Insights into Digital Identity Dynamics through Personal Digital Twins

Pol Hölzmer
University of Luxembourg, SnT
pol.hoelzmer@uni.lu

Muriel Frank
University of Luxembourg, SnT
muriel.frank@uni.lu

Johannes Sedlmeir
University of Luxembourg, SnT
johannes.sedlmeir@uni.lu

Manuel Koschuch
FH Campus Wien, Research Center IT-Security
manuel.koschuch@edu.fh-campuswien.ac.at

Abstract

The digital transformation has resulted in a web of complex systems, creating a tangled mesh of context-specific digital identities for users to manage. Despite the prevalence of privacy risks, users often remain unaware of the extent to which their personal data is used by identity providers and relying parties. Drawing on design science and network analysis techniques, this paper introduces a hierarchical node architecture and data flow model to highlight the interconnectedness of partial digital identities and provides a tool to support the assessment and evaluation process. We used the developed methods to evaluate the digital footprint of ten participants, gaining insights into digital identity dynamics.

Keywords: digital identity, digital footprint, privacy, quantification, network theory

1. Introduction

Digital transformation is ubiquitous in today's world, bringing changes that affect almost every aspect of life (Mahraz et al., 2019). The multifaceted digital landscape pushes individuals to establish several identities to effectively navigate different areas of the digital infosphere (Lengsfeld, 2019). As the digital world is hard for the human eye to observe, many aspects of this transformation and its impact remain invisible (Demirkan et al., 2015). In particular, the digital identity lifecycle (create, use, update, delete, manage) is concealed and intangible, making it challenging for users to identify the risks of something that lacks immediate, visible consequences (Soh et al., 2024). Users often do not realize the extent of data they share with others, how this data is utilized to create detailed digital profiles (Christl, 2017), and especially the risks it poses to their privacy (Liyanaarachchi, 2020). The concept of privacy is defined here as the protection of an individual's rights against the exploitation of (personal) information (Westin, 2003). Even when faced with privacy violations, individuals often put forward the "I've got nothing to hide" argument (Solove, 2007). Furthermore, individuals often state they value privacy, yet their actions indicate a willingness to disclose personal data for little in return or neglect steps to safeguard privacy. This phenomenon is known as privacy paradox, wherein users prioritize ease of use over robust privacy protection (Awad & Krishnan, 2006; Solove, 2020).

We focus on the problem of digital identities being largely invisible to users, resulting in a lack of understanding of corresponding risks and their impact on privacy. Privacy is considered a fundamental human right (United Nations, 1948), whose implementation has evolved significantly, expanding in geographical extent, legal coverage, and enforcement (Seubert & Becker, 2021). Nevertheless, the way current digital technologies and corporations operate requires individuals to provide all kinds of information to gain access to services that can be collected and connected to them (Christl, 2017). Personal data has become the target of modern business models and a product for monetization. Until privacy-enhancing technologies (PETs) mature (Polonetsky & Sparapani, 2021) or there is a shift in business models (Stark et al., 2016), many users will prioritize ease of use and efficiency over robust privacy measures, inadvertently compromising their personal data due to a convenience-privacy tradeoff (Alashoor et al., 2022). Making users' choices about their digital identities visible can be instrumental (Roeber et al., 2015) because it creates awareness, which is key to navigating the current digital landscape and effectively mitigating privacy risks (Story et al., 2021). Therefore, this paper explores the following research question:

How to assess and visualize the dynamics of digital identity interrelations to provide a quantifiable, user-centric measure of privacy?

URI: https://hdl.handle.net/10125/109728
978-0-9981331-8-8

Building on an iterative design-and-evaluate design science research (DSR) process (Peffers et al., 2007), we contribute a hierarchical node architecture and data flow model, employing network theory-based methods for visualizing and quantifying privacy-related metrics. Our goal was to design a user-centric model and tool (Offermann et al., 2010; Winter, 2008) for capturing a representation of digital identities and their connections in a way that is both feasible to produce by and understandable for the generic end-users. We implemented a tool for assessment and evaluation, enabling the analysis of identity attributes while unveiling and contextualizing previously obscured relationships for privacy purposes.

2. Background

The concept of identity is constantly evolving, and different notions have emerged from the philosophical (Noonan, 2019), legal (Leenes, 2007), management (Bélanger & Crossler, 2011) and technical (ITU, 2009) domains. In this paper, we refer to digital identity as the information that identifies an entity in a particular context (Wilson & Hingnikar, 2019) and focus on informational privacy, which refers to individuals' rights to control the collection, use, and dissemination of personal data (A. Westin, 1968) to protect themselves from the collection, processing, storage, and disclosure of (personal) information (Bélanger & Crossler, 2011).

Individuals typically have multiple digital identities linked to different services or identity providers (Pfitzmann & Hansen, 2010). To refer collectively to these (partial) digital identities, researchers use the term digital persona (Clarke, 1994; Roosendaal, 2010), which encompasses all the information about an individual that exists online (Christl, 2017). Digital personas are shaped by the information an individual shares online and interactions with digital service providers.

The context-specific digital identities associated with a digital persona can be separated into two major types. A projected digital identity is created and controlled by individuals according to their preferences, like a personal email or social media account. On the other hand, an imposed digital identity is created by third parties based on information they collect, such as an identity issued by a government or employer (Clarke, 1994).

