ACDC: Anonymous Crowdsourcing using Digital Cash

Luis A. Saavedra(✉)[1][0000−0001−8684−8763] and Alastair R. Beresford[1][0000−0003−0818−6535]

Department of Computer Science and Technology, University of Cambridge, Cambridge, United Kingdom {Luis.Saavedra,Alastair.Beresford}@cl.cam.ac.uk

Abstract. Crowdsourcing applications are vulnerable to Sybil attacks where attackers create many accounts to submit bogus or malicious data at scale. The traditional approach to manage Sybil attacks is privacy invasive since it requires contributors to identify themselves when contributing data. In this paper we present a new reporting protocol which supports the anonymous submission of data to crowdsourcing systems by honest contributors, while identifying malicious individuals who attempt to submit multiple reports. Our approach builds on Chaum's digital cash, and we demonstrate its practicality and deployability on mobile devices based on its low storage, network, runtime, and power requirements.

Keywords: Digital cash · Crowdsourcing · Anonymity · Sybil attacks.

1 Introduction

Crowdsourcing apps outsource the collection and processing of data to the *crowd*, a loose collection of contributors who support the operation of the app by providing or processing data. Crowdsourcing apps can operate entirely in the virtual domain, and include tasks such as labelling photos, and can also cross the boundary into the real-world, using mobile devices owned by individuals to collect information about our physical environment. Such technologies bring many benefits, including cheap and scaleable ways to monitor road-related events such as congestion, accidents and hazards, showing occupancy of public spaces, transport, and private venues (e.g. Waze and Google Maps) or generating temperature and air quality (AQI) maps [11, 39, 41].

Crowdsourcing apps authenticate individuals who provide their data via direct personal identifiers obtained during registration, such as name and email address; or indirect identification methods, such as rate limiting by IP address. Authentication is required to prevent Sybil attacks [19] where a malicious contributor overwhelms a crowdsourcing system with enough bogus data that the app no longer functions correctly [32, 49]. Once data are collected, many crowdsourcing apps apply anonymisation or privacy-preservation algorithms such as differential privacy to aggregate or summarise data safely before publication. This approach requires individuals to trust the app with their data and (where applicable) that the privacy-preserving processing steps are applied correctly.

In this paper we improve the privacy of data contributors in crowdsourcing systems by separating contributor authentication (required to prevent Sybil attacks) from the individual reports submitted by the contributor. Our approach improves the privacy of contributors since reports often contain sensitive information about the person, such as the places they visit in a city.

Our approach builds on Chaum's digital cash [8, 10], allowing authenticated contributors to obtain a limited number of digital tokens which can then be attached to reports in a crowdsourcing app, thereby allowing the app to know such reports come from a genuine contributor, but prevent apps from associating particular contributions with specific individuals. The use of Chaum's protocol allows the system to identify malicious contributors who attempt to contribute multiple data items using a single digital token.

The reported data may itself contain information which allows re-identification of the contributor, so we explore the use of k-anonymity to provide data privacy of the report contents too. We illustrate our approach by focusing on location information, which is required in almost all apps in the real world and also has repeatedly been shown to contain sensitive personal information [16, 26, 51]. We use the running example of crowdsourcing traffic data since the mobility of contributors and frequent reporting requirements make it particularly challenging, however our approach is applicable to a wide range of crowdsourcing scenarios.

Our paper makes the following contributions:

- A new protocol for anonymous data collection that prevents Sybil attacks.
- A method to cryptographically bind data items to digital tokens, allowing data to be shared over untrusted networks, such as peer-to-peer networks.
- Support for client-side k-anonymity, suitable for use in crowdsourcing applications in the physical domain.
- An evaluation of the practicality of our proposal, where we find that one million validated data items can be stored in 3.6 GB; data submission from a smartphone takes less than 34 ms and 6.5 mJ; and signature and verification steps performed by crowdsourcing applications take less than 100 ms.
- An evaluation of our k-anonymity approach to protect the privacy of data items themselves, showing, through traffic simulations, that 20% of cars can always report with k between 22 and 25 in a medium-sized European city.

