An Analysis of Apple's Private Click Measurement

Martin Thomson
Mozilla
mt@mozilla.com

June 23, 2022

Abstract

Apple has proposed Private Click Measurement (PCM) as a means of allowing websites to measure the effect of advertising. A privacy analysis of PCM is provided, examining its claims in terms of what privacy assurances it provides. We also include a discussion of the impact of the privacy protections in PCM on its usefulness for measuring advertising.

1 Introduction

Apple has enabled Private Click Measurement (PCM) [Wil21] for all users in Safari and in iOS applications. The goal of PCM is to support the measurement of advertising, while providing Apple's users with improved privacy. PCM allows a website to request that the browser generate a report about the activity of a user on two sites, by convention the publisher site on which an ad appears, and the advertiser site on which a purchase occurs.

1. A first web site annotates links — which are expected to be outgoing links in advertisements — that it wants to measure, providing a *click identifier* for the link and the identity of a second site that is the target of the link.
2. The user "clicks" on the link from the first site and navigates to the second site. The browser remembers this click for up to 7 days, along with the identity of both sites and the identifier.
3. The user later interacts with the second site — for instance to make a purchase — which requests that the browser generate a report. The site provides a second *trigger identifier* for this event. If a click was recorded by the browser that navigated to this site, the browser adds this identifier to the store and schedules a report to be sent.
4. Between 24 and 48 hours later, the browser submits a report to both sites using an anonymization service. This report includes the identity of both sites and the identifiers provided by each site.

PCM is designed to measure the most immediate form of ad interations: where users click on ads and subsequently purchase a product or perform some action on the destination web site.

PCM also allows for the case where the initial "click" occurs in an iOS application. This analysis mostly concentrates on the Web case, only noting where a click in an app produces notably different properties.

The remainder of this document is structured as follows. In Section 2 we analyze the privacy properties of PCM. In Section 3 we discuss the effect of the privacy protections in PCM on its effectiveness of measuring advertising. Section 4 concludes with an overall assessment of PCM.

Note that this analysis is based on a reading of the PCM specification [Wil21], together with postings on the Webkit blog. A number of details of Apple's implementation of PCM are only available in these blog postings. These are are not always sufficiently detailed to know exactly how the feature is intended to operate.

2 Privacy Analysis

Apple's proposal includes several measures designed to limit the amount of information that sites can gain about user activity on other sites:

- Reports are delayed and routed via an anonymizing proxy to limit the use of timing and IP addresses to match reports to the triggering event.
- The number of identifiers sites can use is limited.
- Reports are only generated when users navigate between sites.

2.1 Reporting Delays

A report that is sent immediately can easily be matched to the event that triggered it. This would give the site that triggered the report precise information about the click, such as whether a user clicked on another site, which site it was, and the click identifier that was used. Immediate reports therefore provide high fidelity information that could be used for tracking. To prevent this sort of information leakage, PCM delays report submissions by 24 to 48 hours. Then, when the report arrives, it could correspond to any triggering event in a 24 hour period.

Apple's implementation of PCM routes the sending of reports through their Private Relay service using fresh connections that do not include any identifying information. This ensures that sites are not able to use IP addresses or any part of the request to match reports to triggering events.

This aspect of the design is effective only if we assume that sites trigger the generation of a large number of reports. If a very small number of reports

are generated in a 24 hour period, it becomes possible to match reports to the triggering event.

Take the case where just two triggering events occur with t hours between each. If the first report arrives less than $24 + t$ hours after the first event, it cannot be for the second event. The potential for delays in reporting makes it possible for the first event to arrive more than 48 hours after the first event, but between 48 and $48 + t$ hours, a second report will most likely be caused by the second event.

This method might be used to match reports with trigger events probabilistically, which is more effective if the click and trigger identifiers are not used for a large number of users. If the same pair of identifiers is reused with larger numbers of users, the amount of information gained diminishes.