Generally, a digital identity consists of identifiers, credentials, and attributes (Pfitzmann & Hansen, 2010). Some of these change over time, while others are static. Digital identity attributes are characteristics associated with an identity. Identifiers (e.g., usernames, email addresses, and phone numbers) are special attributes that uniquely identify an entity in a given context. Lastly, credentials (e.g., passwords and authentication tokens) represent means to prove claims about identity or attributes (Sedlmeir et al., 2021). It's crucial to recognize the importance of identifier attributes as they can pose significant privacy threats if used for deanonymization or profiling across services. (Pfitzmann & Hansen, 2010). However, other attributes may also provide some degree of identifiability if combined, profiled, or if they happen to be unique to a single candidate in a given context (Sweeney, 2000).

A digital twin is an intricate virtual model that mirrors the real-time status of a physical entity, often linked to machinery. Unlike a digital persona, which is limited to identity information, a digital twin covers the entire lifecycle and interactions of the physical entity within a cyber-physical system, capturing its attributes, behaviors, and operational data to provide a comprehensive digital representation (Barricelli et al., 2019).

Data hubs are entities involved in everyday digital interactions that are difficult to avoid and controlled by data processors. Major data hubs include big tech platforms, financial services, telecom providers, and media companies, serving the consumer and analytics industry, creating a pervasive personal data ecosystem that represents hotspots for potential privacy risks (Christl, 2017).

Digital identity dynamics involve the relations and processes related to the identity lifecycle, such as creating, managing, and using digital identities (Bertino & Takahashi, 2011). This includes the social, ethical, and security implications of digital identities, such as privacy concerns and the impact of digital footprints. A crucial example of these dynamics is the flow of data towards data hubs, where personal information is aggregated and processed. Understanding these dynamics is essential for understanding digital identities over time and (digital) space. Moreover, network models can provide a powerful tool for describing and investigating such complex interactions.

Network theory has emerged as a useful approach for understanding the relational characteristics of complex systems by analyzing interconnection patterns among various network elements (Keast & Brown, 2005). Network theory can be applied to study socio-technical networks of nodes using node-link diagrams and matrices (Cutillo et al., 2011; Henry & Fekete, 2006).

The force-directed graph is a type of network visualization where a number of nodes (or vertices) are connected by links (or edges) (West, 2000). The term graph refers to an ordered pair $G = (N, E)$ of a set N of nodes and a set E of edges. In a directed graph $E \subseteq \{(x, y) \,|\, x, y \in N\}$, each link is represented by an ordered pair (x, y), i.e., the link (x, y) is distinct from (y, x) (McGuffin, 2012). Force-directed graphs use physical simulation algorithms to position nodes

based on their relational dynamics, creating intuitive visual representations. The forces are drawing related nodes together while pushing unrelated nodes apart.

Network analysis has introduced different notions to describe parts of a graph with specific properties: Cliques are sets of nodes $C \subseteq N$ in which every node is connected to any other node in this subgraph. Bicliques are (sub)graphs where all nodes of one of the partition subsets are connected to all nodes of the second subset and vice versa. Clusters are more general and can describe any set of nodes forming a subgraph defined by its clustering coefficient, which tells how close that subgraph is to a clique (complete graph) (West, 2000).

3. Research Method

We build on DSR (Gregor & Hevner, 2013; Hevner et al., 2004) to answer our research question. DSR is deeply rooted in information systems (IS) engineering (Winter, 2008) and particularly suitable for investigating interactions between social and technical systems (Baxter & Sommerville, 2011; Carlsson et al., 2011). We follow the DSR process by Peffers et al. (2007), which provides an iterative approach to creating rigorous and useful artifacts. The design of our artifact follows the problem-centered entry point of the DSR process model by Peffers et al. (2007). This involved multiple iterations with ten participants in a partially supervised experiment for data assessment and evaluation.

We identified the problem of insufficient support for individuals in understanding their digital identity dynamics. The motivation stems from the increasing complexity of digital interactions and the need for user awareness of their data usage and associated privacy risks. Our objective was to design a node architecture and data flow model that represents digital identities and their relationships with data hubs. The solution aims to enable users to visualize and quantify privacy-related metrics in an easily understandable, user-centric way.

As an artifact, we designed a new model (Offermann et al., 2010; Winter, 2008) to represent identity attributes and their connections using network theory. We developed a prototype tool implementing interactive processes for assessing, visualizing, and evaluating model instances. The model and tool were iteratively evaluated and refined over six months with ten participants who volunteered to contribute in exchange for learning more about their digital selves following an open talk on data protection at the university. Participants used the prototype tool to assess their digital identities, providing feedback on tool usability and model assessability, comprehensiveness, and accuracy. Each iteration involved participants collecting data about their digital identities, which was then pseudonymized, collected, and evaluated to identify patterns and validate the model's effectiveness. We systematically integrated feedback into each subsequent iteration, focusing on improving the clarity of the model, enhancing the user interface, and ensuring the model captured all relevant identity attributes. Visualization techniques helped identify points of interest within the digital footprints. In particular, the visualizations allow users to accurately represent and communicate their digital identity dynamics, providing insights into privacy risks. By involving users throughout the design process, a useful tool was developed to assess insights into digital identity dynamics that can be leveraged by future research.