2 System overview and security goals

Separating the verification of contributor identities and the submission of reports protects contributor privacy. We assume the presence of an identity authority that authenticates contributors but is not involved in the reporting or app data processing. Apps collect valid reports from contributors registered in the system but do not know the identity of contributors. Therefore, there are three different types of entity: the identity authority, the apps, and the contributors who submit reports. In our traffic reporting example, identity authorities might be car manufacturers or smartphone vendors; apps run by city councils or another app developer; and contributors could be car drivers or public transport passengers.

Identity authorities identify contributors based on credentials already used by crowdsourcing systems, and might include personal data such as IP or email addresses, digital attestation by smartphones, hardware tokens, or other methods. Identity authorities then provide contributors with a limited number of digital tokens which can be attached to reports which contributors can share with apps. Apps can then check that the digital tokens are valid (and therefore come from a genuine contributor) without revealing the identity of the contributor either to the app or the identity authority. Thus our approach meets our primary privacy goal: it allows crowdsourcing platforms to prevent Sybil attacks through contributor authentication, without either the identity authority or the app being able to associate a specific report with a particular contributor. In order to protect apps against malicious contributors, our protocol should ensure that digital tokens cannot be forged or attached to more than one report (i.e. prevent double spending). It should prevent multiple, colluding contributors from send more valid reports to an app than they are authorised to submit.

In order to prevent apps from identifying contributors through their network address, we assume reports are submitted to apps via trusted proxies or anonymous communication networks such as Tor, Loopix, or Nym [18, 42, 17].

3 Authentication protocol

Our system requirements are similar to Chaum's digital cash [8–10], where the *identity authority* is equivalent to the bank, *apps* to merchants, and *contributors* to customers. The blind token generation in this scheme meets our primary security goal, preventing the apps and identity authority from associating a specific token with a contributor, and also prevents double spending.

However, there is an important difference: in the case of digital cash, the token is the entire piece of information sent by the customer to the merchant. In our scheme, the contributor needs to bind the token with a report in a non-invertible way, preventing any intermediary who might receive the report from altering the contents of the report or extracting the token (see §4). It also prevents a malicious app from altering the report or extracting the token in the case where we wish to have an external validator to check the work of an app.

To do so, we follow the token generation protocol as presented by Chaum. The identity authority can use different private keys to sign valid tokens for different crowdsourcing reporting periods as well as limit the number of tokens and reports per contributor. Our protocol follows the payment protocol [10, pp.320–321] except for step 3. Instead of simply answering the app's challenge Z, Alice computes a new challenge $Z' = h(Z, R)$ where R is her report and h is a collision-free one-way cryptographic hash function specified by the identity authority. Subsequently, steps 4–6 are adjusted to include verification and forwarding of the variables Z' and R. By answering challenge Z' instead of Z, Alice has bound her report to the token, preventing their undetectable modification.

Our modified version of Chaum's protocol meets our security goals. Chaum's RSA-based blind signature scheme has been proven secure in the random oracle

model [4]. Digital cash protocols [9, 10] were also proven secure [20], including the inability to forge or alter tokens, link contributors to tokens, and distinguish two tokens submitted by the same contributor.

In the original protocol, banks generate digital cash based on actual Fiat currency deposited. In our use case, the identity authority can decide how many reports each contributor is allowed to submit in a given reporting period. This is a policy question that can be decided via discussion between the identity authority and the apps. For example, the identity authority could issue one token per time period per device; one per contributor; or a different number of tokens for different types of device. Furthermore, the identity authority might have different types of token, or different tokens for different apps, allowing apps to set their own policies. Since the identity authority sees contributor identities, different numbers of tokens can be issued for different types of contributor. These could be made visible to apps by the identity authority by using different different private keys for each type, or not. Our protocol therefore supports a range of policy options. The constraint is each app has to decide how many tokens are issued per-identity and how contributors are grouped into different types (if any).

