FoT-PDS: A User-Centric Paradigm for Privacy-Preserving IoT

George P. Pinto[1,2], **Cássio V. S. Prazeres**[1]

[1]Institute of Computing – Federal University of Bahia (UFBA)
Salvador – BA – Brazil

[2]Federal Institute of Bahia (IFBA)
Salvador – BA – Brazil

georgepacheco@ifba.edu.br, prazeres@ufba.br

Abstract. *The adoption of IoT technologies has amplified concerns regarding personal data privacy, as the continuous collection and processing of personal data expose users to structural privacy risks, including identification, localization and tracking, profiling, and linkage. Current IoT platforms primarily rely on service-centric models, offering limited user control, transparency, and support for informed consent throughout the data lifecycle. This thesis addresses these limitations by proposing FoT-PDS, a user-centric privacy-preserving IoT paradigm that integrates Personal Data Stores and the Fog of Things. The paradigm reassigns data control from service providers to users, enabling decentralized personal data management, fine-grained access control, transparency, and explicit consent. A central component of the paradigm is an AI-assisted consent mechanism that supports users in making informed consent decisions by analyzing their own data. The mechanism leverages clustering methods to evaluate potential profiling risks from users' personal data. FoT-PDS was implemented and evaluated through both technical and empirical studies. Technical experiments demonstrate the feasibility of the consent mechanism under typical fog and IoT constraints, while a user-centric empirical study shows that the paradigm improves perceived data control, transparency, and privacy awareness. Trust is indirectly supported through increased privacy awareness, highlighting its mediating role in privacy-preserving IoT systems. These insights provide evidence that FoT-PDS can effectively mitigate privacy risks in IoT environments.*

1. Introduction and Research Context

The Internet of Things (IoT) technologies have enriched everyday activities by enabling interconnected devices to generate and collect increasing amounts of data from physical environments, thereby supporting the creation of more personalized and valuable services [Ashton 2009]. Different applications, ranging from personal health monitoring to smart urban services, can be developed by collecting, processing, and sharing sensor data across heterogeneous infrastructures. While this data-driven model has expanded both the range and quality of services, it has also amplified concerns regarding the protection of personal data.

IoT platforms typically rely on centralized data management architectures, in which data is transferred to service providers, where storage, processing, and use occur

with limited transparency and user control. This contributes to increasing users' exposure to privacy risks, including identification, localization and tracking, profiling, and data linkage [Pinto et al. 2024]. As a consequence, the benefits offered by IoT services often coexist with privacy risks, reinforcing what has been described in the literature as the *privacy paradox* [Kokolakis 2017].

At the same time, users increasingly express concerns about their data privacy. In response, user control has been increasingly recognized as a key principle for privacy protection. However, existing IoT platforms provide users with limited means to effectively understand or influence how their data is collected, processed, and shared over time. Empirical evidence suggests that insufficient control is often accompanied by reduced transparency and low privacy awareness, which negatively affects users' trust and engagement in data-driven services [CISCO 2019, Mugariri et al. 2022, Zheng et al. 2018]. Despite this recognition, IoT platforms still provide limited practical mechanisms for users to exercise control over their data.

Within this context, Personal Data Stores (PDS) [Verborgh 2023], a user-centric approach, have emerged as a promising direction to address this gap by decoupling data storage from service provision and restoring control to data owners. In parallel, the Fog of Things (FoT) [Prazeres and Serrano 2016] paradigm has been proposed to decentralize computation and data management by bringing processing capabilities closer to data sources. Combining these paradigms offers an opportunity to decentralize data management and reassign control to data owners.

This thesis addresses this gap by proposing FoT-PDS [Pinto and Prazeres 2024], a user-centric paradigm for privacy-preserving IoT that integrates Personal Data Stores within a Fog of Things architecture. By combining decentralized processing with individualized data repositories, FoT-PDS redefines data governance in IoT systems, enabling fine-grained access control, explicit consent enforcement, and transparency throughout the IoT data lifecycle. Privacy is thus operationalized through architectural design rather than treated solely as an external regulatory or policy requirement.

A distinguishing component of the FoT-PDS paradigm is the AI-assisted consent mechanism [Pinto et al. 2025] that supports users during data-sharing decisions. The mechanism analyzes users' personal IoT data using unsupervised learning techniques to estimate potential profiling risks. These risks are then communicated through interpretable metrics, enhancing user awareness without automating consent decisions. This approach directly addresses the limitations of traditional consent models, which often assume high levels of user expertise.

