Interdependent privacy issues are pervasive among third-party applications

Shuaishuai Liu, Barbara Herendi, and Gergely Biczók

CrySyS Lab, Dept. of Networked Systems and Services,
Budapest Univ. of Technology and Economics, Hungary
{sliu,biczok}@crysys.hu

Abstract. Third-party applications are popular: they improve and extend the features offered by their respective platforms, whether being mobile OS, browsers or cloud-based tools. Although some privacy concerns regarding these apps have been studied in detail, the phenomenon of *interdependent privacy*, when a user shares others' data with an app without their knowledge and consent. Through careful analysis of permission models and multiple platform-specific datasets, we show that interdependent privacy risks are enabled by certain permissions in all platforms studied, and actual apps request these permissions instantiating these risks. We also identify potential risk signals, and discuss solutions which could improve transparency and control for users, developers and platform owners.

Keywords: interdependent privacy · third-party apps · permissions · Android · browser extensions · Google Workspace · risk signal

1 Introduction

Third-party applications (apps) occupy a prominent position in the current Internet ecosystem; such apps add extra features and functionality to already popular platforms, e.g., mobile operating systems, social networks, browsers, storage clouds, etc. Data sharing is the foundation of the ever-increasing number of third-party apps and their respective platforms. Faced with the large-scale, diverse and virtually non-stop data exchange, privacy issues have become even more pressing.

Nowadays, various application platforms are stepping up to improve the data privacy mechanisms; e.g., as of iOS v14.5[1], applications must obtain user consent to track user data from other applications and websites. Furthermore, A new privacy dashboard has been added to Android v12[2], which allows users to monitor the usage of app permissions more accurately. These updates are welcome efforts in the quest for enhancing transparency, user control and privacy in general; yet, none of these current developments addresses *interdependent privacy* [2], where

[1] https://developer.apple.com/app-store/user-privacy-and-data-use/
[2] https://www.androidauthority.com/android-privacy-dashboard-1233846/

others might share your data without your knowledge and control. Especially pronounced in the third-party app scenario, data shared with the app by a single user might also contain personal and, potentially, sensitive information on their friends, contacts or colleagues.

The highest profile recent cases featuring interdependent privacy are connected to Facebook. The best-known is the Cambridge Analytica scandal [22], where 87 million Facebook profiles were harvested by an app called "thisisyourdigitallife", then used to build detailed personal psychological profiles, and, consequently, the users were targeted with personalized political ads to affect the outcome of the 2016 US presidential elections. The app in question exploited the *collateral information collection* mechanism on Facebook, where it was installed by 270,000 users but reached tens of millions of friend profiles through the controversially designed permission system [22] that allowed for harvesting friend profiles. More recently, on February 08, 2021, Facebook compensated more than 1.6 million users to the amount of $650 million, one of the largest privacy-related settlements to date, owing to creating and storing scans of their faces without permission. The class action lawsuit was initiated in Illinois in 2015, and involved Facebook's use of facial recognition technology in its photo tagging function. The photo tagging feature allowed users to tag friends in photos they had uploaded to Facebook, creating a personal link to the friends' profile without their consent[3]. In this case, a vivid example of interdependent privacy, the uploading user consciously granted Facebook the permission to display his photo, but the users who were tagged automatically in the photo suffered a privacy loss without even being aware of it.

However, there have been less publicized interdependent privacy incidents, as well. A security bug allowed third-party apps to access Google+ user profile data since 2015 until Google discovered and patched it in March 2018 [4]. When a user gave permission to an app to access their public profile data, the bug also enabled developers pull their and their friends' non-public profile fields; around $500,000$ profiles were affected. To make matters worse, Google decided not to inform the world on this issue. In another unfortunate case, the popular TrueCaller Android app, used for blacklisting spam numbers, came under scrutiny. In addition to uploading the address book of the installing user to its servers (an interdependent privacy issue in itself as noted by the Article 29 Working Party in 2017[5]), TrueCaller allows its users to tag unknown numbers after taking the call, and to upload them to the servers for all other users to see the information. In 2019, this feature blew the cover of an investigative journalist in a hostile country; luckily, no actual harm has been inflicted, but there was a non-negligible threat to the physical well-being of the journalist and her sources[6].

[3] https://www.theguardian.com/technology/2021/feb/27/facebook-illinois-privacy-lawsuit-settlement

[4] https://www.wsj.com/articles/google-exposed-user-data-feared-repercussions-of-disclosing-to-public-1539017194

[5] https://ec.europa.eu/newsroom/article29/items/610173

[6] https://privacyinternational.org/node/2997

The above incidents naturally invite the question *(RQ1)*: *are interdependent privacy issues pervasive among third-party app platforms?* Moreover *(RQ2)*, *do actual apps request the permissions enabling collateral information collection on their respective platforms?* In this paper, we answer *RQ1* and *RQ2* affirmatively. Specifically, the contribution of this paper is threefold. First, we analyze the permission systems of multiple third-party app platforms of different types, and identify permissions for each of them which can potentially cause interdependent privacy issues. Second, through analyzing freshly collected datasets from the respective platforms, we demonstrate and quantify the extent to which real-world apps request these permissions. Pointedly, we show that the type (i.e., category) of the app is a good predictor for the number of potentially interdependent privacy related permissions requested. Finally, we briefly discuss potential measures which can augment and/or improve the current, permission-based access control mechanisms with regard to transparency and control.