4 *K*-anonymous reporting

While our protocol guarantees that reports are submitted anonymously to apps, the contents of those reports may reveal the contributor's identity. One of the biggest risks is a report containing location data, since it is essential in many crowdsourcing scenarios and may also identify individuals uniquely, even when no other information is reported. For example, a sequence of timestamped contributor locations may record a journey from home to work, which has previously been shown to uniquely identify most individuals [30]. In this section we focus on how to protect the privacy of contributors submitting reports which include location data; apps may need to protect other types of potentially identifiable information in a similar way depending on the data collected.

We assume that the electronic devices of contributors can communicate peer-to-peer wirelessly with other devices nearby. This includes smartphones communicating using Bluetooth or WiFi Direct, or vehicles talking with other vehicles and local infrastructure using wireless vehicle communications. Our basic idea is to use the local wireless network to check physical proximity of k contributors who also intend to submit a report to an app. Doing so means that each contributor knows that there will be k reports received by the app from this location, thereby protecting location privacy through k-anonymity. In order to allow contributors to check that $k-1$ others intend to submit reports from their current location, contributors can share their reports via the peer-to-peer network.

Since our protocol in Section 3 binds a report to a token, these reports can safely be shared with otherwise untrusted local contributors over the peer-to-peer network as the report cannot be modified and the token cannot be recovered or attached to a new report. The downside of sharing reports in this way is that there may not be k nearby contributors, and if one or more nearby contributors

are malicious, then the contents of the reports might be shared with the app or another entity without reaching the required threshold to achieve k-anonymity.

To solve this problem, we introduce a three-phase protocol. In the first phase, each contributor, c, creates a hash of their token, $H_c = h(C_c)$ and broadcasts this to all other local contributors. When a contributor receives k or more distinct hashed token values, they enter the second phase by forming a single challenge, Z_c, which is the concatenation of all the hashed token values received along with a hash of their report ($Z_c = H_1|H_2|...|H_k|R_c$); including the hash of the report in Z_c binds the token to the report. The contributor runs through our reporting protocol (see §3) using Z_c as the challenge to produce a response, U_c. This way each contributor can only perform the second phase once, otherwise the contributor will reveal their identity to the identity authority and app in the third phase as the token would be double spent. Then, contributors broadcast their token C_c and values of Z_c and U_c to all other contributors on the local network. Note these values are sufficient to validate the token, challenge and report, but not enough for others to generate a new response [10].

In the third phase, contributors check the validity of all tokens C_{c_1} through C_{c_k} using the challenges Z_{c_1} through Z_{c_k} and their responses U_{c_1} to U_{c_k} respectively. If the challenges are valid, then each contributor knows that there are at least k contributors with valid tokens connected to the local network. Each contributor, c, can now submit their report along with Z_c and U_c to the app.

Applying this protocol to the traffic reporting example, vehicles can detect each other's presence over a local network, run through the three-phase protocol initially and then, if the threshold k is reached, submit their reports to the app.

5 Results, performance and simulations

This section discusses protocol performance in terms of storage, CPU, time, and power requirements, as well as a brief feasibility assessment. We also estimate values of k for the k-anonymous protocol based on traffic simulations.

Storage: before tokens are used, contributors need to store each signed token of size N, together with 5 parameters of size N (a, c, d, x, and y) [10]. Thus $6N$ storage is required for each token. After reporting, the identity authority will need to store the following values: the token, report R, challenge Z, and $Z' = h(Z, R)$, as well as the response U which requires $3N$ storage. In order to detect double-spending, Z and Z' need to be of length N. The report R can be of any arbitrary size $|R|$, thus each report and token requires $6N + |R|$ of storage. Assuming RSA 4 096 is used as the secure signing algorithm $N = 4\,096$ bit. If $|R| = N$, then each transaction and report requires 3.6 KB of storage; this report size is more than enough to store timestamps, position and traffic condition data in our example. This would scale as 179.2 MB, 1 GB, and 3.6 GB to store 50 000, 279 000, and 1 million validated tokens and reports, respectively.