4. Design and Development

This section presents our designed artifact, which we call the Personal Digital Twin (PDT) model (Offermann et al., 2010). The PDT is leaning on the concept of digital twins (Barricelli et al., 2019) but focuses specifically on the representation of individuals' digital identities and the actions that connect and bridge into the digital world. We designed the PDT as an artifact that embodies both model (the 'what') and process (the 'how') related aspects, central to crafting solution-oriented designs (Gregor & Hevner, 2013; March & Smith, 1995).

Several key questions guided our solution design: What abstract node architecture does the model represent? What connections exist between individuals' identity attributes? How do these attributes relate to their environment, particularly concerning privacy hotspots?

4.1. Abstract Node Hierarchy and Data Flows

We present the final version of the abstract node hierarchy and data flows of the PDT in Figure 1. This model depicts a directed network of nodes organized into two major layers: the identity attribute and data hub processor layer. Each layer contains nodes representing different aspects of the user's digital identity, along with their connections, highlighting the flow of data and the level of user control. This model serves as a generic starting point and can be adapted for additional scenarios.

Nodes represent abstract classes of information related to the user's PDT, with arrows indicating their relationships. This approach gives the dataset the characteristics of a directed (non-cyclic) graph and provides context for how data flows from the identity attribute layer at the top to the data hub processor layer at the bottom. The left axis represents the level of control. Users generally have more control over nodes in the identity attribute layer and less control over those in the data hub processor layer. The right axis shows the direction of

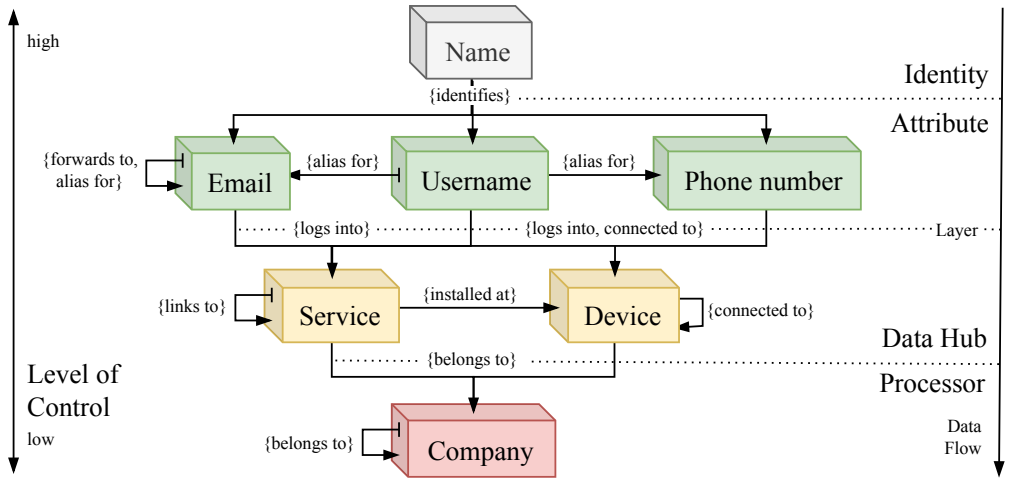

Figure 1: Abstract node hierarchy and data flow for the Personal Digital Twin.

the data flow, where user control decreases and privacy risks increase as data moves downward.

The layers in this system represent broad categories of node classes. The data flow represents the inevitable exchange of identity attributes with data processors. Data hubs are chosen for their potential to collect and process related information, reflecting the complex interactions within the digital ecosystem. This hierarchical node architecture effectively represents the PDT.

In the identity attribute layer, the nodes classes cover: Name, indicative of genuine identity and offering a straightforward method for aggregating nodes into clusters. It enables delineating services associated with an individual's legal name or pseudonyms, particularly in scenarios involving data obfuscation or anonymization strategies. Email, Username, and Phone Number, which serve as prevalent, long-lived, and strong identifiers (Jin et al., 2011), initiating the data flow as symbolic gateways upon identification. These attributes represent essential identifiers in current identity management (IdM) systems (Wilson & Hingnikar, 2019).

The data hub processor layer includes the following types of nodes: Services are of central importance as a major data aggregation point and act as a drain in the data flow. Users can decide which service they use but don't retain control over the data processing. The assessment of the trustworthiness of a service is at the discretion of the respective PDT owner, whereby this study takes an impartial stance towards all services. Devices and their operating systems are major data collection points. Beyond this, devices facilitate integrating multiple services, enabling the identification of usage patterns across different services and devices. As such, they provide a more comprehensive view of user interaction within the digital ecosystem. Company nodes serve to cluster data hubs, highlighting services or devices belonging to or provided by the same company. They reveal the extent of corporate influence on data flows and user interactions within their ecosystem.

Relations between node instances are designed to be chainable but not cyclical. Data can flow sequentially from one node to another but it cannot loop back to a previous node. This acyclic structure maintains the integrity of the data flow, accurately reflecting the unidirectional and non-repetitive nature of interactions within digital identity dynamics.

