by learning from the CTC experiment. We investigate two questions: (i) By pooling data from its community, did the CTC improve coverage of COVID-19-related threats over existing defenses? And (ii) Did publishing threat data in a freely available blocklist improve the ability of defenders to act against threats, compared to the existing abuse mitigation infrastructure? To answer these questions, we first describe the organizational setup of the CTC and observe how the community pooled data and conducted quality assurance. We then evaluate the blocklist through manual classification of a sample and by identifying false negatives for COVID-19-related domains. Next, we conduct longitudinal measurements to infer *who* acted *when* against domains on the CTC blocklist: registry, registrar, browser vendor, or security provider? We end by identifying key lessons for improving the impact of open large-scale data sharing mechanisms.

In sum, our contributions are:

- We present the first empirical evaluation of a large and open threat sharing community, the COVID-19 Cyber Threat Coalition, and describe its mechanisms for producing and vetting threat intelligence.
- We find that user contributions were heavily skewed, with just 10 users making 90% of contributions. Further, the high threshold posed by the indicator vetting process, which relied on VirusTotal, resulted in a mere 5.14% of user-contributed indicators actually being propagated to the blocklist. Instead, most data came from commercial firms, and at least part of this was vetted against a lower threshold.
- We show that the CTC list went well beyond the scope of just COVID-19-related threats. Generic phishing made up a large portion of the blocklist: domains containing the word *whatsapp* (2.8%) outnumbered those containing keywords related to COVID-19 (2.6%).
- Based on longitudinal measurements, we demonstrate that for 58.4% of the domains on the CTC blocklist, existing abuse-mitigation mechanisms were faster: domain-level or client-side interventions had already taken place before the domains appeared on the blocklist. For COVID-19-related domains, this share is smaller: 25.1%, which is evidence of added value of the community. The remaining portion was intervened against later or not at all. This means that the blocklist improved the ability of defenders to protect themselves. Its impact could have been larger, if it had not depended as much on VirusTotal for vetting indicators.

2 THE CTC COMMUNITY

The CTC started as a Slack community on March 19 – a week after the WHO declared COVID-19 a pandemic. Joshua Saxe, a security specialist at Sophos, founded the community out of a “personal sense of alarm”, conceiving it as a “crisis commons model” where “traditional competition and grievance [were] set aside in a moment of exceptional need” [16]. The CTC founders articulated objectives for the community and then set up community support services for volunteers to join and participate. We also joined and revealed ourselves as researchers interested in learning more about how the community was functioning. The CTC’s mission was threefold: fostering collaboration across organizations to uncover otherwise missed threats, producing professional-grade output that the community can rely on, and prioritizing the public good over the interests of individual actors [8].

2.1 Information products

The CTC set out to publish an open high-quality blocklist containing COVID-19-themed abuse to “supplement the existing defensive structure” by setting up a “separate threat intelligence platform dedicated just to pandemic-related cyber activity” [10, 17]. Blocklist data was published on the CTC’s website from March 29, 2020 onwards [18]. In May, it reported that over 60,000 distinct IP addresses had been consuming the blocklist [19]. Initially, they offered two types of lists: one with vetted indicators of compromise (IOCs) and a larger list with unvetted IOCs submitted by the community. The latter was eventually discontinued to “produce the highest quality feeds with the least amount of false positives possible” [20]. Within the vetted category, four blocklist files were available: for domains, URLs, IP addresses and file hashes. The latter two, however, have remained empty at all times. Our analysis focuses on what is arguably the main list, i.e., the domain list.

The technical indicators were supplemented by threat advisories and community meetings. The advisories featured themes like phishing and ransomware related to COVID-19, trends in pandemic-related domain registrations, and the security challenges of remote work. They were published on the CTC website¹ and sent out via a mailing list every week from April 6 until May 26. From April 16 onward, community leaders also organized online community meetings or ‘town halls’ hosted on Zoom and archived on YouTube [21]. These webinars were intended to disseminate findings from the reports, to update the community on changes to the CTC procedures and infrastructure, to interview representatives of security companies, and to allow for Q&A with the community leaders. For the first two months, town halls were organized weekly; after June 11, the CTC planned to organize them every two weeks, but no town halls took place since.

¹<https://www.cyberthreatcoalition.org/advisories>

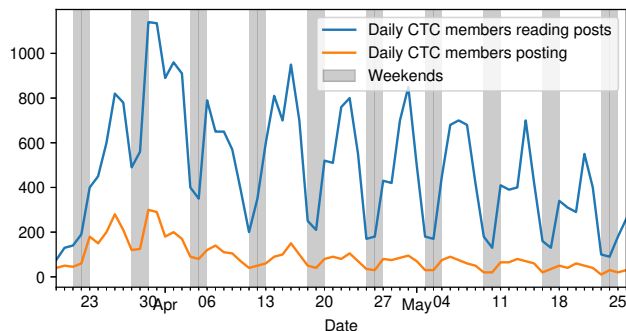


Figure 1: Member activity on the CTC Slack workspace peaked 1.5 weeks after the community was founded [19].

2.2 Community structure

The primary means of collaboration in the Cyber Threat Coalition was a Slack workspace. The aim of this Slack workspace, which anybody could join, was for members of the security community to share information and make the necessary contacts for interventions [22]. After a brief period of unstructured posts, the workspace was organized into themed channels for posting community updates, IOCs per topic (e.g., web, email, malware), and region-based networking. Membership grew to 1,500 in one week [23] and to over 4,000 after 3 months [24]. Members were requested to add their affiliation and job title to their screen name, although this information was not validated. Of the users who indicated a country in their username, around half came from North America, around a third from Europe (mainly the United Kingdom and the Netherlands), 3% from Australia, with the rest scattered around the world. Based on stated affiliations, the CTC counted representatives of security vendors, healthcare providers, CERTs, financial institutions, domain name infrastructure providers, major technology firms and law enforcement agencies among its members. Activity peaked at the end of March and then slowly decreased (Figure 1). By July, the public channels saw much less activity, with daily posts in the single digits, and no official announcements were made except for a single update in February 2021 stating that the CTC was in “low battery mode” but “evolving”.

The community had a simple organizational structure [19]: a ‘steering committee’ functioned as administrators for the community, with its members committed to “sacrificial” time contributions [22], leading teams for various tasks like vetting process development, advisory writing, and media outreach. Within the larger community, over 100 ‘vetted volunteers’ [25] underwent identity verification via their social media profiles to obtain access to private channels, where more sensitive data could be shared [22].




Figure 2: Counts of unique domains newly seen in the CTC AlienVault OTX group and on the CTC blocklist (logarithmic scale), and the proportion of AlienVault OTX domains that were propagated to the CTC blocklist.

3 PRODUCTION OF THE BLOCKLIST

In this section, we describe how the CTC’s prime information product was sourced and vetted. Four sources fed the CTC blocklist: (i) user contributions on Alienvault OTX; (ii) user contributions via a Slackbot; and industry lists from both (iii) named and (iv) unnamed vendors.

Initially, TI sharing was ad hoc, consisting of community members posting free-form IOCs, first in the general Slack channel and soon after in field-specific channels (e.g., for email and domains). After one week, the CTC set up groups for machine-readable TI sharing on AlienVault OTX [23]. Users could submit IOCs to the CTC Slack workspace. After they had been checked for maliciousness, indicators were published on the ‘vetted’ list [22]. In June, the CTC introduced a Slackbot that let users contribute indicators as well as vet indicators supplied by other users [11]. We describe these two vetting mechanisms in subsection 3.2. Further, commercial security providers contributed “hundreds of millions of indicators per day” [22] outside of the AlienVault group, although these were also subject to the same vetting procedure. These indicators appear to represent the vast majority of domains on the final vetted blocklist, as demonstrated in Figure 2. Known industry contributors to the CTC are Symantec [12], Microsoft and Cofense [11], but other vendors have asked to remain anonymous [22].

3.1 User contributions on AlienVault OTX

We analyze the indicators contributed to the CTC group on AlienVault OTX, using the API to gather all submissions over time. By July 2020, 738 users had been accepted into the CTC’s closed AlienVault OTX group. Only 47 users actually contributed IOCs – 10 of whom made 90% of all contributions (Figure 3). The two heaviest contributors did so on a fixed schedule, and described drawing on newly registered domains and certificate transparency for their lists. Others added more opportunistically, with a downward trend over time. In general, recipients were often left guessing as the source or method behind the contributed indicators.




Figure 3: Daily count of unique domain indicators contributed to the CTC AlienVault OTX group, with the top 10 and other contributors ranked on the total number of contributions over time.

In other words, even though the community has over 4000 members, only 1–2% of them contributed indicators, with 10 users contributing the bulk. This skewed distribution resembles the pattern found in the PhishTank community project, where volunteers collect and classify phishing sites [13]: there, the top 10 contributors made 69.9% of submissions and 57.4% of votes [26]. It also resembles free and open source software (FOSS) development, where a tiny fraction of contributors supplies most code and most others only contribute once, typically in the form of a bug report [27]. As we discuss in the next section, most user contributions to the CTC never made it past the vetting stage (see also subsection 3.2). At the outset, a large proportion of AlienVault submissions made it to the CTC blocklist, but after the list was reset on April 12 (subsection 4.2), domains added to the CTC blocklist came mostly from named and unnamed industry sources. Over the entire period of our data collection, just 1.23% of domains on the blocklist came from user contributions on OTX.

