

Investigating Anonymity in Group Based Anonymous Authentication

Daniel Slamanig^{1,2} and Christian Stingl¹

¹ Carinthia University of Applied Sciences,
Medical Information Technology – Healthcare IT & Information Security Group,
9020 Klagenfurt, Austria

² University of Klagenfurt, Computer Science – System Security Group,
9020 Klagenfurt, Austria
{d.slamanig,c.stingl}@cuas.at

Abstract. In this paper we discuss anonymity in context of group based anonymous authentication (*GBAA*). Methods for *GBAA* provide mechanisms such that a user is able to prove membership in a group $\mathcal{U}' \subseteq \mathcal{U}$ of authorized users \mathcal{U} to a verifier, whereas the verifier does not obtain any information on the actual identity of the authenticating user. They can be used in addition to anonymous communication channels in order to enhance user's privacy if access to services is limited to authorized users, e.g. subscription-based services. We especially focus on attacks against the anonymity of authenticating users which can be mounted by an external adversary or a passive verifier when *GBAA* is treated as a black box. In particular, we investigate what an adversary can learn by solely observing anonymity sets \mathcal{U}' used for *GBAA* and how users can choose their anonymity sets in case of $\mathcal{U}' \subset \mathcal{U}$. Based on the information which can be obtained by adversaries we show that the probability of user identification can be improved.

1 Introduction

The Internet is nowadays used by a permanently increasing number of people. Their actions comprise on the one hand private activities, e.g. using it as a source of information, doing online banking, communicating with other persons, reading their newspapers, participating in electronic auctions. On the other hand people use it for business related activities. Obviously, the collection of information divulged and exchanged during these activities may represent an extensive picture of a person and covers many topics related to ones privacy. For example, these information may be highly valuable for providers hosting online services when analyzing user's behavior [20,35]. In this context there are tools available for free, e.g. Google Analytics, which provide a huge set of functionalities for aforementioned purposes, even for unaware and casual users. Nevertheless, users may also benefit from these methods by means of Web personalization, i.e. the customization of delivered Web content with respect to the user's preferences. However, privacy issues are very often neglected, which questions the before

discussed advantages. This can be illustrated by a user who queries a health information service to obtain information on a serious disease. Recent studies show that 80 percent of health searchers seek the information for themselves and 60 to 80 percent of Americans have already used the Internet to find health information [32]. Hence, if any other party is able to link these information to the user, then it may be possible to draw compromising conclusions.

The aforementioned threats are in our opinion highly realistic, since protocols used in Internet communication do not explicitly provide mechanisms to preserve the anonymity of users. Additionally, we are confronted with a phenomenon denoted as privacy myopia [19]. This means, that people often are not aware of dangers related to privacy and sell or give away their data without reflecting on potential negative consequences. For instance, in context of the Internet this means that users reveal IP-addresses which enable third parties to link several actions and may enable third parties to identify the physical users behind their computers. Furthermore, users often easily give away person related information to third parties which exceeds the amount of information necessary. The latter aspect is the subject of privacy enhanced identity management and has experience major research interest in recent years (cf. [5]).

In this paper we will discuss anonymity aspects related to group based anonymous authentication (*GBAA*), which provides anonymity for users if access to services is limited to an authorized set of users. If a user needs to authenticate to a service provider by means of traditional authentication mechanisms, in general the identity of the user is known by the service provider. By means of *GBAA*, the server solely learns the membership of the authenticating user in the set of authorized users, but does not learn the exact identity. This can be valuable for users, if the sole knowledge of service accesses, i.e. the frequency of access, may lead to compromising conclusions. The applications we have in mind for *GBAA* are any kind of Internet services that require user authentication, but users want to hide their behavior from the service provider. Thereby, the main advantage of *GBAA* schemes is that they can be build upon existing and widely deployed public key infrastructures based on X.509 certificates (PKIX) and user registration for services solely requires the user to provide a valid X.509 certificate to the service provider. This results on the one hand in higher security compared to widely used username/password authentication schemes and on the other hand in a privacy improvement for the user. For instance, consider a Internet service which provides access to an electronic health record (EHR) of a person, whereas the EHR represents a life-long documentation of the medical history of a person. In this context, even the knowledge of the frequency of access to the EHR may enable a third party to draw compromising conclusions about the state of health of the person. Additionally, the use of *GBAA* schemes, which are based on public key certificates, prevents users from identity theft by means of password guessing, dictionary attacks or other threats.