Beyond its technical contributions, this work advances a conceptual shift in how privacy is addressed in IoT environments. FoT-PDS positions users as active participants in the data lifecycle by integrating decentralized storage, semantic data representation, and intelligent consent support into a coherent paradigm tailored to distributed IoT systems. The proposed approach is validated through both technical evaluation and user-centered empirical analysis, demonstrating improvements in perceived data control, transparency, and privacy awareness, as well as the feasibility of AI-assisted consent in decentralized settings.

2. Research Problem, Motivation, and Questions

Privacy has become a central topic in IoT research. However, most existing platforms still address privacy protection as an auxiliary concern rather than a core architectural requirement. In practice, privacy is often enforced through isolated mechanisms such as encryption, access control, or regulatory compliance, without reconsidering how data governance is implemented across distributed IoT systems [Pinto et al. 2024].

A fundamental limitation lies in the data management model. Personal data generated at the network edge is transferred to infrastructures controlled by service providers, where storage, processing, and sharing decisions are made without users' knowledge. Once data leaves the device, users typically lose effective control over how their data is reused, combined with other data sources, or retained over time. This situation is particularly critical in distributed IoT environments, where data flows in multiple contexts, computational layers, and domains.

Consent mechanisms are another concern, as they are implemented as static and coarse-grained processes, often based on all-or-nothing agreements that are poorly aligned with the dynamic and continuous nature of IoT interactions. Users are therefore expected to make informed decisions without adequate knowledge about privacy risks. As a result, consent becomes a formal requirement rather than an effective instrument for privacy protection.

From a user perspective, this lack of meaningful control is compounded by limited transparency and low privacy awareness. Empirical evidence suggests that when users are unable to understand how their data is collected, processed, and shared, their perceived control decreases, negatively impacting trust in IoT services.

Motivated by these limitations, the main research question guiding this thesis was: ***Can a user-centric PDS-based approach mitigate personal data privacy risks in IoT environments?***.

To address this central question, the thesis was guided by the following questions. The first three questions adopt a user-centered perspective, focusing on how users perceive and respond to privacy-related dimensions when interacting with the FoT-PDS platform. The last question presents a technical perspective on the AI-assisted consent mechanism proposed in this thesis, focusing on its ability to detect privacy risks through the analysis of personal sensor data.

- **RQ1:** To what extent does the FoT-PDS improve users' perception of data control over collecting, storing, and sharing personal data?
  Answer: The experimental results show that the FoT-PDS significantly improves users' perceived control over their personal data. By decentralizing data storage and enforcing consent-based access through PDSs, users report greater awareness and authority over how their data are collected, stored, and shared.
- **RQ2:** How does the FoT-PDS influence user's perceptions of transparency, privacy awareness, and trust?
  Answer: The results indicate that the FoT-PDS has a positive and direct effect on users' perceptions of transparency and privacy awareness. Additionally, the paradigm indirectly enhances trust in IoT services, demonstrating that user-centric data control mechanisms improve users' confidence in how their data are handled.

- **RQ3:** To what extent do increases in transparency and privacy awareness contribute to users' trust in IoT services?
  Answer: The findings confirm that increases in transparency and privacy awareness significantly contribute to users' trust. Privacy awareness plays a mediating role, indicating that trust is strengthened when users both understand and perceive control over data processing practices.
- **RQ4:** To what extent can the AI-assisted consent mechanism assess the risk of user profiling and support informed consent decisions in IoT environments?
  Answer: The technical evaluation demonstrates that the AI-assisted consent mechanism is capable of assessing profiling risks by analyzing personal sensor data using clustering-based techniques. The generated indicators provide meaningful support for informed consent decisions, enabling users to better understand potential privacy risks before sharing their data.

3. Fog of Things and Personal Data Store Paradigm Overview

The FoT-PDS paradigm represents a structural rethinking of how personal data are governed in IoT environments. Rather than treating privacy as an external requirement enforced through isolated mechanisms, FoT-PDS embeds privacy-preserving principles directly into the architecture of distributed IoT systems. The paradigm is grounded in a user-centric perspective, in which individuals retain control over how their personal data is stored, processed, and shared throughout the data lifecycle.