3.2 Quality assurance

The community admins instituted vetting mechanisms for submitted indicators, in order to “provide reasonable assurance that what we re-share with the public are examples of truly malicious artifacts” [16]. Initially, vetting consisted of a pool of volunteers manually verifying maliciousness of submitted indicators. The workflow was eventually automated using a security orchestration service that integrated VirusTotal: if a domain received 10 or more hits, this would lead to a domain being marked as malicious and propagated to the blocklist. A domain with between 4 and 10 hits would require manual review, while for fewer than 4 hits, it would be marked as “clean” [11] and dropped.

Although automation allowed for higher volumes of indicators to be processed, this workflow left the community with three problems. First, only indicators that were already known to be malicious by many VirusTotal scanners could be added to the blocklist. Earlier research has consistently found

very low overlap among TI sources [1, 4], so requiring that indicators are found by 10 scanners imposes a high threshold to propagation. It meant that just 5.14% of domains from the AlienVault OTX group made it onto the blocklist. Moreover, VirusTotal aggregates labels from 84 established automated scanners, negating the community contribution aspect of the CTC. The second problem: some of the contributed industry sources, in particular Cofense [11], were seen as reliable, yet their indicators did not appear on the vetted blocklist as too few VirusTotal scanners flagged them. As a solution, the CTC admins lowered the thresholds for these trusted sources [11].

The third problem was how to evaluate the indicators that fell in-between the thresholds for ‘clean’ or ‘malicious’. At first, this was done manually by a small pool of volunteers. It is unlikely that they could keep up with the overall volume, so probably these indicators did not make it onto the vetted list. On June 11, the CTC admins announced a new Slackbot that would allow all members to contribute to vetting [11]. Upon request, the bot would serve a domain or URL to an individual user for evaluation. It resembled crowdsourcing mechanisms like PhishTank, albeit without the built-in validation of multiple users checking the same indicator. The bot potentially increased the scalability of manual vetting and would allow indicators to be included long before VirusTotal would provide enough ‘hits’. The downside was that vetting might be done by members with unverified expertise or even potentially adversarial motives. Unfortunately, when the crowdsourcing functionality arrived by June, the peak of user activity had already passed, so it came too late to actually change the vetting process. Despite this process, some false positives slipped through, leading domain owners to join the CTC Slack workspace and request removal from the blocklist.

4 EVALUATION OF THE BLOCKLIST

In this section we address our first research question: *By pooling data from its community, did coverage of COVID-19-related threats improve over existing defenses?* First, we manually label a sample of domains to evaluate the nature of new IOCs appearing on the CTC blocklist. Second, we track the evolution of the blocklist composition, and measure its focus on COVID-19-related abuse. Third, we measure the role and impact of VirusTotal on the vetting process. Fourth, we estimate the coverage of COVID-19-related domains through a comparison with external sources.

4.1 Manual classification of domains

Our first assessment of the blocklist content is to manually inspect a sample of domains. Over the course of 5 days (17–22 May 2020), we took a daily sample of 50 domains that were newly added to the CTC blocklist. We visited those within 4 hours of their appearance on the list, in order to minimize the

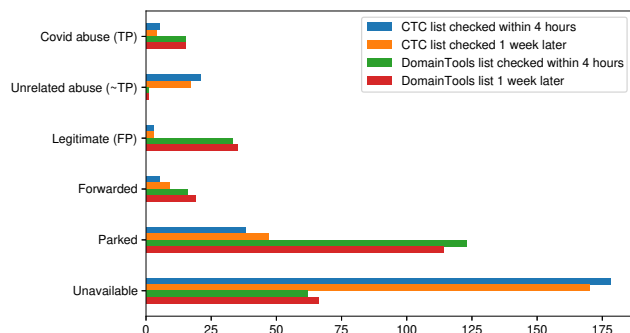


Figure 4: Manual classification of 500 domains, visited at two moments in time.

chance that the domain would already be affected by a countermeasure. We navigated to the site with Microsoft Internet Explorer 11 in a virtual machine on a computer located in the Netherlands. If necessary, we translated page contents using Google Translate. After one week, we visited each site again to see if anything had changed. As a reference, we carried out this process not only for the CTC blocklist but also for a COVID-19 blocklist published by DomainTools [28].

We based our label taxonomy on existing classifications [29, 30]. In many cases, the labeling could be applied in a straightforward manner, such as with parked domains. Some categories contain more ambiguity. When we were uncertain about what the appropriate label was, we took a conservative approach and counted them as true positives, giving the blocklist publishers the benefit of the doubt. As shown in Figure 4, out of 250 domains visited from the CTC list, we encountered just 5 cases of COVID-19 related abuse (2%), including phishing sites themed with the pandemic and sites selling dubious protective materials. The CTC list also included 21 examples (8%) of internet abuse that was not visibly tied to COVID-19, such as generic phishing sites, counterfeit products and pharmaceuticals – although it is conceivable that these domains were in fact used in a campaign that somehow played upon COVID-19, for example in a spam run. We deemed 3 websites as legitimate – and therefore false positives on the list (1%). The majority of domains were unavailable (71%).

The DomainTools list contained more examples of actual COVID-19 related abuse (6%), but at the cost of more false positives (13%), many of which had in common that they contained keywords related to the pandemic, such as a Wuhan-based welding hardware supplier.

When revisiting the domains after one week, we saw minor changes (4% overall), most of which were previously unavailable domains becoming reachable as parking pages. We saw seven examples (1.5% overall) of generic abuse not related to COVID-19 becoming active in this interval of one week.




Figure 5: Evolution of the composition of the CTC vetted domain blocklist.

4.2 Composition of the blocklist

Where subsection 4.1 analyzed a sample, further analysis in this paper is based on the full CTC blocklist of vetted domains from March 31 up to July 1, 2020, during the peak of community activity. We retrieved this list once a day starting March 31, every hour starting April 4, and every five minutes from April 15 onward. Over time, the vetted domain blocklist steadily grew (Figure 5), reaching 46,103 domains on July 1, 2020. Cumulatively, 46,832 fully qualified domain names (FQDNs) on 27,096 unique second-level domain names (SLDs; i.e., domains that can be bought from a registrar) appeared on the vetted list by July 1, 2020. Domain removals were rare, except when the list was reset on April 12 due to the shift towards automated vetting [24] which suggests that stale entries were not purged from the list, potentially generating false positives after cleanup or takedown of a malicious domain². Further, we suspect that back-end changes caused a temporary eight-day blockage in updates, followed by the addition of 5,918 domains on June 22, as users reported download issues around the same time [24]. The list continued to be updated until December 6, 2020, albeit without much community input, after which the list remained available, but no longer changed, indicative of dwindling efforts in the CTC.

Given the CTC’s objective to track COVID-19-related abuse, we examine the coverage of the blocklist in this area. For this purpose, we generated 370 COVID-19-related keywords in 15 languages (see Appendix A). At its peak, over 5,000 domains containing such keywords were registered every day [31]. However, over our measurement period, only 1,229 (2.6%) out of the 46,832 domains seen on the CTC blocklist contain at least one keyword. The share of COVID-19-related domains was stable between 1.4% and 3.6% throughout time, except before the list reset on April 12, when up to 73% of the (much shorter) list was COVID-19-related³.

²1,140 domains were removed; of these domains, 38.8% re-emerged at some later point in time.

³Out of the 1,140 domains removed then, 826 were COVID-19-related, of which 474 would never return.

Provider	Phishing (%)	Malware (%)	Other (%)	Total flagged
GSB*	27 462 (96.8)	1 017 (3.6)	364 (1.3)	28 833
VirusTotal**	40 132 (85.7)	6 680 (14.3)	20 (0.0)	46 832

* Domains may be assigned to multiple abuse types.

** By majority vote over engines with a specific abuse classification.

Table 1: Abuse types for domains on the CTC list as identified by Google Safe Browsing (GSB) and VirusTotal.

Overall, generic phishing was much more frequent: 17.3% of domains matched a brand tracked by PhishTank⁴: the keyword *whatsapp* occurred in 2.8% of domains, making it more prevalent than all pandemic-related keywords combined. Likewise, Google Safe Browsing and VirusTotal classify 96.8% and 85.7% of flagged domains respectively as engaging in phishing (Table 1).

Our findings contrast with the CTC’s stated goal of sharing “high quality threat intelligence related to the COVID-19 pandemic” [32]. The small share of COVID-19-related domains suggests that the collected TI goes beyond COVID-19-specific abuse, and instead captures *any* abuse observed during the pandemic. This could be the result from automated submission processes, with generic TI sources being redirected to the CTC instead of curated and targeted feeds. It also seems directly related to the decision to move to vetting based on VirusTotal. Before that shift, the proportion of COVID-19-related domains was much higher. Afterwards, the list relied on the ability of existing scanners to detect the abuse and therefore it converged on conventional forms of abuse, potentially discarding highly relevant IOCs on new threats. We revisit this issue in subsection 4.4.

4.3 Effect of VirusTotal scanners on vetting

We independently replicated the CTC’s vetting procedure based on VirusTotal, in order to gain insight into the effects on the outcomes. We requested VirusTotal data once a day for all domains that up to that date had appeared on the CTC blocklist, from April 28 until July 1, 2020.

Any domain with more than 10 detections in VirusTotal was automatically considered to be vetted as malicious. We indeed observe that domains meeting that criterion make up the large majority (88.7–97.2%) of domains throughout our measurement period (Figure 6). We also see a smaller set (2.8%–11.3%) of domains with between 4 and 10 detections, which suggests that they either went through the manual review process or they came from trusted industry feeds that were vetted against a lower threshold (subsection 3.2). The share of domains with fewer than 4 detections is negligible, which indicates that the industry sources contributed only do-

⁴https://www.phishtank.com/target_search.php; the list was pruned to reduce the likelihood of false positives, and no lookalike terms were added.




Figure 6: Evolution of the proportion of domains with a given detection count by VirusTotal domain scanners at the time of presence in the CTC blocklist.




Figure 7: Distribution of domains that were (not) detected by Google Safe Browsing (GSB) over the number of detections by VirusTotal engines.

main domains that were already known to multiple security vendors in VirusTotal.

No domain was marked as malicious by more than 27 of VirusTotal’s 84 domain scanning engines; this low ratio is consistent with that found by Peng et al. [15]. Among the 84 scanners, only 23 detected over 5% of vetted domains, and 21 scanners detected over 5% of the domains with COVID-19-related keywords in them. In other words, the CTC threshold for vetting is quite high: 10 detections means that a domain is already known to nearly half of all engines that contribute a non-trivial amount of detections.

As an external corroboration, we determine whether domains on the CTC list were also flagged by Google Safe Browsing (GSB). We find that the more VirusTotal engines detect a domain, the more likely it is to also be flagged by GSB (Figure 7). We confirm with a χ^2 test that the distributions of VirusTotal engine counts for domains that are and are not detected by GSB significantly differ ($\chi^2 = 14595$, critical value at $\alpha = 0.05$: 40.113, $p < 0.0001$). The CTC’s threshold of 10 detections is close to the crossover where more domains are flagged by GSB than are not (11 detections). While this supports a low false positive rate, the threshold also makes it very hard for the CTC to contribute new threat intelligence beyond that of existing anti-abuse infrastructure.




Figure 8: Detection rates for the top 10 engines in VirusTotal sorted by maximum detection rate.

Overall, slightly fewer scanners tend to detect malicious domains as our measurement period progresses (Figure 6). Among the top 10 engines (Figure 8), most have a consistent detection rate and therefore contribution over time. AegisLab WebGuard achieved a high detection rate only after mid May: this is a possible indicator that they may have then started ingesting the CTC blocklist, instead of proactively flagging domains and therefore contributing to the vetting process. Decreasing detection rates for Avira and ESET may translate into a decreased contribution to the vetting process over time. Once a domain was included on the vetted blocklist, few additional engines marked it as malicious (Figure 9), even as its listing duration increased; a larger increase was only visible for the earliest listed domains. This suggests that the vetting decision by the CTC was based on stable detections by the scanners, so domains warrant their near-indefinite presence on the vetted list (subsection 4.2).

In summary, the CTC vetted list comprised (only) domains where a relatively large proportion of security scanners agree on their maliciousness, making false positives unlikely. However, it is also *required* that at least 10 VirusTotal scanners flag a domain. Unless these scanners successfully adapt to the novel COVID-19-related abuse, this strict threshold causes many false negatives. This led to a problem for the CTC: their whole rationale was to “supplement the existing defensive structure”. If the existing scanners were to adapt, then the CTC no longer supplemented them. If they did not adapt, then using them for vetting while maintaining a high threshold

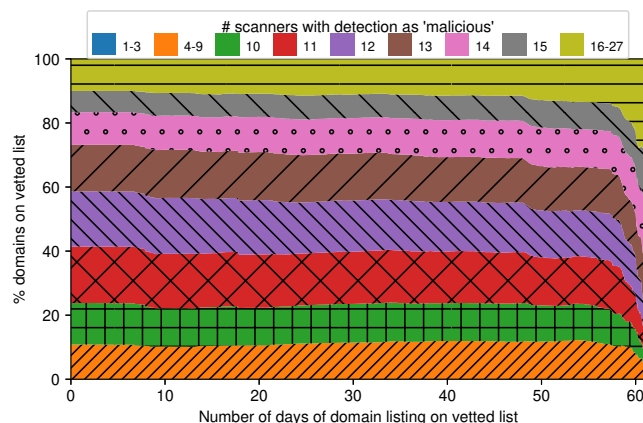


Figure 9: Proportion of domains with a given detection count by VirusTotal domain scanners given the duration a domain has been present on the vetted list.

also meant that the CTC no longer supplemented them. This reaffirms our observation in subsection 4.2 that instead of capturing COVID-19-related abuse not seen in other TI sources, as was the CTC’s primary target, the CTC list rather contained primarily ‘generic’ TI on which many existing sources (here VirusTotal scanners) agreed. While the TI may therefore conform to the CTC’s goal of being of “high quality” (few false positives), the trade-off is that coverage of the threats for which it was set up is undermined (many false negatives). To empirically evaluate this trade-off, we explore the coverage of COVID-19-related domains in the next subsection.

4.4 Coverage of COVID-19 related malicious domains

What COVID-19-related malicious domains were missed by the CTC? For this question, we use passive DNS data from DNSDB by Farsight Security, produced through passive, real-time collection and aggregation of DNS query-response traffic between authoritative servers and recursive resolvers around the world [33]. In particular, we expect higher coverage through DNSDB where discovery of malicious domains may be harder: subdomains as well as top-level domains (TLDs) without available registry zone files (usually ccTLDs).

In order to evaluate the CTC blocklist quality, for each day between March 1 and June 30, 2020, we extracted from DNSDB all FQDNs that match our COVID-19-related keywords from subsection 4.2. In total, we obtain 3,011,717 COVID-19-related domains that contain at least one keyword. We want to know which of these domains are flagged by security vendors as being malicious. In total, 188,305 domains were flagged by at least 1 VirusTotal scanner; however, 162,406 of these are flagged by only one scanner (Fortinet). As presented in Figure 10, 5,767 domains were flagged by 4 to 10 scanners, while only 610 domains were flagged by




Figure 10: Proportion of COVID-19 keyword domains of subsection 4.2 detected by a given number of VirusTotal domain scanners. The CTC vetting process used a threshold of 10.

the CTC’s threshold of 10 or more scanners. Of these, 535 domains were on the CTC blocklist. According to Google Safe Browsing, the majority of these domains were related to social engineering (562 – 92.13%) followed by malware (33 – 5.40%) and unwanted software (15 – 0.24%). Recent work reported similar patterns in keyword-matched COVID-19 domains [34, 35].

We found 75 COVID-19-related domains in DNSDB that had more than 10 detections on VirusTotal and therefore met the inclusion criterion of the CTC’s vetting process, yet were not included on the CTC vetted list. These missed domains are therefore false negatives of the CTC vetted list. In Table 2, we compare how false negatives and domains on the CTC vetted list are distributed across TLDs. While there are no significant differences per TLD type, we can observe differences in how ccTLDs are represented. For instance, .gg is the most prominent ccTLD among the false negatives, whereas it does not appear in the top 10 ccTLDs for domains on the CTC list.

TLD Type	# False Negatives	# CTC domains
gTLD	99 (80.49%)	1,042 (84.78%)
grTLD	3 (2.44%)	3 (0.24%)
ccTLD	21 (17.07%)	183 (14.89%)
gg	3 (14.29%)	ru 41 (22.40%)
ga	2 (9.52%)	cl 20 (10.93%)
eu	2 (9.52%)	cc 18 (9.84%)
ru	2 (9.52%)	tk 14 (7.65%)
pl	2 (9.52%)	br 9 (4.92%)
tk	2 (9.52%)	gd 7 (3.83%)
de	2 (9.52%)	ml 6 (3.28%)
su	1 (4.76%)	gq 6 (3.26%)
co	1 (4.76%)	ca 5 (2.73%)
us	1 (4.76%)	in 4 (2.19%)

Table 2: Number of false negatives versus CTC vetted domains per TLD.

In sum: the CTC blocklist contained false negatives on COVID-19-related abuse, possibly because of the CTC setting a high threshold of 10 detections on VirusTotal. Furthermore, that threshold caused an extreme reduction in how

many COVID-19-related domains made it onto the blocklist, resulting in more false negatives compared to using lower thresholds. Of course, all of this underlines again the more fundamental problem that the CTC’s reliance on VirusTotal undermines its goal to cover threats that are not well covered by existing anti-abuse infrastructure.