The remainder of this paper is organized as follows: In section 2 we discuss aspects of anonymity that are important for Internet based services. In the subsequent section 3 we will briefly introduce *GBAA*, attack models and scenarios

as well as some problems related to *GBAA*. Section 4 discusses the choice of anonymity sets used for authentication and provides a detailed analysis. Finally, section 5 concludes the paper and discusses some future aspects.

2 Different Aspects of Anonymity

Anonymity aspects of users in the context of Internet services are twofold. Firstly, the anonymity of a user may be revealed by the communication channel itself. Consequently, users need to hide their identity when sending messages over the communication channel. This can be achieved by means of anonymous communication channels. Secondly, identities of users may be revealed at higher network layers, i.e. the application layer. This is especially of interest if services require user authentication at the application layer. Subsequently, we will briefly discuss the aforementioned aspects.

2.1 Communication Anonymity

Mechanisms that provide anonymity and unlinkability of messages sent over a communication channel are denoted as anonymous communication techniques and have been intensively studied in recent years, see [12] for a sound overview. There are several implementations available for low-latency services like Web browsing, e.g. Tor [15], JAP [18], as well as high-latency services like E-Mail, e.g. Mixminion [13].

These anonymous communication channels help to improve the privacy of users in context of eavesdroppers and curious communication partners. Especially, regarding the latter one anonymity can be preserved if electronic interaction does not rely on additional identifying information at higher network layers, i.e. the application layer. For example, a user who queries a public web page using an anonymous communication channel may remove all identifying information from higher network layers and thus will stay anonymous.

In our considerations we assume that we have a communication channel that guarantees perfect anonymity and unlinkability. Then a user is connected to a service provider (server) via a kind of *“magic channel”* that leaks no information on the identity of the user at the communication layer. Clearly, this is a somewhat idealized consideration, since real world anonymous communication channels do not realize perfect anonymity resp. unlinkability (cf. [30]) and there may exist additional side channels, e.g. online-behavior of users, which can be used to improve the probability of identification of communicating parties.

2.2 Anonymity at Higher Layers

However, if service providers offer their services only to authorized sets of users, e.g. subscription-based services, closed communities, they require identification of users which in general takes place at higher layers by means of entity authentication mechanisms. In entity authentication or identification protocols the holder of an identity usually claims a set of attributes including an identifier

and interactively proves the possession of the claimed identity to a verifier. This identifier is usually unique within a specific context, e.g. application, but may be a pseudonym, which is not linkable to the physical identity of a person. But there usually exists a party which is aware of this link and additionally, actions conducted under the same pseudonym can be linked. Nevertheless, there exists anonymous credential systems which can be used to anonymously prove the possession of attributes of credentials while preserving unlinkability of different showings of a credential and anonymity of the holders (cf. [4,7,8,24]). These approaches are especially suitable in a multi-provider setting, where users obtain credentials for a pseudonym from one provider and are able to show these credentials under different pseudonyms to other providers. Nevertheless, there are also known attacks (cf. [21,27]) against unlinkability and anonymity of anonymous credential systems when using them in a real world context. We do not consider the aforementioned approaches, since we are interested in a single-provider and “ad hoc” setting. However, the aforementioned mechanisms can also be used to realize anonymous authentication (cf. [6]), but in general they do not provide “ad hoc” mechanisms as discussed below, i.e. they rely on a proprietary setup with every user. Therefore, we will subsequently discuss an alternative approach based on cryptographic primitives like ring signatures [31], which we call group based anonymous authentication (\mathcal{GBAA}), that provides mechanisms to perform anonymous authentications based on “ad hoc” groups, i.e. without relying on interaction with other group members and without any additional proprietary setup.

3 Group Based Anonymous Authentication

Group based anonymous authentication (\mathcal{GBAA}) aims to provide a somewhat paradoxical solution to enhance user’s privacy in context of authentication. It provides mechanisms such that a user is able to prove membership in a group $\mathcal{U}' \subseteq \mathcal{U}$ of authorized users \mathcal{U} , whereas the verifier does not obtain information on the identity of the authenticating user. The set \mathcal{U}' will also be denoted as the anonymity set [29]. Clearly, anonymous communication systems are a prerequisite for providing anonymity in the context of anonymous authentication.

