IDENTITY AND PRIVACY

Michael Sonntag

Johannes Kepler University Linz, Altenbergerstr. 69, A-4040 Linz, Austria
michael.sonntag@ins.jku,at

Keywords

Identity, security, privacy, proof of identity

Abstract

Security depends on permissions, which themselves depend on the identification of the person., This is contrary to another aim, namely privacy. This papers investigates the aspect of controllability of identification, i.e. whether the person to be identified can prevent others from doing this, or restrict the information disclosed to them to what this person provides. A method to "sign" stored messages, like forum posts, is presented, where the author can later on, for a self-determined time, prove authorship to interested parties and then disclose whatever information about her/him as desired.

1. Introduction

What is the "identity" of a user? Apart from the philosophical aspects of the question, this gets ever more complicated especially in IT. The reason is, that practically all security measures depend on first identifying the users and then (in various ways) assigning them permissions. However, the "securely and uniquely identifying a user" is getting more and more difficult. The past has shown that e.g. passwords are problematic: requiring frequent changes leads to sticky notes, and remembering multiple complex passwords is too difficult for users. Password managers can help to a certain degree, but do not solve the basic problem: this tool still has to – this time very – securely identify the user before it discloses the stored passwords. This might mean a reduction in the number of identification acts, but increases their importance. An additional factor is the increasing use of mobile devices: they might contain a camera, 3D-scanner, fingerprint sensor etc, but these are often cheap and not very secure (see Roy/Memnon/Ross, 2017). Also, because of a lack of external interfaces e.g. two-factor identification with additional tokens is problematic (note that NFC might improve this, but such tokens do not seem to be in widespread use at the moment).

This is exacerbated with the increasing problems of biometric identification: deep fakes are now getting trivial to create. Face recognition should therefore be considered completely broken if unable to reliably detect them. 3D scanning might still be quite secure, but whether this remains so, is doubtful. The current approaches require 3D masks to be created – but who can hide the shape of their head in public? Note that faking the voice is today possible too. This approach suffers from the same problem practically all biometric features have: they are public. If they are not public, they typically cannot be used for identification purposes. For example, a picture of the genitals might not be easily available, but unlocking your phone in public would be impossible too. Some feature might be more difficult to obtain (vein or retina scans), but still not completely impossible to obtain and later falsify if a specific person is attacked deliberately. These increases in "hacking

technology" lead to ever more invasive and complicated identification procedures – or a much lower certainty for identification, giving rise e.g. to multi-factor authentication

2. Identity and Privacy

If identification is getting ever more problematic, is this not a good direction regarding privacy? Unfortunately it is not. While identification might not be as reliable any more, this applies against deliberate attacks to impersonate someone. On the contrary, for "normal" use it is getting better. Face recognition might not be extremely reliable, but coupled with voice recognition and perhaps augmented by fingerprint scans and quick automated DNA testing (the last only a future prospect) persons can be identified in public easily. This also means that they can be traced in their movements. Also consider, that it is not always strictly necessary to identify the person – a mobile phone might differ from the identity of its owner, but it is a very good surrogate, as it is rarely given to others and almost continuously carried by the person.

These two aspects lead to a paradox: identification should be better, to improve security, and worse, to enhance privacy. How can those two be combined – if this is possible at all? One option is to increase the controllability of identification, i.e. whether we want to be identified (→ quick and securely) or not (→ impossible to do it). This is obviously a good idea, but the difficulty is in the implementation. We will therefore take a look at the various methods of identification from this point of view and how they support it.

2.1. Controllability of Biometrics

As discussed above, controllability of biometric identification is difficult. Everything that is publicly obtainable (viewable, observable) is hard to constrain. For example, gait recognition might be very secure and for humans almost impossible to replicate, but if we can obtain exact measurements (carried device, long video), a robot could be technically easily built to exactly replicate these motions. This is true at least for all widely-used biometrics: they can be observed and be replicated.