The rest of this paper is organized as follows. Section 2 lays out the related work. Section 3 analyzes the permissions of the studied third-party app platforms, and identifies interdependent privacy invoking permissions. Section 4 introduces our dataset, and quantifies the proliferation of such permissions requested by actual apps from their respective platforms. Section 5 briefly touches upon prospective transparency and control enhancing techniques. Finally, Section 6 reflects on our contributions and concludes the paper.

2 Related work

Here we briefly summarize related work in the intersection of interdependent privacy and third-party apps.

Interdependent privacy captures the networked characteristics of privacy-related decisions. Owing to this networked nature, the privacy of individuals are bound to be affected by the actions of others, e.g., Facebook users sharing the data of their friends' [2]. In economic terms, unaware fellow users fall victim to a negative externality. Extending this interpretation, a data entry, seemingly concerning a single individual, may actually be also related to (multiple) others because of data correlation [15]. Note that the same concept is known under different monikers, such as collective privacy [20], networked privacy [3] and multiple-subject privacy [7], among others. For a comprehensive overview, we refer the interested reader to [10].

Interdependent privacy affects different types of data and data sharing scenarios. A subset of attributes from the profile of a social network user may be harvested [2]. The location privacy of certain individuals may be threatened by sharing co-location information [16]. Photo sharing may affect the privacy of friends and bystanders captured in the photo [15]. Even the genetic profile of an individual and associated inferrable medical information might get exposed by an eager relative (i.e., kin genomic privacy) [9]. A common trait among the aforementioned scenarios is that all of them could be instantiated through a variety of third-party apps.

General privacy considerations regarding third-party apps, platforms, permissions and ecosystems have been a strong focus area of researchers in the last decade. We do not even attempt to give a comprehensive overview here, rather, we highlight a few studies with close relations to this paper. Wang et al. studied the data collection practices of Facebook third-party apps and proposed control mechanisms which can increase transparency [23]. King et al. conducted an exploratory survey on how Facebook users interact with apps, and how much they understand the privacy implications of such interaction [14]. Androidleaks uncovered how sensitive data is used once the user gave the required permissions and the Android app was installed [6]. Chia et al. studied app permissions, privacy risk signals and community ratings on multiple app platforms [4]. FlowDroid and its follow up works provided taint analysis for Android apps that sheds light on potential unintended and malicious data leaks [1]. Reardon et al. explored the many ways apps can circumvent the Android permission system [19]. Finally, Kelley et al. (and many others building on their study) showed that users actually factor in their privacy concerns when choosing between apps if they are presented with easy-to-understand privacy facts before installation [13]. The above selection of studies clearly demonstrate that i) permission models are imperfect, ii) various privacy leaks do occur in apps, and iii) users act on their concerns when presented with tractable information on app privacy.

Yet, there are only a handful of scientific studies dealing explicitly with interdependent privacy situations regarding third-party apps. Biczok and Chia showed that the personal, relational and spatial privacy of Facebook users are threatened by their friends [2]. Pu and Grossklags investigated the effect of selfish and other-regarding preferences in social app adoption [18]. Harkous and Aberer analyzed Google Drive apps, and pointed out that users suffer more privacy loss owing to their collaborators than their own actions [8]. Finally, Symeonidis et al. presented a comprehensive data analytics, modeling and legal study on the *collateral information collection* practices of Facebook apps affecting the friends of the user [22]. While both these studies and further anecdotal evidence suggest that interdependent privacy issues might be the norm rather than the exception on most third-party app platforms, the research community lacks a data-driven study for available, active, but previously uncharted platforms, such as Android, Google Chrome and other browsers, and cloud services. This paper aims at filling this gap.

3 Platforms, permissions and interdependent privacy

3.1 Permissions and interdependent privacy

Permission-based access. Third-party app platforms share a common security model which is based on requesting and granting permissions. App permissions guard the access to i) restricted data, such as location or contact information, and ii) restricted actions, such as taking photos or connecting to the Internet. Generally, the main objectives of app permissions include: i) enabling user control over data shared, ii) achieving transparency, so that the user

understands what data an app is using and why, and iii) promoting data minimization, so that the app accesses and utilizes only the data absolutely required for a specific task the user invokes.