2.1.1 Sleeping Browsers

The analysis above assumes that browsers are constantly active. However, browsers that are not active cannot send reports. Reports will therefore be submitted during the hours that users are active. This makes report timing less predictable in the sense that reports might be submitted later than planned.

This might have some privacy benefit because late arriving reports could obscure when they were triggered. However, the condition that a browser is active potentially makes it much easier to link reports to trigger events. Sites are likely to know when users are active, simply by observing visits, which only happen during active periods. Sites can correlate user activity with report submission times in order to match reports to trigger events.

This might make it easier to identify users with less common browsing hours, such as those in time zones that are less popular with visitors of the site or those who keep less common hours.

As an aside, it is not clear from PCM documentation how long the browser will retain a report if it is not active for an extended period and therefore unable to send it. On the basis that the PCM design might want to allow for repurposing identifiers on a predictable time scale, in subsequent analysis we will assume that unsubmitted reports will be discarded 9 days after the original click.

2.2 Identifiers in PCM

Apple aims to reduce the ability to track using reports by making identifiers relatively scarce. Each click can be assigned one of 256 identifiers; each trigger event can be assigned one of 16 identifiers.

The implied goal with PCM is to force sites to assign the same identifier to multiple users by making the space of identifiers very small. PCM is unable to force sites to comply as the allocation of identifiers is discretionary.

Sites are not prevented from tracking specific individuals; thus PCM only offers mild incentive to avoid tracking by making measurement less effective identifiers are reserved for tracking. PCM therefore only limits the number of people that can be tracked.

2.2.1 Tracking with PCM

The way an API like PCM might be used for tracking is that a pair of sites agree to track a particular user.

To begin, we have both sites with some number of users that might visit those sites. Each site gets a different perspective on each user. The same person might visit both sites, but they can interact with each site differently. The sites might not be able to connect the user identity used by a person on one site to an identity on the other site; see Figure 1 for an illustration of this.

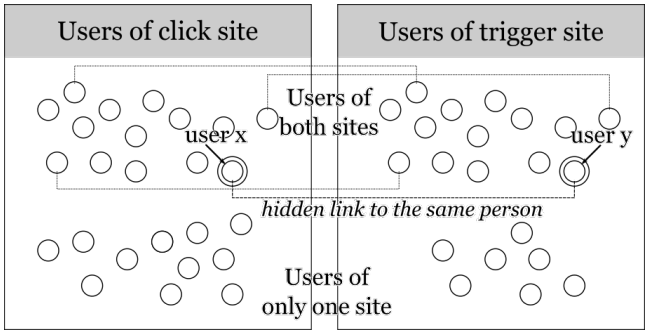

Figure 1: Linking user identities across sites

The two sites might have some idea about the correspondance between user identities. They might use IP addresses, fingerprinting, or just observations about user activity to make some guesses. The sites want to confirm a guess that a user on one site (we might say user x) is the same person as a user on the other site (user y).

To link these identities to the same person, the two sites need to convince the user to navigate from one site to the other. The first site creates content that it only shows to user x. This content links to the second site and includes a PCM click identifier that is specific to user x.

The second site then creates a trigger event for user y. If user x and user y are the same person, PCM will generate an attribution report that both sites will receive 24 to 48 hours later, confirming this guess. If the guess is incorrect and user x and user y are different people, no report is created.

Once two sites have confirmed that their different user identities correspond to the same person, they can exchange information about the activities of that person as they choose. They can track this person across the two sites without any further indication that the tracking is happening.

2.2.2 Linking Multiple Users

If the second site is uncertain about which of its users are the same person, it can use the trigger identifier to select between different options. The second site assigns a different trigger identifier to up to 16 users. The trigger identifier in the report will show which guess was correct; see Figure 2.