Traditional IoT architectures are predominantly service-centric, relying on centralized data storage and unclear data flows that limit user control, awareness, and trust. In contrast, FoT-PDS shifts data management toward individuals by integrating PDS into a FoT architecture. This integration enables decentralized data management at the network edge while preserving interoperability and scalability across heterogeneous IoT domains.

At a conceptual level, FoT-PDS combines three fundamental principles. First, personal data is managed within individualized repositories, ensuring that data ownership and decision-making remain aligned. Second, data processing is decentralized and positioned as close as possible to data sources, reducing dependence on centralized cloud platforms. Third, all data access and sharing operations are mediated through explicit and fine-grained user consent.

From an architectural perspective, FoT-PDS extends the original FoT model by introducing PDS instances at both local and global network layers (Figure 1). The local network comprises smart devices, FoT gateways, and FoT-PDS servers. Smart devices are responsible for sensing and data generation, typically operating under strict resource constraints. FoT gateways act as intermediaries, performing lightweight preprocessing tasks such as aggregation, metadata enrichment, and temporary storage before data is forwarded. FoT-PDS servers provide more robust storage and processing capabilities and host individualized PDS instances, where semantically enriched personal data is persisted under user control.

Data management within FoT-PDS follows a semantic approach to enhance interoperability and machine interpretability. Sensor observations are annotated using domain ontologies and represented as structured triples, enabling queries, reasoning, and integration across heterogeneous applications. This semantic enrichment supports consistent

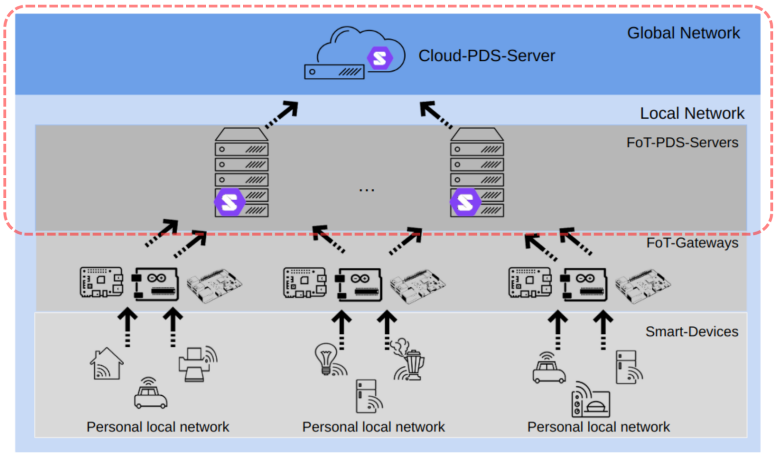

Figura 1. FoT-PDS paradigm.

data interpretation while preserving user-defined access constraints.

In addition to local data management, FoT-PDS supports selective synchronization with cloud-based PDS instances. Users may choose to replicate specific data items to remote repositories hosted in cloud infrastructures, either managed directly by the user or delegated to trusted third parties. This selective synchronization mechanism reinforces data minimization principles and enables cross-domain data portability without compromising user control.

A defining characteristic of the FoT-PDS paradigm is its focus on the entire IoT data lifecycle. Privacy risks are not limited to the moment of data generation but often emerge or intensify during storage, processing, and sharing stages. By enforcing explicit consent and decentralized control at each of these stages, FoT-PDS mitigates key IoT privacy threats, including identification, localization and tracking, profiling, and data linkage.

While FoT-PDS enforces user control through architectural mechanisms, it also recognizes the cognitive challenges associated with managing privacy decisions in data-intensive environments. For this reason, the paradigm incorporates an AI-assisted consent mechanism, which supports users by assessing potential profiling risks associated with data sharing requests. This mechanism complements, rather than replaces, user decision-making and is discussed in detail in the following section.

Overall, FoT-PDS establishes a coherent paradigm for privacy-preserving IoT that aligns distributed system design with user-centric data governance. By integrating decentralized storage, semantic data management, and consent enforcement into a unified architecture, the paradigm provides a foundation for trustworthy IoT ecosystems.

3.1. AI-Assisted Consent Mechanism

The FoT-PDS paradigm enforces explicit consent as a core requirement for accessing personal data in IoT environments, positioning users as active decision-makers in the data sharing. While this approach reinforces data sovereignty, it also introduces practical challenges. Granting meaningful consent requires users to understand not only the purpose

of data usage but also the potential risks associated with them. Prior studies indicate that relying solely on users to manage such complex decisions is ineffective in practice, particularly in data-intensive environments characterized by continuous data generation and opaque analytics [Acquisti et al. 2020].