An individual's manifestation of the abstract model depicted in Figure 1 is (a subset of) their personal digital twin. Each PDT manifests the presented node architecture and, therefore, represents the base set of identity attributes that play a major role in digital identity management. A concrete example of node relations are: $\{Email\text{: }john@example.com\}$ *logs into* $\{Service\text{: }Facebook\}$; $\{Email\text{: }john@example.com\}$ *logs into* $\{Service\text{: }Instagram\}$; $\{Service\text{: }\{Facebook, Instagram\}$ *belongs to* $\{Company\text{: }Meta\}$; $\{Service\text{: }Instagram\}$ *installed at* $\{Device\text{: }Smartphone\}$. In this context, manifestation refers to the (somewhat) unique assortment of nodes and relations in a PDT. The abstract architecture provides guidelines on the types of nodes most commonly assessed and their relationships. Initially, the design included more nodes, such as passwords and hardware identifiers. However, we simplified the model using generalizations and focused on concepts that are more easily understood by users. Most internet users can relate to logging into a service with an email address and password, making these key elements of the model. While hardware identifiers and other technical details are crucial for tracking and privacy violations, assessing them requires high technical literacy and considerable time.

The final model balances completeness, accessibility, and the time required to create the PDT. Our goal was to support privacy quantification research while ensuring that the artifact remains accessible to the average internet user. We included only the most commonly assessed nodes and links, focusing on a minimal, overlapping subset to ensure participant comparability. However, the model remains extendable, allowing for the addition of nodes and links as needed. We aimed to make the process straightforward so that individuals with limited knowledge of digital identities could complete it independently, using information typically known to end-users. This approach makes the representation of identity dynamics easy to understand and assess.

4.2. Assessment and Evaluation Process

The assessment and evaluation processes revolve around an abstract architectural model. The assessment process involves creating an instance of the PDT model by individually evaluating the specific information outlined in the model. This requires more complex methods than using common survey techniques. The evaluation process uses advanced visualization and statistical methods derived from network theory to analyze the data and draw conclusions.

The assessment process begins with the identification and categorization of various identity attributes. Participants are guided to collect information about their PDT. This process is a highly manual task and requires a lot of reflection on what system each individual is using to be as complete as possible. Data collection is guided by the hierarchical node architecture, ensuring a comprehensive and systematic capture of the identity attributes and their relationships. Participants are supported by a custom-built tool that facilitates the entry and management of this data. The tool allows users to input their data through an intuitive interface that organizes information into node-relations. This method not only simplifies data entry but also the evaluation process. The evaluation process utilizes advanced visualization techniques and statistical methods derived from network theory. The data is analyzed and visualized primarily using force-directed graphs, which intuitively display the relationships and data flows between different nodes. These visualizations help uncover patterns, such as clusters and cliques and highlight potential privacy hotspots within the digital footprint.

4.3. Prototype Instantiation

The designed architecture model and processes have been implemented as `Python3`-based cross-platform graphical user interface using the `Qt` framework. The visualizations are written in JavaScript using `d3.js`. Furthermore, the tool stores the data in a local SQLite database file. The application is called `pdt` and is designed to provide transparent and offline data processing under the user's sole control. The source code and pre-built executables, as well as demo images, are available open-source[1]. By default, the prototype includes a pseudonymized PDT, named Alpha, further used in Section 7 to demonstrate the utility of the artifact without first performing the data collection for their own PDT.

5. Data Collection

We systematically evaluated the completeness and usefulness of our designed model, process, and tool. Over several months, we engaged ten individuals in assessing their PDT. We used selective sampling to recruit participants during an open university lecture on privacy risks. This allowed us to target individuals who were inherently interested in the topic. Anyone who expressed interest was given a personal introduction to our method. The participants were split into five groups of two participants each. Hence, we conducted five iterations, plus an initial ex-ante iteration performed by the authors. For every iteration, we ensured that the prototype application was installed correctly on the participants' devices and that they understood how to use it. Since the evaluation of the PDT required them to share some (pseudonymized) results and insights with us, we assured the participants that their responses were kept confidential, highlighting the importance of trust between us and the participants to conscientiously evaluate and discuss identity dynamics (Podsakoff et al., 2003). We merely used samples of the resulting network graph to discuss and communicate the results, which cannot be linked to the participants. The participants helped us evaluate the overall model by examining and contemplating their digital identities, giving insight into typical internet users' understanding of their digital footprint.

The participants completed the assessment process individually and by themselves over several days or even weeks. Individuals were encouraged to take all the time they needed to assess and reflect on their digital footprint to ensure that results would comprehensively represent their PDT. We were available to answer any questions and provided support during this period. We gathered the collected data in a way that protected the participants' privacy: After finishing their assessment, the participants used a function of the prototype tool to pseudonymize their data by replacing their personal information with generic node names (besides others) while still maintaining the relation. Afterward, the

[1] https://gitlab.com/personal-digital-twin/pdt

participants shared an exported version of their local database file with us for further evaluation.

The process of replacing each assessed piece of data with a generalized group representation and data recovery prevention in SQLite is what we refer to as pseudonymization. We obtained feedback on the node groups, relation types, usability, bugs, and the participants' motivation and understanding while assessing a complete PDT. Initially, we did not inquire about the details and dynamics of their PDT. Instead, we collected and analyzed every PDT and then provided feedback to the participants. This enabled us to further assess our design by analyzing network patterns of the pseudonymized PDT and asking participants about the accuracy of our observations and hypotheses.

The iterative DSR process allowed us to manage feedback effectively (Peffers et al., 2007), which greatly influenced the final design, particularly in the assessment's abstract node model and user interface. The ending conditions for each iteration included achieving a comprehensive representation of digital identities, ensuring the model's completeness, and confirming the tool's usability in enabling users to assess their data individually and accurately. We received positive feedback from the majority of participants, indicating they left the study with an overview and a deeper understanding of their digital footprint. Three participants expressed being shocked about the extent of their PDT and asked for advice on improving their privacy. The remaining participants had a neutral stance, acknowledging the state of their partially concerning digital footprints but not actively seeking ways to improve.

6. Data Analysis

During the evaluation phase, we collected ten different PDTs that provide valuable comparisons and insights into the uniqueness of identity dynamics. The results of all assessed PDTs indicate that, on average, each participant assessed 107 nodes spread across 6 node classes. The average number of email addresses and devices is 5 and 4, respectively, excluding Alpha with its 45 emails addresses. The count of assessed relations averages 158, with relations and services averaging 5 and 70, respectively. Email-related metrics show an average of 9 links per address, peaking at 46, while devices have an average of 8 links and a maximum of 21. These averages indicate a varied and interconnected structure within the PDT ecosystem, highlighting significant diversity in node and link distributions.

In this section, we focus on insights from three PDTs (*Alpha*, *Delta*, and *Eta*), demonstrating the artifact's usefulness for visualizing partial digital identities. Note that

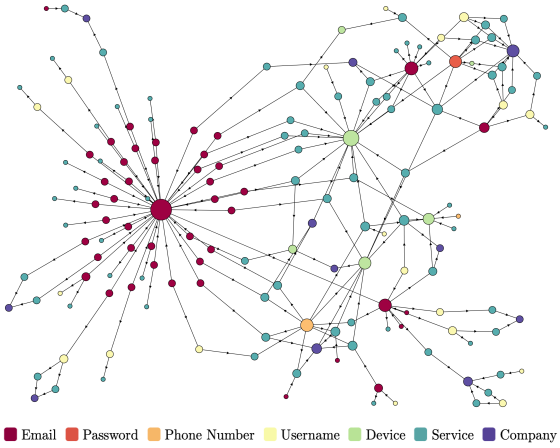

Figure 2: Alpha's unfiltered Personal Digital Twin.

there are other aspects of the PDTs that could be demonstrated. Furthermore, our tool can visualize eight more types of graphs from the same dataset and measure 137 different features for the study of the structure, dynamics, and functions of complex networks. We provide in-depth insights into what we consider the most important aspects in the following.

Figure 2 illustrates an unfiltered and cluttered force-directed graph of PDT *Alpha*. Clutter, characterized by overlapping links in the visualization, introduces noise and must be filtered to prevent distorting statistics and focusing exclusively on certain subgraphs. Subgraphs are subsets of nodes and edges, and filters are necessary to streamline points of interest during analysis, as demonstrated in Figure 3.

Network structures like clusters and cliques are crucial for evaluating identity dynamics. Here, partial identities are represented as clusters of interconnected nodes with any identity attribute connected to one or more data hubs in the middle. A highly interconnected cluster with a strong centrality factor and large cliques indicates poor privacy properties. Cliques should occur in a few cases and be relatively small. Bicliques are more common and indicate redundancy and multiple uses of the same attribute. These structures and more can be detected using network visualization and statistics.

We figured that it is possible to identify various types of partial identities easily through visual analysis. In Figure 3a, filtered by service, email address, and username, we pinpoint clusters of digital identities represented as star graphs, where an identity attribute is the central force of gravity. In the following, we present our categorization of partial identities based on this analysis.

The Personal Identity represents a critical subgraph, prominently at the center of force. We observed that most participants used a single email address linked to

many services, which is, therefore, easy to spot as a large star graph in the center. Alpha highlights a privacy focus through additional segregation into a public and private identity. Its public identity consists of a single mailbox with many aliases that forward to the private identity that is only used with trusted services.

The Business Identity is an essential subgraph representing an imposed identity since the employer primarily collects, enters, and manages the corresponding information. Many individuals receive at least one business email account or other identification means to access and work with the business's infrastructure and services. Alpha has strictly separated business and personal identity, while other candidates have entangled. This may pose a risk for businesses if personal behavior leads to, e.g., a data breach.

We identified three additional cluster types of identities: a legacy identity, which refers to remnants of past identities that were not adequately deleted; a local identity, which relates to an account that operates on a local device; and a fake identity, which is created with fabricated information to conceal one's real identity.

Delta, filtered by service, email address, username, and password, is illustrated in Figure 3b. This graph shows the value of an expandable model. The candidate added the relations {*Password*} *logs into* {*Service, Device*}. By adding this attribute, we can expand our analysis by security considerations visible through cliques between password usage and identities described before. Using only two unique passwords across multiple online and offline (see username links) services poses a considerable risk, as a breach of one service could potentially compromise multiple accounts. A corporate account could be compromised due to a personal data leak, especially in this scenario where the same passwords are shared between the personal and business identities. Thus, Delta's PDT underscores the importance of using unique passwords per service to enhance security.

Eta is depicted in Figure 3c and filted by email addresses and services. Eta consists of a massive and tangled web of services and email addresses, which is typical for many users. Eta has multiple accounts with numerous services and uses various email addresses. The high level of interconnectedness among different services and email addresses indicates a higher potential for privacy risks and a potential lack of privacy risk awareness. The dense network of relationships also suggests that losing access to one email account could have cascading effects and breach many other services. Eta's structure highlights a challenge many users face in managing their digital identities and being prone to extensive profiling, especially when creating accounts for various web services, is unavoidable.

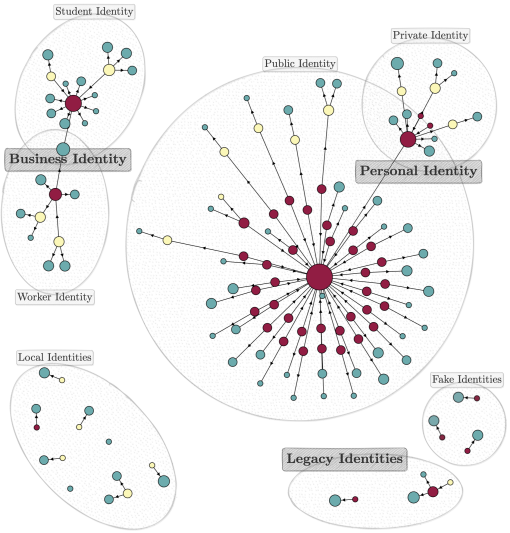

(a) Alpha sub-graph (Service, Email, and Username)

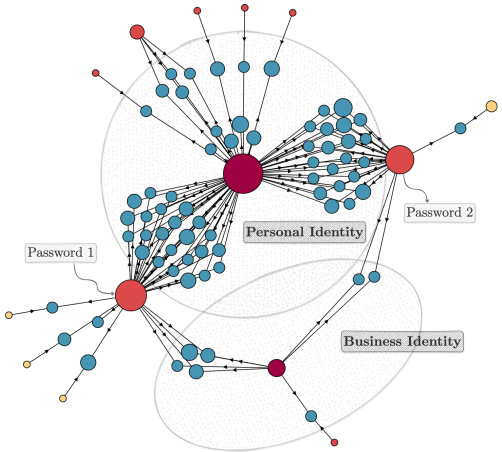

(b) Delta sub-graph (Service, Email, and Password)

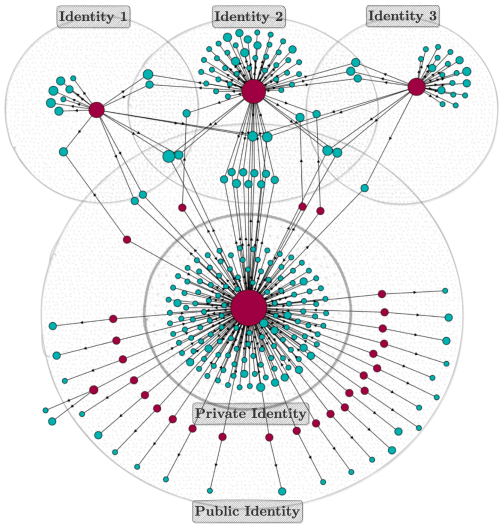

(c) Eta filtered by Service and Email

Email Password Phone Number Username Device Service Company

Figure 3: Demonstration of PDTs Alpha, Delta and Eta.

The analysis of these three PDTs – Alpha, Delta, and Eta – demonstrates the varied approaches and challenges in managing digital identities. Alpha exemplifies a highly organized and privacy-focused structure, Delta reveals potential security vulnerabilities related to password reuse, and Eta represents the typical complexity many users face when creating many accounts and losing track of their privacy and identity management.

7. Discussion

The PDT model introduced in this study offers a new approach to managing digital identities. It visually demonstrates how these identities are connected and measures privacy-related metrics. Our findings emphasize the significant diversity in digital identity dynamics among users, illustrating the complexity and uniqueness of individual digital footprints. An empirical evaluation with ten participants revealed diverse structures, ranging from highly organized and segregated identities to intricate webs of interconnected services and credentials. This diversity highlights the need for user-centric digital identity management solutions. The PDT model effectively increased participants' privacy awareness by making relationships and privacy hotspots visible. The visualizations help users gain insights into their digital identities, promoting better privacy and security practices. Overall, the PDT model successfully addresses the challenges identified in the design phase, demonstrating its utility and effectiveness.