- Fingerprints: Fingerprints are left everywhere, and there is absolutely no practical way to keep them secret. Also note that e.g. driving licenses or passports might require/produce large libraries of fingerprints.
- Face recognition (2D): Unless continuously wearing a mask, there are just too many possibilities of obtaining a photo of your face. See also Facebook, where a picture of any person might be uploaded and tagged, even though that person might not be a member at all.
- Face recognition (3D): Same a 2D pictures, videos of a face from multiple angles are trivial to obtain from afar and without notice by the victim. Even deliberate and detailed IR scans (exactly like phones use) are not noticeable to the victim.
- Iris scans: As has been demonstrated, a high-resolution picture from several meters away is enough (see also Fancourt et al 2005 for recognition via long distance). Therefore this must be considered as impossible to keep secret.
- Retina scans: Controllability exists only partially. The information might be obtainable, but only difficult (e.g. during an eye examination), especially from afar (flooding with infrared light, zoomed high-res video of eyes to obtain picture of whole retina).

- Hand geometry: If not permanently wearing with gloves, the geometry of a hand can be deduced from pictures easily. Also, prints on surfaces can be helpful.
- Vein scans (hand, finger): This can be more difficult, as it required infrared light and the inside of the hand/fingertip. Still, if deliberately targeting a person e.g. traps can be created, where a scanner is mounted beneath a surface where that person is going to put their hands.
- Gait recognition: The way of walking cannot be kept secret, as creating a video of someone walking is easily possible for as long as needed.
- Keystroke dynamics: Speeds, delay, etc of typing cannot be kept secret and might even be deduced from merely videotaping or listening to it (Asonov/Agrawal 2004).
- Heart pattern: Heart beat patterns have been monitored from afar with precision distance measurements. However, at the moment it is unclear whether this is enough to fool identification algorithms. But see also smart watches: if you hack them, you definitely get all the data necessary, potentially even in real-time.
- Brain pattern: Probably at least at the moment impossible it is impossible to detect brain waves at a distance, but verification is therefore similarly intrusive. This does not seem to be currently in practical use.
- Voice analysis: The voice is impossible to keep secret and it is today easy to synthesize arbitrary texts in a different voice.
- Signature recognition: Your manual signature is trivial to keep secret in theory, but easy to obtain if targeting a specific person, as they exist at many places. Advanced signature recognition not only takes the picture into account, but also dynamics (speed, pressure etc); these are more difficult to obtain, as only monitoring of an actual signing discloses these values (e.g. videos of a signature will not provide pressure information). Replication is probably easy with advanced technology readily available (simple robot with lots of degrees of freedom).
- DNA: There is absolutely no way to keep this secret at all. You distribute your DNA all over the area whenever being merely present.

Typical countermeasures against attacks on physiological biometrics (behavioral biometrics has this already built-in) are liveness checks. These usually can be countered relatively easily, as the degrees of freedom are limited, especially as the person still has to be recognized and therefore the biometric feature itself cannot change too much. This is more a question of “what exact feature is the verifier looking for?” than impersonating a whole human being.

The best biometrics regarding controllability require strong physical closeness to obtain the data. The natural tendency of persons to keep a distance helps there. This is reduced with features that are regularly used with extreme closeness, i.e. by contact, like hands.

Controllability has a second dimension: are fakes easy to “present” under observation? For example a 2D picture with the eyes cut out might fool lots of mobile phones into unlocking themselves, but if a human is observing the interaction, it will be practically impossible to not raise suspicion. A further aspect here is, whether the observer can detect manipulations, but does not receive identity information (e.g. using a fake external eye to fool a retina scanner is detectable, but neither in this manipulation attempt nor on real verifications the observer obtains the identity of the person authenticating). Similarly, while special makeup or clothing can prevent facial recognition by public video surveillance, this will easily be apparent and can probably also be trivially identified as

"problematic" by an algorithm i.e., the aim for privacy would be not for the cameras to not identify the person, but not detect the presence of a person (which then needs to be identified) at all.

The ideal biometric identification would be a kind of permanently closed eye that looks like normal skin (no recognition that it is there), that will solely be opened for identification (no chance of observing it during other activities) at very close range (no obtaining by observing an identification), can be opened only deliberately (no accidental or prompted disclosure outside of identification) – and is not used for anything else (no possibilities for re-purposing "normal" equipment). While this does not exists, its features are important for assessing identification in general, and biometrics specifically.