To address these limitations, FoT-PDS incorporates an AI-assisted consent mechanism designed to support informed decision-making. The mechanism provides users with interpretable indicators that estimate the potential exposure to profiling risks before data disclosure occurs. This approach complements explicit consent by increasing privacy awareness at the moment data access is requested.

The mechanism operates on users' data stored within their respective PDS and is deployed at both the FoT-PDS-Server and Cloud-PDS-Server layers. To preserve user privacy, all analyses are performed on isolated, per-user datasets, ensuring that the system itself introduces no cross-user aggregation or external profiling. The goal is not to infer or label specific user profiles, but to assess whether the structure of the available data may facilitate profiling by third parties.

From a technical perspective, the mechanism relies on unsupervised learning techniques, specifically clustering algorithms, to analyze patterns in users' sensor data. Clustering is well-suited to this context, as personal IoT data typically lacks ground truth labels and exhibits high dimensionality and variability. By identifying cohesive and well-separated data groupings, the mechanism estimates the possibility that distinct behavioral or contextual profiles could be inferred from the data.

To evaluate clustering quality and derive a profiling risk indicator, the mechanism employs internal validation metrics that assess intra-cluster cohesion and inter-cluster separation. Among these metrics, the Silhouette index is adopted as the primary basis for the Profile Metric (PM) estimation due to its interpretability and suitability for per-user data analysis. Complementary other indices (Davies-Bouldin and Calinski–Harabasz) are used during algorithm selection to improve robustness, but the final PM is derived from the normalized Silhouette score of the best-performing clustering configuration.

The resulting PM is translated into qualitative risk levels and presented to users through the consent interface. Higher values indicate that the data exhibits clear and separable structures, suggesting a greater potential for profiling if shared, whereas lower values indicate more diffuse data distributions. Importantly, the metric does not measure profiling risk in an absolute sense; instead, it acts as an indicator of potential exposure based on structural properties of the data. This distinction reinforces the role of the mechanism as a decision-support tool rather than a deterministic privacy assessor.

By integrating AI-assisted consent into the FoT-PDS architecture, the paradigm addresses a key limitation of traditional consent models: the assumption that users possess sufficient knowledge to anticipate privacy risks. The proposed mechanism enhances transparency and privacy awareness while preserving user control, thereby operationalizing informed consent as a practical component of privacy-preserving IoT systems.

4. Research Methodology

This research adopts a mixed-methods approach that combines architectural design, platform implementation, empirical investigation, and technical evaluation. The methodology

was structured to assess both the feasibility of a user-centric IoT paradigm and its effects on user perception and behavior.

First, we conducted a systematic mapping of the literature to identify prevailing privacy threats in IoT systems and assess existing solutions based on Personal Data Stores, fog computing, and related paradigms. The findings of this mapping informed the design requirements of the FoT-PDS architecture and the consent mechanism [Pinto et al. 2024], [Pinto and Prazeres 2025b].

Subsequently, the FoT-PDS paradigm was specified following a design-oriented research approach and realized as a functional platform [Pinto and Prazeres 2025a]. This implementation includes mechanisms for decentralized data management, semantic data enrichment, fine-grained access control, and AI-assisted consent support. A technical evaluation was performed to assess the feasibility of the consent mechanism, considering computational constraints typical of fog and IoT environments.

Finally, a user-centered empirical study was conducted to evaluate the impact of FoT-PDS on key privacy-related constructs, including data control, transparency, privacy awareness, and trust. Quantitative data were collected through controlled experiments and analyzed using statistical techniques to validate the proposed research model and examine the relationships among these constructs [Pinto et al. 2026].

5. Scientific Contributions

This thesis advances the state of the art in privacy-preserving IoT by proposing, formalizing, and empirically validating a user-centric paradigm that redefines how personal data is managed in distributed environments. Rather than addressing privacy through isolated technical mechanisms, the contributions of this work lie in establishing privacy as an architectural and technical property of IoT systems.

The primary scientific contribution is the FoT-PDS paradigm, an original integration of PDS into the FoT architecture. This paradigm bridges two previously independent research lines, user-centric data management and distributed IoT infrastructures, by demonstrating that decentralized processing can coexist with continuous user control over personal data. By shifting data governance from service-centric to user-centric models, FoT-PDS establishes a new architectural framework for privacy-preserving IoT systems, promoting control, transparency, privacy awareness, and trust throughout the data lifecycle.

