Trilobyte: Measurements of Censorship Evasion Potential for Single-Player Games

Anonymous Author(s)

ABSTRACT

Censorship evasion solutions that rely on communication services popular in Western countries may invite closer scrutiny into the activities of their users in censored countries. This paper investigates the ability of popular and ubiquitously available single-player games to enable users to evade censorship. We find that most investigated games allow users to communicate keywords considered sensitive in China, when compressed, encrypted or hidden in game chat channels. We introduce Trilobyte, a system that further hides data in game state generated opportunistically during user regular game-playing activities, and shares data-hiding state through accounts on gaming platforms. We find that for many games, Trilobyte can store megabytes of data in a single game state, and evade current censorship in China. Even in the presence of future, more rigorous censors, Trilobyte can hide up to 5.3 MBs of data in game state saved in a one hour gaming session.

KEYWORDS

plausible deniability, censorship evasion, video games

1 INTRODUCTION

Secure and private communication and data sharing tools (e.g., E2EE apps, file sharing sites, social networks) have become almost ubiquitous. However, powerful state-level censors ban access to such technologies [4] and replace them with approved alternatives, e.g., [7, 8, 12], that monitor and remove communications and content deemed sensitive [21, 22]. Further, while a censor may not completely prevent users from acquiring secure communication tools (e.g., WhatsApp, Telegram, Signal apps), the mere use of such apps in the censored area is sufficient to invite closer scrutiny of user activities. This leads to a need for secure plausible deniable communication and data sharing solutions. We posit this can be best achieved by leveraging services that are already tolerated (and sometimes even developed) in censored regions.

Previous work has used play-time covert channels in online games, to embed information into user commands [19, 20, 23], in-game state (e.g., light updates) [31], avatar movements [29, 31], gaming packets [24], or game locations [17, 26]. Such solutions either make strong adversary assumptions, e.g., that game operators do not collude with the adversary who also cannot inspect the content of game communications, or require specialized software to be installed on game servers.

Trilobyte [1] is a plausible deniability communication system built on single-player gaming platforms with billions of players worldwide [6], many located in censored countries [16]. Trilobyte uses shared single-player gaming accounts as a cover for storing and accessing data to evade monitoring by adversaries who control gaming servers. By leveraging covert channels within locally saved game states, Trilobyte facilitates delay-tolerant communications for censored game players. This approach hides encrypted data within game state generated during regular gameplay, achieving plausible deniability by separating gaming and data hiding activities. Experiments conducted with accounts on 46 games(30 online games and 16 single-player games), including those available in China, demonstrate the efficacy of this method, with potential to hide tens to hundreds of KB in a single game state. Additionally, Experiments with tens of rented accounts, surveys conducted with Chinese users who rent gaming accounts (n = 114). These experiments further reveal that accounts often store tens of game states, resulting in the ability to hide up to 5.31 MB of data in a one-hour gaming session.

2 RELATED WORK

Early work employed steganography in computer board games to conceal information within player moves or strategies [20, 23]. For example, StegoGo [20] conceals data within Go game moves, but its altered game states can be distinguished from regular ones. FPSCC [29] utilizes first-person shooter games to hide information in slight player character movements, yet it's vulnerable to adversaries aware of players' positions. Castle [19] embeds data into real-time strategy (RTS) game commands, assuming operators and censors lack collusion or the ability to inspect game states or communications. In contrast, Rook [24] assumes unencrypted game traffic and decentralization, hiding data within mutable fields of game packets. Telepath [31] conceals data in Minecraft network traffic, mimicking regular patterns, and relies on

[1] Trilobites are marine arthropods that have evaded the censorship of evolution for 270 million years. Their fossils hid in shale and limestone for a further 250 million years.

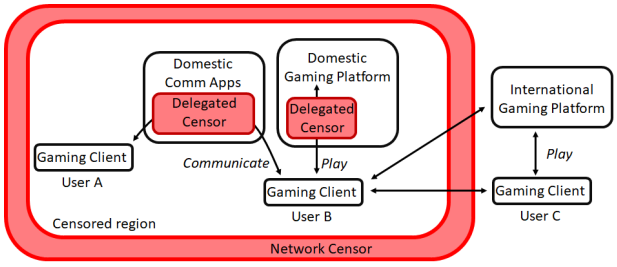

Figure 1: System model. Users install gaming clients which they use to connect to gaming platforms. Censorship is implemented at all levels of the Internet infrastructure. The censor delegates censorship responsibilities to content providers, e.g., gaming, communication platforms.

encrypted client-server communications. Similarly, Wajid et al. [26] propose visual communications in shooter games, vulnerable to in-game monitoring.

Trilobyte hides data within covert channels of game state thumbnails, shared through gaming platform accounts. It defends against adversaries inspecting both game states and user communications. Unlike previous systems focus on a specific game or game genre, it's adaptable to any single-player game providing cover thumbnails.

3 MODEL AND BACKGROUND

We assume users located in a censored area, who need to share content with other users who are either in or outside this censored area, see Figure 1. Users consist of both *writers* who hide data and *readers* who access the hidden data. We consider a system that provides such read and write functionality through single-player games accessible in the censored area. Both writers and readers install and run a system client on their devices. This section details the threat model, provides background on the single-player game ecosystem, and specifies system requirements.