A naive approach to realize \mathcal{GBAA} would be to give a copy of a secret k to every user $u \in \mathcal{U}$, which could be used in conjunction with a traditional authentication scheme. Obviously, the revocation of a single user u_i would result in a reinitialization and thus in reissuing a fresh secret k' to every remaining user $u \in \mathcal{U} \setminus u_i$. Hence, this approach is far from being practical. Improved techniques for \mathcal{GBAA} were explicitly treated in [3,23,28,33,37] and additionally with special properties like being anonymous as long as the number of authentication is beyond a threshold [36], the ability to detect fraudulent users [6,11] and with the ability to revoke the anonymity of users [3,22]. They can be realized by means of group signatures [1,9, etc.], witness indistinguishable signatures [10], ring signatures [16,31, etc.] or similar concepts as (deniable) ring authentication [26].

The latter two classes of signature and authentication schemes are preferable to group signatures in the context of large and dynamic groups, as it is the case

with Internet services, since they can be generated “ad hoc” without depending on an explicit setup phase or reinitialization in case of dynamic groups. Thereby “ad hoc” means that an authenticating user does not need the knowledge, consent or assistance of the remaining members of an ad hoc group to perform an authentication. Furthermore, in general they do not require a proprietary setup and do only rely on standard public key certificates, i.e. X.509 certificates, which are widely deployed and available. It must be mentioned that there are already attempts to integrate group, ring and traceable signatures, which can be used for \mathcal{GBAA} , into the PKIX framework [2].

There are three important properties that \mathcal{GBAA} mechanisms need to provide (cf. [23,33]):

1. **Anonymity:** The verifier is not able to determine the identity of an authenticating user with probability higher than $1/|\mathcal{U}'|$.
2. **Unlinkability:** It is impossible to link k , $k > 1$, instances of the \mathcal{GBAA} protocol of one (anonymous) user $u_i \in \mathcal{U}'$.
3. **Security:** Only authorized users $u \in \mathcal{U}$ should be able to pass the \mathcal{GBAA} .

The properties we are focusing on in this paper are anonymity and unlinkability, and in particular we investigate strategies to construct groups used for \mathcal{GBAA} . This is especially of interest in context of large groups, since the computational effort in \mathcal{GBAA} protocols usually grows (linearly) with the size of the anonymity set, i.e. the cardinality of \mathcal{U}' . Thus, a large set of authorized users may force a user to prove his membership using a subset of all authorized users for efficiency purposes. It must be mentioned that we do not explicitly discuss technical details on the construction of methods for \mathcal{GBAA} and will treat them as a black box in the remainder of this paper.

The question that comes up is, whether a verifier or even an observer is able to reduce the anonymity and consequently unlinkability by continuously observing anonymity sets, although the underlying \mathcal{GBAA} method and communication channel provides perfect anonymity and unlinkability.

3.1 Attacker Model

As mentioned above, we are not considering anonymity provided by the \mathcal{GBAA} itself and the communication channel. Consequently, we assume that the \mathcal{GBAA} methods provide perfect anonymity, unlinkability and security and the communication channel provides perfect anonymity and unlinkability (“magic channel”). Clearly, these assumptions are very strong with respect to the real world and thus the results presented in this paper, i.e. the reduction of anonymity of users, may even be improved enormously by substituting the perfect \mathcal{GBAA} and communication channel by actually deployed methods.

The attack model used in this paper considers the following adversaries.

- **Honest but curious (passive) verifier:** An insider who is able to monitor all actions inside the verifier’s system, but does not actively manipulate messages which are exchanged during the \mathcal{GBAA} .

- **Eavesdropper:** Anyone who is able to monitor the inbound traffic of the verifier. As above, the eavesdropper solely behaves passive, i.e. does not manipulate exchanged messages.