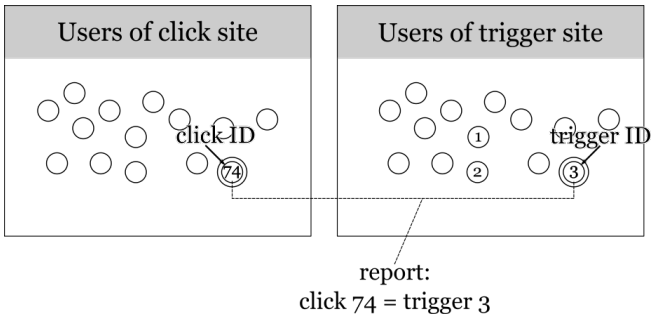

Figure 2: Resolving uncertainty about user identity

Sites that are more confident about how user identities match can instead use the combination of click and trigger identifier to confirm guesses about multiple users. This allows for more efficient use of a limited number of identifiers; see Figure 3.

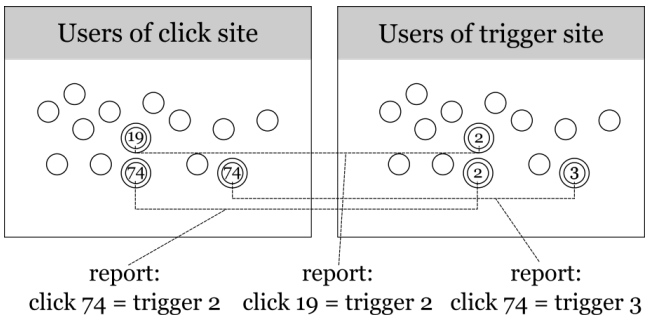

Figure 3: Using identifiers to match multiple user identities

Matching multiple users adds a risk of a bad match. If the sites guess incorrectly about which user identities correspond to a person, but that person has matching identity in the set of users that have been assigned the same click identifier, a report will be generated that appears to confirm a match.

For example, sites might choose to assign two users, x_1 and x_2, the same click identifier of 74. The sites hypothesize that users $y_1 = x_1$ and $y_2 = x_2$. The second site assigns trigger identifiers of $y_1 = 3$ and $y_2 = 2$. If the sites have guessed incorrectly and user x_1 is really the same person as user y_2, a report might be generated for $(74, 2)$, which appears to confirm the guess

that $x_2 = y_2$.

The risk of this sort of false match increases if all 16 trigger identifiers are used. Sites could manage this risk by using fewer trigger identifiers or by using other information to isolate linking attempts. For instance, by choosing users with a history of very different active hours to share a click identifier. Additional matching attempts might be used to increase confidence in matches.

2.2.3 Ideal Tracking Rate

Using this approach, two sites can link the identity of up to 4096 users at a time in one direction. If the sites guess correctly, then the reports PCM generates will confirm each of those guesses in 36 hours on average. This leads to an expected rate of approximately 2419 identities matched per day in one direction if sites are able to guess correctly for each user. This is doubled to 4838 identities per day if users navigate in both directions.

However, sites might not receive reports within 48 hours, either because they guessed incorrectly, or because the user's computer is not running when the report is scheduled to be sent.

If no report is received for a given click and trigger identifier, then that pair of identifiers cannot be safely reused until the browser stops tracking the click and any report that might be generated from that click. As clicks are remembered for 7 days and reports are delayed for a further 2 days, this means that reuse of those identifiers any sooner than 9 days from the time of the click might result in users a false match.

Failed guesses produce a probabilistic negative confirmation after those 9 days. This negative result needs to allow for the possibility that a positive report was lost. This reduces the rate at which users can be matched when sites are less able to guess correctly.

2.2.4 Audience Partitioning

Confirming a guess about user identity across sites is unlikely to be feasible for many sites. Sites might not be able to use fingerprinting or other techniques to reduce the set of candidate user identities to just 16, which would allow for PCM to be used directly.

PCM can be used to progressively partition users that visit both sites. Over multiple iterations of this process, user identities can be matched with no prior knowledge.