Platforms, e.g., Google's Android, have evolved significantly since their inception to achieve these objectives. Android has introduced install-time and run-time permissions; the latter group includes all individual permissions deemed *dangerous* by the platform. Run-time permissions can be explicitly granted (or denied) by the user through a dedicated pop-up window, shown when the execution of the app reached a state where the permission is required. On top of this, very recently, Android has included a *privacy dashboard* that shows which apps had used sensitive permissions and for how long in the last 24 hours; also, with easy access to revoke said permissions if so desired. Despite all these improvements in Android and other permission mechanisms, there are still no specific (neither transparency, nor mitigation) measures targeted at interdependent privacy.

Rubbing salt into the wound, app platforms' definition of certain permissions are vague, as to what extent the app will obtain and use sensitive private information. Combining this more general transparency issue with the specific flaws mentioned above, two sub-optimal privacy outcomes emerge. First, the user does not have sufficient knowledge on the scope of the information to be shared: others' private data might be transferred to the app without even their knowledge. Second, the user might grant excessive permissions to the app to preserve full functionality. Although this latter has been shown to be an issue with respect to one's own sensitive attributes, it could induce an even more negative impact in the context of interdependent privacy.

Permissions related to interdependent privacy. Corresponding to the above two points, when the permission involved is ambiguous, users pay more attention to protecting their own privacy while ignoring the privacy of their friends [18]. When the number of permissions granted by users to apps becomes larger, interdependent privacy issues often emerge. An obvious example is a top-rated Firefox extension called *AdBlocker Ultimate*. The permission-related warnings of this app are the following: W1) "Access browser tabs", W2) "Store unlimited amount of client-side data", W3) "Access browser activity during navigation", and W4) "Access your data for all websites". Plausibly, the combination of W1, W3 and W4 enables the extension to read the website, detect ads, and replace them with blank boxes. However, the same permissions enable the app to collect, e.g., messages sent to and received from a web-based chat; an outcome that could cause privacy loss to the communication partners of the user, an obvious interdependent privacy scenario, no user would prefer to experience. Furthermore, W2 enables the storage of unlimited personal data collected through W1, W2 and W4; this can allow for observing, e.g., personal communications over a longer period of time. Yet, not granting these requested permission makes it impossible to install and use the app.

As we would like to quantify the extent to which interdependent privacy issues are present in third-party app platforms, we classify permissions into three pre-defined categories: invoking interdependent privacy (IDP), potentially invoking interdependent privacy (PIDP), and not invoking interdependent privacy (NIDP). If a permission *directly* enables access to private data related to a natural person other than the user herself, it is in IDP; e.g., the `READ_CONTACTS` permission in Android. If a permission *potentially* enables access to private data related to a natural person other than the user herself, it is in PIDP. Such risk can be realized through i) accessing data that *may* implicate multiple parties, such as photos or documents (e.g., `READ_EXTERNAL_STORAGE` in Android); ii) enabling a restricted action that *may create* multi-party data, such as photos or audio recordings (e.g., `RECORD_AUDIO` in Android); and iii) enabling *inference* of other's private data with reasonable effort, such as location via co-location information from other sources (e.g., `ACCESS_FINE_LOCATION` in Android). Note that granting a PIDP permission does not automatically constitute privacy loss for a third party; the loss is context-dependent and may require additional effort from the app developer or an adversary. If a permission does not belong either to IDP or PIDP, then it is in NIDP, and not in our focus.

3.2 Platform specifics

Here we briefly introduce the app platforms we investigated. For practical data availability reasons, we targeted the most popular mobile app platform Android, two well-known browsers providing an API for third-party extensions (Mozilla Firefox and Opera), and Google Workspace, a cloud-based enterprise collaboration tool bundle. Although these 4 platforms vary greatly in both their functionality and technical mechanisms, all of them offer the equivalent of an app store, where the access control of apps is based on the user granting permissions.

Android. Android users can download and install more than 3 million apps from the Google Play store, making Android the largest third-party app platform, both by user base and the set of available apps. This popularity has made the platform's permission model change continuously over time, while trying to keep a balance between being appealing to both users and third-party developers alike. The current stable OS is v11 (with v12 right around the corner), while the API version, also defining the current permission model, is level 30. Android has evolved into a general purpose OS with plenty of protected data objects and actions; this amounts to 91 permissions in total, offered to third-party apps). We make 91 our baseline for the total number of relevant permissions. Out of these 91, there are 4 which explicitly and 16 which potentially interfere with others' personal data instantiating interdependent privacy, see Table 1. Note that the pop-up messages, appearing when installing an app from Google Play, contain warnings which can be mapped directly to API-level permissions with reasonable effort.