2.2. Controllability of Possession

Possession typically means a token, which produces some sort of code (numeric, picture etc), either displayed visually or communicated directly to the other side (via e.g. cryptographic protocols). As long as the physical device is in your possession, controllability is typically good. But most of these tokens do not have any further security measures, so as soon as an attacker got hold of them, there exists no further controllability at all. The only possibility is discovering the loss and still possessing enough information to lockout the device for all entities that will attempt verifications in the future – which is only possible if these potential verifiers are know or check at a specific instance for any updates of exclusions.

While the secret of these tokens is most of the time not absolutely safe, i.e. with enough investment it might be possible to extract it, this takes a long time and changes/destroys the token. Therefore "undetectable copying", like in biometrics, is not possible absent security issues (see e.g. Charette 2011 regarding the RSA tokens, where probably the seeds were stolen).

Regarding privacy tokens are usually very advantageous. They will be connected to a certain person, but this information does not have to be apparent from the device at all. This means, stealing "some" device does not automatically disclose its owner's identity. Similarly, using a token for logging in does not provide information for identification (separately from the rest of the login process, e.g. an username), as tokens are typically at least outwardly identical. Additionally, the verification protocol should not disclose the identity of the verifier, to listeners, but this is not always guaranteed. Moreover, even for tokens with NFC capabilities, many of them incorporate a physical button, meaning they cannot be incited to perform an unnoticed authentication without the knowledge of the owner. Therefore large-scale identification of users is completely impossible.

2.3. Controllability of Knowledge

Knowledge itself is controllable very well, as extraction from the brain does not work. However, extortion, violence etc are of course still possible. However, as soon as the information has been disclosed, no further controllability exists at all, as there is no influence over any further dissemination. Additionally, as the knowledge originally only exists within humans, it must be disclosed in some way to serve as identification. This act may be, and typically is, observable. Contrary to e.g. fingerprints this is usually also easy to "copy" during this act, as knowledge is almost always typing some kind of "code" into a version of a "keypad" (selecting specific symbols/ pictures with a mouse, entering numbers on a keyboard etc). Compared to biometrics therefore the act of identification is here a bigger problem of controllability.

Knowledge does not enable any kind of identification (absent an explicit act of identification) through mere observation. The knowledge (or a matching counterpart) must be known to the verifier too, who additionally associates it to an identity. Therefore privacy is very well protected.

Moreover, while e.g. password are often not unique, knowledge elements are changeable and not directly tied any person. i.e. someone can "change" their identity or "share" it with other persons or entities, e.g software programs or agents, effectively creating a surrogate themselves (which is impossible with biometrics and usually extremely hard with tokens).

2.4. What is identifying you?

Another aspect of "identification" is the question, what identifies you. E.g. an IP address might individualize you and allow certain person/institutions to discover your name, but your passport number – while much more reliable and unique – can often be more anonymous as fewer persons have access to it in ordinary business or access to the special database storing them all. Essentially, "proving the identity" always takes place in relation to a *specific* other entity. That entity has certain other knowledge, access to some databases etc. The identity data disclosed could therefore be specifically tailored to the current counterpart to provide them only with as much information as desired. As long as that information is not passed on - and deleted soon after - this would work. But the knowledge of an entity may increase (or decrease) over time, and data can be passed on to others with a different set of related databases and more possibilities of deriving further information through combining them.

It is therefore better to take one step back and think about the *need* for identification. Mostly there is no actual need to know the identity, merely to ascertain certain properties. For instance, when attempting to cross a border the identity of the person is of no interest to the border police: they want to ensure the person is allowed to cross the border, not on a list of fugitives/wanted person, old enough etc. This can be reduced to the following classes of properties:

- (Not) Being on a list: e.g. persons granted access or possessing a driving license. The exact position on the list as well as any other data (columns, date/time of being added to it etc) should remain unavailable. Example for not being on a list are not being wanted or having had the driving license revoked.
- Properties of the person, for instance age or social security number.
- Knowledge of some data, e.g. an access token
- Possession of abilities or ressources, e.g. a minimum of available computation power

Note that there are different ways of achieving this. For instance, "being on the list" could be confirmed by a trusted third person (e.g. the issuer of the list), or a proof that a list provided by the verifier does include you (trivially: everyone on the list knows a shared secret value – proving the knowledge of the value is evidence of being on the list without disclosing the position on it). The latter is preferable from the point of privacy, as "confirmation by a third party" requires identification to that party. If unavoidable, these should be offline proofs (with all their disadvantages, like the revocation problem), so that the third party need not be involved in the verification (and therefore at least know this fact and the time of it taking place).