This requires selecting a set of 256 users of the click site, which are each allocated a unique click identifier. At each iteration, users of the trigger site uniformly partitioned into 16 different groups, each with a different trigger identifier. Each iteration allows the sites to learn which of the 16 groups each user falls in. If there are N users of the trigger site, after

$\lceil \log_{16}(N) \rceil = \lceil {}^{log_2(N)}/_4 \rceil$ iterations of this process, the identity of each of these 256 users can be matched.

This results in matching the identity of at most ${}^{604}/_{log_2(N)}$ users per day[1].

The same process can be applied in reverse with unique identifiers being allocated on the trigger site, with users at the source site being progressively partitioned. This is 8 times slower for sites with the same number of users. Applying the partitioning at the source site would only become more feasible if the trigger site has more than 256 times more users than the click site, due to the smaller number of trigger identifiers.

Information obtained using fingerprinting or other tracking techniques might be used to reduce the number of iterations needed to match each pair of user identities.

2.2.5 Multiple Sites

The rate at which user activity can be linked across sites applies to each pair of "sites"[2]. Each site that performs this sort of tracking can independently perform this user identity matching process with any other site.

Site operators can increase the number of users they can track using PCM by creating additional sites. Though cross-site navigation like this has a user-visible effect, this might go unnoticed. Sharing information between these sites likely requires the use of navigation tracking techniques; see Section 2.3. For closely coordinated sites like this, this sort of information sharing is presently viable.

Apple has stated that if it detects abuse of PCM they might penalize sites by blocking access to the API. Use of multiple sites might be readily detectable and qualify as abuse. The less overt forms of abuse such as those described in Sections 2.2.1 through 2.2.4 might be difficult to detect.

The rate at which user identity can be linked across sites is directly proportional to the number of sites. If multiple sites are used on each end, the increase is quadratic; ten sites on one side of the interaction can track 10 times as many users, but ten sites on both ends can track 100 times as many. For clicks that originate in iOS applications this cannot be used on the source side.

[1] This follows the same expected rate of confirmations as in previous sections for the 256 users that are targeted $({}^{256}/_{(1 - \log(0.5))})$. This is then divided by the number of iterations needed to identify each user on the trigger site $({}^{log_2(N)}/_4)$. This only depends on the 256 *targeted* users all submitting reports within their scheduled 24–48 hours.

[2] PCM defines a site as origins that share the *registrable domain* (or *eTLD+1*) portion of a domain name. For example, "example.com" is the registrable domain for both "www.example.com" and "mail.example.com". This is the same boundary that is used for sharing cookies across origins. This makes sense in that cookies can be used to share information about users between different origins that share a registrable domain.

2.3 Only Navigation

Limiting the API to measure only navigations ensures that PCM only provides information under tightly-controlled conditions. PCM also requires that users engage with sites to navigate, either through a click or key press. Most importantly however is that navigation can already be used to transfer information between sites about users.

Transferring information between sites when following a click is called navigation tracking [SY21]. The most obvious way form of navigation tracking is link decoration. The linking site modifies links to include parameters that — rather than aiding in the identification of the destination page — pass information to the destination site. This might include information about the user, the ad they clicked on, or any other information the linking site wants to include.

Tracking using navigation is something that browsers are in the process of trying to understand and prevent [SY21]. However, it is generally recognized that fully preventing navigation tracking is an unsolved problem. This is partly because it is hard to distinguish between the information in a link that is necessary to identify the target page and other information that might be used to identify a user. It is also because there are timing side channels that allow sites to correlate user activity across sites without relying on signals that users — or browsers — can observe.

Limiting measurement to navigation means that PCM doesn't introduce new ways for sites to communicate about the people that visit them, which limits the additional privacy loss the proposal might produce. PCM can therefore be framed as an improvement over the use of navigation tracking for measuring clicks.

Measuring only navigation makes PCM less usable for measurement, as we will see in Section 3.3.