Table 1. Android permissions: IDP and PIDP

IDP	PIDP
read call log_Phone	read the contents of your USB storage_Photos/Media/Files
read your contacts_Contacts	modify or delete the contents of your USB storage_Photos/Media/Files
modify your contacts_Contacts	approx. location (network-based)_Location
read your text messages (SMS or MMS)_SMS	precise location (GPS and network-based)_Location
	access extra location provider commands_Location
	take pictures and videos_Camera
	read sensitive log data_Device & app history
	read your Web bookmarks and history_Device & app history
	record audio_Microphone
	read the contents of your USB storage_Storage
	modify or delete the contents of your USB storage_Storage
	find accounts on the device_Contacts
	read cal events plus confidential information_Calendar
	add or modify cal events and send email to guests w/o owners' knowledge_Calendar
	read cell broadcast messages_SMS
	find accounts on the device_Contacts

Browser extensions: permissions and warnings. Although referred to differently, browser extensions are very similar to apps. Extensions usually expand browser functionality, and manage user operations. Owing to their objectives and architecture, browser extensions are all about interacting with their respective platforms, often times resulting in obtaining large amounts of information about user operations in the browser in real time, but also about content downloaded by the browser. Note that browsers are also used to access intranets and other non-public resources, therefore, they might leak a variety of personal (and other confidential) information if something goes wrong. Both Firefox and Opera are based on Chromium, therefore their APIs and permission models facing third-party extensions are all based on the Chrome API (along with Chrome, Edge, Brave and Safari, to be correct). Both browsers have their own extension store.

Albeit they are based on the same APIs, Firefox and Opera have some unique characteristics. They both support the majority of permissions but not all[7], and

[7] `https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/manifest.json/permissions#browser_compatibility`

they both define their own warning messages that users can see before/when they install an extension[8]. In fact, Opera does not show these warnings when installing; they are only visible on their dedicated page in the extension store. Making things more complicated, i) not all permission requests generate warning messages, and ii) warning messages and API-level permissions are not totally consistent: the platforms have decided to simplify warnings for the sake of clarity to the average user. While this is laudable from one aspect, these explanations sometimes do not fully reflect the risks of granting the requested permissions. Exact mappings between permissions and user warnings are hard to find, but may be extrapolated from Chrome's official documentation[9]. Since it is only feasible to scrape the extension stores for per app warnings (and not for API-level permissions), we base our analysis on these. Note that our datasets contain information on Manifest V2 extensions; however, the changes introduced in Manifest V3 do not have a significant impact on interdependent privacy[10].

Firefox has 26 different warning messages, 19 of which are potential culprits for interdependent privacy violations, see Table 2. Opera extensions make use of 20 types of warning messages, 13 of which pose a potential threat owing to privacy interdependence, see Table 2. Note that all affected warnings (and their corresponding permissions) are in PIDP, and we omit NIDP due to space constraints. Also, note that all Opera warnings start with "This extension (can/will)".

Google Workspace (formerly GSuite). GSuite is a collaborative enterprise office platform launched by Google in September 2016, that was rebranded to Google Workspace in 2020, when it already had more than 2 billions users. Users do not need to download applications, they only need to edit and share files in the cloud, realizing remote collaboration. The platform has an app store, the Marketplace[11], where business, productivity and educational tools are offered by third-party developers. One of Workspace's subsystems, Google Drive, has already been shown to leak others' personal information through apps owing to its collaborative nature [8]; however, its permission model has changed completely due to the integration of various Google subsystems into the Workspace. The platform has many specialized permissions catering to its intended usage as a collaborative office productivity solution. Specifically, there are 87 different permissions, out of which 3 are IDP ("See and download your contacts", "View customer related information" and "View, edit, or permanently delete contacts" and 71 are PIDP (which we omit due to the lack of space). Note that, although Workspace is a subscription-based service for enterprises and universities, it hosts huge amounts of private data. What is more, if an employee (usually a system

[8] e.g. Firefox: `https://support.mozilla.org/en-US/kb/permission-request-messages-firefox-extensions`

[9] `https://developer.chrome.com/docs/extensions/mv2/permission_warnings/#permissions_with_warnings`

[10] `https://developer.chrome.com/docs/extensions/mv3/intro/mv3-overview/`

[11] `https://workspace.google.com/marketplace`

Table 2. Browser extension permissions: PIDP

Firefox	Opera
Access browser tabs	Access your data on all websites
Access browser activity during navigation	Access your tabs and browsing activity
Access your data for named site	Access your data on some websites
Exchange messages with programs other than Firefox	Exchange messages with programs other than Opera
Download files and read and modify the browser's download history	Capture the content of the entire screen or of individual tabs and windows
Access your location	Access data you copy and paste
Access recently closed tabs	Allows other installed extensions and web pages to communicate with this extension
Access your data for all websites	Detect your physical location
Store unlimited amount of client-side data	Manipulate privacy-related settings
Access your data for sites in the named domain	Know which sites you're visiting most often
Read and modify bookmarks	Read and modify bookmarks
Access your data on # other sites	Store an unlimited amount of client-side data
Get data from the clipboard	Read and modify your browsing history
Extend developer tools to access your data in open tabs	
Read the text of all open tabs	
Access browsing history	
Access your data in # other domains	
Access browsing history	
Read and modify browser settings	

administrator) installs a third-party app resulting in a privacy violation for other natural persons, the company can be held responsible as per the GDPR.