A second contribution is the conceptualization and realization of an AI-assisted consent mechanism that operationalizes informed consent in IoT environments. Unlike traditional consent models based on static policies or predefined rules, the proposed mechanism leverages unsupervised learning techniques to estimate profiling risks from users' personal data. This work makes a scientific contribution, utilizing artificial intelligence as a tool for privacy risk awareness and mitigation, thereby enabling context-aware and user-informed consent decisions.

From an architectural and data management perspective, the thesis contributes a fully specified and validated model for decentralized personal data management in IoT. This includes the integration of semantic data annotation, individualized data repositories, fine-grained access control, and selective synchronization between local and cloud-based

PDS instances.

Beyond architectural innovation, this thesis makes a methodological contribution by combining technical evaluation with a user-centered empirical study. The empirical analysis provides evidence that data control directly influences transparency and privacy awareness, which in turn strengthens user trust.

Collectively, these contributions advance the scientific discourse on privacy-preserving IoT systems by demonstrating that adequate privacy protection requires the integration of distributed architectures, user-centric data management, and intelligent consent support. The results of this thesis provide both conceptual and empirical foundations for future research on privacy in IoT ecosystems.

6. Related Works

Privacy is a significant challenge in IoT ecosystems due to the pervasive collection, processing, and sharing of personal data. Several studies concentrated on identifying privacy threats such as identification, localization and tracking, profiling, linkage, and information leakage, as well as on proposing mitigation strategies grounded in security mechanisms or regulatory compliance [Pinto et al. 2024].

Other studies have explored privacy-enhancing technologies (PETs) in IoT, including anonymization, encryption, access control, and policy-based enforcement mechanisms [Safa et al. 2022]. While these solutions help protect data during specific stages of the IoT data lifecycle, they typically rely on service-centric architectures that collect and process data outside the users' control. As a result, users often lack transparency and control over how their data is stored, processed, and shared, negatively impacting their privacy awareness and trust in service providers.

To overcome these limitations, PDS has emerged as a user-centric approach that emphasizes data control as a foundational principle for privacy protection. In this context, PDS has been proposed as a paradigm that promotes user control. Several works have demonstrated the feasibility of applying PDS platforms, particularly Solid, to IoT-related scenarios, including industrial environments [Bader and Maleshkova 2020], healthcare systems [Ghayvat et al. 2022], and smart environments [Boi et al. 2023].

Despite these advances, existing PDS-based IoT solutions primarily focus on access control and data management, often assuming that privacy is ensured once user authorization is granted. Consent is typically treated as a static, binary decision, implicitly handled through Solid's access control mechanisms, without supporting users in understanding the potential privacy risks associated with data disclosure [Fries et al. 2023].

In contrast to existing work, this thesis advances the state of the art by integrating FoT and PDS to establish a user-centric paradigm in which privacy is treated as an architectural property of IoT systems. FoT-PDS enables privacy protection closer to data sources, promoting user control over personal data, and enhancing transparency, privacy awareness, and trust. Besides that, unlike relying on implicit or static consent mechanisms, our paradigm introduces an AI-assisted consent mechanism that supports informed data-sharing decisions by estimating potential profiling risks from users' own data and presenting them through interpretable indicators.

7. Conclusion

This thesis examined how personal data privacy risks in IoT environments can be mitigated through a user-centric approach. To address this challenge, the FoT-PDS paradigm combines FoT and PDS to redefine the data management. By shifting control from service-centric platforms to individual users, FoT-PDS embeds privacy-preserving principles directly into its architecture, mitigating risks such as identification, profiling, localization, tracking, and linkage.

A central contribution of the proposed paradigm is the integration of an AI-assisted consent mechanism that supports informed consent decisions. We move from static or policy-based consent models to a model of real informed consent based on interpretable indicators of potential profiling risks derived from users' own data.

The technical evaluation of the paradigm demonstrates that this approach is computationally feasible in distributed and fog-based environments. At the same time, the empirical analysis reveals that FoT-PDS enhances users' perceived control over data, transparency, and privacy awareness. Trust is indirectly supported through increased privacy awareness, reinforcing its role as a key mediator in privacy-preserving IoT systems.