3 Functionality

The ways in which PCM attempts to protect privacy limit the usefulness of the information it generates. PCM has several obvious limitations:

- Reports are delayed
- Identifiers are deliberately scarce
- Only the last click is measured
- Only involved sites receive reports
- Additional fraud mitigation is necessary

There is some amount of natural tension between privacy and utility goals; so it is not surprising that privacy safeguards have some adverse effects on measurement. A privacy-preserving system cannot maintain the

fidelity of a system that tracks user activity. However, these limitations would seem to be quite significant when taken together.

3.1 Delays

The 24 to 48 hour delay between a triggering event and the submission of a report delays any response to measurement of a campaign. Actionable information might not be available until at least 24 hours after starting or changing a campaign.

3.2 Scarcity

Identifier scarcity is an obvious limitation. PCM permits the use of just 256 identifiers for click events and 16 identifiers for trigger events. Even for a smaller advertiser, this can mean that multiple campaigns need to share identifiers. Advertisers routinely divide run variations on campaigns to observe how different approaches are received. This might involve tuning a number of variables such as the creative (the image, video, or audio) and targeting or bid conditions. Moreover, it they run multiple concurrent campaigns, choosing how to assign from the limited set of 256 identifiers presents a difficult challenge. The 16 identifiers for trigger events are even more scarce.

Larger advertisers will have greater problems. Advertisers might find that 256 identifiers is not sufficient to even allocate a unique identifier to each campaign.

Over time, reuse of identifiers is possible, but as browsers remember click events for up to 7 days, this can only be done safely — that is, without risk of capturing reports that use the old meaning — after 9 days[3].

Limitations on the number of identifiers has created some awkward externalities for the API, such as pressure on the public suffix list[4]. Registering a domain name on the public suffix list allows an advertiser to allocate multiple subdomains, for which PCM will allow a separate identifier space. This workaround arbitrarily increases number of identifiers, but at a cost of increasing the number of public suffixes.

3.3 Attribution Models

it might be possible to register different events, such as impressions, PCM deliberately only counts clicks at the source site. This means that PCM is

[3]For trigger identifiers, this is not obvious. It is because a trigger identifier that is used immediately after a click event might be retained until the click expires if the browser is not able to submit a report.

[4]`https://github.com/publicsuffix/list/issues/1245`, `https://github.com/privacycg/private-click-measurement/issues/78`

unable to support the range of measurements that the advertising industry currently relies on.

Advertising is used for a variety of purposes. Not all advertising exists to drive something as direct as a click. Some ads exist to improve brand awareness. Some ads are placed with no expectation of immediate action. Measuring only direct clicks does not always capture all of the ways in which ads might be effective.

Existing ad measurement systems recognize this by enabling a range of attribution models. These models aim to apportion credit for the display of ads in multiple places. Where there are multiple ads involved, each model places different value on views and interactions with ads, plus the time that these different events occur.

Some attribution models place high value on the most recent direct interaction, excluding older interactions and any views. For this sort of attribution model, PCM might be suitable. Attribution models that value other events are not supported by PCM.

3.4 Same-Browser, Same-Device Matching

PCM only matches events that occur within Safari, or clicks in iOS applications that ultimately convert in Safari on the same iOS device. Should another browser adopt the API, only events that occur in that single browser could result in a report. Events that occur on different devices cannot be recognized as contributing to a single report.

Users that view advertisements on one device (or browser) but subsequently convert in another won't be recognized by PCM. This means that the values PCM produces will miss a non-trivial proportion of attributions.

The same criticism applies to many of the proposals in this space.

3.5 Reporting Constraints

The PCM specification [Wil21] states reports are provided only to the site where the click occurred. However, a recent blog posting[5] notes that both source and trigger site receive these reports.

For sites where advertising is not central to their business, they might rely on other parties to operate those parts of their business. Sites that show ads often use one or more external Supply Side Platforms (SSPs); sites that place ads will use one or more agencies or Demand Side Platforms (DSPs). PCM requires that sites receive reports directly. Sites cannot delegate responsibility for handling reports to another site.