It is straightforward to see that each platform has a significant proportion of its permissions and warnings connected to interdependent privacy; see Table 3. This answers *RQ1* affirmatively: *interdependent privacy issues are indeed pervasive among third-party app platforms.*

4 Application-level statistics

4.1 Data collection

We collected datasets by scraping the app stores of 4 different third-party app platforms in late 2020 and early 2021: Android (10,589 apps), Mozilla (16,546), Opera (1,682)) and Google Workspace (882). Each record contains all available meta-data, e.g., app name, category, permissions/warnings, number of users, rating, etc., depending on the actual platforms. Due to the app stores' protection against scraping i) we did not manage to collect enough data for Chrome

Table 3. Summary: IDP and PIDP permissions/warnings

Platform	No. of permissions/warnings	IDP + PIDP	Ratio
Android	91	20	21.98%
Firefox	26	19	73.08%
Opera	20	13	65%
Google Workspace	87	74	85.06%

Table 4. Number of apps with IDP/PIDP

Platform	Apps with IDP	Apps with PIDP	Total Apps	Ratio
Android	1029	8307	10589	78.66%
Firefox	0	13704	16546	82.82%
Opera	0	1421	1682	84.48%
Google Workspace	29	845	882	97.62%

and Edge extensions, therefore we omit these platforms from our analysis; ii) our Android dataset contains only a fragment of the millions of available apps (yet, large and random enough to be significant). To the best of our knowledge, we collected complete datasets for Firefox, Opera and Google Workspace. Note that automatic scraping was infeasible for Google Workspace; we collected information on all available apps manually. Both datasets and scraping scripts are available for download[12].

4.2 Do real apps request IDP/PIDP permissions?

Table 4 shows the number of apps that requested at least one IDP or PIDP permission. The last column calculates the proportion of the union of these apps versus all the apps in the dataset. It is clear that the vast majority of apps are affected as evidenced by proportion larger than 80% for all platforms. Note that the browser platforms offer only PIDP permissions.

To further study the privacy protection permissions of apps, we calculated the permissions requested by each app. Regarding Android (Google Play store), 17.2% of apps requested the permission "Contacts", which means that their users have shared their contact list with the third-party developer, directly exposing others' personal data without their knowledge. Besides, 78.66% of apps have the potential to leak private information owing to interdependent privacy. On average, each app requests 11.21 permissions, out of which 4.4 are IDP or PIDP.

Mozilla and Opera extensions, despite their similar architecture, differ significantly in terms of PIDP warning types (0.83 vs. 3.93 on average) and total warning types (0.85 vs. 4.63) displayed. Note that, although here we observe warning instead of permissions, the difference holds, as warning-permission mappings are alike on both platforms. One reason could be that more Mozilla extensions make use of the `active_tab` permission (which does not generate a warning) instead of

[12] `https://www.dropbox.com/s/iz9kedsbzaw2vn1/liu_dpm2021_data.zip?dl=0`

Table 5. Average number of IDP&PIDP permissions per app

Platform	IDP/PIDP permissions	Total permissions	Proportion
Android	4.40	11.21	39.3%
Firefox	0.83	0.85	97.6%
Opera	3.93	4.63	84.9%
Google Workspace	2.04	2.42	84.3%

`_url` type permissions[13]. Some other permissions also do not generate warnings, therefore the total number of permissions requested in Table 5 is underestimated, while the proportion in the last column is overestimated for browser extensions.

Google Workspace is dedicated to collaborative enterprise features with a lot of PIDP permissions, therefore we expected a high proportion of those requested by apps. Indeed, 85% of total permissions requested (2.04 out of 2.42) are IDP or PIDP. We also observed that a majority of permissions are requested only by a few apps. This can be explained by the relatively low number of available apps, and the fact that permissions are very specific (especially compared to browser extensions), e.g., "View and manage your Google Slides presentations" instead of "View and manage your documents".

Based on the results above, we can also answer *RQ2* affirmatively: *actual apps do request IDP/PIDP permissions enabling collateral information collection on all studied platforms.*

4.3 Risk Signals

Users can obtain limited information when deciding upon installing third-party apps, such as category, number of users, user ratings, and permission types. Taking the Google Play store as an example, here we investigate whether the user can interpret these pieces information as risk signals towards interdependent privacy. Previous studies found that neither popularity (number of users), nor community ratings (stars) are good indicators for privacy-conscious app behavior [4]. We also found evidence supporting this hypothesis. In fact, community ratings show a weak positive correlation with both the number of total permissions and the number of IDP/PIDP permissions requested: favorable ratings are mostly based on advanced functionality requiring more permissions.

The only promising indicator for an app to enable collateral information collection is its category. In order to demonstrate this, we select $2,043$ apps from the Google Play dataset randomly, with the constraint of around 200 samples should belong to each of the 10 major categories. The average number of total permissions (left) and IDP/PIDP permissions (right) can be seen in Figure 1.