5 IMPACT OF THE CTC BLOCKLIST

In this section we address our second research question: *Did publishing threat data in a freely available blocklist improve the ability of network defenders to act against threats, compared to the existing abuse mitigation infrastructure?* We combine various data sources to understand which, how and when actors intervene to take down domains on the blocklist, and produce a longitudinal measurement of these countermeasures. We look specifically at *domain-level* interventions and *client-side* interventions. Registrars, registries and hosting providers are responsible for domain-level interventions. This method protects all users, as it prevents them from accessing the domain, but is relatively invasive, as the intervention cannot be circumvented; it may therefore be applied more cautiously [36]. Client-side interventions inherently only protect those clients who enable the intervention, but may be able to respond more quickly and aggressively to threats. For the blocklist of the CTC to improve the ability of defenders to act against COVID-19-related threats, it should flag domains that existing defenses do not act against or it should flag domains more quickly than existing defenses.

5.1 Domain-level interventions

A core countermeasure is the takedown of a website. This is requested by law enforcement agencies, by targeted organizations, or specialized – e.g., brand-protection – companies acting on their behalf [37–39]. The takedown can subsequently be implemented at (sub)domain level by registries, registrars and/or hosting providers. We measure takedowns by registrars and registries primarily through ‘Extensible Provisioning Protocol’ (EPP) status codes [40, 41] within the WHOIS domain registration data: we consider a domain as taken down when the EPP status codes CLIENTHOLD for registrars and SERVERHOLD for registries respectively are set, which indicates that the domain is not delegated, i.e., activated in the DNS [40, 42]⁵. We retrieve historical WHOIS domain registration data for all vetted domains through VirusTotal, on July 6 and 7.

In total, 6,635 (30.8%) out of the 21,524 distinct second-level domains where we could obtain WHOIS data saw either

⁵As noted by Alowaisheq et al. [38] and confirmed by our own observations, other status codes (in particular *PROHIBITED) are at best unreliable indicators of takedown, and often reflect registry- or registrar-specific behavior (e.g., default configurations).

TLD type	# domains	Intervention by			Any ∇
		Registry	Registrar	p -value	
New gTLDs	4 714	45.2%	20.6%	< 0.0001	56.3%
Legacy gTLDs	14 604	2.7%	23.9%	< 0.0001	26.0%
ccTLDs	2 206	3.9%	4.7%	0.181	8.3%
Total	21 524	12.1%	21.2%	< 0.0001	30.8%

Table 3: Registry and registrar interventions on second-level domains from the CTC blocklist, grouped by TLD type, with p -value given for χ^2 test.

a registrar or a registry intervention⁶. The coverage and actors of domain-level interventions depends on the type of TLD (Table 3). The intervention rate was the highest among new gTLDs (56.3%), with most enacted by registries as they must comply with the most stringent requirements [43]. Meanwhile, on legacy gTLDs (com/net/org), intervention is less prevalent (26.0%) and rather fell to registrars, as this has historically been their responsibility [44, 45]. Finally, ccTLDs saw relatively little intervention (8.3%), potentially owing to their independence in setting (abuse) policies [46–48]. Consequently, assuming the CTC list contained useful TI, it could have acted as a complementary defense where other interventions were less common, in this case especially for ccTLDs.




Figure 11: Delay between the first appearance of a domain on the CTC blocklist and interventions by registries, registrars or Google Safe Browsing (GSB).

Next, we measure if appearance on the CTC list is more timely than these interventions, as this would provide those using the blocklist with an advance warning of a live threat before it is intervened upon. We find that if an intervention did take place, it was usually faster than the CTC: 61.3% of registry and 77.1% of registrar interventions already occurred before the domain appeared on the vetted CTC list (Figure 11). In cases where the CTC blocklist does predate the intervention, the delay tends to be small. This therefore suggests that for the share of domains that registrars and registries do intervene upon, their malicious contents were mostly already unavailable by the time the domains appeared on the CTC vetted list, so end users would not need the CTC’s blocklisting to be protected from those domains.

⁶For 35 domains, we ignore the intervention as it occurred before January 1, 2020 and is therefore unlikely to be COVID-19- or CTC-related.




Figure 12: Empirical survival function for domains on the CTC blocklist where WHOIS registration data is available. We measure lifetime from domain creation until registry or registrar intervention or first appearance on the CTC list.

Another way to see this effect is that for newly registered domains, interventions by registries and in particular registrars tend to occur closer to the registration time than CTC blocklisting (Figure 12). This may indicate close monitoring of new and suspicious domains by registrars and registries that results in more immediate action than the publication of the malicious domain on the CTC blocklist. However, the CTC list covers many more domains, suggesting that registrars and registries may be more careful in taking action against domains, especially once they are older, while the CTC captures TI more broadly. Moreover, the CTC list may also include more novel threats that registrars and registries are not well able to detect and take down. This is further supported by interventions across domains on the blocklist that contain any of the COVID-19 keywords from subsection 4.2: we see registry or registrar interventions for 185 out of 821 COVID-19-related second-level domains (22.5%), lower than the 30.8% seen across the whole list ($\chi^2 = 32.591$, $p < 0.0001$) even though multiple registries and registries had subjected COVID-19-related domains to additional verification [49–51]. The CTC list is therefore even more comprehensive for this novel abuse type: 68.3% and 49.3% of domains appear on the CTC blocklist before a registry or registrar intervention respectively, meaning the CTC blocklist is also more proactive and therefore more useful in flagging COVID-19 domains than domains overall.

5.2 Client-side interventions

Client-side solutions such as domain scanning engines, firewalls, DNS-based filters and browser interstitials provide a complementary countermeasure by blocking access to malicious content, although only for their users. These solutions typically generate or ingest threat intelligence – such as the CTC blocklist – in order to determine if a resource should be considered as malicious [1, 4]. Note that VirusTotal, which the CTC used to confirm maliciousness, and its constituent scanners also serve as client-side solutions. We discussed in subsection 4.3 how only 23 of its 84 engines succeed at detecting at least 5% of domains on the CTC blocklist, rein-




Figure 13: NXDOMAIN responses (i.e., blocking interventions) by the Quad9 DNS resolver, indicating ingestion of the CTC blocklist on May 6.

forcing the blocklist’s utility, as it can be ingested separately to complement the other engines’ protection.

Browser-based intervention Major browsers check every URL that a user navigates to against the Google Safe Browsing (GSB) service, and display a warning interstitial when the URL is known to be malicious [52], which therefore has the potential to protect a large user base. We use the Google Safe Browsing API⁷ to receive hash prefixes of detected malicious URLs⁸. We retrieved an initial state on April 17, 2020, and afterwards collected updates every half hour until July 15, which allowed us to determine when a domain was first flagged. Throughout our measurement period, 28,383 domains on the CTC list (60.6%) were flagged at some point, meaning that the CTC list provides some complementary protection. However, as discussed in subsection 4.3, the domains that GSB fails to flag tend to have fewer detections by VirusTotal engines (Figure 7), suggesting that their malicious status is less agreed upon. By the end, 20,962 domains (44.8%) remained flagged: it is unclear whether GSB removes domains because they are no longer considered malicious, or automatically after a certain delay. Importantly, GSB performs worse on COVID-19-related domains, at some point flagging 302 out of 1,229 domains (24.6%, $\chi^2 = 713.858$, $p < 0.0001$). This might reflect that these COVID-19-related domains contain more scams and forms of abuse outside of the normal scope of GSB, which is focused on the conventional abuse categories of phishing, malware and spam. The CTC list may therefore provide greater benefits for these domains. While GSB does not achieve full coverage, if it flags domains, it does so before the CTC in 96.3% of general cases, and 98.0% of COVID-19 domains (Figure 11), with remaining domains being added quickly to GSB after their first appearance on the CTC list. The CTC list therefore leaves users vulnerable for longer across the domains that GSB detects, and would only be useful if its additional domains consisted of novel TI.

DNS-based intervention Quad9 is a public DNS resolver that blocks malicious domains by responding to queries with

⁷<https://developers.google.com/safe-browsing/v4>

⁸Discrepancies exist between the output of the Google Safe Browsing API and actual browser interventions, a.o. due to “data sharing restrictions” [53]. Our data are therefore an approximation of the latter interventions.

NXDOMAIN [54]. We retrieved DNS records for all vetted domains from Quad9 [55] once a day from April 10 until June 21. Quad9 relies on threat intelligence from at least 18 providers [55] and was reported to include the CTC blocklist from May onward [19]. This inclusion proved beneficial: before May 6, Quad9’s detection rate was 30% at its lowest point (70% for COVID-19-related domains), but from then onward Quad9 included almost all of the blocklist (Figure 13). This shows that the CTC blocklist managed to incorporate threat intelligence that was unknown to at least some security service providers.