7.1. Limitations

Considerations were made regarding whether and which PDT artifacts could be assessed through automated processes, which was the most frequently received feedback. The most straightforward answer is found in the definition of the PDT, which states that it is composed of information only known to the individual. We could not identify a way to reliably and holistically collect this information. Some options were considered, but none could outperform the manual assessment after weighing the costs and benefits. Complexity, limited coverage, and high error rates limit the usefulness of automated processes. After all, individuals should also learn and reflect during manual processing, which is crucial for increasing awareness. Nevertheless, we identified the following starting points for (machine-supported) assessment guidance: A *password manager* can be a helpful starting point for assessing an individual's data management practices. It stores structured data such as usernames, email addresses, passwords, and services. Example relations are $\{$*Username, Email*$\}$ *logs into* $\{$*Service*$\}$. Generating a list of *owned devices*, including smartphones, laptops, desktop PCs, wearables, and IoT devices, is helpful to identify further which applications are permanently linked to a service account. This allows us to get the relation between services and the devices, like $\{$*Service*$\}$ *installed at* $\{$*Device*$\}$. The browser or email history can be searched for used services, saved credentials, and registration emails. This approach can assist individuals in identifying used services and exposed information that may not be obvious or long forgotten. Furthermore, the manual assessment was meant also to help the user become aware of the full extent of the PDT. Automated methods may jeopardize the individual's highly valued privacy and control and reduce trust in the tool.

Our node model's design is built for flexibility and scalability, allowing for the seamless integration of new nodes and relationships without disrupting the existing structure. Each node represents a unique aspect of digital identity and is connected through well-defined relationships reflecting real-world interactions. The model's modularity enables the easy integration of new digital services and identity attributes by adding new nodes and defining their connections. Its flexibility in categorizing and grouping nodes makes it a versatile tool for a wide range of privacy assessment and management applications. Some evaluation participants leveraged the extendable abstract architecture of our model. For example, users added "password" as an identity attribute (i.e., credential), highlighting security risks associated with leaked passwords and identifying affected related services and identity attributes. This flexibility future-proofs the model and enhances its utility and relevance in the rapidly evolving digital landscape.

7.2. Future Work

This study has laid the groundwork for deploying the PDT model at a larger scale, increasing its generality by acquiring and assessing more and higher-quality data. While current research provides initial insights, launching new studies to collect more qualitative data is essential for creating universally applicable conclusions. We aim to streamline the analysis of PDTs and provide automated conclusions on their completeness, privacy, and security maturity while also giving guidance on how individuals can learn to improve on these aspects. To achieve this, we target to use features from advanced statistics and machine learning to augment the analysis, moving beyond the simplicity of initial statements. Automated data analysis and advanced analytics will be crucial in scrutinizing these metrics, pushing the boundaries of how privacy can be quantified and apprehended in the context of PDTs.

Our research on PDTs aims to further improve our understanding of privacy risks in digital identity management through advanced processing techniques that provide deeper insights into privacy maturity at scale. While the presented visualization method is useful for context-based manual evaluation, it still requires a lot of understanding of how to interpret it, and there is room for improvement. By broadening our focus beyond the individual level and focusing on visualization and automated statistical network analyses, we can use additional metrics (e.g., degree centrality or clustering coefficient (Cutillo et al., 2011)) that offer additional features for the quantification of privacy and security risks.

8. Conclusion

The PDT node architecture presents a structured, user-centric, and extensible approach to visualizing digital identities. This architecture utilizes a network of interconnected nodes representing various facets and their distinct role in mapping out the digital footprint. By employing graph theory, specifically directed graphs, the architecture enables the visualization of complex relationships and data flows between these nodes. This visualization enhances our understanding of individuals' digital identities and facilitates the identification of potential privacy risks and areas requiring better management. The feasibility of this method is underscored by its ability to transform abstract digital data into tangible, user-friendly visual representations, making it an effective tool for awareness and proactive privacy and identity management.

Acknowledgments

This work was funded by Luxembourg's National Research Fund (FNR) and PayPal (PEARL grant ref. 13342933/Gilbert Fridgen, and PABLO grant ref. 1632675), and supported by Luxembourg's Ministry for Digitalisation. For open access purposes, the authors have applied a CC BY 4.0 license to any Author Accepted Manuscript arising from this submission.

References

Alashoor, T., Keil, M., Smith, H. J., & McConnell, A. R. (2022). Too Tired and in Too Good of a Mood to Worry About Privacy: Explaining the Privacy Paradox Through the Lens of Effort Level in Information Processing. *Information Systems Research*.

Awad, N. F., & Krishnan, M. S. (2006). The Personalization Privacy Paradox: An Empirical Evaluation of Information Transparency and the Willingness to Be Profiled Online for Personalization. *MIS Quarterly*, *30*(1).

Barricelli, B. R., Casiraghi, E., & Fogli, D. (2019). A Survey on Digital Twin: Definitions, Characteristics, Applications, and Design Implications. *IEEE Access*, *7*.

Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, *23*(1).

Bélanger, F., & Crossler, R. E. (2011). Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. *MIS Quarterly*, *35*(4).

Bertino, E., & Takahashi, K. (2011, September). *Identity Management: Concepts, Technologies, and Systems*. Artech House Publishers.

Carlsson, S. A., Henningsson, S., Hrastinski, S., & Keller, C. (2011). Socio-technical IS design science research: Developing design theory for IS integration management. *Information Systems and e-Business Management*, *9*(1).

Christl, W. (2017, June). *Corporate Surveillance In Everyday Life. How Companies Collect, Combine, Analyze, Trade, and Use Personal Data on Billions* (tech. rep.).