3. Identification on synchronous communication

Identification over synchronous remote communication is very problematic, and getting increasingly so. For example, with deep fakes and voice faking the co-called "nephew trick" of impersonating a relative to obtain money, is going to be even more problematic. And in times of a virus outbreak where old persons should not be visited in person – and simultaneously ever more

old persons are tech-savvy and use electronic communication means – this could produce huge problems. If the person looks like the nephew (one photo from Facebook may be enough), sounds like the nephew (calling him under some pretense), and merely has a slightly different (but with a good explanation) online address, convincing someone to part with money or valuables can become much more likely and therefore dangerous. Note also the methods for video identification, e.g. when remotely opening a bank account, are going to be either more complicated or unreliable. But at the same time, we can also pretend to be at multiple locations. E.g. with a good chat-bot and a not-so-important-role, meetings do not necessarily have to be attended personally.

Consequently, how do we know the person on the other side of a call is who we think they are? This has always been a problem for strangers, i.e. people we do not know very well. But now this is getting problematic even for familiar people. If this becomes a larger issue, demands for mandatory identification/registration/official communication accounts will increase and lead to a swift decline of privacy. Even then, taking over a Skype-/E-Mail/... account will only get more interesting for attackers, as these become more valuable and simultaneously harder to replace. Therefore it will be necessary to strongly increase the identification before using them (additionally a very nice position for logging; position/time/device, but also communication partners/habits/ etc).

What is therefore needed is identification *inside* a synchronous communication (like video conference, chat, or audio call). This should be separate from the communication itself, so establishing the connection does not give any listener or intermediary (like to service provider) any information or proves the identity. Therefore end-to-end encryption is a necessary, but not sufficient, requirement. Nobody outside the call can obtain any provable information, as the verification takes place only inside. This preserves privacy, as solely the actual communication partners securely obtain the identity of the other side, but requires this additional verification step. This does not completely solve the problem, as a connection still has to be established, i.e. some "number" has to be "dialed" to reach the intended communication partner by the initiator (the respondent does not necessarily need any information on the other side, which might remain anonymous to him or others). This "number" need not be a direct identification (could e.g. also be a tor hidden service), but the caller must be able to store it for later communication. Because of this it must be stable, i.e. not change over time (a lookup does not help – then the "number" merely transfers into the "name" to lookup). Therefore privacy is necessarily reduced at least to the point of linkability: this communication is to the same person as the one in the past, even if we don't know exactly who is calling, and whom is being called.

4. Identification on asynchronous communication

Asynchronous communication has the same problem: how to securely identify the origin of a message. Of course this is possible via the hosting service, i.e. the operator of the forum, blog, newsgroup etc. But regarding privacy that is not ideal. The same problem exists for electronic signatures – everyone can verify them – and in this way they disclose the identity to whatever degree is contained in the certificate. Therefore the possibility to post anonymously, but the author still being able to securely identify themselves – if necessary – is desirable. Note that in this model it is completely up to the creator of the content to voluntarily disclose her/his identity, and to which degree. This means that a short (or longer) time after posting someone may try to verify the author of a message or obtain some other information (e.g. age, nationality, location) from them, but this cannot be enforced (fallback to anonymity). There exist several problems that need to be solved for an implementation of this approach:

- The post must contain some contact information, but this information may neither allow to identify the author, nor provide linkability, i.e. the possibility to declare that this article must be from the same person as arbitrary other postings. Anonymity also extends to the process of identification: starting the process should not reveal any information about the identity of the author. Ideally the author can stop the process of identification at any point before the successful end without giving away *any* information.
- The owner of the forum cannot identify the author from any stored information, i.e. the posting is anonymous or all tracing information is removed. This is not absolute, as e.g. the identification feature could also be used to securely prove the identity as opposed to some claimed identity (service operators often identify users merely via an E-Mail address).
- The process of identification should not involve the hosting service, i.e. they should know that someone looked at the posting, but not whether this person tried to identify the author or what the identity of the author is. Of course it may attempt the identification itself, but only just like any third party.
- The author should be able to arbitrarily decline the request. This could be e.g. based on elapsed time ("will prove for 1 month only, then it should have been deleted anyway") or conditional to the asking person willing and able to proove their identity (and being "acceptable" according to the user).
- No central registry exists, i.e. there may be collisions in the contact information of an article. In case there are several posts with identical data, multiple persons must be able to be contacted. It must always be possible to reach the "correct" person for identification - as long as this person is willing to identify. This simultaneously produces plausible deniability.

It should be noted that this system supports identification in synchronous calls to: communicating a "verification code" will allow the other side to initiate an identity verification.

A possible implementation for this problem could consist of the following elements:

1. To ensure asynchronicity, it cannot depend on the author being logged in, reachable at the moment etc. Therefore a surrogate of the user in the form of a software agent is needed. This can execute and remain active around the clock. As long as no authentication takes place, it needs very few resources, as there is nothing to do for it. If this is not needed (e.g. for synchronous communication, some software is still needed, but it need not be active without the user).
2. When a new message is created, the author tells its own agent to "sign" it. The agent then creates a new public/private keypair and uses this to create a Tor hidden service. As this is separate for every signature, i.e. post, no linkability exists through them (if desired, the same key can be used for all posts on a single site or based on any other distinction too, reducing linkability). Simultaneously it prevents identification of or recognizing the agent e.g. through its IP address. This hidden service is the "entrance" point for verification. Whenever the author intends to "abandon" the post temporarily or permanently , the hidden service is simply deregistered/not started again. Verification will then always fail.
3. The agents signs the message with a secret value (which could be the same for all instances or again individually – here this would only improve security). The exact algorithm depends on the kind of verification attempted later (see next element).
4. If asked to verify, the agent uses a Zero-Knowledge proof to prove to the client that he indeed knows this secret value. It then provides any identification information it deems

appropriate for this verifier (which might have to identify themselves). As this is a Zero-Knowledge proof, the inquiring person/agent can rely (to an arbitrary degree) that this response comes from the real author's agent, but third parties cannot. If both collude, they can falsify the protocol in a way undetectable, and therefore still believable, to such third parties. For them it therefore looks like someone prooved authroship of a post, presented an identity, and someone else believes this. But they cannot verify it themselves (but can reasonably depend on the result, because mostly verifiers will be independent and not lie). Note that third parties can obtain the identification information only from the verifier, not from listening in on the protocol, which is encrypted because of the Tor connection.

Such a Zero-Knowledge proof can consist e.g. of the discrete logarithm (as described by Chaum/Evertse/van de Graaf 2000): $g^x \bmod p = y$. The message (or its hash value) is used as the base g and the secret number (perhaps randomly generated for each post or always identical for this user) as the exponent x. The modulo number p can be publicly known or, similar to the result y, be added as the signature to the post (in addition to the Tor hidden service descriptor). The agent can then prove the knowledge of this x without actually disclosing it by (multiply) correctly answering one of two possible challenges by the verifier. As the message is also part of the equation this will only work if it is unmodified, effectively creating a "signature". Therefore even a successful verification does not allow a malicious verifier to impersonate the author later on or convince third parties that the identification is correct (or wrong).

There remains only a single problem: as each agent individually and randomly generates a new hidden service identifier for each post, collisions may occur. While these are highly unlikely (tor hidden services are essentially public-private key pairs and a collision means that two agents accidentally generated exactly the same keys), they cannot be ruled out. Additionally, this possibility even has some advantages, as an agent can then plausibly decline that this is not "his" hidden service, even if it should be proven that he employs it for some message. As long as it cannot be forced to show this alternative message, tracing a hidden service only gives a high probability that this agent is responsible for it. As the Tor system is designed for only a single hidden service behind every descriptor, if two servers exist with the same, usually the later one to register wins, but there seems to be no quick deterministic result – some clients could still see (at least for some time) the old descriptor. To correctly support the approach described here, an "enumeration" of all servers would be necessary, i.e. connecting to the first, if this one declines connecting to the second and so on. Load balancers for hidden servers do exists, but these work by knowing all the servers they are balancing – something that cannot be assumed in this context.