In summary, did the CTC improve the ability of defenders to act against threats, compared to existing abuse mitigation infrastructures? We find that for 58.4% of the FQDNs on the CTC blocklist, existing abuse mitigation pipelines at the domain level or in the browser were all faster than the CTC at intervening. For these domains, the pooling and sharing of data in a public blocklist then provided little additional value. For the remaining 41.6%, defenders – such as public and private organizations or managed security service providers – who ingested the open CTC blocklist did however improve their ability to defend themselves, compared to relying on existing anti-abuse pipelines. This advantage was even more sizeable for COVID-19-related domains, at 74.9% additional coverage, showing once again that the CTC was more effective when focusing on its original goal of collecting COVID-19-related abuse. One additional indication of the CTC’s utility occurred in May 2021, when the DNS provider Quad9 started ingesting the CTC list. This created an almost complete overlap between the CTC blocklist and Quad9’s client-side interventions, leading to better protection for its users. To put it briefly: if actors in abuse mitigation intervened, they tended to do so more quickly than the CTC, but often they didn’t intervene, so the list improved the ability of defenders to act, especially on COVID-19-specific threats.

6 ETHICS

Our study describes a threat information sharing community that anyone could join [32] but was nevertheless not public: information was TLP:GREEN, i.e., restricted to the community [56]. Throughout this work, we adhere to the CTC’s Code of Conduct [57]. In accordance with the cited “Chatham House Rule”, we take great care to prevent identifying any community members because this might impact them professionally, except for the founder, who has publicly stated his role. Further, we do not cite from observed conversations, as this might impede trust in the confidentiality of this and future threat information sharing initiatives. To still provide a rich image of the community, we rely on public sources such as the CTC’s webinars [21] and interviews with steering committee members [10, 16, 17, 58, 59]. The blocklist that we analyze in-depth has been made fully public by the CTC

[18]. The lead author's institutional review board (IRB) has approved this study design.

7 DISCUSSION

We have described collaboration in the Cyber Threat Coalition community, which removed many of the entry barriers normally present on threat information sharing arrangements. The CTC relied on Slack and AlienVault OTX groups to pool threat information, then vetting it through VirusTotal before making the resulting blocklist freely available. To learn whether this unique arrangement led to improved threat defenses, we investigated two questions.

First, *by pooling data from its community, did the CTC improve coverage of COVID-19-related threats over existing defenses?* We find that the CTC primarily consolidated existing sources, rather than produce or propagate new threat intelligence. Just 10 users contributed most of the domains in the CTC's AlienVault group and over time, user contributions on the blocklist were outstripped by data from named and unnamed security firms. The community drifted away from its initial goal of tracking pandemic-related abuse: on the CTC blocklist, domains with COVID-19-related keywords (2.6%) were overshadowed even by those with simply the keyword *whatsapp* (2.8%). General phishing domains eventually made up most of the CTC blocklist. This drift was caused by the reliance of the CTC on VirusTotal for vetting domains to be propagated on its blocklist: it required 10 detections by VirusTotal engines. This high threshold meant that the resulting list tended to reproduce the detection of conventional abuse by antivirus engines, rather than contribute new threat intelligence focused on COVID-19. In terms of coverage of these latter threats, we found 75 false negatives – domains that should have been on the list according to the standards of the CTC itself, yet were missing. Thousands more COVID-19-related domains might have been missed by the CTC, with the exact number depending on what threshold one chooses in terms of VirusTotal detections.

Second, *did publishing threat data in a freely available blocklist improve the ability of network defenders to act against threats, compared to the existing abuse mitigation infrastructure?* Our analysis presents a dual view on who acted on domains listed by the CTC and when. On the one hand, no actor achieved full coverage in their interventions – be it at domain-level or client-side – meaning that the CTC vetting process succeeded in delivering an aggregated, more complete set of malicious domains warranting action. This is perhaps best demonstrated by its inclusion in the Quad9 DNS service. On the other hand, where actors in abuse mitigation did intervene, they were usually faster. For 58.4% of the FQDNs on the blocklist, the CTC lagged in incorporating the indicators in its list. Here, they provided little added value – in particular because other interventions such as domain take-downs or browser interstitials typically have a much wider

reach. For the small fraction of COVID-19-related domains, the CTC blocklist was more effective in terms of coverage and improving existing defenses. Here, the apparent lack of focus of the CTC on COVID-19-related abuse impaired the overall utility of its blocklist.

Based on our findings, we draw three lessons for future open source threat information sharing initiatives. Our first lesson is that *scaling up the community does not automatically lead to better pooling of threat information*. In just a few weeks' time, the CTC managed to set up a pipeline for collecting, vetting, and disseminating threat intelligence. Throughout its progression, CTC admins repeatedly pressed members to “signal boost our social media posts”, inviting more people to join the community [22]. Possibly, the admins had assumed that network effects would only increase as the community grew, and did not anticipate the dynamics of a volunteer organization. The CTC did not capitalize as much as it could have on its pool of 4,000 volunteers: contributions on AlienVault were made by just a fraction of community members. This is in line with earlier research on open source software development (subsection 3.1). However, that research also found that open-source communities have a long-tail of members with small contributions, like a bug report. Along the same lines, the CTC could have benefited from indicator vetting by its members (our second lesson). Scaling up the CTC community may also have disincentivized threat information sharing, because it exacerbated the free-rider problem: having a large number of untrusted participants may have discouraged some contributors from sharing indicators that were sensitive or that they feared would be used commercially by other firms.

The second lesson is that *openness of the community requires a scalable quality assurance process for the contributed indicators*. The CTC chose to fulfill that need via VirusTotal, but thereby undid some of the benefits of pooling new threat information in the first place. As described in subsection 3.2, indicators had to meet a threshold level of 10 scanners in VirusTotal before they were propagated to the CTC blocklist. Only 21 scanners were able to detect more than 5% of the indicators on the blocklist, so a threshold of 10 means that half of the dominant scanners need to already detect the indicator before it was shared and published. Although this workflow is scalable and produced a blocklist with a low number of false positives, it also meant that valuable indicators that were not yet known to the dominant VirusTotal engines were discarded because they did not meet the threshold. In light of the CTC's mission, it was ironic that this particularly affected indicators related to COVID-19. Had the CTC's solution for manual vetting of indicators, the Slackbot, come earlier, then it might have prevented the impact of relying on VirusTotal. Crowdsourced vetting of indicators can be successful, as PhishTank has shown [13]. User participation had already tapered off, however, when the Slackbot was

introduced, so the transition away from VirusTotal was not successful.

The CTC was founded on the premise that existing abuse and threat information sharing mechanisms were unsuitable or unprepared for the risks posed by COVID-19-related abuse [10] and that a new response was needed. Our third lesson is that *existing threat intelligence and abuse mitigation structures are actually quite resilient and able to adapt to 'new' types of threats*. Where the existing anti-abuse pipelines intervened, they did so faster than the CTC could detect the domain and include it on its blocklist. For example, 96% of the client-side interventions through Google Safe Browsing occurred before the domains were featured on the CTC list (subsection 5.2). Another way to read these findings is that due to its loss of focus on COVID-19-related domains, the CTC was forced to ‘compete’ with general-purpose abuse sharing mechanisms, a battle it was unlikely to win.

8 LIMITATIONS

The fast-paced evolution of the CTC and its operational processes introduces inherent limitations for our study. Changes to the CTC vetting process and our own data collection, as well as blocklist hosting issues, caused temporal gaps in our data. Moreover, where we do not have access to historical data, in particular for VirusTotal detections and active DNS records, we only collect data starting from the moment when a domain was first included on the blocklist. Finally, certain improvements to the vetting process, such as the Slackbot or the preferential treatment of curated third-party sources, were only introduced by the time community participation had dwindled. Given the low number of contributions, we could therefore not investigate if these would have had a significant impact on the blocklist.

As VirusTotal is the main driver of the CTC vetting process, our analyses are inherently biased by its classifications. Figures 6-10 provide some insight into how these classifications are distributed for our data. However, researchers have questioned the reliability of VirusTotal [15, 60] and other (phishing) blocklists [61]. These shortcomings may be exacerbated by the novelty of COVID-19-related abuse, as well as the semantic discussion on whether tactics such as scams or price gouging constitute ‘maliciousness’ at all. Nonetheless, VirusTotal provides us with objective and independent detection metrics across a large set of domain scanning engines, serving as a strong signal for maliciousness, meaning that domain-level and client-side interventions can be expected. We attribute these interventions based on indicators that carry a level of uncertainty, as also observed in previous work [38]. In particular, the availability of sufficiently detailed WHOIS data is skewed towards gTLDs [39], and we assume correct parsing of its non-standard format by VirusTotal. Our results in section 5 therefore serve as a lower bound to actual interventions. Similarly, the count of COVID-19-specific abu-

sive domains is a lower bound, as we assume the presence of pandemic-related keywords. However, other domains in section 4 may have only carried COVID-19-related content on their web page, or have been propagated within a COVID-19-related context (e.g., a spam email). Where we quantify potential false negatives, we equally did not contend to do so exhaustively.