Passive attacks conducted by an eavesdropper can easily be prevented by means of encrypted communication, i.e. a communication channel which provides confidentiality and integrity of transmitted messages. However, it must be mentioned that an external adversary may run a denial of service (DoS) attack against the verifier’s system in order to deter authentications of users anyway. We do not consider active attackers, since there exist measures incorporated into *GBAA* protocols to detect a cheating verifier (cf. [23]), which are outside the scope of this paper. Furthermore, in practice an actively cheating verifier may leak out some day and will consequently not be trustworthy anymore.

An adversary may mount the subsequent attacks, whereas we focus on the first one in this paper and the latter one will only be stated for the sake of completeness.

- **Anonymity sets only:** An adversary is clearly able to record all information which are shown to him during any instance of a *GBAA*. Thus he can count the occurrences of users in anonymity sets. The adversary will try to reduce the anonymity of single users solely by means of the aforementioned information.
- **Behavioral heuristic:** Since unlinkability is a required property, every action inside the system requires a single *GBAA* protocol. Thus, authentications of a single user are likely to occur cumulative since in general at least a few operations are conducted within the verifier’s system.

3.2 Some Problems Related to *GBAA*

One inherent problem in “ad hoc” *GBAA* is, that the physical person which holds a digital identity, irrespective of the representation, e.g. X.509 certificates, is not directly known to a user. Consequently, a user may not be able to distinguish “real” from “fake” identities. Especially in large groups, a verifier may be able to forge identities of “authorized users” which look valid to all other users. This is crucial, if verifiers set up their system (parameters) autonomous, i.e. issue credentials or certificates on their own, and do not involve a commonly trusted party, e.g. a trusted certification authority which issues public key certificates. It must be stressed, that this attack is an active one which can be conducted by malicious verifiers to reduce the anonymity of users. However, we assume that the verifier is honest but curious in our attack model. This fake-user insertion attack can be somewhat compared to the sybil attack [17], which has been investigated in a somewhat similar context [25]. But, in our case the verifier creates a set of forged identities on his own and integrates them into the system. Consequently, the effective anonymity set for *GBAA* will be reduced by users unawarely including fake identities in their anonymity sets.

Another problem in this context is that the authenticating user needs to be sure, that all actually chosen users are indeed authorized users at the point

of time of authentication. We want to emphasize, that the task of determining authorized users, i.e. to check if a user is authorized and the respective certificate is valid, is a time consuming and non-trivial task, but is inherent to all certificate based \mathcal{GBAA} protocols. However, we will not consider this problem in detail in this paper since it does not affect our investigations.

4 Analysis

In this section we firstly analyze strategies to construct anonymity sets and secondly propose methods that can be used by an adversary to improve the probability of identification of users.

4.1 Group Construction Strategies

In the following we will discuss the strategies to construct anonymity sets for \mathcal{GBAA} . Thereby, we consider the two possible scenarios, i.e. on the one hand the entire group and on the other hand a subgroup of authorized users.

Entire Group. If a user chooses the entire group \mathcal{U} for \mathcal{GBAA} , the probability of user u_i being the one who actually authenticates in any anonymity set is $p_{u_i} = 1/|\mathcal{U}|$. Hence, this approach guarantees perfect anonymity [14,34]. This strategy is immune against fake-user insertion attacks, since always all users are chosen. However, it must be emphasized that for actual available protocols for \mathcal{GBAA} the computational effort grows at least linearly with the size of the anonymity set. Hence, in case of a large set of authorized users, this approach is impractical.




Fig. 1. Static subgroup approach from the point of view of the prover

Subgroup. This alternative approach is characterized by choosing a subset $\mathcal{U}' \subset \mathcal{U}$ for \mathcal{GBAA} , whereas we assume that $|\mathcal{U}'| \ll |\mathcal{U}|$. Therefore, users need to construct subgroups following some specific strategy. The obvious method for a user u_i to construct an anonymity set of size k is, to independently choose $k - 1$ users uniformly at random from \mathcal{U} and to subsequently integrate himself into the anonymity set. This approach is prone to a fake-user insertion attack, since

the verifier may include faked “authorized” users into the set of all authorized users. Consequently, the level of security depends on the fraction of “fake” users.

Considering the subgroup-approach we distinguish between static subgroups and dynamic subgroups.