Controllability of this approach rests in the hands of the author: she can instruct the agent to abandon individual posts, or simply destroy the agent, thereby rendering everything anonymous. A further advantages of this approach is, that even later advances in cryptography will not help. Should the secret for the Zero-Knowledge proof be discovered, other persons can claim authorship too, but this does not allow anyone to identify the author or proof such a suspicion. Similarly breaking the hidden service key would only allow impersonation, but not attributing it to a specific agent. Privacy is practically complete, as apart from any registration/login requirements to be able to communicate, no identifying information about the author is disclosed at all and there is no way to link it to another communication act.

The biggest drawback is the mirror image of the advantages: the author of a post cannot be discovered by anyone, not even the police. But if a specific suspicion for a certain person exists, first the agent of that person would have to be found/identified. This should normally be possible (note that it could also be created/contacted/… via Tor, rendering this at a minimum extremely difficult). After analyzing the agent, it can be proven that the secret key for the hidden service is

present, and that the secret for the Zero-Knowledge proof is available in it too. Essentially the agent can be "forced" (or a similar agent provided with the necessary data) to perform the validation of the signature. As the agent must be able to provide the identity to be useful in some form, this must be accessible to it too (optionally only an encrypted version could be stored, with the matching key being part of the post – only this post can then unlock at most that data). Essentially, verifying the authorship of an already known person with a known post remains possible and discloses at least all the identity information stored in the agent for this post. As technically the verification is easy, testing all the contained data against a large number of posts is feasible too. Securing the agent against unauthorized access to its data is therefore paramount. It should therefore be executed only a a trustworthy server.

The only countermeasure if the author does not provide her identity or refuses to participate in identification is deleting the post (terminating the phone/video call etc) – or trying to get hold of the (respectively all existing) agent somehow and the try the verification as explained above. It should be noted, however, that at the moment the situation is very much the same: obtaining an E-Mail address via Tor might be getting ever more difficult (e.g. Google requires verification via a telephone number), but remains possible. Registering for most platforms only requires an E-Mail address and no further proof of identity. Therefore all communications are anonymous, although potentially traceable (IP addresses in logs, ISP records of IP assignments etc). But contrary to the scheme presented here proof of authorship is very hard or impossible (especially without giving away information about other posts or platforms), as only current (i.e. not necessarily when the act took place) control of the E-Mail address (which is still some way from an identification) can potentially be proven.

5. Summary

A method to prove the identity of a person by generating a privacy-friendly signature was prevented and its implementation discussed. It could also be used in synchronous communication, where the implementation through an independent software agent would allow it to take place automatically in the background. Generally identification is the natural enemy of privacy – which no longer applies if nobody can identify the entity with reasonable means – it is therefore very important put identification into the hands of the person. It should not be possible without their consent (controllability). Especially in time with a virus epidemic and increasing demands for smartphone Apps logging all interactions/meetings between persons, restrictions on identification by others become supremely important.

6. References

Asonov, D., Agrawal, R. (2004). Keyboard Acoutic Emanations. *IEEE Symposium on Security and Privacy, 2004.* Berkeley, 2004, 3-11.

Charette, R. (2011). SecurID - Used by 30,000 Organizations and 40 Million People - Sees Security Reliability Partially Compromised. IEEE Spectrum, 18.3.2011, https://spectrum.ieee.org/riskfactor/telecom/internet/securid-used-by-30000-organizations-and-40-million-people-sees-reliability-partially-compromised

Chaum, D., Evertse, J.-H., van de Graaf, J. (2000). An Improved Protocol for Demonstrating Possession of Discrete Logarithms and Some Generalizations. *Proceedings of Advances in Cryptology — EUROCRYPT' 87*, Springer, Berlin, Heidelberg, 2000, ISBN 978-3-540-39118-0, 127-141

Fancourt, C., Bogoni, L., Hanna, K., Guo, Y., Wildes, R., Takahashi, N., Jain, U. (2005). Iris Recognition at a Distance. *Audio- and Video-Based Biometric Person Authentication*, LNCS 3546, 2005, ISBN 978-3-540-27887-0, 1-13

Roy, A., Memon, N., Ross, A. (2017). MasterPrint: Exploring the Vulnerability of Partial Fingerprint-Based Authentication Systems, *IEEE Transactions on Information Forensics and Security*, 12(9), 2013-2025.