Privacy Therapy with Aretha: What If Your Firewall Could Talk?

William Seymour
University of Oxford
Oxford, UK
william.seymour@cs.ox.ac.uk

ABSTRACT

The rapid adoption of smart home devices has brought with it a widespread lack of understanding amongst users over where their devices send data. Smart home ecosystems represent complex additions to existing wicked problems around network privacy and security in the home. This work presents the Aretha project, a device which combines the functionality of a firewall with the position of voice assistants as the hub of the smart home, and the sophistication of modern conversational voice interfaces. The result is a device which can engage users in conversation about network privacy and security, allowing for the forming and development of complex preferences that Aretha is then able to act upon.

CCS CONCEPTS

• **Security and privacy** → **Usability in security and privacy**; • **Human-centered computing** → *Visualization systems and tools*; *Sound-based input / output*.

KEYWORDS

Voice Assistants, Firewall, Smart Home, Voice Interfaces, Sensemaking

CHI'19 Extended Abstracts, May 4–9, 2019, Glasgow, Scotland UK

ACM ISBN 978-1-4503-5971-9/19/05.
https://doi.org/10.1145/3290607.3308449

ACM Reference Format:
William Seymour. 2019. Privacy Therapy with Aretha: What If Your Firewall Could Talk?. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI'19 Extended Abstracts), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3290607.3308449

INTRODUCTION

Maintaining privacy and security in the networked home is a tremendous challenge, with well explored problems such as the privacy paradox and the marketing of end user security products (such as antivirus and firewalls) as the first and last step required to achieve domestic cyber security. This comes at a time of increasing reports of corporate data breaches and a gradual shift in public opinion over the sharing of personal data. But as the number of 'smart' devices in the home continues to rise, consumers continue to lack situational awareness of their home networks.

At the same time, voice assistants such as the Amazon Alexa and Google Assistant have quickly risen in popularity. Driven by accurate voice recognition, these devices have become the centre of the smart home, able to coordinate other automation devices in addition to functioning as assistants. With the exception of small open source projects[1], commercially available voice assistants are geared towards novelty functionality and the creation of large machine learning datasets.

But what if firewalls and voice assistants were combined to create a product that *truly* worked on behalf of its users? Being at the centre of the smart home, voice assistants are perfectly placed to monitor other connected devices. In addition, the conversational nature of today's top assistants could allow for something that is sorely missing from the modern smart home: *conversations* about security and privacy.

Continuing previous research on the concept of 'respectful' smart home devices [14, 17], this work describes an exploratory project called Aretha: a prototype voice assistant that captures and analyses data from a user's home network in order to support conversations about security and privacy. After detailing relevant prior work, the remaining sections of this paper describe the structure of Aretha, as well as the novel interaction modalities by which it presents analysis.

BACKGROUND

Voice Interfaces

Pioneering work by Nass et. al. showed that a number of phenomena normally associated with human interaction also apply to interactions with computer voice interfaces (e.g. that computer generated voices are gendered social actors, and social responses to computers are automatic and unconscious) [8, 13]. Speech activates the same centres in the brain regardless of whether it originates from a home assistant or another person, presenting an array of functional and ethical challenges when designing voice controlled systems.

[1] Such as:

- Jasper (https://jasperproject.github.io)
- Home Assistant (https://www.home-assistant.io)
- Mycroft (https://mycroft.ai/)