Clarke, R. (1994). The digital persona and its application to data surveillance. *The Information Society*, *10*(2), 77–92.

Cutillo, L. A., Molva, R., & Onen, M. (2011). Analysis of Privacy in Online Social Networks from the Graph Theory Perspective. *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, 1–5.

Demirkan, H., Bess, C., Spohrer, J., Rayes, A., Allen, D., & Moghaddam, Y. (2015). Innovations with Smart Service Systems: Analytics, Big Data, Cognitive Assistance, and the Internet of Everything. *Communications of the Association for Information Systems*.

Gregor, S., & Hevner, A. R. (2013). Positioning and Presenting Design Science Research for Maximum Impact. *MIS Quarterly*.

Henry, N., & Fekete, J.-d. (2006). MatrixExplorer: A Dual-Representation System to Explore Social Networks. *IEEE Transactions on Visualization and Computer Graphics*, *12*(5).

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *Management Information Systems Quarterly*, *28*(1), 75.

ITU. (2009, January). Y.2720 : NGN identity management framework. Retrieved November 10, 2023, from https://handle.itu.int/11.1002/1000/9574

Jin, L., Takabi, D., & Joshi, J. (2011). Analyzing Security and Privacy Issues of Using E-mail Address as Identity. *International Journal of Information Privacy, Security and Integrity*, *1*.

Keast, R., & Brown, K. (2005). The network approach to evaluation: Uncovering patterns, possibilities and pitfalls.

Leenes, R. E. (2007). Do They Know Me? Deconstructing Identifiability.

Lengsfeld, J. (2019). *Digital Era Framework*.

Liyanaarachchi, G. (2020). Online privacy as an integral component of strategy: Allaying customer fears and building loyalty. *Journal of Business Strategy*.

Mahraz, M.-I., Berrado, A., & Benabbou, L. (2019). A Systematic literature review of Digital Transformation.

March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*.

McGuffin, M. J. (2012). Simple algorithms for network visualization: A tutorial. *Tsinghua Science and Technology*, *17*(4), 383–398.

Noonan, H. (2019, February). *Personal Identity* (3rd edition). Routledge.

Offermann, P., Blom, S., Schönherr, M., & Bub, U. (2010). Artifact Types in Information Systems Design Science – A Literature Review. In R. Winter, J. L. Zhao, & S. Aier (Eds.), *Global Perspectives on Design Science Research* (pp. 77–92). Springer.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, *24*(3), 45–77.

Pfitzmann, A., & Hansen, M. (2010, August). A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management.

Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *The Journal of Applied Psychology*, *88*(5), 879–903.

Polonetsky, J., & Sparapani, T. (2021). A Review of the Privacy-Enhancing Technologies Software Market. *IEEE Security & Privacy*, *19*(6), 119–122.

Roeber, B., Rehse, O., Knorrek, R., & Thomsen, B. (2015). Personal data: How context shapes consumers' data sharing with organizations from various sectors. *Electronic Markets*, *25*(2), 95–108.

Roosendaal, A. (2010). Digital Personae and Profiles as Representations of Individuals. In M. Bezzi, P. Duquenoy, S. Fischer-Hübner, M. Hansen, & G. Zhang (Eds.), *Privacy and Identity Management for Life* (pp. 226–236). Springer.

Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering*, *63*, 603–613.

Seubert, S., & Becker, C. (2021). The Democratic Impact of Strengthening European Fundamental Rights in the Digital Age: The Example of Privacy Protection. *German Law Journal*, *22*(1), 31–44.

Soh, S., Talaifar, S., & Harari, G. M. (2024). Identity development in the digital context. *Social and Personality Psychology Compass*, *18*(2), e12940.

Solove, D. J. (2007, July). 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy.

Solove, D. J. (2020). The Myth of the Privacy Paradox. *GW Law Faculty Publications & Other Works*.

Stark, L., King, J., Page, X., Lampinen, A., Vitak, J., Wisniewski, P., Whalen, T., & Good, N. (2016). Bridging the Gap between Privacy by Design and Privacy in Practice. *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 3415–3422.

Story, P., Smullen, D., Yao, Y., Acquisti, A., Cranor, L. F., Sadeh, N., & Schaub, F. (2021). Awareness, Adoption, and Misconceptions of Web Privacy Tools. *Proceedings on Privacy Enhancing Technologies*, *2021*(3), 308–333.

Sweeney, L. (2000). Simple Demographics Often Identify People Uniquely. *Health*, *671*.

United Nations. (1948). Universal Declaration of Human Rights. Retrieved November 17, 2023, from https : / / www . un . org / en / about - us / universal - declaration-of-human-rights

West, D. B. (2000, January). *Introduction to Graph Theory*. Pearson College Div.

Westin. (2003). Social and Political Dimensions of Privacy. *Journal of Social Issues*, *59*(2), 431–453.

Westin, A. (1968). Privacy And Freedom. *Washington and Lee Law Review*, *25*(1), 166.

Wilson, Y., & Hingnikar, A. (2019). *Solving Identity Management in Modern Applications: Demystifying OAuth 2.0, OpenID Connect, and SAML 2.0*.

Winter, R. (2008). Design science research in Europe. *European Journal of Information Systems*, *17*(5), 470–475.