The number of permissions varies greatly across categories. Apps belonging to "Business" and "Communication" request an average of 14.74, 19.92 permissions, while "Art&Design" and "Comics" only have 8.26, 6.71. A reason-

[13] `https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/manifest.json/permissions#activetab_permission`

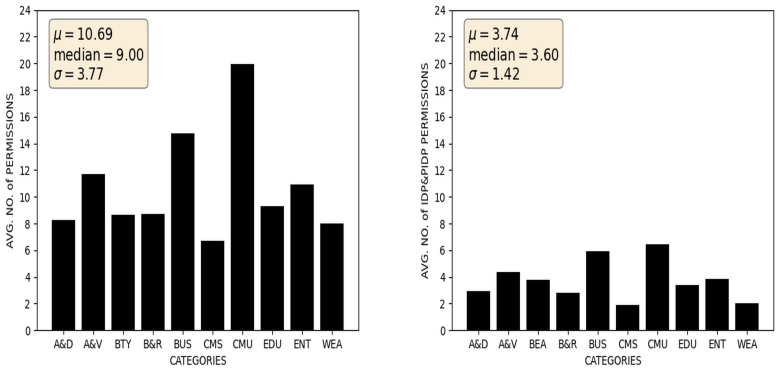

Fig. 1. Average number of permissions per app with different categories; Average number of IDP/PIDP permissions per app with different categories. Art&Design, Auto&Vehicles, Beauty, Books&Reference, Business, Comics, Communication, Education, Entertainment, Weather

able explanation to this result is that communications and business apps have more advanced features requiring more permissions. Interestingly, the same result holds for IDP/PIDP permissions. Categories with a high number of total permissions have a high number of IDP/PIDP permissions, and vice versa. The reasoning above can explain this result partially, but we argue that it is also in the characteristics of communication/business apps to involve more collaboration and multi-party interaction, a main theme behind permissions invoking interdependent privacy.

To illustrate this observation, we turn to the distribution of the number of IDP/PIDP permissions across all apps in a given category. Figure 2 shows the histogram for this metric for the categories "Art&Design" (left) and "Communication" (right). The difference between the two plots are striking: both the average number of IDP/PIDP permissions (2.95 vs. 6.41), and the shape of the histograms (top-heavy vs. normal-like) are very different. These patterns are mostly consistent for categories with a low and high number of permissions, respectively. This corroborates our previous observation, as more interactive/collaborative categories have more apps requesting a large number of IDP/PIDP permissions.

Naturally, our observations on risk signals can be Android-specific. It constitutes important future work for us to investigate these signals with respect other platforms.

5 Discussion: avoidance, transparency and control

In Section 3 and 4 we observed that i) all observed platforms offer permissions potentially invoking interdependent privacy, and ii) real apps do request a number of these permissions. Users are not particularly aware of interdependent privacy risks [22], and app platforms neither i) do a good job of informing the installing user and other persons affected by this issue, nor ii) offer control levers

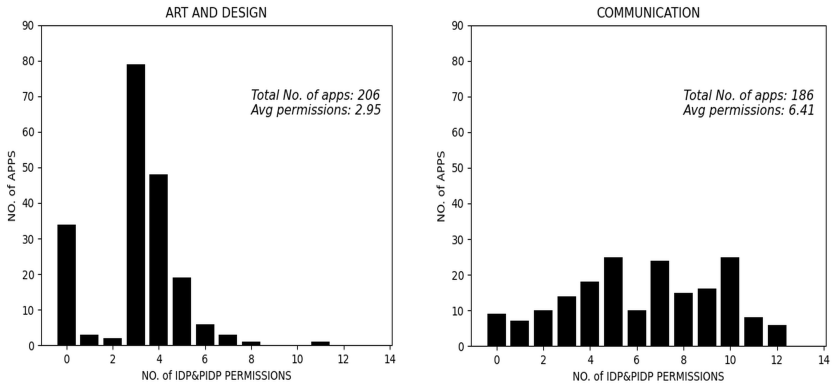

Fig. 2. Number of apps with different number of IDP/PIDP permissions in ART&DESIGN; Number of apps with different number of IDP/PIDP permissions in COMMUNICATION.

to influence such sharing. Therefore, it is up to the individual privacy awareness of installing users (acting as "amateur data controllers" [22]) and blind luck, that none of these platforms will experience its own Cambridge Analytica moment. Naturally, these options are neither satisfactory, nor systemic; in the following we discuss potential mitigation mechanisms, promoting risk avoidance, transparency and control.

Risk avoidance. A visceral response by app platforms to avoid exposed interdependent privacy risks could be to banish (most) IDP/PIDP permissions from their API. In fact, this is exactly what Facebook did in 2018, in response to the Cambridge Analytica scandal: it gutted its API for third-party apps, and introduced strict manual app review (hiring thousands of new employees)[14]. As evidenced by the declining popularity of Facebook apps, this might not be the most efficient way to deal with such risks. Indeed, the strong two-sided network effects characterizing app platforms require catering for both users and developers [17].

Transparency. Inspired by the GDPR and defined eloquently by Kamleitner's 3R insight framework [12], the sharing party (i.e., the amateur controller) can take 3 steps to reduce interdependent privacy risks: realize that there is a data transfer, recognize others' rights and respect others' rights. It is clear that transparency enhancing technologies can facilitate the first two steps. A potential way to make the sharer aware of interdependent privacy is to add a special warning sign to the already existing permission notification dialogues. Such a solution has to be platform-specific, and needs the co-operation of the platform owner. If such co-operation is unlikely, a dedicated interdependent privacy dashboard app can be implemented, in the manner of proposed dashboard designs for Facebook

[14] `https://about.fb.com/news/2018/04/restricting-data-access/`

apps [22]. Note that an exact public mapping of API-level permissions to user warnings could also improve awareness (especially for browser extensions).

Following the opinion of the Article 29 Working Party and the subsequent recommendations of Privacy International[15] in the TrueCaller case, affected data subjects (i.e., "others") should/could also be notified by the app developer using SMS, using the very data it acquired unlawfully (i.e., contact list). Such notification, however, is not a general possibility: it depends on the platform and the actual data collected.

Control. There are some privacy best practices that, when adhered to, would improve the situation on the developer and the platform owner side. These include requesting the exact minimum privileges an app needs (developer), and introducing well-defined, fine-grain permissions to enable asking for the minimum privilege (platform owner, especially for browser extensions).

Best practices aside, there is potential for interdependent privacy specific solutions that can enable better control of personal information both for affected users, privacy-conscious apps and platforms. Notifying affected data subjects and asking for their consent can be feasible for i) specific data types (e.g., contact list) or closed platforms such as Facebook (where all data are connected to other users of Facebook). In cases, where a certain data object is clearly connected to multiple natural persons (e.g., photos, messages, collaborative documents, calls), sharing mechanisms tailored to multi-party data may be utilized [21, 15]. It remains to be seen whether these can be incorporated efficiently into a third-party app platform. Another way to go is to combine permissions with enforceable policies (in the manner of [5] but regarding privacy), and control the information flow in run-time [11] (not between components, but among platform, developer, users and others affected). An interesting restriction would be to keep data acquired through IDP/PIDP permissions locally on the user device, enabling computation (if needed for full app functionality) but restricting data transfer. There are many challenges for such a solution, starting with non-structured data that is hard to label as "multi-party".

Indeed, we can make a case for interdependent privacy being inherently present in current app platforms. A radical solution to mitigate this situation would be to completely redesign the currently widespread permission-based access for app platforms, and try different alternatives.

6 Conclusion

In this paper, we have investigated whether interdependent privacy issues are present in popular third-party application platforms, such as Android, browser extensions and Google Workspace. Specifically, we have shown that there exist a significant number of permissions in all platforms that, directly or with reasonable probability, invoke interdependent privacy (RQ1). Moreover, via datasets

[15] `https://privacyinternational.org/node/2997`

collected from multiple platforms, we have demonstrated that actual apps do request these permissions (RQ2). We also found that the category of apps can be used as a risk signal for interdependent privacy. Finally, we have discussed potential solutions which could help in enhancing transparency and control. In our future work, we will aim at i) a comprehensive analysis of permission models in browser platforms, and ii) implementing the most promising of the potential transparency and control enhancing solutions discussed above.

References

1. Arzt, S., Rasthofer, S., Fritz, C., Bodden, E., Bartel, A., Klein, J., Le Traon, Y., Octeau, D., McDaniel, P.: Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. Acm Sigplan Notices **49**(6), 259–269 (2014)
2. Biczók, G., Chia, P.H.: Interdependent privacy: Let me share your data. In: Sadeghi, A. (ed.) Financial Cryptography and Data Security - 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers. Lecture Notes in Computer Science, vol. 7859, pp. 338–353. Springer (2013). https://doi.org/10.1007/978-3-642-39884-1_29, `https://doi.org/10.1007/978-3-642-39884-1\_29`
3. Boyd, D.: Networked privacy. Surveillance & society **10**(3/4), 348 (2012)
4. Chia, P.H., Yamamoto, Y., Asokan, N.: Is this app safe?: a large scale study on application permissions and risk signals. In: Mille, A., Gandon, F., Misselis, J., Rabinovich, M., Staab, S. (eds.) Proceedings of the 21st World Wide Web Conference 2012, WWW 2012, Lyon, France, April 16-20, 2012. pp. 311–320. ACM (2012). https://doi.org/10.1145/2187836.2187879, `https://doi.org/10.1145/2187836.2187879`
5. Fragkaki, E., Bauer, L., Jia, L., Swasey, D.: Modeling and enhancing android's permission system. In: European Symposium on Research in Computer Security. pp. 1–18. Springer (2012)
6. Gibler, C., Crussell, J., Erickson, J., Chen, H.: Androidleaks: Automatically detecting potential privacy leaks in android applications on a large scale. In: International Conference on Trust and Trustworthy Computing. pp. 291–307. Springer (2012)
7. Gnesi, S., Matteucci, I., Moiso, C., Mori, P., Petrocchi, M., Vescovi, M.: My data, your data, our data: Managing privacy preferences in multiple subjects personal data. In: Preneel, B., Ikonomou, D. (eds.) Privacy Technologies and Policy - Second Annual Privacy Forum, APF 2014, Athens, Greece, May 20-21, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8450, pp. 154–171. Springer (2014). https://doi.org/10.1007/978-3-319-06749-0_11, `https://doi.org/10.1007/978-3-319-06749-0\_11`
8. Harkous, H., Aberer, K.: "if you can't beat them, join them": A usability approach to interdependent privacy in cloud apps. CoRR **abs/1702.08234** (2017), `http://arxiv.org/abs/1702.08234`
9. Humbert, M., Ayday, E., Hubaux, J., Telenti, A.: Addressing the concerns of the lacks family: quantification of kin genomic privacy. In: Sadeghi, A., Gligor, V.D., Yung, M. (eds.) 2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013. pp. 1141–1152. ACM (2013). https://doi.org/10.1145/2508859.2516707, `https://doi.org/10.1145/2508859.2516707`