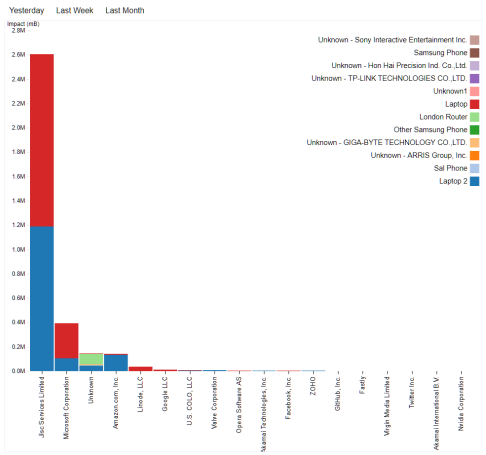

Figure 1: Data destination profile in the Aretha web interface

[2]https://solid.inrupt.com/

These theories are validated in practice—while some of the social rules users apply to assistants might simply be 'over-learned social behaviours' [7]—trends associated with much more 'human' interaction have also been documented. Similar to interpersonal conversations, interactions with voice assistants can often be positive even when failing to fulfil their functional objectives (i.e. the interactions themselves are satisfying) [6, 12], and participants in my own research (currently under submission) sometimes describe voice assistants as having the same physical 'presence' as a person would. So while conversation structure with voice interfaces might be fundamentally different to speech between humans [11], these interactions come with many of the same risks and benefits.

A major problem with smart home devices is that computational modes of thought often translate poorly to the medium of speech, leaving exchanges with machines at risk of feeling more mechanical than normal conversation [11]; the move from information dense media (e.g. text, images) to the relatively low bandwidth of speech can lead to metallic sounding soliloquies [5] or a loss of precision.

Security & Privacy in the Smart Home

Privacy as applied to technology usage is a well developed field, with widely used models of privacy as a dynamic, dialectic process [9] and as 'contextual integrity' [2], the latter considering privacy as based on norms surrounding the appropriateness of information in context, as well as flows of information between different parties. Smart home devices are often able to observe users across these contexts, leading to violations of informational boundaries due the fact that devices lack contextual awareness of the environments in which they operate.

Contributions exploring user understanding around smartphone apps [15] and smart homes [18] show a generally poor level of understanding amongst end users about where their data is going, coincident with deep seated problems with current privacy controls such as privacy policies[10]. The resultant problems are often conceptualised as ones of poor situational awareness, but even with perfect information, it is unlikely that the users of *any* skill level would have the time or inclination to successfully process it [1], leading Van Kleek to conclude that the most effective tools in this space are likely to be those that *support* user decision making with analysis and automation, rather than simply providing information [16].

Other contributions show the second order effects that smart devices can have on relationships in the home; Choe et al. demonstrate how continuous sensing can cause tension between cohabitants [3], and Zeng highlights how smart home devices can disrupt or exacerbate power balances within the home [18]. These 'wicked' problems with smart home privacy have led to speculative works promoting a more decentralised internet architecture, such as Databox and Solid. Databox provides users with a container residing on their local network for storing personal data, rather than a remote server [4]. The Solid framework[2] provides a similar model, with applications accessing a user controlled 'POD' which can be hosted anywhere on the web.

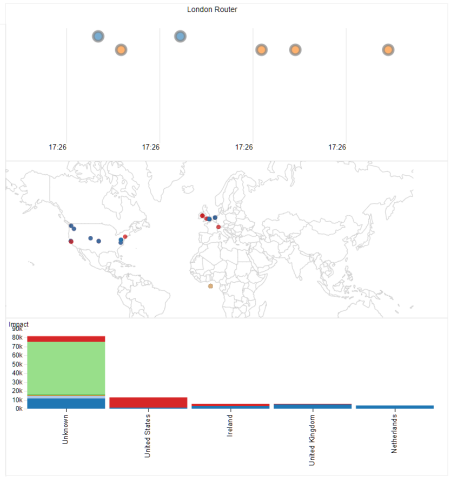

Figure 2: Behaviour and geographic profiles in the Aretha web interface

CONVERSATION: SMART FIREWALL

*Featuring the **U**ser and **A**retha*
U: Aretha, do you have any updates for me?
A: *This week the fridge has started sending data to a new company in Croatia called DataCorp. Is this expected, or would you like me to block traffic to that destination?*
U: Aretha, who owns that company?
A: *DataCorp is owned by EvilCorp.*
U: I don't understand why it needs to send data there. Aretha, please block that traffic.
A: *Will do. Please let me know if the fridge stops working correctly because of this.*

EXPLORING DATA FLOWS IN THE HOME

The focus of the Aretha project is to explore how the position of the voice assistant at the centre of the smart home can be leveraged to help users maintain situational awareness of their home networks. Aretha achieves this by capturing traffic meta data from devices connected to its WiFi hotspot, collating these with information about the companies and countries that data is transferred to (based on IP address). Time series analysis is used to detect uncharacteristic behaviour, and, for supported devices, re-identify behaviour from encrypted network traffic to show users information their devices might inadvertently be leaking about them.

Analysis from Aretha is accessed via an API which supports a web interface and companion Alexa skill. The web sensemaking interface is based on X-Ray Refine [16], and displays information about network flows grouped by destination company (see Figure 1), traffic category, or geographic location (see Figure 2). The Alexa skill supports verbal reasoning about the data in human terms, as well as providing novel ways to respond to suspicious device behaviour (see Conversation 1 in sidebar). In line with research on the Databox and Solid projects, data from Aretha does not leave the home without good reason and user consent.

In doing this, Aretha is able to to 'retrofit' existing smart home devices for more accessible interaction, allowing users to make assisted security and privacy decisions about their devices through Aretha. Letting users decide who their devices can send data to is an important step towards promoting user-centric privacy and security interactions in the home, and Aretha provides this in a way that is compatible with existing devices without requiring extensive configuration.

WHAT IF YOUR FIREWALL COULD TALK?

A major challenge with conventional firewalls is that they are notoriously difficult to use and configure by non-experts. The information required to form a mental model is complex, with unclear mappings between data firewalls understand (e.g. ports and IP addresses) and data users require to make decisions (e.g. company names). To address this, when users interact with Aretha (via the companion Alexa skill) they make requests in terms of *who* should be allowed to receive their data, rather then *where* that data should be allowed to go. By using rules associated with richer, human motivations, Aretha offers meaningful decisions to users, pairing computational scale with human creativity in an effort to reduce the cognitive friction of making security and privacy decisions; "stop sending data to Google" is much easier to reason about than "`iptables -A INPUT -d 64.233.160.0/19 -j DROP`".

Beyond this, the combination of a voice interface with the semantic functionality of Aretha allows for a number of novel interactions that transcend the abilities of conventional firewalls. The lack of conversations about security and privacy in the smart home leaves users with a poor understanding of

what their home is doing with their data. To this end, Aretha is designed to assist users in developing their understanding through conversation—a kind of Socratic dialogue about data privacy and network security (see Conversation in sidebar) that reflects the nature of privacy as a dynamic, dialectic process [9].

Interactions like these allow for the idiosyncratic nature of security and privacy preferences by providing users with the questions they need to form their preferences and develop them over time, as well as the analysis and tools they need to act on them. In this way, Aretha handles the repetitive tasks required to perform analysis, such as aggregating logs or effecting decisions, leaving users to exercise their creativity when dealing with unusual scenarios as determined by each device's data flow model. For convenience, users can also specify a list of blacklisted companies or countries that Aretha will automatically block when detected.

EVALUATION & FUTURE WORK

As part of the PETRAS IoT in the Home demonstrator (see acknowledgements), Aretha will be installed in a fully equipped smart home at the Building Research Establishment in Watford, UK. This is expected to take place in Q1 2019, and will be paired with a user study designed to investigate the ability of tools such as Aretha to help users develop their understanding of network security and privacy in the home. A usability evaluation of the Refine interface itself is provided in [16].

Planned future work involves the expansion of Aretha to assist users in exercising their rights under the European General Data Protection Regulation (GDPR, see sidebar). Aretha provides the perfect platform to *automate* the usage of these rights. By identifying which companies are processing data from an individual's IoT devices, the platform will be able to generate emails invoking the rights to access and erasure using templates and publicly available contact email addresses. These requests will be sent, monitored, and stored from within the web interface, dramatically lowering the barrier for entry of GDPR data rights.

CONCLUSION

The poor level of understanding amongst users of smart home devices, combined with the rapid adoption of voice assistants as the hub of the smart home has produced an opening for a new type of device that assists users in developing complex and highly personal preferences around home security and privacy.

Initial exploration through the Aretha project in engaging users in conversation about their preferences demonstrate that shifting the design calculus of these devices can lead to developments in terms of security and privacy in the home. Making use of modern voice interfaces allows devices such as Aretha to empower users by building situational awareness and control over their home networks rather than patronise them with inaccessible or overly simplified representations of their technology.