Static Subgroup. In case of static subgroups, a user $u_i \in \mathcal{U}$ initially chooses $k - 1$ users uniformly at random from \mathcal{U} and forms his static anonymity set by adding himself to this set. Subsequently, he uses his initial chosen anonymity set for every \mathcal{GBAA} . If \mathcal{U} is large, e.g. $|\mathcal{U}| = 200$, and the size of the anonymity set is smaller than the size of \mathcal{U} ($|\mathcal{U}'| \approx 100$), it is very unlikely that two distinct users choose exactly the same anonymity set, i.e. $\approx 1/\binom{|\mathcal{U}|}{|\mathcal{U}'|}$. Hence, if a user applies this strategy all \mathcal{GBAA} s are in general linkable. Additional, with side channel information, e.g. user’s behavior, it may be easier to identify a single user. As a consequence we want to point out that this approach is in our opinion not appropriate.

Dynamic Subgroup. In case of dynamic subgroups a user $u_i \in \mathcal{U}$ constructs his anonymity set \mathcal{U}' independently for every single authentication. Thus unlinkability is guaranteed and with respect to the above strategies in our opinion it is the preferred strategy in context of large sets of authorized users. Nevertheless,




Fig. 2. Dynamic subgroup approach from the point of view of the prover

we will subsequently examine potential weaknesses of this dynamic subgroup approach which can be used by an adversary to improve the probability of identifying authenticating users.

4.2 Anonymity Sets Only Attack

As mentioned above, we are now focusing on anonymity sets independent of the protocol used for \mathcal{GBAA} . Furthermore, we assume that these information can be monitored by an adversary, e.g. the verifier or an eavesdropper. In particular, we introduce methods to analyze the anonymity sets and derive measures to improve attacks against anonymity. In order to compute these measures a $|\mathcal{U}| \times N$ history




Fig. 3. The setting for this example is: $|\mathcal{U}| = 1000$, $N = 10^8$. Subfigure (a) illustrates the uniform distribution of the passive frequencies. In this example it was assumed that the active frequencies are Gaussian distributed (b). Subfigure (c) shows that the distribution of the active frequencies is still reflected in the global frequency. Furthermore, it can be conjectured that the value of the global frequency is directly “connected” to the value of the active frequency and vice versa.

matrix \mathcal{H} will be used, where \mathcal{U}'_j , $1 \leq j \leq N$, is the j -th anonymity set and N is the overall number of \mathcal{GBAA} protocol runs.

$$\mathcal{H}(i, j) = \begin{cases} 1, & \text{if } u_i \in \mathcal{U}'_j, \\ 0, & \text{else.} \end{cases}$$

Put differently, the matrix represents the collection of all anonymity sets which were used in \mathcal{GBAAs} and the element $\mathcal{H}(i, j)$ contains the value 1 if and only if user u_i occurred in the respective anonymity set \mathcal{U}'_j . Based on this matrix, we are defining the global frequency $\nu^G_{u_i}$ of user u_i which is the sum of the i -th row. The global frequency of a user u_i itself consists of an active part $\nu^G_{u_i, \mathcal{A}}$, i.e. the number of actual authentications of the user, and a passive part $\nu^G_{u_i, \mathcal{P}}$, i.e. other users choose u_i in their anonymity sets. Obviously, an adversary can solely determine the sum $\nu^G_{u_i} = \nu^G_{u_i, \mathcal{A}} + \nu^G_{u_i, \mathcal{P}}$ of the users frequency from the history matrix. The two subsequent facts can easily be obtained.

$$\sum_{i=1}^{|\mathcal{U}|} \nu^G_{u_i, \mathcal{A}} = N \tag{1}$$


$$\sum_{i=1}^{|\mathcal{U}|} \nu_{u_i, \mathcal{P}}^G = \sum_{i=1}^N (|\mathcal{U}'_i| - 1) \tag{2}$$