10. Humbert, M., Trubert, B., Huguenin, K.: A survey on interdependent privacy. ACM Comput. Surv. **52**(6), 122:1–122:40 (2020). https://doi.org/10.1145/3360498, https://doi.org/10.1145/3360498
11. Jia, L., Aljuraidan, J., Fragkaki, E., Bauer, L., Stroucken, M., Fukushima, K., Kiyomoto, S., Miyake, Y.: Run-time enforcement of information-flow properties on android. In: European Symposium on Research in Computer Security. pp. 775–792. Springer (2013)
12. Kamleitner, B., Mitchell, V.: Your data is my data: a framework for addressing interdependent privacy infringements. Journal of Public Policy & Marketing **38**(4), 433–450 (2019)
13. Kelley, P.G., Cranor, L.F., Sadeh, N.: Privacy as part of the app decision-making process. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 3393–3402 (2013)
14. King, J., Lampinen, A., Smolen, A.: Privacy: Is there an app for that? In: Proceedings of the Seventh Symposium on Usable Privacy and Security. pp. 1–20 (2011)
15. Olteanu, A., Huguenin, K., Dacosta, I., Hubaux, J.: Consensual and privacy-preserving sharing of multi-subject and interdependent data. In: 25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018. The Internet Society (2018), http://wp.internetsociety.org/ndss/wp-content/uploads/sites/25/2018/07/ndss2018_06B-1_Olteanu_paper.pdf
16. Olteanu, A., Huguenin, K., Shokri, R., Humbert, M., Hubaux, J.: Quantifying interdependent privacy risks with location data. IEEE Trans. Mob. Comput. **16**(3), 829–842 (2017). https://doi.org/10.1109/TMC.2016.2561281, https://doi.org/10.1109/TMC.2016.2561281
17. Parker, G.G., Van Alstyne, M.W.: Two-sided network effects: A theory of information product design. Management science **51**(10), 1494–1504 (2005)
18. Pu, Y., Grossklags, J.: Towards a model on the factors influencing social app users' valuation of interdependent privacy. Proc. Priv. Enhancing Technol. **2016**(2), 61–81 (2016). https://doi.org/10.1515/popets-2016-0005, https://doi.org/10.1515/popets-2016-0005
19. Reardon, J., Feal, Á., Wijesekera, P., On, A.E.B., Vallina-Rodriguez, N., Egelman, S.: 50 ways to leak your data: An exploration of apps' circumvention of the android permissions system. In: 28th {USENIX} Security Symposium ({USENIX} Security 19). pp. 603–620 (2019)
20. Squicciarini, A.C., Shehab, M., Paci, F.: Collective privacy management in social networks. In: Quemada, J., León, G., Maarek, Y.S., Nejdl, W. (eds.) Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009. pp. 521–530. ACM (2009). https://doi.org/10.1145/1526709.1526780, https://doi.org/10.1145/1526709.1526780
21. Such, J.M., Porter, J., Preibusch, S., Joinson, A.: Photo privacy conflicts in social media: A large-scale empirical study. In: Proceedings of the 2017 CHI conference on human factors in computing systems. pp. 3821–3832 (2017)
22. Symeonidis, I., Biczók, G., Shirazi, F., Pérez-Solà, C., Schroers, J., Preneel, B.: Collateral damage of facebook third-party applications: a comprehensive study. Comput. Secur. **77**, 179–208 (2018). https://doi.org/10.1016/j.cose.2018.03.015, https://doi.org/10.1016/j.cose.2018.03.015
23. Wang, N., Xu, H., Grossklags, J.: Third-party apps on facebook: privacy and the illusion of control. In: Proceedings of the 5th ACM symposium on computer human interaction for management of information technology. pp. 1–10 (2011)