DATA SUBJECT RIGHTS UNDER THE GDPR

The GDPR gives every citizen of the EU a number of rights in relation to how their personal data is processed. These rights are far reaching, but have been criticised for not being accessible to regular users (i.e. the majority of uses are from academics, lawyers, and journalists). The rights provided are:

- Right to be informed about collection and use of personal data
- **Right of access to personal data**
- Right to have inaccurate personal data rectified
- **Right to have personal data erased**
- Right to request restrict the processing of personal data
- Right to have personal data provided in a machine readable format
- Right to object to processing of personal data
- Rights in relation to automated profiling and and decision making

ACKNOWLEDGEMENTS

This work is supported by ReTiPS: Respectful Things in Private Spaces, a project funded through the UK Engineering and Physical Sciences Research Council PETRAS IoT Hub under grant number N02334X/1. The author is supported through the EPSRC Centre for Doctoral Training in Cyber Security under grant number P00881X/1.

REFERENCES

[1] Alessandro Acquisti and Jens Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE security & privacy* 3, 1 (2005), 26–33.

[2] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. 2006. Privacy and contextual integrity: Framework and applications. In *Security and Privacy, 2006 IEEE Symposium on*. IEEE, 15–pp.

[3] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, Shwetak N Patel, and Julie A Kientz. 2012. Investigating receptiveness to sensing and inference in the home using sensor proxies. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 61–70.

[4] Andy Crabtree, Tom Lodge, James Colley, Chris Greenhalgh, Richard Mortier, and Hamed Haddadi. 2016. Enabling the new economic actor: data protection, the digital economy, and the Databox. *Personal and Ubiquitous Computing* 20, 6 (2016), 947–957.

[5] Joseph Galen Lindley, Paul Coulton, Haider Akmal, and Brandin Hanson Knowles. 2017. Anticipating GDPR in Smart Homes Through Fictional Conversational Objects. (2017).

[6] Irene Lopatovska, Katrina Rink, Ian Knight, Kieran Raines, Kevin Cosenza, Harriet Williams, Perachya Sorsche, David Hirsch, Qi Li, and Adrianna Martinez. 2018. Talk to Me: Exploring User Interactions with the Amazon Alexa. *Journal of Librarianship and Information Science* (2018).

[7] Irene Lopatovska and Harriet Williams. 2018. Personification of the Amazon Alexa: BFF or a Mindless Companion. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval (CHIIR '18)*. ACM, 265–268.

[8] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are Social Actors. In *Proceedings of the CHI conference on Human factors in computing systems (CHI '94)*. ACM, 72–78.

[9] Leysia Palen and Paul Dourish. 2003. Unpacking privacy for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 129–136.

[10] Irene Pollach. 2007. What's wrong with online privacy policies? *Commun. ACM* 50, 9 (2007), 103–108.

[11] Martin Porcheron, Joel E Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, 640.

[12] Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. 2017. Alexa Is My New BFF: Social Roles, User Satisfaction, and Personification of the Amazon Echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI '17)*. ACM, 2853–2859.

[13] Byron Reeves and Clifford Ivar Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places.* Cambridge University Press.

[14] William Seymour. 2018. How Loyal is Your Alexa?: Imagining a Respectful Smart Assistant. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*.

[15] Irina Shklovski, Scott D Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, 2347–2356.

[16] Max Van Kleek, Reuben Binns, Jun Zhao, Adam Slack, Sauyon Lee, Dean Ottewell, and Nigel Shadbolt. 2018. X-Ray Refine: Supporting the Exploration and Refinement of Information Exposure Resulting from Smartphone Apps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, 393.

[17] Max Van Kleek, William Seymour, Reuben Binns, and Nigel Shadbolt. 2018. Respectful things: Adding social intelligence to 'smart' devices. In *Living in the Internet of Things: Cybersecurity of the IoT*. IET.

[18] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End user security & privacy concerns with smart homes. In *Symposium on Usable Privacy and Security (SOUPS)*.