3.1 Threat Model

We consider an adversary with control over domestic services and employs tools like China's Great Firewall (GFW) to censor internet activities, even at the ISP level, see red outer boundary in Figure 1. we assume that the adversary *delegates* content moderation and censorship responsibilities to services [21, 28], e.g., gaming platforms, using filters and staff to censor user communications, removing sensitive content, and specific keywords [21], see delegated censor blocks in Figure 1. While current gaming platforms primarily filter plaintext communication for sensitive content, our experiments assume the adversary can inspect all communications and data on both domestic and international platforms, including game states, enabling them to verify validity and search for hidden data.

Adversary Limitations. The adversary can install and use the Trilobyte client on any number of devices. However, it does not have access to the state of Trilobyte clients installed by other users, has no control over the software they install and execute on their devices, and has no access to game behaviors they perform on their local devices and the game state generated and saved locally.

3.2 The Single-Player Game System Model

The following discusses the main components of the single-player gaming ecosystem used in this work to provide plausibly deniable data sharing and communication solutions for censored gamers.

Gaming Platforms and Clients. Gaming platforms like Steam [9], EA App [2], WeGame [11], GOG [5], Epic Game Store [3], host video games uploaded by game developers, and provide functionality for users to search, access and play video games. Users access gaming platforms through clients which they install on their devices. Gaming platform clients connect to the gaming platform server, and provide a user interface to authenticate, and to search, purchase, manage, install, and uninstall video games.

Game Platform Accounts. Users need to create accounts on gaming platforms, in order to acquire, purchase, manage and access games. In most cases, users create a username and password, and verify their phone number during the account creation process. Several games in countries like China require additional identity verification, such as providing government-issued identification.

Single-Player Game State. To minimize disruptions during player disconnections, most single-player games provide status-saving functionality, allowing users to save game state on their device. Each game stores state files in dedicated directories, and each game state is stored in a different file or sub-directory. To enable users to play on multiple devices, major gaming platforms [2, 3, 5, 9, 11] provide game state online storage services for supported single-player games. Upon user launch and exit game events, the gaming platform client synchronizes saved game state between the user device and gaming platform-provided online storage.

3.3 System Requirements

Let Γ denote an environment used to provide cover to hide data. Further, let $\perp$ denote the null message. Then, a system hiding data into a cover environment Γ should satisfy the following requirements.

Plausibly Deniable Data Storage. A system provides plausibly deniable data storage (PD-DS) if each action performed

to hide data can be explained to the adversary through public, gaming-related actions. We formalize this requirement by extending a plausible deniability (PD) game [18]. The PD-DS game involves a challenger C and an adversary $\mathcal{A}$ (i.e., the censor) and takes place over r rounds. The game uses an $Hide(\Gamma, M, K)$ function that given the cover environment Γ, a message M and a key K, returns an updated environment that embeds M obfuscated, e.g., encrypted, using K. The game also uses a $canHide(\Gamma, M, K)$ predicate that returns true only if the cover environment Γ can hide the input message M obfuscated with K. Note that $canHide(\Gamma, \perp, K) = True$. The $Hide$ and $canHide$ functions are instantiated later in the paper. The PD game proceeds through the following steps:

- C chooses a secret symmetric encryption key K using a security parameter s. C further selects a bit b.
- $\mathcal{A}$ and C engage in the following r rounds:
  - $\mathcal{A}$ generates a message M and sends it to C.
  - C generates cover Γ s.t. $canHide(\Gamma, M, K) = True$.
  - C sets $M_0 = \perp$ and $M_1 = M$. It then computes and sends $\Gamma' = Hide(\Gamma, M_b, K)$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs b', its guess of b.

The advantage of the adversary $\mathcal{A}$ in the PD-DS game is $Adv(\mathcal{A}) = |P(b' = b) - P(b' \neq b)|$. A system is said to provide plausible deniability if any probabilistic polynomial time (PPT) adversary $\mathcal{A}$ has only negligible advantage in the PD-DS game.

Note that the original cover environment Γ is created by the challenger and is not available to the adversary. Thus, the game models the inability of the adversary to distinguish a game state that hides its message ($\Gamma' = Hide(\Gamma, M, K)$ from a game state that that does not ($\Gamma' = \Gamma = Hide(\Gamma, \perp, K)$).

Readability of Hidden Data. Given a cover environment Γ and a key K, a user should be able to determine if Γ hides data encrypted with K. Given K and a data-hiding cover environment $\Gamma' = Hide(\Gamma, M, K)$ the reader should be able to identify and reconstruct the hidden data M.

Game State Validity. Trilobyte also needs to satisfy the following *Game State validity* requirements: **Loading Validity**. Any data-hiding game state can be used to load in the game. **Reachability**. The game state loaded from a data-hiding game state is reachable via regular user play. **Sequence Validity**. When hiding data in multiple game states from the same game session, a game state S_2 captured at time T after a previous state S_1 needs to be reachable within time T after the game state captured in S_1.