Considering the above mentioned method to create subgroups one can conclude that the passive frequencies are uniformly distributed and the average of all passive frequencies is $\bar{\nu}_{\mathcal{P}} = \sum_{i=1}^N (|\mathcal{U}'_i| - 1)/N$. In contrast, the distribution of the active frequencies is in general unknown, but it is very unlikely that the distribution is uniform in real world scenarios. At this point the following question arises: What kind of information can be obtained about the global frequency? A first observation is, that the sum of the passive frequencies is much greater than the sum of the active frequencies. For instance, if the size of the anonymity set is constant then $\sum \nu_{\cdot, \mathcal{P}} / \sum \nu_{\cdot, \mathcal{A}} = |\mathcal{U}'| - 1$. Secondly, we know that the passive frequencies are uniformly distributed and thus we are able to derive a confidence interval α for the expected value. Hence, all passive frequencies will lie in the confidence interval with probability p (see table 1). Based on the global frequency of a single user u_i it is possible to derive an interval for the active frequency of this user (see figure 4). The parameter (α) determines the lower and upper bound of the confidence interval. By subtracting these bounds from the global frequency one obtains an interval for the active frequency which holds with probability p .


$$\max(\nu_{u_i}^G - \bar{\nu}_{\mathcal{P}}(1 + \alpha), 0) \leq \nu_{u_i, \mathcal{A}}^G \leq \max(\nu_{u_i}^G - \bar{\nu}_{\mathcal{P}}(1 - \alpha), 0) \tag{3}$$

Table 1. Confidence interval α , probability p and number of outliers

Authorized users $ \mathcal{U} $	Anonymity set $ \mathcal{U}' $	Number of auth. N	α	p	Outliers
1000	100	1000.000	0.005	≈ 0.8863	≈ 113
1000	100	1000.000	0.01	≈ 0.9984	≈ 2
1000	100	1000.000	0.02	≈ 1	≈ 0



(a)



(b)

Fig. 4. In subfigure (a) an example confidence interval is shown. Based on this confidence interval ranges for the active frequency of two users are derived (b).

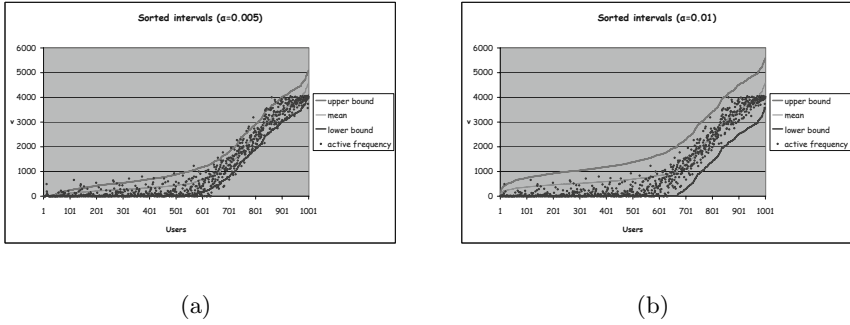


Fig. 5. Upper and lower bounds for active frequencies for two choices of the parameter α ; (a): $\alpha = 0.005$; (b): $\alpha = 0.01$

Note that the lower α the more precise are the lower and the upper bound for the active frequency. But, the number of passive frequencies which are not inside the confidence interval will grow and consequently the number of active frequencies which do not satisfy equation (3). From equation (3) it is possible to derive the maximum size of the interval $\delta_{\mathcal{A}}$, whereas $\delta_{\mathcal{A}} \leq 2\alpha\nu_{\mathcal{P}}$. This is also reflected in figure 5 where $\nu_{\mathcal{P}} = 100.000$ and $\alpha = 0.005$ ($\alpha = 0.01$). Consequently, $\delta_{\mathcal{A}} = 1000$ (2000) which can also be seen in figures 5. Considering two users u_i and u_j where the upper bound of user u_i is significantly smaller than the lower bound of the user u_j , then $\nu_{u_i, \mathcal{A}}^G$ is also significantly smaller than $\nu_{u_j, \mathcal{A}}^G$. These information can prospectively be used to improve the probability of identification of u_j in comparison to u_i .

It must be mentioned, that this estimation is independent of the distribution of the active frequencies. Furthermore, we have evaluated a number of random number generators (RNG) provided by standard libraries of different programming languages and most of them behave as the probability theory predicts and clearly was the basis for our investigations. However, we have also encountered a few RNGs that provide “better” results than expected and consequently more precise bounds could be obtained. Put differently, RNGs that behave “better” than the theory predicts, i.e. the passive frequency $\nu_{u_i, \mathcal{P}}^G$ of every user u_i will be very close to the mean passive frequency $\bar{\nu}_{\mathcal{P}}$, the active frequency of every user $\nu_{u_i, \mathcal{A}}^G$ can be determined precisely.

5 Conclusion

In this paper we have briefly discussed group based anonymous authentication (\mathcal{GBAA}) and strategies to construct anonymity sets. Furthermore, we have discussed attacks which can mainly be conducted by passive adversaries and finally we have pointed out how to estimate the number of authentications per user. This result can be used to reduce the anonymity of authenticating users. Additional side-channel information, e.g. user’s behavior, can be used to further improve the efficiency of the proposed approach. We conclude, that \mathcal{GBAA} , even

considered as a black box, leaks information on authenticating users over a period of time. One important fact is, that the approximated active frequencies of users are more precise the greater the number of protocol runs. In order to counter this kind of attack we recommend to significantly reduce the number of *GBAAs*. This can be achieved by a combination of *GBAA* and token based anonymous transactions, which is topic to current and future research.

Acknowledgements

The authors would like to thank the anonymous referees for their valuable comments and Michael Reiner for providing the simulation results to us.

References

1. Ateniese, G., Camenisch, J., Joye, M., Tsudik, G.: A Practical and Provably Secure Coalition-Resistant Group Signature Scheme. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 255–270. Springer, Heidelberg (2000)
2. Benjumea, V., Choi, S.G., Lopez, J., Yung, M.: Anonymity 2.0 - X.509 Extensions Supporting Privacy-Friendly Authentication. In: Bao, F., Ling, S., Okamoto, T., Wang, H., Xing, C. (eds.) CANS 2007. LNCS, vol. 4856, pp. 265–281. Springer, Heidelberg (2007)
3. Boneh, D., Franklin, M.: Anonymous Authentication with Subset Queries. In: Proc. of the 6th ACM conference on Computer and communications security, pp. 113–119 (1999)
4. Brands, S.: Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy. MIT Press, Cambridge (2000)
5. Camenisch, J., Shelat, A., Sommer, D., Fischer-Hübner, S., Hansen, M., Krasemann, H., Lacoste, G., Leenes, R., Tseng, J.: Privacy and Identity Management for Everyone. In: DIM 2005: Proceedings of the 2005 workshop on Digital identity management, pp. 20–27. ACM, New York (2005)
6. Camenisch, J., Hohenberger, S., Kohlweiss, M., Lysyanskaya, A., Meyerovich, M.: How to Win the Clone Wars: Efficient Periodic n -Times Anonymous Authentication. In: Proceedings of the 13th ACM conference on Computer and communications security, CCS 2006, pp. 201–210. ACM, New York (2006)
7. Camenisch, J., Lysyanskaya, A.: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Heidelberg (2001)
8. Chaum, D.: Security without identification: transaction systems to make big brother obsolete. *Commun. ACM* 28(10), 1030–1044 (1985)
9. Chaum, D., van Heyst, E.: Group signatures. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 257–265. Springer, Heidelberg (1991)
10. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (1994)
11. Damgård, I., Dupont, K., Pedersen, M.Ø.: Unclonable Group Identification. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 555–572. Springer, Heidelberg (2006)

12. Danezis, G., Diaz, C.: A Survey of Anonymous Communication Channels. Technical Report MSR-TR-2008-35, Microsoft Research (January 2008)
13. Danezis, G., Dingledine, R., Mathewson, N.: Mixminion: Design of a Type III Anonymous Remailer Protocol. In: SP 2003: Proceedings of the 2003 IEEE Symposium on Security and Privacy, Washington, DC, USA, pp. 2–15. IEEE Computer Society, Los Alamitos (2003)
14. Diaz, C., Seys, S., Claessens, J., Preneel, B.: Towards Measuring Anonymity. In: Dingledine, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 54–68. Springer, Heidelberg (2003)
15. Dingledine, R., Mathewson, N., Syverson, P.: Tor: The Second-Generation Onion Router. In: Proceedings of the 13th USENIX Security Symposium, p. 21 (2004)
16. Dodis, Y., Kiayias, A., Nicolosi, A., Shoup, V.: Anonymous Identification in Ad Hoc Groups. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 609–626. Springer, Heidelberg (2004)
17. Douceur, J.R.: The Sybil Attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)
18. Federrath, H.: Privacy Enhanced Technologies: Methods, Markets, Misuse. In: Katsikas, S.K., López, J., Pernul, G. (eds.) TrustBus 2005. LNCS, vol. 3592, pp. 1–9. Springer, Heidelberg (2005)
19. Froomkin, M.: The Death of Privacy? *Stanford Law Review* 52(5), 1461–1543 (2000)
20. Joshi, A., Joshi, K., Krishnapuram, R.: On Mining Web Access Logs. In: Proceedings of the 2000 ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery, pp. 63–69. ACM, New York (2000)
21. Kesdogan, D., Pham, V., Pimenidis, L.: Information Disclosure in Identity Management. In: Proceedings of 12th Nordic Workshop on Secure IT-Systems, Reykjavik, Iceland, October 11–12 (2007)
22. Kilian, J., Petrank, E.: Identity Escrow. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 169–185. Springer, Heidelberg (1998)
23. Lindell, Y.: Anonymous Authenticaion. Whitepaper Aladdin Knowledge Systems (2007), <http://www.aladdin.com/blog/pdf/AnonymousAuthentication.pdf>
24. Lysyanskaya, A., Rivest, R.L., Sahai, A., Wolf, S.: Pseudonym Systems. In: Heys, H.M., Adams, C.M. (eds.) SAC 1999. LNCS, vol. 1758, pp. 184–199. Springer, Heidelberg (2000)
25. Martucci, L.A., Kohlweiss, M., Andersson, C., Panchenko, A.: Self-Certified Sybil-Free Pseudonyms. In: Proceedings of the first ACM conference on Wireless network security, WiSec 2008, pp. 154–159. ACM, New York (2008)
26. Naor, M.: Deniable Ring Authentication. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 481–498. Springer, Heidelberg (2002)
27. Pshalidis, A., Meyer, B.: Linking Anonymous Transactions: The Consistent View Attack. In: Danezis, G., Golle, P. (eds.) PET 2006. LNCS, vol. 4258, pp. 384–392. Springer, Heidelberg (2006)
28. Persiano, P., Visconti, I.: A Secure and Private System for Subscription-Based Remote Services. *ACM Trans. Inf. Syst. Secur.* 6(4), 472–500 (2003)
29. Pfitzmann, A., Köhntopp, M.: Anonymity, Unobservability, and Pseudonymity - A Proposal for Terminology. In: Federrath, H. (ed.) Designing Privacy Enhancing Technologies. LNCS, vol. 2009, pp. 1–9. Springer, Heidelberg (2001)
30. Raymond, J.-F.: Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In: Federrath, H. (ed.) Designing Privacy Enhancing Technologies. LNCS, vol. 2009, pp. 10–29. Springer, Heidelberg (2001)

31. Rivest, R.L., Shamir, A., Tauman, Y.: How to Leak a Secret. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 552–565. Springer, Heidelberg (2001)
32. Sarasohn-Kahn, J.: The Wisdom of Patients: Health Care Meets Online Social Media (April 2008), <http://www.chcf.org>
33. Schechter, S., Parnell, T., Hartemink, A.: Anonymous Authentication of Membership in Dynamic Groups. In: Franklin, M.K. (ed.) FC 1999. LNCS, vol. 1648, pp. 184–195. Springer, Heidelberg (1999)
34. Serjantov, A., Danezis, G.: Towards an Information Theoretic Metric for Anonymity. In: Dingledine, R., Syverson, P.F. (eds.) PET 2002. LNCS, vol. 2482, pp. 41–53. Springer, Heidelberg (2003)
35. Srivastava, J., Cooley, R., Deshpande, M., Tan, P.-N.: Web Usage Mining: Discovery and Applications of Usage Patterns from Web Data. SIGKDD Explor. Newsl. 1(2), 12–23 (2000)
36. Teranishi, I., Kurukawa, J., Sako, K.: k-Times Anonymous Authentication. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 308–322. Springer, Heidelberg (2004)
37. Tzeng, W.-G.: A Secure System for Data Access Based on Anonymous Authentication and Time-Dependent Hierarchical Keys. In: Proc. of the ACM Symp. on Information, computer and communications security, pp. 223–230. ACM, New York (2006)