



REPEL: A UTILITY PRESERVING PRIVACY SYSTEM FOR IOT BASED ENERGY METER

¹ Mr. RAMAKRISHNA DOREDLA, ² Ashritha Pullag, ³ Belige Navya, ⁴ Ch. Sneha

¹Assistant Professor, Department of ECE, Mallareddy Engineering College for Women, Hyderabad

^{2,3,4}Students, Department of ECE, Mallareddy Engineering College for Women, Hyderabad

ABSTRACT

Modern Internet of Things (IoT) applications transmit sensor data to the cloud where it is subjected to analytics to provide useful services to users. Unfortunately, IoT sensor data often embeds sensitive private information that is vulnerable to leakage when sent to the cloud. Prior work on preserving IoT data privacy, particularly in the energy domain, focuses on obfuscating data to prevent extraction of private information. However, unless done carefully, data obfuscation significantly reduces the ability to extract useful but non-private information from the data. As a result, these existing techniques also reduce much of the utility derived from deploying IoT devices. In this paper, we address this problem by designing RepEL, a new utility-preserving privacy technique, which intelligently obfuscates smart energy meter data to prevent leaking a home's private occupancy information, while retaining the ability to perform useful energy disaggregation analytics. To preserve energy data's utility, our approach creates a randomized permutation of actual device usage via load replay while suppressing private user behavior information (such as occupancy) in the original data. We implement our algorithm on an embedded gateway node to demonstrate its feasibility and empirically evaluate our approach using real energy traces from homes. Our results show that the privacy leak rate for nearly two-thirds of the homes is below 10%, with four homes having no privacy leak. At the same time, the change in device usage for these homes is less than 3%. Further, we also demonstrate that RepEL has the flexibility to randomly replay loads, which can prevent adversaries from inferring behavioral patterns from device usage or use the information to determine occupancy

I.INTRODUCTION

Recent advances in embedded systems hardware and wireless networking have led

to the emergence of the Internet of Things (IoT) that is increasingly monitoring all aspects of our lives. Many consumer IoT

products are now available for applications such as smart home automation, smart health, and others. Current IoT products use a cloud-based architecture, depicted in Figure 1, where the IoT device sends sensor data to the cloud, where it is subjected to analytics. Further, as shown in the figure, users often use a mobile app (or a voice assistant, such as Alexa) to interact with the cloud service and send commands to the IoT device via the cloud. Such a cloud-based architecture for IoT devices, while commonplace, raises numerous privacy concerns. Since IoT sensor data often embeds private information about the user, sending this data to the cloud service implies the cloud analytics service now has the ability to extract this private information. In many cases, the embedded private information is orthogonal to the original purpose for which the data was collected by the IoT device, and can result in privacy leakage and unintended consequences. As an example, consider the privacy problems that arose when Strava, a popular fitness tracking app used by runners, published anonymized heat maps of popular running routes. Even

though the data was anonymized by removing user identities, the location information embedded in the running routes inadvertently revealed the locations of secret military bases in remote regions [30]. Studies have shown that similar privacy leakages can occur from the use of smart home IoT products [19, 24, 25, 31, 32]. For example, when users operate their smart lights or smart appliances, the cloud service can infer a home's occupancy, i.e., whether it is occupied or vacant. Researchers [26] have shown that inferring occupancy is the first step towards launching various types of sophisticated privacy attacks (see Table I) such as when residents take vacations, when they eat out at restaurants, etc. While the issue of IoT privacy has broad applicability to multiple application domains, in this paper, we focus on privacy in the domain of smart homes and energy. Maintaining user privacy in the face of cloud-based IoT services is a challenging problem. One privacy preserving approach is to use IoT devices "locally" without using any backend cloud services. However, this prevents the user from taking advantage of many useful features that cloud analytics services can provide. For example, such services use sophisticated machine learning to provide tailored recommendations for

making the home more energy efficient or to identify anomalous energy usage based on comparisons across peer groups. Users can use these insights and recommendations to reduce energy waste and cut their energy bills. Running such analyses locally is often infeasible due to the computational needs of these analytics algorithms or the inability of a standalone device to perform comparative analytics across users. Thus, not sharing IoT data with a cloud backend implies that the user has to forgo useful insights provided by these services. Another privacy-preserving approach is to use data obfuscation where the data is transformed by adding noise, or other means, prior to sending it to the cloud service [25, 36]. However, data obfuscation methods are a blunt instrument that, while enhancing privacy, destroy all useful information embedded in the data, whether private or not. As a result, running cloud analytics on this transformed data will produce erroneous results and is no longer useful to the user. Thus, traditional privacy-preserving techniques are often not compatible with cloud-based IoT services. To address these drawbacks, we argue for a utility preserving privacy approach where the data produced by the IoT device is intelligently transformed to strip out private information while retaining

other useful information embedded in the data. Such an approach will continue to allow cloud services to perform analytics that is useful to the user while preventing “mining” of private information. Figure 2 compares such as an approach to the current (non-private) cloudbased IoT services and recent privacy preserving methods for energy data. As shown, current IoT services preserve utility and sacrifice privacy; privacy-preserving data transformations based on obfuscations preserve user privacy but destroy utility. Our approach provides utility to users and respects privacy. While other recent efforts have examined utility-privacy tradeoffs [14, 16], this prior work has focused on theoretical issues, such as information theoretic analysis of this tradeoff or mathematically analyzing the impact of adding noise.

II. METHODOLOGY

A) System Architecture




Fig1 .Block Diagram

The architecture for the REPEL system focuses on ensuring privacy while monitoring energy usage in an IoT-based energy meter. It involves energy meters equipped with IoT sensors that collect real-time data. This data is processed locally or at edge devices to ensure privacy, while a secure communication protocol (e.g., encrypted data transfer) is used to send data to a central cloud server. The server performs analysis and generates reports for users, enabling them to monitor energy consumption efficiently without compromising their privacy.

B) Proposed Raspberry pi

The Raspberry Pi Pico is an affordable microcontroller board created by the Raspberry Pi Foundation. Unlike full-fledged computers, microcontrollers are small and have limited storage and peripheral options, such as the absence of devices like monitors or keyboards. However, the Raspberry Pi Pico is equipped with General Purpose Input/Output (GPIO) pins, similar to the ones found on Raspberry Pi computers, allowing it to connect with and control a variety of electronic devices. Introduced in January 2021, the Raspberry Pi Pico is based on the

RP2040 System on Chip (SoC), which is both cost-effective and highly efficient. The RP2040 SoC includes a dual-core ARM Cortex-M0+ processor that is well-known for its low power consumption. The Raspberry Pi Pico is compact, versatile, and performs efficiently, with the RP2040 chip as its core. It can be programmed using either Micro Python or C, providing a flexible platform for users of various experience levels. The board contains several important components, including the RP2040 microcontroller, debugging pins, flash memory, a boot selection button, a programmable LED, a USB port, and a power pin. The RP2040 microcontroller, custom-built by the Raspberry Pi Foundation, is a powerful and affordable processor. It features a dual-core ARM Cortex-M0+ processor running at 133 MHz, 264 KB of internal RAM, and supports up to 16 MB of flash memory. The microcontroller provides a wide range of input/output options, such as I2C, SPI, and GPIO. The Raspberry Pi Pico has 40 pins, including ground (GND) and power (Vcc) pins. These pins are grouped into categories such as Power, Ground, UART, GPIO, PWM, ADC, SPI, I2C, System Control, and Debugging. Unlike the Raspberry Pi computers, the GPIO pins on the Pico can

serve multiple functions. For instance, the GP4 and GP5 pins can be set up for digital input/output, or as I2C1 (SDA and SCK) or UART1 (Rx and Tx), though only one function can be used at a time.

C) Design Process

The design of embedded systems follows a methodical, data-driven process that requires precise planning and execution. One of the core elements of this approach is the clear separation between functionality and architecture, which is crucial for moving from the initial concept to the final implementation. In recent years, hardware-software (HW/SW) co-design has gained significant attention, becoming a prominent focus in both academia and industry. This methodology aims to align the development of software and hardware components, addressing the integration challenges that have historically affected the electronics field. For large-scale embedded systems, it is essential to account for concurrency at all levels of abstraction, impacting both hardware and software components. To facilitate this, formal models and transformations are employed throughout the design cycle, ensuring efficient verification and synthesis. Simulation tools are vital for exploring design alternatives and confirming

the functional and timing behavior of the system. Hardware can be simulated at different stages, including the electrical circuit, logic gate, or RTL level, often using languages like VHDL. In certain setups, software development tools are integrated with hardware simulators, while in other cases, software runs on the simulated hardware. This method is generally more suited for smaller parts of an embedded system. A practical example of this methodology is the design process using Intel's 80C188EB chip. To reduce complexity and manage the design more effectively, the process is typically divided into four main phases: specification, system synthesis, implementation synthesis, and performance evaluation of the prototype.

APPLICATIONS

Embedded systems are being increasingly incorporated into a wide range of consumer products, such as robotic toys, electronic pets, smart vehicles, and connected home appliances. Leading toy manufacturers have introduced interactive toys designed to create lasting relationships with users, like "Furby" and "AIBO." Furbies mimic a human-like life cycle, starting as babies and growing into adults. "AIBO," which stands for Artificial Intelligence Robot, is an advanced robotic

dog with a variety of sophisticated features. In the automotive sector, embedded systems, commonly referred to as telematics systems, are integrated into vehicles to offer services like navigation, security, communication, and entertainment, typically powered by GPS and satellite technology. The use of embedded systems is also expanding in home appliances. For example, LG's DIOS refrigerator allows users to browse the internet, check emails, make video calls, and watch TV. IBM is also developing an air conditioner that can be controlled remotely via the internet. Given the widespread adoption of embedded systems across various industries.

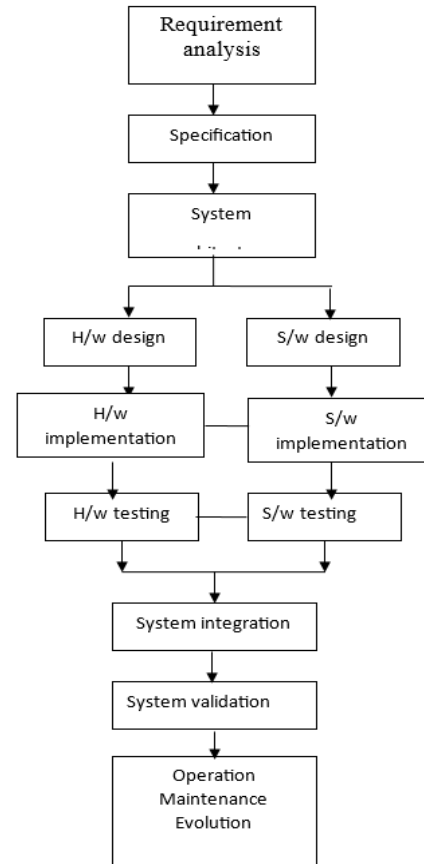


Fig 2. Embedded Development Life Cycle

III.CONCLUSION

In this paper, we considered the privacy of energy data from smart meters and argued that while existing techniques enhance privacy, they do not preserve utility as they destroy all useful information embedded in the data, whether private or not. We presented the notion of utility-preserving privacy, which prevents privacy leakage while retaining the ability to perform energy-efficiency analytics. We designed and

implemented RepEL, a utility-preserving privacy approach, that maintains appliance usage information while thwarting privacy attacks by randomly replaying energy loads to suppress private information leaks. We implemented our algorithm on a Raspberry Pi to demonstrate its feasibility and used real energy traces to analyze its performance. Our results showed that the privacy leak rate for nearly two-thirds of the homes is below 10%, with four homes having no privacy leak. At the same time, the change in device usage for these homes is less than 3%. Further, we demonstrate that RepEL has the flexibility to randomly replay loads, which prevents adversaries from inferring behavioral patterns from device usage.

IV.FUTURE SCOPE

The future scope of the REPEL system includes enhancing data security through advanced encryption algorithms and integrating AI-powered analytics for better energy usage predictions. The system could also incorporate machine learning techniques for real-time anomaly detection, helping users optimize energy consumption. Additionally, the adoption of blockchain technology could further ensure the integrity

and privacy of the data, fostering trust in energy monitoring solutions.

V.REFERENCES

- [1] G. Acs and C. Castelluccia. I Have a DREAM! (Differentially private smart Metering). In International Workshop on Information Hiding, 2011.
- [2] G. Acs, C. Castelluccia, and W. Lecat. Protecting against physical resource monitoring. In Proceedings of the 10th annual ACM workshop on Privacy in the electronic society, 2011.
- [3] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng. Occupancy-driven energy management for smart building automation. In Proceedings of the 2nd ACM workshop on embedded sensing systems for energyefficiency in building, 2010.
- [4] M. Backes and S. Meiser. Differentially private smart metering with battery recharging. In Data Privacy Management and Autonomous Spontaneous Security. 2014.
- [5] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava. Nilmtk: an open source toolkit for non-intrusive load monitoring. In

Proceedings of the 5th international conference on Future energy systems, 2014.

[6] N. Batra, A. Singh, and K. Whitehouse. Neighbourhood nilm: A big-data approach to household energy disaggregation. arXiv preprint arXiv:1511.02900, 2015.

[7] C. Beckel, W. Kleiminger, R. Cicchetti, T. Staake, and S. Santini. The eco data set and the performance of non-intrusive load monitoring algorithms. In Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings, 2014.

[8] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, and P. Shenoy. Non-intrusive occupancy monitoring using smart meters. In Proceedings of the Workshop on Embedded Systems For Energy-Efficient Buildings, 2013.

[9] D. Chen and D. Irwin. Sundance: Black-box behind-themeter solar disaggregation. In the Eighth International Conference on Future Energy Systems, May 2017.

[10] D. Chen, D. Irwin, P. Shenoy, J. Albrecht, et al. Combined heat and privacy: Preventing occupancy detection from smart meters. In 2014 IEEE International Conference on Pervasive Computing and Communications (PerCom), 2014.

[11] Dataport. pecan street., 2018.
<http://www.pecanstreet.org/>.