Reports for iOS applications can be sent to an arbitrary URL.

[5] `https://webkit.org/blog/11940/pcm-click-fraud-prevention-and-attribution-sent-to-advertiser/`

3.6 Fraud Mitigation

Fraud is a major problem in online advertising. Advertisers spend large amounts of money to place advertisements. Some sites might try to take advantage of poor accountability and traceability to defraud advertisers of that money by claiming to have displayed ads.

The prevalence of fraud is particularly relevant in a system like PCM, where the goal is to remove any ability to link different events: the click, the trigger, and the report. The limited information that is carried across into the attribution report makes it harder to use contextual information to determine if a click was fraudulent.

PCM aims to address this by attaching unlinkable tokens to reports[6]. These tokens are generated by the click source when a click occurs and then included in the report. These tokens are made unlinkable through the use of a technique called "blinding" [Cha83]. Blinding prevents the issuer of the token being able match the value they receive with the one they issued, they can only learn that they issued it.

The use of tokens goes some way to addressing the fraud problem. If sites are honest, only a genuine click will result in a valid token being issued. Sites can use the token to identify which reports are genuine as opposed to which were generated by dishonest clients when clicks did not occur.

However, one of the primary sources of fraud are sites that try to claim that ads are displayed when they were not. These sites are responsible for creating tokens. This gives a potential fraudulent site good opportunities to insert false reports to inflate the number of reports that lead to conversions.

As reports are also sent to advertisers, an advertiser is able to bound the total number of reports by the page views they receive from a site. However, bots or other fraudulent practices might be used to inflate this number. It also does not address other tricks that drive clicks, which might then be misrepresented, either as being from an advertisement — when they are not — or for a more valuable advertisement than was clicked.

It appears as though sites are unable to delegate the creation of unlinkable tokens to another site. However, use of HTTP redirects might allow this, depending on how PCM is implemented.

4 Conclusions

PCM relies on the limited availability of identifiers being sufficient encouragement for advertisers to aggregate events by sharing identifiers across users. The idea seems to be that obtaining information about individual users using PCM is sufficiently impractical that sites will not bother using the feature for tracking.

[6]This information is also based on blog postings (ibid.) rather than the specification. This format lacks some of the detail that might be necessary to fully analyze the design.

At a fundamental level, measuring advertising means that sites will learn something about how people use other sites. Ads are shown on different sites with a goal of encouraging actions on another site. Thus, producing any sort of measurement naturally reveals something about how people interact with multiple sites.

Any proposal in this space needs to guarantee that sites receive less information about people and their browsing activities than would be required to link the identity of users on different sites. PCM provides accurate, cross-site information about specific people, making such guarantees difficult or impossible to provide.

PCM implies a simple trade-off between utility for advertisers and privacy for users:

- Fewer identifiers means worse utility and better privacy.
- More identifiers means better utility and worse privacy.

Choosing a small number of identifiers in PCM suggests that Apple has chosen to favor privacy. This analysis shows that there is no limit on identifiers that avoids potential privacy problems. The scarcity of identifiers only disincentivizes tracking by honest sites that seek to use the API for its intended purpose. For malicious sites, the privacy controls in PCM only limit the rate at which additional users can be tracked.

Overall, PCM does not provide users with any guarantee that sites are unable to use the information it provides for tracking. For sites and apps that seek to measure their advertising, the measurement capabilities PCM provides comes with functional limitations that appear to make it difficult to use effectively.

References

[Cha83] David Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology*, pages 199–203, Boston, MA, 1983.

[SY21] Pete Snyder and Jeffrey Yasskin. Navigational-Tracking Mitigations. `https://privacycg.github.io/nav-tracking-mitigations/`, Oct 2021.

[Wil21] John Wilander. Private Click Measurement. `https://privacycg.github.io/private-click-measurement/`, Apr 2021. Draft Community Group Report.