Our research focused on the CTC blocklist. We evaluated it through its fit for purpose to inform real-time enforcement actions, in line with the CTC’s mission, but the list of COVID-19 related abuse material could conceivably also be used in retrospect to evaluate security controls or even as training data for machine learning purposes. Finally, although the blocklist was an important and visible outcome of the community’s efforts, it was not the only one. The CTC produced threat advisories and community meetings, and facilitated communication between members of the security community. We refrain from analyzing the posts in the CTC Slack workspace due to our adherence to the CTC code of conduct [57]. These conversations likely continued in private messages between members, where we would not have visibility on their outcomes, such as takedowns or law enforcement intervention. Indeed, an important added value of the CTC may be not in the data that they output, but in the network of peers that they managed to bring together at short notice, and as a showcase for what open source threat sharing might look like.

9 RELATED WORK

The performance of blocklists as sources of threat intelligence has been the topic of previous studies, which raised questions about the coverage of relevant threats by open sources [1, 2, 62, 63] as well as closed, commercial threat intelligence sources [4]. Peng et al. [15] found that even the best engines on VirusTotal missed 30% of submitted phishing sites and Oest et al. [64] managed to evade being blocklisted for 55% of phishing domains using simple cloaking techniques. Metcalf and Spring [3] hypothesized such problematic coverage to be an artefact of the collection method of abuse infrastructures, with each using disparate methods to detect threats from their specific vantage points. Notably in the blocklist literature, Li et al. [1] proposed metrics by which to understand threat intelligence and calculated those metrics for a set of open source blocklists. We draw on its coverage and timeliness metrics and our study also describes a blocklist, that of the CTC, but we go beyond descriptive metrics and measure the blocklist’s ability to inform countermeasures. More generally, where Li et al. and other studies take blocklists as-is for their analysis, we conduct measurements on the CTC blocklist to shed light on the open source threat sharing process by which it was produced.

The open source threat sharing model is not entirely new: security information has always propagated through “informal networks of trusted security professionals that exist across

[organizations]]” [65] and earlier research has described initiatives that share the CTC’s objective to pool threat information, but do not match it in terms of access and scale. Thomas et al. described aggregating data from within Google services, in a lab environment and with access restricted to the authors [14]. Another aggregator, Facebook ThreatExchange, was included in the set of feeds analyzed by Li et al. [1], who described it as a “closed-community platform” of “hundreds of companies and organizations”. Here, access is based on being involved in software development for the Facebook platform. Outside of academic research, we see the Cyber Threat Alliance, an industry partnership of 33 firms that share reports and indicators with each other ahead of publication on the basis of quid pro quo, and therefore not available to the public [6]. VirusTotal combines scan engine logic to classify files and URLs [60], and ingests external data feeds as inputs [15]. Despite the fact that users can also contribute binaries and URLs as inputs on a limited free plan, VirusTotal is not a community but a commercial service. Probably most approximate to the CTC, PhishTank is an aggregator of domains and URLs suspected of phishing that provides a practical, real-world example of the promise of collaboration to identify threats. It has a crowdsourcing capability to let users validate maliciousness and like the CTC blocklist, its feed is shared free of charge. It is an older initiative, with the last study of it from 2008, which did not evaluate the ability to inform countermeasures [5]. New user registration on PhishTank has been closed since 2020. Three properties set the experiment of the CTC in threat information sharing apart from these initiatives. First, the low barriers to entry. Formalizing a threat information sharing community has not been attempted before at the scale of thousands of volunteers [23]. Second, the community’s efforts were documented in open sources such as webinars and interviews [10, 16, 17, 21, 58, 59]. And third, it produced information goods that were made freely available – potentially magnifying the impact of the threat information [18]. Because of these properties, the CTC approached what Benkler [66] called peer production: open creation and sharing performed in online groups. In this case, peer production of security information. We conclude that the scale at which threat information sharing occurred in the CTC in response to the pandemic offered a unique opportunity to investigate the collaborative model.

More generally, authors have recently described internet abuse related to COVID-19, which the CTC also tracked. They have pointed to signs of coordinated campaigns [67, 68] and drew attention to the risk of overzealous filtering of COVID-19-related material [34, 69]. Bitaab et al. investigated examples of phishing related to COVID-19 and concluded that the existing anti-phishing ecosystem fell short, based on the sheer volume of COVID-19 related fraud reported on by the FTC [70]. Our measurements support their finding that for the niche of COVID-19 related material, an organization like the CTC could play a valuable supplementary role.

10 CONCLUSIONS

The Cyber Threat Coalition had the aim to “break down traditional barriers to intelligence sharing [and] produce a professional-quality threat feed that the broad IT security public [could] rely upon” [8]. We found that by pooling data from its community, the CTC managed to improve coverage of threats related specifically to COVID-19 over that of existing defenses, and we found evidence that the CTC was faster than other defenses to list such domains. Therefore the community improved the ability of network defenders to take action by publishing its threat data in a freely available blocklist. Over time the CTC lost focus, until it aggregated mostly generic abuse information. We described how this can be traced back to choices that the community made in scaling up its quality assurance processes using VirusTotal. This prevented the CTC from delivering some of its value, as it relied on a threshold number of scanners to recognize a domain as malicious, dropping valuable new indicators in the process and causing it to lag behind established defense mechanisms.

Looking back on his experiences, founder Joshua Saxe said: “We had a lot of [volunteering] energy, but we didn’t have the right organizational machinery to funnel that energy.” [19]. Given that no open source threat sharing community of this kind previously existed, it is unsurprising that the Cyber Threat Coalition went through growing pains. Although we have examined its impact critically, it is only as a result of the hard work of the members of this community that we have been able to investigate the principles of open source threat information sharing at all – and such a community may have impact in subtle ways that are not easily observed or quantified, such as building trust and facilitating contact. We dedicate this paper to the volunteers of the CTC and hope that the lessons discussed in this work may contribute to future threat information sharing.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their comments and advice which has helped to bring focus to the paper. This research was partially funded by the Ministry of the Interior and Kingdom Relations of the Netherlands, Delft University of Technology (M75B07), the Research Fund KU Leuven, and the Flemish Research Programme Cybersecurity. Victor Le Pochat holds a PhD Fellowship of the Research Foundation Flanders - FWO (11A3421N). This research was supported by academic licenses from VirusTotal and Google Cloud.

REFERENCES

- [1] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M. Voelker, Stefan Savage, and Kirill Levchenko. “Reading the Tea leaves: A Comparative Analysis of Threat Intelligence”. In: *28th USENIX Security Symposium*. USENIX Security ’19. 2019, pp. 851–867. URL: <https://www.usenix.org/conference/usenixsecurity19/presentation/li>.
- [2] Harm Griffioen, Tim M. Booij, and Christian Doerr. “Quality Evaluation of Cyber Threat Intelligence Feeds”. In: *18th International Conference on Applied Cryptography and Network Security*. ACNS ’20. 2020, pp. 277–296. DOI: 10.1007/978-3-030-57878-7_14.
- [3] Leigh Metcalf and Jonathan M. Spring. “Blacklist Ecosystem Analysis”. In: *2nd ACM Workshop on Information Sharing and Collaborative Security*. WISCS ’15. 2015, pp. 13–22. DOI: 10.1145/2808128.2808129.
- [4] Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. “A different cup of TI? The added value of commercial threat intelligence”. In: *29th USENIX Security Symposium*. USENIX Security ’20. 2020. URL: <https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman>.
- [5] Tyler Moore and Richard Clayton. “The Consequence of Non-Cooperation in the Fight Against Phishing”. In: *2008 APWG eCrime Researchers Summit*. 2008. DOI: 10.1109/ECRIME.2008.4696968.
- [6] Cyber Threat Alliance. 2020. URL: <https://www.cyberthreatalliance.org>.
- [7] *Anti-Phishing Working Group (APWG)*. URL: <https://apwg.org/membership>.
- [8] COVID-19 Cyber Threat Coalition. *Our mission*. May 2020. URL: <https://www.cyberthreatcoalition.org/about-us/our-mission>.
- [9] Microsoft. *Microsoft report shows increasing sophistication of cyber threats*. Sept. 2020. URL: <https://blogs.microsoft.com/on-the-issues/2020/09/29/microsoft-digital-defense-report-cyber-threats>.
- [10] Jeremy Kirk. “New Platform Collects COVID-19 Threat Intelligence”. In: *BankInfoSecurity* (Apr. 22, 2020). URL: <https://www.bankinfosecurity.com/collecting-covid-19-threat-intelligence-a-14154>.
- [11] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 008 June 11, 2020*. YouTube. June 11, 2020. URL: <https://www.youtube.com/watch?v=0aBcwNqJWDA>.
- [12] Chris Larsen. *Sample Results From Processing a Large Feed of Shady Covid-Type Domains*. Symantec. Apr. 22, 2020. URL: <https://symantec-enterprise-blogs.security.com/blogs/expert-perspectives/sample-results-processing-large-feed-shady-covid-type-domains>.
- [13] Tyler Moore and Richard Clayton. “Evaluating the Wisdom of Crowds in Assessing Phishing Websites”. In: *12th International Conference on Financial Cryptography and Data Security*. 2008, pp. 16–30. DOI: 10.1007/978-3-540-85230-8_2.
- [14] Kurt Thomas, Rony Amira, Adi Ben-Yoash, Ori Folger, Amir Hardon, Ari Berger, Elie Bursztein, and Michael Bailey. “The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges”. In: *19th International Symposium on Research in Attacks, Intrusions, and Defenses*. 2016, pp. 143–164. DOI: 10.1007/978-3-319-45719-2_7.
- [15] Peng Peng, Limin Yang, Linhai Song, and Gang Wang. “Opening the Blackbox of VirusTotal: Analyzing Online Phishing Scan Engines”. In: *2019 Internet Measurement Conference*. IMC ’19. 2019, pp. 478–485. DOI: 10.1145/3355369.3355585.
- [16] Lennart Maschmeyer. *Why the Covid-19 Crisis Offers an Opportunity for Lasting Change in Cybersecurity*. The Center for Security Studies at ETH Zurich, Apr. 28, 2020. URL: <https://isnblog.ethz.ch/cyber/why-the-covid-19-crisis-offers-an-opportunity-for-lasting-change-in-cybersecurity>.
- [17] Arielle Waldman. “Volunteers join forces to tackle COVID-19 security threats”. In: *SearchSecurity* (May 11, 2020). URL: <https://searchsecurity.techtarget.com/news/252482930/Volunteers-join-forces-to-tackle-COVID-19-security-threats>.
- [18] @ThreatCoalition on Twitter. Mar. 29, 2020. URL: <https://twitter.com/ThreatCoalition/status/1244340555007614976>.
- [19] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 007 May 28, 2020*. YouTube. May 28, 2020. URL: https://www.youtube.com/watch?v=cMeDAiG10_4.
- [20] @ThreatCoalition on Twitter. Apr. 8, 2020. URL: <https://twitter.com/ThreatCoalition/status/1247958809856937985>.
- [21] COVID-19 Cyber Threat Coalition. YouTube channel. Apr. 19, 2020. URL: <https://www.youtube.com/channel/UCHfhxcqhQADRV2h5hFggAww>.
- [22] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 001 April 16, 2020*. YouTube. Apr. 20, 2020. URL: <https://www.youtube.com/watch?v=2euXCNUMhww>.

- [23] @ThreatCoalition on Twitter. Mar. 28, 2020. URL: <https://twitter.com/ThreatCoalition/status/1243711825927118853>.
- [24] COVID-19 Cyber Threat Coalition. *Slack workspace of the Cyber Threat Coalition*. Access upon request through <https://www.cyberthreatcoalition.org/join-us>. 2020.
- [25] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 003 April 30, 2020*. YouTube. Apr. 20, 2020. URL: <https://www.youtube.com/watch?v=3EdzxEltbzk>.
- [26] *Statistics about phishing activity and PhishTank usage*. PhishTank. 2021. URL: <https://www.phishtank.com/stats.php>.
- [27] Josh Lerner and Jean Tirole. “Some Simple Economics of Open Source”. In: *The Journal of Industrial Economics* 50.2 (2002), pp. 197–234. ISSN: 1467-6451. DOI: 10.1111/1467-6451.00174.
- [28] *DomainTools COVID-19 Threat List*. Mar. 2020. URL: <https://covid-19-threat-list.domaintools.com>.
- [29] SANS Internet Storm Center. *Covid 19 Domain Classifier*. 2020. URL: <https://isc.sans.edu/covidclassifier.html>.
- [30] Giovane C. M. Muora, Thymen Wabeke, and Christian Hesselman. *Coronavirus and DNS: view from the .nl ccTLD*. Tech. rep. TR-2020-01. SIDN Labs, Mar. 26, 2020. URL: <https://www.sidnlabs.nl/downloads/SzbDWlJkFgwEv9K8pwSCy/9443d3e686ef9a64f7674b0f05b9fd8d/SIDN-LABS-TR-2020-01.pdf>.
- [31] COVID-19 Cyber Threat Coalition. *2020-05-26 Weekly Threat Advisory*. May 26, 2020. URL: <https://www.cyberthreatcoalition.org/advisories/2020-05-26-weekly-threat-advisory>.
- [32] COVID-19 Cyber Threat Coalition. *Join us*. 2020. URL: <https://www.cyberthreatcoalition.org/join-us>.
- [33] Pawel Foremski, Oliver Gasser, and Giovane Moura. “DNS Observatory: The Big Picture of the DNS”. In: *2019 Internet Measurement Conference*. 2019, pp. 87–100. DOI: 10.1145/3355369.3355566.
- [34] Ryo Kawaoka, Daiki Chiba, Takuya Watanabe, Mitsuaki Akiyama, and Tatsuya Mori. “A First Look at COVID-19 Domain Names: Origin and Implications”. In: *2021 Passive and Active Measurement Conference*. PAM ’21. 2021. DOI: 10.1007/978-3-030-72582-2_3.
- [35] Sean McGrath. *COVID-19 Malicious Domain Report: Mapping malicious activity throughout the pandemic*. ProPrivacy. June 16, 2020. URL: <https://proprivacy.com/privacy-news/covid-19-malicious-domain-report>.
- [36] Maciej Korczyński, Samaneh Tajalizadehkhoob, Arman Noroozian, Maarten Wullink, Cristian Hesselman, and Michel van Eeten. “Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs”. In: *2017 IEEE European Symposium on Security and Privacy*. EuroS&P ’17. 2017, pp. 579–594. DOI: 10.1109/EuroSP.2017.15.
- [37] Alice Hutchings, Richard Clayton, and Ross Anderson. “Taking down websites to prevent crime”. In: *2016 APWG Symposium on Electronic Crime Research*. eCrime ’16. 2016. DOI: 10.1109/ECRIME.2016.7487947.
- [38] Eihal Alowaisheq, Peng Wang, Sumayah Alrwais, Xiaojing Liao, XiaoFeng Wang, Tasneem Alowaisheq, Xianghang Mi, Siyuan Tang, and Baojun Liu. “Cracking the Wall of Confinement: Understanding and Analyzing Malicious Domain Take-downs”. In: *26th Annual Network and Distributed System Security Symposium*. NDSS ’19. 2019. DOI: 10.14722/ndss.2019.23243.
- [39] Victor Le Pochat, Tim Van hamme, Sourena Maroofi, Tom Van Goethem, Davy Preuveneers, Andrzej Duda, Wouter Joosen, and Maciej Korczyński. “A Practical Approach for Taking Down Avalanche Botnets Under Real-World Constraints”. In: *27th Annual Network and Distributed System Security Symposium*. NDSS ’20. 2020. DOI: 10.14722/ndss.2020.24161.
- [40] Scott Hollenbeck. *Extensible Provisioning Protocol (EPP)*. STD 69. IETF, Aug. 2009. URL: <https://tools.ietf.org/rfc/std/std69.txt>.
- [41] S. Hollenbeck. *Domain Registry Grace Period Mapping for the Extensible Provisioning Protocol (EPP)*. RFC 3915. IETF, Sept. 2004. URL: <http://tools.ietf.org/rfc/rfc3915.txt>.
- [42] Internet Corporation for Assigned Names and Numbers. *EPP Status Codes | What Do They Mean, and Why Should I Know?* June 16, 2014. URL: <https://www.icann.org/resources/pages/epp-status-codes-2014-06-16-en>.
- [43] Maciej Korczyński, Maarten Wullink, Samaneh Tajalizadehkhoob, Giovane C. M. Moura, Arman Noroozian, Drew Bagley, and Cristian Hesselman. “Cybercrime After the Sunrise: A Statistical Analysis of DNS Abuse in New gTLDs”. In: *13th ACM Asia Conference on Computer and Communications Security*. ASIACCS ’18. 2018, pp. 609–623. DOI: 10.1145/3196494.3196548.

- [44] Natasha Tusikov. “Access Chokepoints”. In: *Chokepoints: Global Private Regulation on the Internet*. 1st ed. University of California Press, 2017. Chap. 4, pp. 116–155. ISBN: 9780520291218.
- [45] Kevin Murphy. “VeriSign demands website take-down powers”. In: *The Register* (Oct. 11, 2011). URL: https://www.theregister.com/2011/10/11/verisign_asks_for_web_takedown_powers/.
- [46] Kim G. von Arx and Gregory R. Hagan. “SOVEREIGN DOMAINS: A Declaration of Independence of ccTLDs from Foreign Control”. In: *Richmond Journal of Law & Technology* 9.1 (Fall 2002), p. 4.
- [47] Peter K. Yu. “The Never-Ending ccTLD Story”. In: *Addressing the World: National Identity and Internet Country Code Domains*. Ed. by Erica Schlesinger Wass. Rowman & Littlefield Publishers, Inc., 2003. Chap. 1, pp. 1–16. ISBN: 9780742528109.
- [48] National Research Council. *Signposts in Cyberspace: The Domain Name System and Internet Navigation*. The National Academies Press, 2005. ISBN: 9780309096409. DOI: 10.17226/11258.
- [49] EURid. *COVID-19 and extraordinary measures on .eu domain names*. Apr. 8, 2020. URL: <https://eurid.eu/en/news/doteu-covid19-measures/>.
- [50] Eleanor Bradley. *Keeping a close watch on coronavirus domains*. Nominet. Mar. 26, 2020. URL: <https://www.nominet.uk/keeping-a-close-watch-on-coronavirus-domains/>.
- [51] Frank Bajak. “Internet Firm Restricts Virus-Themed Website Registrations”. In: *ABC News* (Mar. 26, 2020). URL: <https://abcnews.go.com/Technology/wireStory/internet-firm-restricts-virus-themed-website-registrations-69825166>.
- [52] Devdatta Akhawe and Adrienne Porter Felt. “Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness”. In: *22nd USENIX Security Symposium*. USENIX Security ’13. 2013, pp. 257–272. URL: <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe>.
- [53] Sergio Villegas et al. *sbserver differs from online/browser lookup?* GitHub repository google/safebrowsing. Issue 30. Aug. 31, 2016. URL: <https://github.com/google/safebrowsing/issues/30>.
- [54] Sean Gallagher. “New “Quad9” DNS service blocks malicious domains for everyone”. In: *Ars Technica* (Nov. 16, 2017). URL: <https://arstechnica.com/information-technology/2017/11/new-quad9-dns-service-blocks-malicious-domains-for-everyone/>.
- [55] Quad9. *Quad9 DNS: Internet Security and Privacy in a Few Easy Steps*. 2019. URL: <https://www.quad9.net/>.
- [56] Forum of Incident Response and Security Teams. *Traffic Light Protocol (TLP)*. Aug. 2017. URL: <https://www.first.org/tlp>.
- [57] COVID-19 Cyber Threat Coalition. *Code of Conduct*. URL: <https://www.cyberthreatcoalition.org/about-us/code-of-conduct>.
- [58] Tammy Xu. *Volunteer Devs Around the World Fight Pandemic-Inspired Cybercrime*. Built In, June 9, 2020. URL: <https://builtin.com/cybersecurity/coronavirus-cyber-threat-coalition>.
- [59] Brian Krebs. “COVID-19 Has United Cybersecurity Experts, But Will That Unity Survive the Pandemic?” In: *Krebs on Security* (Apr. 15, 2020). URL: <https://krebsonsecurity.com/2020/04/covid-19-has-united-cybersecurity-experts-but-will-that-unity-survive-the-pandemic>.
- [60] Ravindu De Silva, Mohamed Nabeel, Charith Elvitigala, Issa Khalil, Ting Yu, and Chamath Keppitiyagama. “Compromised or Attacker-Owned: A Large Scale Classification and Study of Hosting Domains of Malicious URLs”. In: *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/desilva>.
- [61] Adam Oest, Yeganeh Safaei, Adam Doupé, Gail-Joon Ahn, Brad Wardman, and Kevin Tyers. “PhishFarm: A Scalable Framework for Measuring the Effectiveness of Evasion Techniques against Browser Phishing Blacklists”. In: *2019 IEEE Symposium on Security and Privacy*. 2019, pp. 1344–1361. DOI: 10.1109/SP.2019.00049.
- [62] Marc Kührer, Christian Rossow, and Thorsten Holz. “Paint It Black: Evaluating the Effectiveness of Malware Blacklists”. In: *17th International Symposium on Research in Attacks, Intrusions and Defenses*. RAID ’14. 2014, pp. 1–21. DOI: 10.1007/978-3-319-11379-1_1.
- [63] Thomas Vissers, Peter Janssen, Wouter Joosen, and Lieven Desmet. “Assessing the Effectiveness of Domain Blacklisting Against Malicious DNS Registrations”. In: *2019 IEEE Security and Privacy Workshops*. 2019, pp. 199–204. DOI: 10.1109/SPW.2019.00045.

- [64] Adam Oest, Penghui Zhang, Brad Wardman, Eric Nunes, Jakub Burgis, Ali Zand, Kurt Thomas, Adam Doupé, and Gail-Joon Ahn. “Sunrise to Sunset: Analyzing the End-to-end Life Cycle and Effectiveness of Phishing Attacks at Scale”. In: *29th USENIX Security Symposium*. USENIX Security ’20. 2020. URL: <https://www.usenix.org/conference/usenixsecurity20/presentation/oest-sunrise>.
- [65] Michel J. G. van Eeten and Johannes M. Bauer. “Economics of Malware: Security Decisions, Incentives and Externalities”. In: *OECD Science, Technology and Industry Working Papers* 2008/01 (May 2008). DOI: 10.1787/241440230621.
- [66] Yochai Benkler. “Peer Production and Cooperation”. In: *Handbook on the Economics of the Internet*. Ed. by J. M. Bauer and M. Latzer. Edward Elgar, 2013.
- [67] Sameera Horawalavithana, Ravindu De Silva, Mohamed Nabeel, Charitha Elvitigala, Primal Wijesekara, and Adriana Iamnitchi. *Malicious and Low Credibility URLs on Twitter during COVID-19*. 2021. arXiv: 2102.12223 [cs.SI].
- [68] Pengcheng Xia, Mohamed Nabeel, Issa Khalil, Haoyu Wang, and Ting Yu. “Identifying and Characterizing COVID-19 Themed Malicious Domain Campaigns”. In: *11th ACM Conference on Data and Application Security and Privacy*. CODASPY ’21. 2021, pp. 209–220. DOI: 10.1145/3422337.3447840.
- [69] Anjali Vyas, Ram Sundara Raman, Nick Ceccio, Philipp Lutscher, and Roya Ensafi. “Lost in Transmission: Investigating Filtering of COVID-19 Websites”. In: *Financial Cryptography and Data Security*. FC ’21. 2021.
- [70] Marzieh Bitaab, Haehyun Cho, Adam Oest, Penghui Zhang, Zhibo Sun, Rana Pourmohamad, Doowon Kim, Tiffany Bao, Ruoyu Wang, Yan Shoshitaishvili, Adam Doupé, and Gail-Joon Ahn. “Scam Pandemic: How Attackers Exploit Public Fear through Phishing”. In: *2020 APWG Symposium on Electronic Crime Research (eCrime)*. eCrime 2020. 2020.
- [71] Yi-Min Wang, Doug Beck, Jeffrey Wang, Chad Verbowski, and Brad Daniels. “Strider Typo-Patrol: Discovery and Analysis of Systematic Typo-Squatting”. In: *2nd Workshop on Steps to Reducing Unwanted Traffic on the Internet*. 2006, pp. 31–36.
- [72] Tobias Holgers, David E. Watson, and Steven D. Gribble. “Cutting Through the Confusion: A Measurement Study of Homograph Attacks”. In: *USENIX Annual Technical Conference*. 2006, pp. 261–266.

A COVID-19 KEYWORD LIST

We selected COVID-19-related keywords starting from a set of English keywords, which we then translated into 14 major languages: Arabic, Bengali, Mandarin Chinese, Dutch, French, German, Hindi, Italian, Japanese, Malay, Portuguese, Russian, Spanish and Turkish. Finally, we generated lookalike terms using techniques from typosquatting [71] and homographs [72].

self-isolation	eovid	covic1	cocvid
corona	co-vid	covid	st1mulus
pandemic	fovid	pandemic	pandemic
mask	eovib	timulus	pandemic
ncov	vovid	pandemic	panndemic
vaccine	pademic	pandemic	panndemic
virus	covix	cov8d	pandemic
hydroxychloroquine	landemic	vaccine	pandemic
quinacrine	covoid	covid	vacvine
chloroquine	ocvid	tsimulus	chloraquine
remdesivir	pandemic	stimulus	cokvid
plaquenil	xovid	vaccine	eovicl
azithromycin	stimulis	covid	coviud
metformin	cuid	copvid	vaccine
favipiravir	stimulas	vaccine	pandemic
interferon	stimulus	vacne	pandemic
lopinavir	covid	cpovid	stimulus
ritonavir	stemulus	covid	stimulus
arbitol	stimulu	pandemic	stimulus
stimulus	clovid	vacne	stimulus
infection	cov-id	pandemic	covid1
n95	stimulus	stimul-us	cov9d
respirator	covid	covid	covid
testkit	cpvid	pandemic	covid
distance	covid	pandemic	covid
quarantine	covid	vacne	covid
lockdown	baccine	vacne	covid
covis	vaccine	chloraquine	covid
covic	clvid	caccine	covid
covid	covid	stimulus	covid
covid	covid	chloraquine	covid
covid	vaccine	stimulus	covid
covid	covid	panemic	pandemic
covid	covid	vacne	covid
covid	coovid	colvid	covid
covid	chloraquin	panemic	sars-cov
covid	pandemic	pandemic	sarscov
covid	mandemic	stimulus	ivomec
covid	covid	stikulus	ivermectin
covid	covid	covid	mectizan
covid	stimulus	pandemic	iver-dt
covid	vaccine	covid	ivexterm
accine	pandemic	pandemic	scabo-6
vacine	covid	stimulus	sklice
andemic	stimulus	vaccine	stromectol
covid	pandemic	stipulus	soolantra
pandemic	vaccine	pandemic	mk-933
ftp2	ftp3	c19	

Table 4: List of English keywords and homographs