Overall, the results validate that a PDS-based, user-centric paradigm can effectively address privacy challenges in IoT environments. By combining architectural decentralization with consent support, FoT-PDS demonstrates that privacy-aware system design is a viable alternative to data management models.

Despite these contributions, this research also presents limitations. The empirical evaluation was conducted in controlled experimental settings, which may not fully capture the complexity of real-world IoT deployments. Although multiple data domains were considered, further studies are needed to evaluate behavioral effects and scalability in heterogeneous environments.

Regarding the AI-assisted consent mechanism, the profiling risk estimation is influenced by both the quality of the available data and the selection of clustering algorithms. While the proposed approach demonstrates feasibility, future work may explore alternative learning techniques, the usage of contextual information, and complementary strategies, such as federated learning.

Referências

Acquisti, A., Brandimarte, L., and Loewenstein, G. (2020). Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4):736–758.

Ashton, K. (2009). That 'Internet of Things' thing. RFID Journal (Expert Views). Published June 22, 2009.

Bader, S. R. and Maleshkova, M. (2020). SOLIOT—decentralized data control and interactions for IoT. *Future Internet*, 12(6).

Boi, B., De Santis, M., and Esposito, C. (2023). A decentralized smart city using solid and self-sovereign identity. In Gervasi, O., Murgante, B., Rocha, A. M. A. C., Garau, C., Scorza, F., Karaca, Y., and Torre, C. M., editors, *Computational Science and*

Its Applications – ICCSA 2023 Workshops, pages 149–161, Cham. Springer Nature Switzerland.

CISCO (2019). Consumer privacy survey: The growing imperative of getting data privacy right.

Fries, J., Freund, M., and Harth, A. (2023). A solid architecture for machine data exchange with access control. *Proceedings of the 1st Semantic Web on Constrained Things, Hersonissos, Greece*, 3412:74–81.

Ghayvat, H., Sharma, M., Gope, P., and Sharma, P. K. (2022). SHARIF: Solid pod-based secured healthcare information storage and exchange solution in Internet of Things. *IEEE Transactions on Industrial Informatics*, 18(8):5609–5618.

Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134.

Mugariri, P., Abdullah, H., García-Torres, M., Parameshchari, B. D., and Abdul Sattar, K. N. (2022). Promoting information privacy protection awareness for internet of things (iot). *Mobile Information Systems*, 2022(1):4247651.

Pinto, G. P., Donta, P. K., Dustdar, S., and Prazeres, C. (2024). A systematic review on privacy-aware iot personal data stores. *Sensors*, 24:2197.

Pinto, G. P. and Prazeres, C. (2024). Towards data privacy in a fog of things. *Internet Technology Letters*.

Pinto, G. P. and Prazeres, C. (2025a). A User-Centric IoT Platform for Privacy With AI-Assisted Consent. *IEEE Open Journal of the Computer Society*, 6:1834–1846.

Pinto, G. P. and Prazeres, C. (2025b). Data Privacy in the Internet of Things: A Perspective of Personal Data Store-Based Approaches. *Journal of Cybersecurity and Privacy*, 5(2).

Pinto, G. P., Sousa, N. R., Da Silva, C. N., Peixoto, M. L., Figueiredo, G. B., and Prazeres, C. V. (2025). Enhancing IoT data privacy: AI-assisted consent mechanism in a PDS-based solution. *Internet of Things*, 34:101807.

Pinto, G. P., Sousa, N. R., and Prazeres, C. V. S. (2026). My data, my rules: an experimental study on a user-centric approach to data privacy in the internet of things. *Computing*, 108(3):33.

Prazeres, C. and Serrano, M. (2016). SOFT-IoT: Self-Organizing FOG of Things. In *2016 30th International Conference on Advanced Information Networking and Applications Workshops*, pages 803–808.

Safa, N. S., Mitchell, F., Maple, C., Azad, M. A., and Dabbagh, M. (2022). Privacy enhancing technologies (pets) for connected vehicles in smart cities. *Transactions on Emerging Telecommunications Technologies*, 33(10).

Verborgh, R. (2023). Re-decentralizing the Web, for good this time. In Seneviratne, O. and Hendler, J., editors, *Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web*, pages 215–230. ACM.

Zheng, S., Apthorpe, N., Chetty, M., and Feamster, N. (2018). User perceptions of smart home IoT privacy. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW).