TOWARDS ESTABLISHING THE LINK BETWEEN A PERSON'S REAL-WORLD INTERACTIONS AND THEIR DECENTRALIZED, SELF-MANAGED DIGITAL IDENTITY IN THE DIGIDOW ARCHITECTURE

Tobias Höller

Institute of Networks and Security
Johannes Kepler University Linz
tobias.hoeller@ins.jku.at

Keywords

IDIMT, Digidow, digital identity, biometrics

Abstract

The Digidow architecture is envisioned to tie digital identities to physical interactions using biometric information without the need for a central collection of biometric templates. A key component of the architecture is the distributed service discovery, for establishing a secure and private connection between a prover, a verifier and a sensor, if none of them knows the others ahead of time. In this paper we analyze the requirements of the service discovery with regard to functionality and privacy. Based on typical use-cases we evaluate the advantages and disadvantages of letting each of the actors be the initiator of the discovery process. Finally, we outline existing technologies could be leveraged to achieve our requirements.

1. Digidow vision

Digital identity will be a central requirement for many future applications. There is ample research under many different aspects on how to create, assign, and verify an individual person in a digital world based on interactions in the physical world. The Digidow project (Institute of Networks and Security, 2019) envisions a trustworthy infrastructure enabling biometric authentication without central databases. Such an infrastructure would provide two new capabilities:

- Individuals should no longer need to carry a physical token to prove a digital identity. Providing their biometric attributes to a sensor should be sufficient.
- Individuals should remain in full control of their personal information, the data should remain decentralized and offer extended privacy guarantees.

Of course, these new capabilities must not compromise the security of the infrastructure. Otherwise, it would no longer be useful for critical applications, like i.e. passport checks.

1.1. Security Benefit

Digidow's decentralized approach provides an important security improvement over current implementations. Passports provide a very nice example:

Currently, passports contain a chip with the biometric information of their holders. This information can be used to verify the identity of a person by comparing the measured biometric values with the data stored on the passport's chip. But that also creates a security issue because it means that biometric information could be extracted from stolen passports and used for identity theft (Vijayakrishnan, 2008). With Digidow there is neither a physical token that could be lost, nor a centralized database, which would be hacked, as seen in countries like India (Khaira, 2019). The important biometric information remains under full control of the owner at all times. This also has another advantage over current biometric passport implementations: If a foreign government wants to read out a passport, they will be able to do so, in Digidow the individual user still has the power to prevent that (or at least notice it happening).

1.2. Envisioned architecture

The Digidow architecture is illustrated by Figure 1. The communication always flows between three actors. The personal agent is a piece of software under full control of an individual person and manages all personal information. Verifiers can be operated by everyone who wants to use the Digidow infrastructure to interact with the digital identity of people. The most common scenario will probably be matching an individual with his/her identity. That identity will usually come from a nation state, but of course a personal agent could be linked to other identities as well.

Sensors can be operated by everyone and are responsible for collecting and providing biometric information to personal agents. They must be available in every location where a verifier wants to interact with a personal agent. But that does not imply that sensor and verifier have to be operated by the same entity.

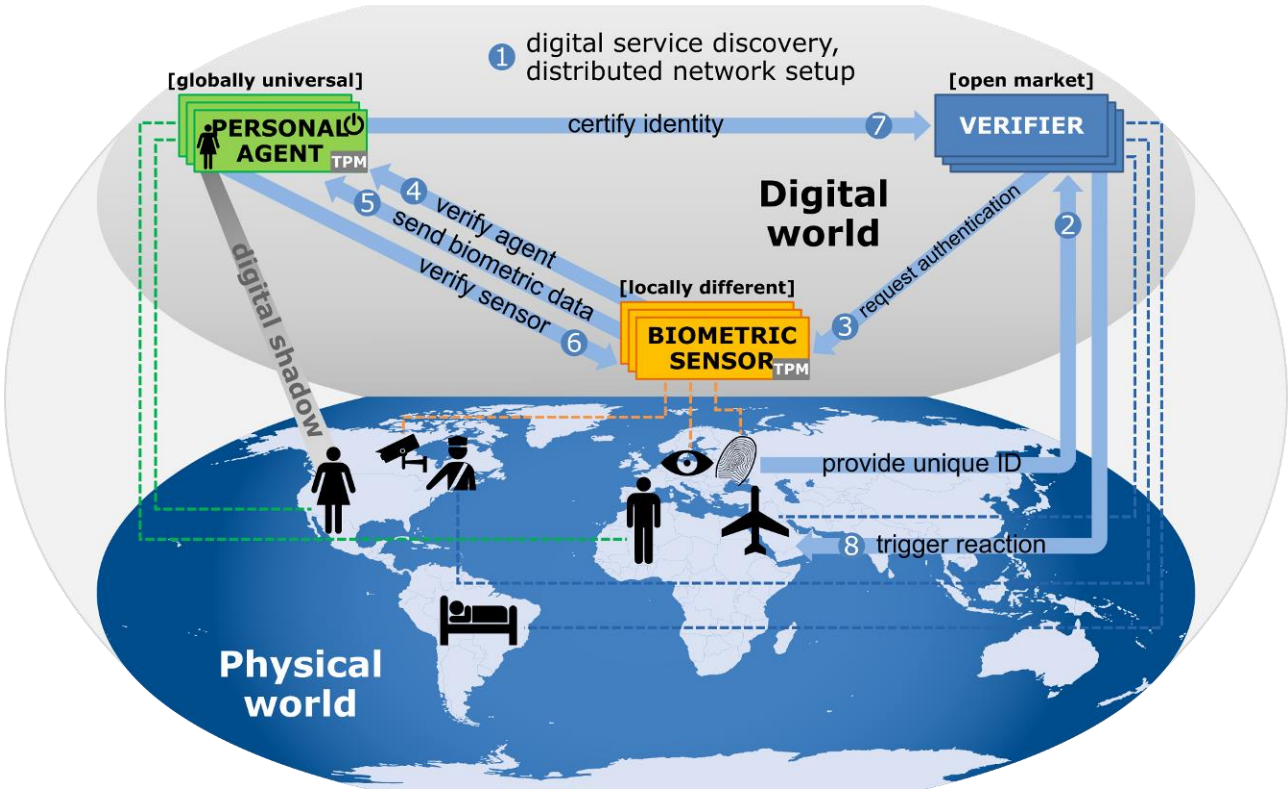

Figure 1: Overview of Digidow Architecture (INS, 2019)

2. Service discovery requirements

Figure 1 has a rather inconspicuous step one called ``digital service discovery''. The main objective of this step is the identification and localization of all parties required for the transaction. A set of functional and privacy requirements has been identified for this initial step.

2.1. Functional requirements

- Working on a global scale: In the digital world there should be no need to restrict a distributed system to a specific reason. Therefore, the service discovery should function equally, independent of the location of the parties involved in the transaction.
- Latency: In order to be practical, the entire process of verifying an ID should take no longer than alternative authentication procedures. The first step must therefore be as time efficient as possible. It is hard to pin down exact timings, but transit fare providers for example have historically required transaction times below 500~ms (Smartcard alliance, 2011). We believe that the generic and decentralized approach of Digidow justifies longer transaction times, but use the 500~ms as a goal for now. More research has to be done on acceptable average and maximum transaction times for Digidow in the future.

2.2. Privacy requirements

- A global listening adversary should not be able to correlate which personal agent is talking to which verifiers. We assume that this meta information could be used to infer the individual behind a personal agent without permission.
- A verifier should not receive any information about the user identity until the personal agent decides to actually reveal it. Otherwise the user would not retain full control over his personal information.

3. Use-cases

In order to identify the challenges presented by this project, two possible Digidow use cases have been selected, which will be used to identify requirements in more detail. For every use case there is a set of questions to be asked:

- Who will be operating the verifier?
- Who will be operating the sensor?
- Who should initiate the transaction?
- Which information is available at sensor, verifier and agent?
- Which additional information must be provided?

3.1. Digital passport

The digital passport would enable citizens to legally prove their identity without a physical identification token (the classical passport). This means that the personal agent needs to be confirmed by the issuing country of the passport. This will most likely work with technologies and procedures similar to the ones currently in place for the European E-ID (European Parliament, 2014).

Fortunately the implementation details are not relevant to the question, which should be answered here. Passports can be used in multiple different scenarios, but the by far most common one is providing identification when traveling in foreign countries. Passport verification usually happens at borders or airports. In both cases the verifiers would be operated by the border guard of a nation. Assuming that the digital passport should allow for the passport check to work as before, it is a fair assumption that the sensors will also be operated by the border guard. In this arrangement the verifier wants to identify one specific individual at a time based on the biometric data provided. It would make sense to expect the verifier to initiate the transaction by trying to find the responsible personal agent.

3.2. Ticketing for public transport

Ticketing is an area where a smooth transition between the digital and the physical world could provide significant benefits. The general trend towards mobile payment (European Parliament, 2017) also applies to the purchase of tickets for public transportation. More current trends try to remove the need for purchasing tickets in advance entirely. Instead users just have to identify themselves when entering (for example public transport) and leaving (Rhein-Main-Verkehrsverbund, 2019). The selection of the ideal ticket as well as the actual payment can be handled automatically in the digital world. A good example for such an approach would be GPS ticketing using smart phones (FairTiq, 2019).

Such a system could also be based on the Digidow architecture. It would allow users to board simply by interacting with a sensor, without any need for physical devices or special applications. Furthermore, it would no longer be necessary to fully track users via GPS in order to select the ticket. Instead only personal agents would be tracked, without any knowledge about the person in the physical world running the personal agent.

4. Open questions

The goal of this research is to identify all criteria the service discovery step needs to satisfy and provide a potential implementation. A core issue is how the entire service discovery process should be structured.

4.1. Who initializes the service discovery?

4.1.1. The sensor

The currently suggested infrastructure requires a human interacting with a sensor in order to work. Considering that most users would not interact with a biometric sensor, unless they wish to authenticate themselves, this event makes for a good starting point for the digital service discovery. At this point the system has very little information available. The sensor only knows the biometric information it has collected and may reasonably assume that there is a personal agent somewhere on the network, which feels responsible for the individual with the collected biometric features. Detecting a personal agent with only that information without compromising privacy and security is a currently unanswered question.

Most likely the sensor would have to provide an additional interface enabling users to provide a specific ID unique to their personal agent (and maybe a second one for the verifier). This idea is reflected in the current Digidow architecture shown in figure 1 where the verifier receives a unique ID from the user.

A problem to keep in mind here is that there are also biometric measurements, which can easily be extracted without consent. Face recognition using surveillance cameras would be the prime example. If the sensor acts as the starting point for service discovery and no additional user input is required, measures must be taken to prevent abuse of the Digidow infrastructure to track users against their will.

4.1.2. The personal agent

An alternative approach towards service discovery might be to start the process earlier. If we assume that a personal agent has extensive information about every real-life interaction of its owner, it might be able to predict necessary interactions even before the user interacts with a sensor. Imagine a scenario where you leave your house, activating the alarm system via your personal agent. Your personal agent now knows your location and can make reasonable assumptions about your next actions. Most likely you will go to your own car, call a taxi (or self-driving car in the future) or use public transport. So the personal agent could proactively contact the very limited set of likely sensors the user might interact with and establish multiple probable connections.

Such a system would be nice while it works, but it would need a backup mechanism where the user can manually tell the personal agent to prepare a connection to a sensor in order to ensure availability. And that in turn would require a physical device, which would not be an issue due to the widespread use of smart phones, but it would compromise on Digidow's goal to work without physical tokens.

4.1.3. The verifier

The last idea would be to split the responsibility for service discovery. What would be the case if the sensor was only responsible for detecting the verifier (which it will often be paired with anyway) and leave the remaining work there?

If the personal agent has already been linked with the verifier, this approach would work very reliably. In the case of passports that might have happened during the visa application process or in other scenarios during an account creation process. If the biometric information of everyone with a valid visa is available to the verifier along with a unique ID of their personal agents, the verifier could easily take care of service discovery.

Again this approach has the disadvantage that it does not work for the initial identification procedure and gives the verifier access to sensor data, which it would not need otherwise.

4.2. Link personal agent to sensor data

The most critical piece of information within the system is the biometric data collected by the sensor. An ideal goal would be to discover the responsible personal agent only based on the measured biometric information, without revealing that information to anyone but the corresponding agent. If that is not possible, which additional information must be provided by the user in order to enable that detection? The classic approach of assigning a unique identifier to a service available on the internet is insufficient, because that would either require a central name service or reveal the names of the personal agents, a sensor is communicating with, to a passive attacker.

5. Potential directions

5.1. Lessons from OAuth

The OAuth framework provides a similar functionality as Digidow. It also defines communication between three parties: A client, an authorization server and a resource holder. The client is verified by the authorization server and receives a token to make requests to the resource holder (Hardt, 2012). A personal agent acts like an authorization server, verifiers are just another word for resource holders and sensors are only the tool a client uses for interactions. For the relevant research question the functionality itself is not too important, but the data exchanged between the parties and the experiences made in the real world will be able to provide valuable input.

5.2. TOR Hidden Services

TOR (The Tor Project, 2019) is a project to enable safe and private routing of traffic on the internet. One of the capabilities of TOR lies in running hidden services. They are designed in a way that their public IP address remains hidden. The only way to access them is via their unique hostname. That improves privacy by making schemes like IP localization impossible, but the individual personal agent can still be tracked reliably.

The protocols of the Digidow project should work independent of TOR, but the service discovery process has to be aware of the additional indirection steps required to reach a hidden service. Can service discovery reliably detect personal agents within the defined performance requirements without compromising the privacy offered by the TOR network? Can the privacy of personal agents improved by replacing the static unique hostname with a temporary one? Otherwise a global passive attacker could still track connections based on who accesses which hidden service when.

6. References

European Parliament and the council. Regulation (EU) No 910/2014 of the european parliament and of the council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC. url: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv%3AOJ.L_.2014.257.01.0073.01.ENG (visited on 06/21/2019).

European Payment Council. White Paper Mobile Payments. Version 5. EPC492-09. 2017.

FairTiq. How it works. url: https://fairtiq.com/en-ch/fairtiq-app/how-it-works (visited on 06/21/2019).

D. Hardt. The OAuth 2.0 Authorization Framework. RFC 6749. RFC Editor, Oct. 2012. url: https://www.rfc-editor.org/rfc/rfc6749.txt.

Institute of Networks and Security. Digidow. url: https://ins.jku.at/research/projects/digidow (visited on 05/11/2019).

Rachna Khaira. Rs 500, 10 minutes, and you have access to billion Aadhaar details. url: https:// www.tribuneindia.com/news/nation/rs-500-10-minutes-and-you-have-access-to-billion-aadhaar-details/523361.html (visited on 06/21/2019).

Rhein-Main-Verkehrsverbund. RMVsmart: Meine Strecke. Mein Preis. url: https://www.rmv.de/c/de/fahrkarten/die-richtige-fahrkarte/rmvsmart-das-neue-tarifangebot/(visited on 06/26/2019).

Smart Card Alliance. Transit and Contactless Open Payments: An Emerging Approach for Fare Collection. TC-11002. 2011.

The TOR project. url: https://2019.www.torproject.org/docs/documentation.html.en (visited on 06/24/2019).

Vijayakrishnan P, Josef Pieprzyk, and Huaxiong Wang. “Formal Security Analysis of Australian e-Passport Implementation”. In: Proceedings of the Sixth Australasian Conference on Information Security - Volume 81. AISC ’08. Wollongong, NSW, Australia: Australian Computer Society, Inc., 2008, pp. 75–82.