4 TRILOBYTE

Trilobyte, a data storage system for single-player gamers, covertly stores data within shared gaming accounts by utilizing game state channels. Its architecture involves a client wrapping around the game platform client, facilitating both gaming and data hiding. Employing two threads—one for gaming and one for data hiding—Trilobyte ensures behavioral independence [25]. The gaming thread relays user actions to the game platform client, generating game state used as cover for hiding user data, which is then shared through gaming accounts.

4.1 Setup

Each user needs to install the Trilobyte client. The client is responsible for managing access to gaming accounts used for communication, and embedding and extracting data from game state stored in those accounts. To communicate through Trilobyte two users need to share key material, i.e., symmetric encryption and authentication keys K_e and K_a, and information about an account *acc* they share, e.g., an account id on a gaming platform site or with a game developer, or the id of a listing in an account rental site, see S 4.3. While outside the scope of this paper, this bootstrapping step could take place over a low-bandwidth out-of-band communication channel.

During setup, the user needs to provide her Trilobyte client with this information, $[K_e, K_a, acc]$, for each of the user's contacts. Further, the Trilobyte client initializes an empty hide data list HD. During operation, the list will store tuples of format $[fid, t, pos, seq]$ for each content file fid to be sent to the contact. t is the timestamp when the user has initiated the sending of the content, pos points to the position in the file up to where data has already been hidden, and the sequence number seq counts how many game state files were used to hide content from fid up to pos.

4.2 Hiding Data in Game State Thumbnails

Trilobyte hides data in game snapshot thumbnails of single-player games. Many video games utilize image thumbnails for their saved game states. The content of these thumbnails can vary depending on the game. Some may display a screenshot of the current gameplay, showcasing the player's character, and surroundings. Since games can have highly dynamic and diverse environments, character actions, and events, each screenshot can capture a unique moment in the gameplay. This introduces a significant amount of entropy. Trilobyte hides data on thumbnail images by using state-of-the-art image steganography approaches SteganoGAN[30], which can achieve a payload of 4.4 bits per pixel.

Data Storage Format Trilobyte organizes data into segments. Each segment is encrypted separately, and stored in a different game state. The length of a segment depends on the size of thumbnail image that is used to hide the data and the threat model. For instance, games from developers that do not collude with the adversary and that encrypt game traffic, offer larger segment sizes than games from colluding

or insecure providers. Section 5 discusses the segment sizes for several single-player games and threat models.

The format of a hidden data segment is $(enc, auth)$, where $enc = E_{K_e}(type, hidden\ data)$ and $auth$ is the first 4 bytes of $E_{K_a}(enc)$. The $type$ field specifies the segment type, e.g., DAT segments to organize hidden data into files, or REQ/REP segments for client/server communications (Appendix A).

For a file with unique identifier fid and whose data needs to be stored in multiple game states, Trilobyte needs to keep track of the segment's sequence number. More specifically, Trilobyte stores fid's data into DAT segments whose field enc is $E_{K_e}(DAT, fid, last, seq, data)$, where the $last$ bit flag indicates if this is the last segment of fid, and $data$ is the actual file content.

Trilobyte adds game thumbnails to a hide list (HL), including only pointers to fields that differ significantly from the previous thumbnail in HL, enabling data hiding, and readers reconstruct HL by comparing game states and extracting hidden data segments.

4.3 Sharing Game State

Trilobyte clients can synchronize game state via shared gaming platform accounts id and key K_a, requiring a secure, low-bandwidth out-of-band communication channel for exchanging account details and login-time SMS codes. The following describes alternative game state sharing techniques.

Account Rental. Trilobyte clients use platforms like Zuhaowan for sharing personal or rented gaming accounts (§5.3). These platforms allow renting accounts for short durations (hourly, daily, or weekly) at low rates (e.g., a few USD cents per hour), with a platform fee of 15-25%. The platform handles authentication and session management, requiring lenders to provide login details. Renters select, pay, and receive credentials, with the platform holding lenders' credentials until payment. Clients can rent accounts, enabling encrypted game state sharing (§4.4). During rental, the account is delisted to prevent simultaneous access, while incentivizing participation through payments.

Cross-platform Progression. Some single-player games like Cyberpunk 2077 and The Witcher III offer cross-platform progression, letting players sync their game state across different platforms. By registering an extra account with the game developer and linking it to their other platform accounts, players can seamlessly continue their game from any device and platforms without switching accounts. For example, two accounts on different platforms (like Steam and GOG) with the same game can share states through this feature, eliminating repetitive logins and logout actions.

Game State Sharing Sites. The gaming ecosystem also contains sites dedicated to sharing game state. Such sites are popular among gamers in censored countries, and some

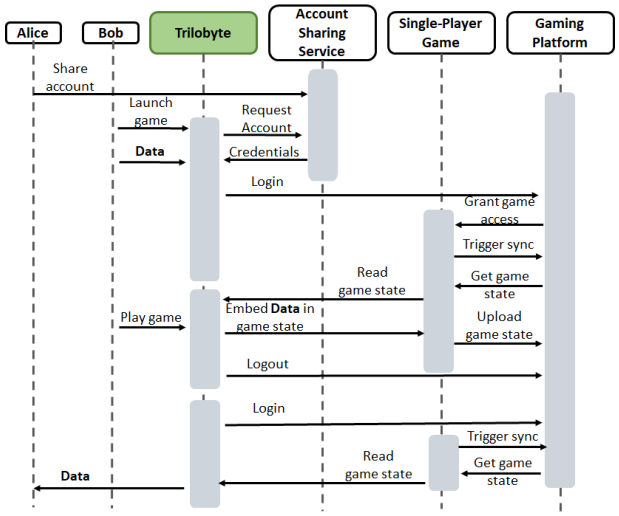

Figure 2: Data communication sequence diagram between a Trilobyte writer Bob (B) and reader Alice (A) that share an account owned by A through a rental platform.

are hosted in censored countries, e.g., 3DMGame [2], GamerSky [3]. Most game state sharing sites require users to register accounts in order to upload game state. While sites like MCBBS and 3DMGame provide social network functionality, i.e., allow users to follow, befriend or subscribe to game state posted by other accounts, most also allow visitors to download any game state without registering an account. This allows Trilobyte users who also use such sites to share game state with plausible deniability.

4.4 Hidden Data Communication

The following details the data communication procedure between a sender (B) and receiver (A) that share a gaming account acc. Figure 2 illustrates this protocol, when B launches a game, his client logs into acc and downloads game data. B's client maintains a hide list (HD) for data transfer. When B creates a new file (fid) to send to A, it adds an entry to HD. During gameplay, B's client reads from HD, retrieves data chunks, and encrypts them for A. Encrypted data is hidden in the game state. Upon logout, B's client uploads game states. When A logs in, her client syncs game states and retrieves hidden data. If verification succeeds, data is decrypted and files are reconstructed.

5 TRILOBYTE EVALUATION

Trilobyte Implementation. We implemented Trilobyte using Python and the SteganoGAN package[30].

[2] 3DMGame: 3dmgame.com

[3] GamerSky: https://www.gamersky.com/

Game	Platform	Licensed	Paid	Online Storage Size	Game State Size	Thumbnail Image
A Perfect Day	WeGame	●	●	100MB	100KB	○
Baldurs Gate 3	Steam	○	●	18.55GB	16MB	●
Battlefield	Origin	○	●	1GB	2MB	○
Cyberpunk 2077	Steam/GOG	○	●	2GB	6MB	●
Divinity: Original Sin	Steam/GOG	○	●	200MB	9MB	●
Dragon Age Origins	Steam	○	●	100MB	16MB	●
Loop Hero	Steam/Epic	○	●	20MB	30KB	●
Minecraft	NetEase	●	○	1GB	<100GB	○
Mount & Blade	Steam	○	●	953MB	5MB	○
Moncage	Steam CN	●	●	10MB	400KB	○
No Man's Sky	Steam	○	●	210MB	256KB	○
Rimworld	Steam	○	●	92.95GB	16MB	○
The Witcher 3	Steam	○	●	953.67MB	2.5MB	●
Watch Dogs Legion	Steam	○	●	1GB	78KB	●
Yakuza: Like a Dragon	Steam	○	●	66.18MB	1MB	●

Table 1: Data sharing through game states over 15 single-player games. All evaluated games and platforms allowed uploading and downloading plaintext and encrypted sensitive keywords appended to or replacing game state.

Experimental setup. Experiments were conducted between devices in China and the US. The setup included a device in China (i5 2.7 GHz CPU, 8 GB RAM, 200 Mbps downlink, 15 Mbps uplink) and a device in the US (i7 3.6 GHz CPU, 8 GB RAM desktop, 1,200 Mbps downlink, up to 35 Mbps uplink). Gaming accounts controlled by the paper authors were used across both devices, with one account logged in on a computer in China and the other on a US computer. Both accounts were utilized for sending and receiving messages.
Experimental Data. We gathered a list of Chinese sensitive keywords from a public dataset spanning 2004 to 2014 [13]. We also added newly reported censored keywords from online forums and gaming chats. The dataset includes 12,858 sensitive keywords. We employed the Baidu censorship API [1] to detect keywords filtered by key Chinese game operators such as Baidu, Tencent, and NetEase. Baidu API identified 3,919 keywords out of the above dataset, with some tagged with multiple labels: 1,925 as politically sensitive, 723 as terrorism-related, 1,000 as sex-related, 125 as curse words, and 41 as related to malicious ads. We evaluated in-game censorship with collected sensitive keywords in China, Appendix B includes the report of measurement of keywords filtering over 30 Chinese online games.

5.1 Game State Storage

We explored using game state files for data storage, detailed in Table 1. Our evaluation covered games from various platforms, allowing two users to log in to the same account, save, and download states. Most games require a one-time purchase, but platforms also offer free games with cloud storage for states i.e., Steam (2034 such games), GOG (22) Epic (5), Tencent WeGame (5). For the evaluated games and platforms, the average game state size ranges from tens of KB to megabytes. However, platforms impose a limit on total game state size that ranges from 10 MB to 18.55GB. Once the total size of the game state files reaches the limit, the oldest state is removed until there is enough space to store the new one.

Table 1 also shows the average game state size and the per-game limit on total game state can be uploaded to online storage.All tested games and platforms allowed injecting new files, plaintext or encrypted, into game state directories, including images, videos, and compressed files. They also permitted injecting plaintext and encrypted keywords into existing game state files or replacing them entirely, across all file types. This was possible from both the US and China, with downloads accessible through the same account in the other country.

5.2 Data Hiding Experiments

We evaluated the ability of Trilobyte to hide data in single-player game state, see also Table 1. Several games do not generate state with thumbnails (Moncage, A Perfect Day, Battlefield, Loop Hero, etc.). Such games cannot be used to hide data, and can only be used to embed data by appending or replacing existing state files, like shown in § 5.1. To investigate the number of single-player game state stored in game platform accounts, we rented or borrowed multiple accounts for nine distinct single-player games featuring saving thumbnail capabilities across various platforms.

Table 2 illustrates the number of investigated accounts, and maximum frequency of game saves within a one-hour session for each game. Our evaluation employs the state-of-the-art image steganography algorithm SteganoGAN, known for its capability to conceal 4.4 bits per pixel. By analyzing the thumbnail size of different games, we derive the maximum data hidden rate in a one-hour gameplay session for each evaluated game. For instance, Baldur's Gate 3, a recent popular game, can conceal 5.31MB of data in a one-hour session with plausible deniability.

5.3 Account Acquisition Experiments

We conducted a survey to investigate the gaming account renting and lending experiences and perceptions of Chinese users. The survey, in Simplified Chinese, was distributed through the Wenjuan.com crowdsourcing site. Appendix C includes the English translation of the questions. We recruited participants through the nga.cn forum, a general gaming community for online game video gamers, and chat groups in WeChat and QQ dedicated to popular video games.
Ethical Considerations. The survey was approved by the institutional review board at our university. Participants received a consent form before starting, with no personal info collected. Responses were stored securely on a university

Game	Hidden Data per State	# of Account Investigated	Max States # in hour	Hidden Data per Session
Baldurs Gate 3	495KB	5	11	5.31MB
Cyberpunk 2077	62.7KB	5	15	940.5KB
Divinity: Original Sin	92.81KB	2	1	92.81KB
Divinity: Original Sin 2	69.61KB	3	3	208.83KB
Dragon Age Origins	70.4KB	3	6	422.4KB
Loop Hero	25.14KB	3	1	25.14KB
The Witcher 3	15.68KB	2	7	109.73KB
Watch Dogs Legion	62.7KB	3	2	125.4KB
Yakuza: Like a Dragon	15.68KB	3	1	15.68KB

Table 2: Data sharing evaluation through game state thumbnails over nine single-player games.

Linux server and accessed only through encrypted channels from password-protected laptops.

Findings. 393 users took part in the survey. We analyzed responses from 285 participants over 18 years old who played video games (127 female, 163 male; 130 employed, 132 students, 17 unemployed, 1 self-employed). All respondents spent at least 30 seconds on the survey. Of the 285, 114 rented game accounts, and 33 lent their accounts to others. Survey findings include: **In-game communications are popular**. Twenty-nine of the 33 account lenders said that renters can use their accounts to communicate; 100 out of 114 account renters have used rented accounts to communicate with other players. **Data storage**. Of the 33 account lenders, 30 said that renters are allowed to upload content to their accounts; 71 out of the 114 account renters have uploaded content in accounts they rented. **Cryptocurrency use**. Six of the 33 account lenders accept cryptocurrency payments for their accounts, including Bitcoin (4), Ethereum (2), and Tether (3). Conversely, ten account renters have used cryptocurrencies to rent accounts, including Bitcoin (all 10), Tether (4), and Ethereum (3).

This suggests that in-game communications and data upload activities are popular among people who rent accounts, and lenders are aware of renters' use of this functionality. Further, the use of cryptocurrencies is moderately popular, suggesting the possibility to remain anonymous during the payment step.

6 DISCUSSION

Readability of Hidden Data. The format $(enc, auth)$ of Trilobyte data segments hidden in individual game state ensures that only the reader with access to the authentication key K_a can determine if a segment hides data or not. Further, only a reader with access to the encryption key K_e can decrypt *enc* and recover the segment's payload.

Game State Validity. Trilobyte hides data in game state thumbnails by image steganography approach SteganoGAN [30]. The thumbnails will remain the same image format as the original image. This ensures that Trilobyte data-hiding thumbnails are valid, and can be loaded. Thus, Trilobyte-edited game state thumbnail files can be used to load game states that are indistinguishable from game states loaded from snapshots generated after regular game-playing sessions.

Trilobyte Provides PD-DS. The adversary can observe sequences of user actions(account login, download game states, game play interval, upload game state, account logout), i.e., the set Σ of the cover environment Γ. However, they cannot access the actual gameplay activities since they occur on the user's device. The behavior of a Trilobyte user appears the same as a regular gaming user to the adversary. Trilobyte utilizes typical gaming activities are identical for Trilobyte and regular users. Additionally, the edited game state by Trilobyte is indistinguishable from regular game state to the adversary, and the number of uploaded game states remains consistent regardless of Trilobyte usage.

Real vs. Synthetic Game Playing Behaviors. The Trilobyte client relies on the user's natural game playing behavior to access and hide data within saved game states. Future exploration may involve integrating Trilobyte with tools like Voyager [27] to simulate user actions for game exploration and play without user involvement. Additionally, the client can learn user behavior models to optimize data storage and sharing, considering factors such as login and logout times, game selections, play durations, etc. These models can be installed on multiple devices, and shared among contacts to enhance data hiding through diverse behavioral patterns.

7 CONCLUSIONS

This paper introduces Trilobyte, a system that hides data with plausible deniability in game state thumbnails generated opportunistically during regular game-playing activities in single-player games. Trilobyte users communicate data-hiding game state through accounts shared on gaming platforms. Experiments reveal that Trilobyte can embed hundreds of megabytes of data in state saved for games currently licensed in China whith one-hour gaming sessions. Surveys with Chinese lenders and renters of gaming platform accounts, and experiments with renting gaming accounts in China, reveal the ease of acquiring and sharing such accounts.

REFERENCES

[1] [n. d.]. Baidu Content Censoring Service API. Baidu, https://ai.baidu.com/tech/textcensoring.

[2] [n. d.]. EA App Game Platform. https://www.ea.com/ea-app.

[3] [n. d.]. Epic Game Store Platform. https://store.epicgames.com/.

[4] [n. d.]. Free privacy and security tools: Blocked In ... Comparitech, https://www.comparitech.com/privacy-security-tools/.

[5] [n. d.]. GOG Game Platform. https://www.gog.com/.

[6] [n. d.]. How Many Gamers Are There? (New 2023 Statistics). Exploding Topics, https://explodingtopics.com/blog/number-of-gamers.

[7] [n. d.]. QQ. https://mail.qq.com/.
[8] [n. d.]. Sina Weibo. https://my.weibo.com/.
[9] [n. d.]. Steam Game Platform. https://steampowered.com.
[10] [n. d.]. Unistego.
[11] [n. d.]. WeGame Game Platform. https://www.wegame.com.cn/.
[12] [n. d.]. Weixin/WeChat. https://weixin.qq.com/.
[13] 2014. Chinese Keywords GitHub Repository. GitHub, https://github.com/jasonqng/chinese-keywords.
[14] 2017. China's great firewall gives rise to a robust industry of information smugglers. https://hongkongfp.com/2017/04/02/chinas-great-firewall-gives-rise-robust-industry-information-smugglers/.
[15] 2022. Most Popular Searched Online Games in China. Baidu, https://lewan.baidu.com/rankland?idfrom=1043&gameSource=client&gameType=0.
[16] 2022. Top 10 Countries/Markets by Game Revenues. newzoo, https://newzoo.com/insights/rankings/top-10-countries-by-game-revenues.
[17] 2023. Finnish newspaper hides Ukraine news reports for Russians in online game. The Guardian, https://www.theguardian.com/world/2023/may/03/finnish-newspaper-hides-news-reports-for-russians-in-online-game.
[18] Chen Chen, Anrin Chakraborti, and Radu Sion. 2019. PD-DM: An efficient locality-preserving block device mapper with plausible deniability. Proc. Priv. Enhancing Technol. 2019, 1 (2019), 153–171.
[19] Bridger Hahn, Rishab Nithyanand, Phillipa Gill, and Rob Johnson. 2016. Games without Frontiers: Investigating Video Games as a Covert Channel. In 2016 IEEE European Symposium on Security and Privacy (EuroS&P). 63–77.
[20] Julio C Hernandez-Castro, Ignacio Blasco-Lopez, Juan M Estevez-Tapiador, and Arturo Ribagorda-Garnacho. 2006. Steganography in games: A general methodology and its application to the game of Go. computers & security 25, 1 (2006), 64–71.
[21] Jeffrey Knockel, Lotus Ruan, and Masashi Crete-Nishihata. 2017. Measuring Decentralization of Chinese Keyword Censorship via Mobile Games. In 7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17).
[22] Jason Q. Ng Lotus Ruan, Jeffrey Knockel and Masashi Crete-Nishihata. 2016. One App, Two Systems: How Wechat Uses One Censorship Policy in China and Another Internationally. 84 (2016).
[23] Steven J Murdoch and Piotr Zieliński. 2005. Covert channels for collusion in online computer games. In Information Hiding: 6th International Workshop, IH 2004, Toronto, Canada, May 23-25, 2004, Revised Selected Papers 6. Springer, 355–369.
[24] Paul Vines and Tadayoshi Kohno. 2015. Rook: Using Video Games as a Low-Bandwidth Censorship Resistant Communication Platform. In Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society. 75–84.
[25] Ryan Wails, Andrew Stange, Eliana Troper, Aylin Caliskan, Roger Dingledine, Rob Jansen, and Micah Sherr. 2022. Learning to Behave: Improving Covert Channel Security with Behavior-Based Designs. Proceedings on Privacy Enhancing Technologies 3 (2022), 179–199.
[26] Abdul Wajid, Nasir Kamal, Muhammad Sharjeel, Raaez Muhammad Sheikh, Huzaifah Bin Wasim, Muhammad Hashir Ali, Wajahat Hussain, Syed Taha Ali, and Latif Anjum. 2021. A First Look at Private Communications in Video Games using Visual Features. Proceedings on Privacy Enhancing Technologies 2021, 3 (2021), 433–452.
[27] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An Open-Ended Embodied Agent with Large Language Models. arXiv preprint arXiv:2305.16291 (2023).
[28] Diwen Xue, Benjamin Mixon-Baca, Anna Ablove, Beau Kujath, Jedidiah R Crandall, and Roya Ensafi. 2022. TSPU: Russia's decentralized censorship system. In Proceedings of the 22nd ACM Internet Measurement Conference. 179–194.
[29] Sebastian Zander, Grenville Armitage, and Philip Branch. 2008. Covert channels in multiplayer first person shooter online games. In 2008 33rd IEEE Conference on Local Computer Networks (LCN). IEEE, 215–222.
[30] Kevin Alex Zhang, Alfredo Cuesta-Infante, Lei Xu, and Kalyan Veeramachaneni. 2019. SteganoGAN: High Capacity Image Steganography with GANs. arXiv:1901.03892 [cs.CV]
[31] Sun Zhen and Shmatikov Vitaly. 2023. TELEPATH: A Minecraft-based Covert Communication System. In Proceedings of the IEEE Symposium on Security and Privacy.

A CONTENT-FETCHING TRILOBYTE SERVER

Trilobyte provides support for building a content-fetching server with plausible deniability: servers receive requests to fetch data for censored clients to access. Applications include *information smugglers*, who package overseas news and sell them to a domestic audience [14].

The following considers therefore a Trilobyte server S with unrestricted Internet access, e.g., located outside of the censored region. The server operator controls a list L_S of gaming platform accounts, and lists (a subset of) these accounts on an account rental platform, to be rented for an hourly fee. We further consider a Trilobyte client C that also controls a list L_C of gaming platform accounts. We assume that the client shares a key K_{CS} with the server. Further, that the client knows the id of at least one account $A_S \in L_S$ controlled by the server. However, the server does not need to know of any client accounts in L_C.

To issue a request to the server, the client first selects an account $A_C \in L_C$ and sets its password to the Base58 encoding of $E_{K_{CS}}(A_C)$. It then generates a REQ block that consists of (1) a unique request id req_id, (2) an identifier for the content to be retrieved (e.g., URL, keyword to search, social network account id) and (3) the id of account A_C. Then, the client rents account A_S. This allows the server to charge a fee for its subsequent work. The client then uses the protocol outlined in § 4.4 to embed his request into game state in account A_C.

When the client's session expires, the server logs back into its account A_S, and uses the protocol outlined in Section 4.4 to recover the game state and read the request, account A_C and its password. The server then fetches the requested content and splits it into DAT blocks, each encrypted with the key K_{CS}. It then uses the protocol of Section 4.4 to embed the DAT blocks into game state on account A_C. If the DAT blocks cannot be stored in the game state for a single account, S creates a REP-type block that contains (1) the req_id value, (2) the number of remaining content blocks, (3) the fee required to communicate them, and (4) the identifier of an account $A_S^\star$.

The client accesses account A_C and reads the DAT blocks. If it identifies a REP block and decides to pay the fee for the

Game Name	Genres	Developer	Licensed	Platform	Paid	Size	Chat Channel	Compressed	Encrypted	Stego-Hidden
Genshin Impact	ARPG	Domestic	●	PC	○	64.5GB	G,P	○	○	○
Honor of Kings	MOBA	Domestic	●	Mobile	○	3.7GB	W,G,P	○	○	○
Sausage Man	Shooter	Domestic	●	Mobile	○	3.5GB	W,G,P	○	○	○
Mini World	Shooter	Domestic	●	Mobile	○	592.3MB	W,G,P	○	○	○
PUBG	Shooter	International	○	PC	○	33.86GB	W,G,P	○	○	○
Cross Fire	Shooter	Domestic	●	PC	○	6.6GB	W,G,P	○	○	○
Game of Peace	Shooter	Domestic	●	Mobile	○	1.89GB	W,G,P	○	○	○
Sanguosha	Card	Domestic	●	PC	○	79MB	W,G,P	○	○	○
League of Legends	MOBA	International	●	PC	○	21.6GB	W,G,P	○	○	○
Fantasy Westward Journey	MMORPG	Domestic	●	PC	○	4.92GB	W,G,P	○	○	○
Naraka: Bladepoint	ARPG	Domestic	●	PC	●	32.33GB	W,G,P	○	○	○
Final Fantasy XIV	MMORPG	International	●	PC	●	35.4GB	W,G,P	○	○	○
Justice	MMORPG	Domestic	●	PC	○	9.61GB	W,G,P	○	○	○
MapleStory	MMORPG	International	●	PC	○	12.9MB	W,G,P	○	○	○
World of Warcraft	MMORPG	International	●	PC	○	75.8GB	W,G,P	○	○	○
Teamfight Tactics	Strategy	International	●	PC	○	2.13GB	W,G,P	○	○	○
Audition Dance Battle Online	Rhythm	Domestic	●	PC	○	1.43GB	W,G,P	○	○	○
AssaultFire	Shooter	Domestic	●	PC	○	21GB	W,G,P	○	○	○
Warframe	ARPG	International	●	PC	○	32.42GB	W,G,P	○	○	○
Hearthstone	Card	International	●	PC	○	2.93GB	W,P	○	○	○
Wulingwaizhuan OL	MMORPG	Domestic	●	PC	○	1.7GB	W,G,P	○	○	○
JX3	MMORPG	Domestic	●	PC	○	7.78GB	W,G,P	○	○	○
Lost Ark	MMORPG	International	○	PC	○	72.47GB	W,G,P	○	○	○
Path of Exile	ARPG	International	●	PC	○	30.01GB	W,G,P	○	○	○
Escape from Tarkov	Shooter	International	○	PC	○	12.1GB	W,G,P	○	○	○
DOTA 2	MOBA	International	●	PC	○	43.6GB	W,G,P	○	○	○
QQ Dance	Rhythm	Domestic	●	PC	○	2.4GB	W,G,P	○	○	○
Honkai Impact 3rd	Action	Domestic	●	Mobile	○	7.6GB	W,G,P	○	○	○
Run Juveniles	Action	Domestic	●	Mobile	○	1.2GB	W,G,P	○	○	○
The Elder Scrolls Online	MMORPG	International	○	PC	●	117.5GB	W,G,P	○	○	○

Table 3: Censorship of communication channels in both domestic (developed in China) and international games, both licensed and unlicensed for operation in China. Empty circles in the last three columns show that Baidu API-labeled sensitive keywords that were either Zlib-compressed, AES256-encrypted or steganographycally hidden in cover text were not censored on any of the available channel types (public, group, world).

remaining content, it repeats the above process with two exceptions. First, it rents the server's account $A_S^\star$ at the fee specified in the REP block. Second, the REQ block hidden in game state in account $A_S^\star$ includes a list of accounts from L_C where the server needs to embed the remaining content. The client and server take turns to access these accounts, where the server hides content in game state and the client reads and erases the state.

Content Caching. Since the client controls the accounts where the content is stored, once the client pays the required fee, it "owns" the content. In an alternative approach, the server can maintain control over the content by embedding it in game state stored in accounts it controls. This enables the server to *cache* popular content (e.g., news, posts of influential social network users) and share access to the accounts storing it, upon request from other clients.

Plausible Deniability. The server's list L_S of controlled accounts introduces a weakness similar to Tor's bridges: an adversary can eventually learn the ids of all the accounts, then coerce the account rental platform to provide the history of all the users who rented accounts in L_S. However, unlike bridges, these accounts provide plausible deniability grounds: users can claim genuine interest in renting those accounts to play games.

B ONLINE GAME MEASUREMENT

To determine the validity of the adversary model defined in Section 3.1, we evaluated in-game censorship in China. We focused first on keyword filtering on communication channels in online games. While initial keyword filtering solutions notified users when their messages contained sensitive keywords, Ruan et al. [22] revealed that platforms like WeChat have started to implement a form of silent censoring.

The evaluation was performed on 30 online games selected based on Baidu's search index [15]: the 15 highest ranked Chinese online games (top section of Table 3), 10 randomly selected mid-ranked games (mid section), and the 5 lowest ranked games (bottom section). All but three games (Naraka

and Final Fantasy XIV) are (mostly) free. They include diverse genres, e.g., massively multiplayer online role-playing games (MMORPG), Shooter, Strategy, Rhythm, Multiplayer Online Battle Arena (MOBA), and Cards. Most investigated games are developed domestically in China, or are licensed for operation in China. However, we also included international unlicensed games (PUBG, Lost Ark, Escape from Tarkov, Elder Scrolls), whose servers are accessible from China. Five of the games are for mobile devices, the others are for PCs.

The clients of 17 games are over 5GB, and seven are over 30GB. All the game clients were available for download in China. For each game we used two accounts controlled by paper authors.

All games support private communication channels. Only one game (Genshin Impact) does not support world channels, and one game (Hearthstone) does not support group communication channels. We have evaluated communications on each type of communication channel available in each game. Only one group channel was evaluated for games supporting multiple group channels.

Obfuscated Communications. Given the 3,919 Baidu API-flagged sensitive keywords, we further compressed them using the Zlib library, encrypted them using AES with 256-bit keys, or steganographically hidden them in benign cover messages using Unistego [10]. None of the evaluated games censored any of these messages, on any available channel types, shown with empty disks in Table 3.

C SURVEY

1. How old are you now?
2. What is your current occupation?
3. What is your gender?
4. Do you play online video games? If the answer is “No” stop the survey.
5. Have you ever rented any game account on an account rental service? (If “No” skip to Q9). (If “Yes”, ask) What account renting platforms have you used?
6. Did you ever rent an account that allowed you to communicate with other accounts on the gaming platform? (If yes, then ask) What platforms provide such accounts?
7. Did you ever rent an account where you were able to upload content to the account? Content includes photos, videos, or any other media. (If yes) What games provide such accounts?
8. What payment methods do you use to rent accounts? (If answer includes “cryptocurrency”, ask) What cryptocurrency have you used to rent accounts?
9. Have you ever advertised and lent any of your own game accounts on an account rental service? (If “No” end the survey). (If “Yes”, ask) What account renting platforms have you used?
10. Did you ever lend an account that allow customers to communicate with other accounts on the gaming platform? (If yes, then ask) What platforms provide such accounts?
11. Did you ever lend an account that allows customers to upload content to the account? Content includes photos, videos, or any other media. (If yes) What games provide such accounts?
12. What payment methods do you accept to lend your accounts? (If answer includes “cryptocurrency”, ask) What cryptocurrency have you used to rent accounts?