CPU performance: we measured the performance of a Python implementation of Chaum's digital cash protocol [37] using random contributor identifiers and token values, resulting in 1 024 bit tokens. The timing data were obtained

using a laptop with an Intel Core i7 (I7-9750H) running Python 3.9.2 and macOS 10.15.7. After running warm up operations, the timing results for 100 000 operations are summarised by their mean and standard deviation, since they present normal distributions. The only exception is the signing operation by the identity authority, which presents a binormal distribution, possibly due to occasionally hitting an uncommon (uncached) code path. However, the frequency of the second peak is negligible and of the order of ten times lower than the most frequent values in the histogram. The creation of the token candidates by contributors takes a mean of 30.3 ms with a standard deviation of 4.2 ms (30.3 ± 4.2 ms), and payment takes 183.5 ± 35.1 ms. The identity authority signs each token in 20.2 ± 1.8 ms, verifying each in 92.6 ± 15.8 ms. The verification and its timings depend heavily on the size of the database in the implementation tested. A distributed system or simply reordering tokens could optimise the results.

Power consumption: we evaluate the feasibility of our approach on mobile devices by measuring the power consumption on a smartphone. We place a power sensor between the battery and the smartphone, similar to previous approaches [35, 43]. For the smartphone we use a Motorola Moto E6 Plus (XT2025-2) as it comes with a removable battery. We measure an upper bound of the total power consumed, obtaining readings at a rate of 2 KHz [48]. The device runs on battery power with display and Wi-Fi off during the measurements. Each operation is run 100 times in order to time and measure their power consumption.

Table 1. Android operation timings and power consumption

	Mean	Std Deviation	Power Consumption	Std Deviation
SHA-256	1.54 ms	0.21 ms	0.40 mJ	0.05 mJ
SHA-512	1.54 ms	0.21 ms	0.40 mJ	0.06 mJ
Blinding step	16.31 ms	3.20 ms	3.44 mJ	0.78 mJ
Verifying step	33.19 ms	2.41 ms	6.49 mJ	4.52 mJ

These measurements show hashing operations are very fast and have negligible impact on power consumption (<0.5 mJ). The blinding and verifying steps require more computational effort but little power: 3.44 mJ and 6.49 mJ. Furthermore, token generation could be run while connected to power overnight. These numbers are very small: previous research estimates 4G/LTE network transmission to consume 12.8 J to send a single packet [33]. Thus, the network cost necessary for all crowdsourcing applications is going to dominate over any energy consumption associated with the CPU requirements of the protocol.

5.1 *K*-anonymity

We used Veins, a vehicular network simulation software which combines traffic and network communication simulations, to perform a brief feasibility assessment and estimate realistic values of k in the domain of urban traffic congestion [46].

Our simulations are based on the real road map from the Aarhus (Denmark) city centre, obtained from OpenStreetMap. We simulate traffic volumes by using data from the CityPulse Aarhus traffic dataset [13], together with speed limits, junctions and traffic light simulations in Veins. We assume vehicles can communicate using a peer-to-peer protocol with a broadcasting range of 300 m as proposed by the U.S. Department of Transportation [40].

We measure the number of vehicles that are in range of each other for 5 seconds, enough time to exchange messages and reports. We run two simulations to cover volumes of traffic representative of peak (5 000 vehicles) and off-peak (2 000 vehicles) periods. The results were plotted and form long-tailed distributions in both scenarios, indicating there are areas where only low values of k are possible. The values of k obtained for indicative percentiles are summarised in the Table 2.

Table 2. Values of k for indicative percentiles

Percentiles	0 – 40	50	60	65	70	75	80	85	90	95
k for 2 000 vehicles	1	3	5	9	11	14	22	30	38	47
k for 5 000 vehicles	1	2	4	5	6	8	25	50	80	112

With 2 000 vehicles we find that at any time, one in five vehicles has 21 or more vehicles nearby, and one in ten has 37 or more. When simulating 5 000 vehicles our results show one in five vehicles has 24 or more nearby, and one in ten has 79 or more at any time.

In other words, half of all vehicles at any point in time have no anonymity ($k = 1$) and are driving around with no other vehicles within 300 m. Conversely, 20% of the time, vehicles are able to communicate with over 20 other vehicles within 300 m. This is likely to be acceptable for many crowdsourcing applications as occasional reports are all that is required. For example, in the domain of traffic monitoring, sporadic reports are useful, and in particular, reports will be produce in all of the busy areas since these will have high values for k.

6 Related work

Previous research has proposed addition of noise and homomorphic encryption for privacy-preserving measurement aggregation by untrusted entities [23, 38, 45], and identity-based encryption to allow anonymous crowdsourcing [14, 15]. However, these assume arbitrary, unauthenticated contributors, making it impossible to identify Sybil attacks and penalise dishonest contributors.

Pseudonym-based approaches to location privacy include mix zones [6, 7], silent period [34], and some blind signature schemes [1]. However, all of these rely on middleware systems, or have no ability to limit user contributions, thus being vulnerable to Sybil attacks. The SPPEAR approach [29] uses pseudonyms and authorisation tokens, but cannot guarantee contributor privacy when some

of its entities collude. It is also very inefficient compared to our approach, requiring more than 20 s to issue 50 pseudonyms, and results in significant CPU use in mobile devices. Blind signature schemes have been proposed for anonymous authentication of vehicles, but require extra trusted infrastructure to be deployed [3, 50]. Another approach is k-anonymity, which ensures that any particular record is indistinguishable from at least $(k-1)$ other records [44]. Spatial and temporal cloaking techniques proposed for k-anonymity assume trusted service providers [31], and are vulnerable to subtraction and negation attacks [5].

Differential privacy enables the publication of aggregate statistics of a dataset via the addition of noise by the aggregator [21, 22]. However, it is unsuitable for small amounts of data and relies on a trusted third-party data collector to apply the differential privacy algorithms. Local differential privacy approaches like RAPPOR ensure responses sent to aggregators are already differentially private [25], but lose detail quickly and degrade their privacy protection much faster than the central model with periodic reports [24, 36, 47]. Another approach is shuffle model differential privacy, with results similar to the central model for secure aggregations [2, 12, 27, 28] but combines limitations of central and local approaches: degraded data quality and granularity and requiring trusted entities.

7 Conclusion

In this paper we improved on the privacy of data contributors in a crowdsourcing system by devising a protocol which separates contributor authentication (required to prevent Sybil attacks) from the individual reports submitted to apps. Since crowdsourcing apps operating in the real-world typically gather location data as part of any report, we proposed a protocol to protect the location privacy of all contributors. Our approach does not need a trusted app, trusted middleware systems nor additional infrastructure. Instead, we rely on an identity authority to oversee contributor identification and registration, adapting Chaum's digital cash scheme to blindly sign and limit the number of reports a contributor can submit over a defined period of time.

The protocols have modest implementation requirements: one million validated reports and tokens fit within 3.6 GB; computation time typically takes less than 34 ms; and power consumption in handheld devices is very modest at 6.5 mJ. The practical feasibility of the k-anonymous protocol implementation was shown via simulation of vehicles driving around in a medium-sized European city: we found that 20% of the vehicles could submit reports at any time with a k between 22 and 25 for typical traffic volumes during peak and off-peak times.

Acknowledgements. We thank Daniel Hugenroth for his help in obtaining mobile device power consumption measurements. We are grateful to anonymous reviewers for their feedback. We acknowledge and thank Fundación Mapfre Guanarteme and Nokia Bell Labs for their generous financial support. The views, opinions and findings in this paper are those of the authors and not necessarily those of our funders.

References

1. Aslam, B., Zou, C.C.: One-way-linkable blind signature security architecture for vanet. In: 2011 IEEE Consumer Communications and Networking Conference (CCNC). pp. 745–750 (2011). https://doi.org/10.1109/CCNC.2011.5766590
2. Balle, B., Bell, J., Gascón, A., Nissim, K.: Private summation in the multi-message shuffle model. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security. p. 657–676. CCS '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3372297.3417242
3. Baza, M., Nabil, M., Mahmoud, M.M.E.A., Bewermeier, N., Fidan, K., Alasmary, W., Abdallah, M.: Detecting sybil attacks using proofs of work and location in vanets. IEEE Transactions on Dependable and Secure Computing **19**(1), 39–53 (2022). https://doi.org/10.1109/TDSC.2020.2993769
4. Bellare, M., Namprempre, C., Pointcheval, D., Semanko, M.: The one-more-rsa-inversion problems and the security of chaum's blind signature scheme. Journal of Cryptology **16**(3) (2003). https://doi.org/10.1007/s00145-002-0120-1
5. Beresford, A.R.: Location Privacy in Ubiquitous Computing. Ph.D. thesis, University of Cambridge (2004)
6. Beresford, A., Stajano, F.: Location privacy in pervasive computing. IEEE Pervasive Computing **2**(1), 46–55 (2003). https://doi.org/10.1109/MPRV.2003.1186725
7. Beresford, A., Stajano, F.: Mix zones: user privacy in location-aware services. In: IEEE Annual Conference on Pervasive Computing and Communications Workshops, 2004. Proceedings of the Second. pp. 127–131 (2004). https://doi.org/10.1109/PERCOMW.2004.1276918
8. Chaum, D.: Blind signatures for untraceable payments. In: Chaum, D., Rivest, R.L., Sherman, A.T. (eds.) Advances in Cryptology. pp. 199–203. Springer US, Boston, MA (1983). https://doi.org/10.1007/978-1-4757-0602-4_18
9. Chaum, D.: Online cash checks. In: Quisquater, J.J., Vandewalle, J. (eds.) Advances in Cryptology — EUROCRYPT '89. pp. 288–293. Springer Berlin Heidelberg, Berlin, Heidelberg (1990). https://doi.org/10.1007/3-540-46885-4_30
10. Chaum, D., Fiat, A., Naor, M.: Untraceable electronic cash. In: Goldwasser, S. (ed.) Advances in Cryptology — CRYPTO' 88. pp. 319–327. Springer New York, New York, NY (1990). https://doi.org/10.1007/0-387-34799-2_25
11. Chen, M., Yang, J., Hu, L., Hossain, M.S., Muhammad, G.: Urban healthcare big data system based on crowdsourced and cloud-based air quality indicators. IEEE Communications Magazine **56**(11), 14–20 (2018). https://doi.org/10.1109/MCOM.2018.1700571
12. Cheu, A., Smith, A., Ullman, J., Zeber, D., Zhilyaev, M.: Distributed differential privacy via shuffling. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2019. pp. 375–403. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_13
13. CityPulse EU FP7 project: Smart City Datasets (2016), http://iot.ee.surrey.ac.uk:8080/datasets.html
14. De Cristofaro, E., Soriente, C.: Short paper: Pepsi—privacy-enhanced participatory sensing infrastructure. In: Proceedings of the Fourth ACM Conference on Wireless Network Security. p. 23–28. WiSec '11, Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/1998412.1998418
15. De Cristofaro, E., Soriente, C.: Extended capabilities for a privacy-enhanced participatory sensing infrastructure (pepsi). IEEE Transactions on Information Forensics and Security **8**(12), 2021–2033 (2013). https://doi.org/10.1109/TIFS.2013.2287092

16. De Montjoye, Y.A., Hidalgo, C.A., Verleysen, M., Blondel, V.D.: Unique in the crowd: The privacy bounds of human mobility. Scientific reports **3**(1), 1–5 (2013). https://doi.org/10.1038/srep01376
17. Diaz, Claudia and Halpin, Harry and Kiayias, Aggelos: The Nym Network (2021), https://nymtech.net/whitepaper, [Online]
18. Dingledine, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. Tech. rep., Naval Research Lab Washington DC (2004)
19. Douceur, J.R.: The sybil attack. In: International workshop on peer-to-peer systems. pp. 251–260. Springer (2002). https://doi.org/10.1007/3-540-45748-8_24
20. Dreier, J., Kassem, A., Lafourcade, P.: Formal analysis of e-cash protocols. In: 2015 12th International Joint Conference on e-Business and Telecommunications (ICETE). vol. 04, pp. 65–75 (2015)
21. Dwork, C.: Differential privacy: A survey of results. In: International conference on theory and applications of models of computation. pp. 1–19. Springer (2008). https://doi.org/10.1007/978-3-540-79228-4_1
22. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science **9**(3–4), 211–407 (2014). https://doi.org/10.1561/0400000042
23. Erkin, Z., Troncoso-pastoriza, J.R., Lagendijk, R., Perez-Gonzalez, F.: Privacy-preserving data aggregation in smart metering systems: an overview. IEEE Signal Processing Magazine **30**(2), 75–86 (2013). https://doi.org/10.1109/MSP.2012.2228343
24. Erlingsson, Ú., Feldman, V., Mironov, I., Raghunathan, A., Talwar, K., Thakurta, A.: Amplification by shuffling: From local to central differential privacy via anonymity. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2468–2479. SIAM (2019). https://doi.org/10.1137/1.9781611975482.151
25. Erlingsson, U., Pihur, V., Korolova, A.: Rappor: Randomized aggregatable privacy-preserving ordinal response. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. p. 1054–1067. CCS '14, Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2660267.2660348
26. Gambs, S., Killijian, M.O., Núñez del Prado Cortez, M.: De-anonymization attack on geolocated data. Journal of Computer and System Sciences **80**(8), 1597–1614 (2014). https://doi.org/https://doi.org/10.1016/j.jcss.2014.04.024, special Issue on Theory and Applications in Parallel and Distributed Computing Systems
27. Ghazi, B., Kumar, R., Manurangsi, P., Pagh, R., Sinha, A.: Differentially private aggregation in the shuffle model: Almost central accuracy in almost a single message. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 3692–3701. PMLR (18–24 Jul 2021), https://proceedings.mlr.press/v139/ghazi21a.html
28. Ghazi, B., Manurangsi, P., Pagh, R., Velingker, A.: Private aggregation from fewer anonymous messages. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology – EUROCRYPT 2020. pp. 798–827. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_27
29. Gisdakis, S., Giannetsos, T., Papadimitratos, P.: Sppear: Security & privacy-preserving architecture for participatory-sensing applications. In: Proceedings of the 2014 ACM Conference on Security and Privacy in Wireless & Mobile Networks. p. 39–50. WiSec '14, Association for Computing Machinery, New York, NY, USA (2014). https://doi.org/10.1145/2627393.2627402

30. Golle, P., Partridge, K.: On the anonymity of home/work location pairs. In: Tokuda, H., Beigl, M., Friday, A., Brush, A.J.B., Tobe, Y. (eds.) Pervasive Computing. pp. 390–397. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01516-8_26
31. Gruteser, M., Grunwald, D.: Anonymous usage of location-based services through spatial and temporal cloaking. In: Proceedings of the 1st International Conference on Mobile Systems, Applications and Services. p. 31–42. MobiSys '03, Association for Computing Machinery, New York, NY, USA (2003). https://doi.org/10.1145/1066116.1189037
32. Hern, A.: Berlin artist uses 99 phones to trick Google into traffic jam alert. The Guardian (2020), https://www.theguardian.com/technology/2020/feb/03/berlin-artist-uses-99-phones-trick-google-maps-traffic-jam-alert
33. Huang, J., Qian, F., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O.: A close examination of performance and power characteristics of 4G LTE networks. In: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services. p. 225–238. MobiSys '12, Association for Computing Machinery (2012). https://doi.org/10.1145/2307636.2307658
34. Huang, L., Matsuura, K., Yamane, H., Sezaki, K.: Enhancing wireless location privacy using silent period. In: IEEE Wireless Communications and Networking Conference, 2005. vol. 2, pp. 1187–1192 Vol. 2 (2005). https://doi.org/10.1109/WCNC.2005.1424677
35. Hugenroth, D.: Measuring energy consumption of privacy-preserving protocols for fun and profit (2021), https://mobiuk.org/2021/abstract/S5-P2_Hugenroth_MeasuringEnergyConsumptionPrivacyPreserving.pdf, mobiUK 2021
36. Joseph, M., Roth, A., Ullman, J., Waggoner, B.: Local differential privacy for evolving data. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 31. Curran Associates, Inc. (2018)
37. Kamthe, N.S., Kamala, K.K., Husaini Basha, A.S., Sharma, P.: Project_DigitalCash (06 2020), https://github.com/koushik-kumar/Project_DigitalCash
38. Liu, Y., Guo, W., Fan, C.I., Chang, L., Cheng, C.: A practical privacy-preserving data aggregation (3pda) scheme for smart grid. IEEE Transactions on Industrial Informatics **15**(3), 1767–1774 (2019). https://doi.org/10.1109/TII.2018.2809672
39. Meier, F., Fenner, D., Grassmann, T., Otto, M., Scherer, D.: Crowdsourcing air temperature from citizen weather stations for urban climate research. Urban Climate **19**, 170–191 (2017). https://doi.org/https://doi.org/10.1016/j.uclim.2017.01.006
40. National Highway Traffic Safety Administration: NHTSA: Vehicle-to-Vehicle Communication (2021), https://www.nhtsa.gov/technology-innovation/vehicle-vehicle-communication
41. Overeem, A., R. Robinson, J.C., Leijnse, H., Steeneveld, G.J., P. Horn, B.K., Uijlenhoet, R.: Crowdsourcing urban air temperatures from smartphone battery temperatures. Geophysical Research Letters **40**(15), 4081–4085 (2013). https://doi.org/https://doi.org/10.1002/grl.50786
42. Piotrowska, A.M., Hayes, J., Elahi, T., Meiser, S., Danezis, G.: The loopix anonymity system. In: 26th USENIX Security Symposium (USENIX Security 17). pp. 1199–1216. USENIX Association, Vancouver, BC (Aug 2017), https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/piotrowska

43. Rice, A., Hay, S.: Decomposing power measurements for mobile devices. In: 2010 IEEE International Conference on Pervasive Computing and Communications (PerCom). pp. 70–78 (2010). https://doi.org/10.1109/PERCOM.2010.5466991
44. Samarati, P., Sweeney, L.: Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Tech. rep., SRI International (1998)
45. Shi, E., Chan, T.H., Rieffel, E., Chow, R., Song, D.: Privacy-preserving aggregation of time-series data. In: Proc. NDSS. vol. 2, pp. 1–17. Citeseer (2011)
46. Sommer, C., German, R., Dressler, F.: Bidirectionally coupled network and road traffic simulation for improved ivc analysis. IEEE Transactions on Mobile Computing **10**(1), 3–15 (2011). https://doi.org/10.1109/TMC.2010.133
47. Tang, J., Korolova, A., Bai, X., Wang, X., Wang, X.: Privacy loss in Apple's implementation of differential privacy on macOS 10.12. arXiv preprint arXiv:1709.02753 (2017)
48. Texas Instruments: INA219 data sheet, product information and support (2021), https://www.ti.com/product/INA219, [Online]
49. Wang, G., Wang, B., Wang, T., Nika, A., Zheng, H., Zhao, B.Y.: Ghost riders: Sybil attacks on crowdsourced mobile mapping services. IEEE/ACM Transactions on Networking **26**(3), 1123–1136 (2018). https://doi.org/10.1109/TNET.2018.2818073
50. Wang, X., Jiang, J., Zhao, S., Bai, L.: A fair blind signature scheme to revoke malicious vehicles in vanets. Computers, Materials & Continua **58**(1), 249–262 (2019). https://doi.org/10.32604/cmc.2019.04088
51. Zang, H., Bolot, J.: Anonymization of location data does not work: A large-scale measurement study. In: Proceedings of the 17th annual international conference on Mobile computing and networking. pp. 145–156 (2011). https://doi.org/10.1145/2030613.2030630