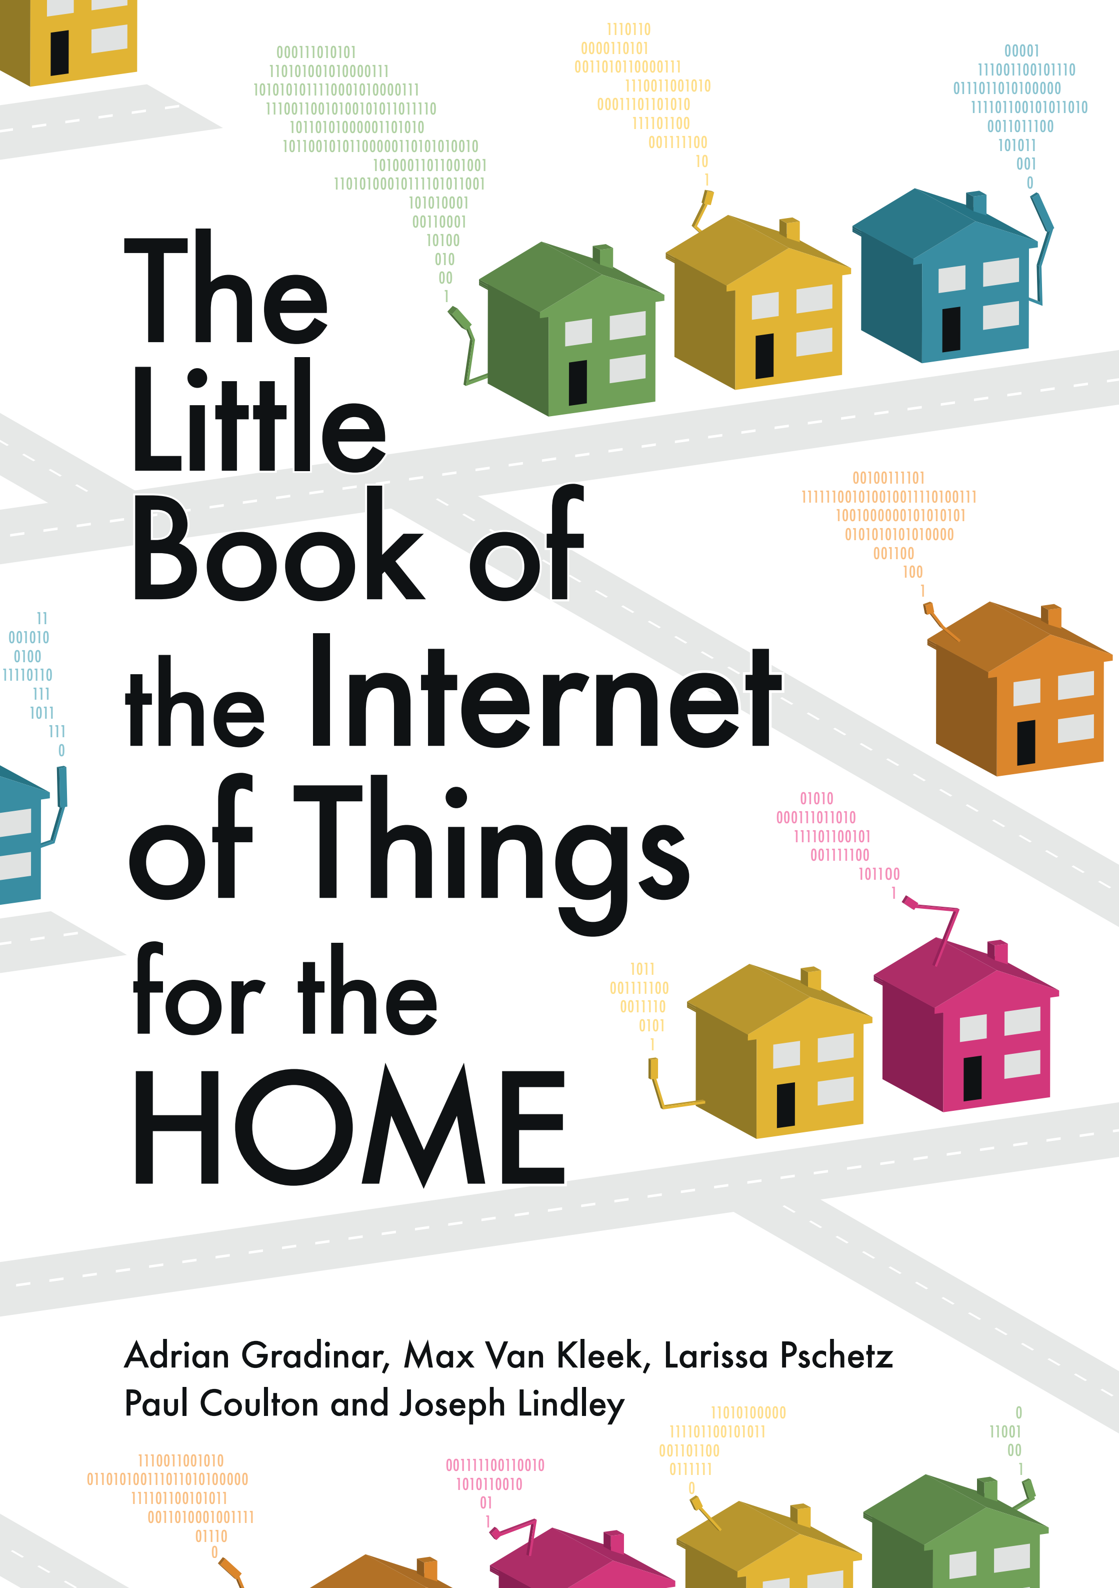

The Little Book of the Internet of Things for the HOME

Adrian Gradinar, Max Van Kleek, Larissa Pschetz
Paul Coulton and Joseph Lindley

Editor of the PETRAS Little Books series:

Dr Claire Coulton
ImaginationLancaster, Lancaster University

With design by Michael Stead, Roger Whitham and Rachael Hill
ImaginationLancaster, Lancaster University

ISBN 978-1-86220-370-9

The Little Book of the Internet of Things for the HOME

Adrian Gradinar, Max Van Kleek, Larissa Pschetz, Paul Coulton and Joseph Lindley

Acknowledgements

This book is based on our research conducted for the Acceptability and Adoption theme of the PETRAS Cybersecurity of the Internet of Things Research Hub funded by the UK Engineering and Physical Sciences Research Council under grant EP/N023234/1, EP/N023013/1, and EP/N02317X/1. The aim of PETRAS is to explore critical issues in privacy, ethics, trust, reliability, acceptability, and security and is undertaken by researchers from UCL, Imperial College London, Lancaster University, University of Oxford, University of Warwick, Cardiff University, University of Edinburgh, University of Southampton, University of Surrey, University of Bristol and Newcastle University. We would also like to thank the Building Research Establishment (BRE) at Watford, where a number of the projects featured in this book were demonstrated.

Contents

What this Little Book tells you

From doorbells to lightbulbs, and thermostats to ventilation, Internet of Things (IoT) enabled devices are increasingly available to buy for our homes or are being embedded within new builds. In this Little Book, we discuss the 'smart' products and services that are being designed for our homes, which are part of the emerging IoT, and highlight potential flaws in the design practices that have produced them. In particular, we discuss two new perspectives about how the design process could demonstrate how complex the IoT is, while also showing users what the potential issues relating to privacy, trust and security might be.

Based on our research conducted for the Acceptability and Adoption theme of the PETRAS Cybersecurity of the Internet of Things Research Hub, this Little Book explains:

- What we mean by the IoT and Smart Homes
- How the home has evolved as a site of technology innovation
- Why the IoT could be called the Internet of Superfluous Things
- How IoT issues give rise to new IoT design strategies, such as:
    - Human Data Interaction (HDI)
    - More-Than Human Centred Design
- Practical examples of new design strategies that address IoT challenges, such as:
    - Data legibility
    - User Agency
    - User/Device Negotiation
    - How all of this connects together and what it all means

What is the Internet of Things?

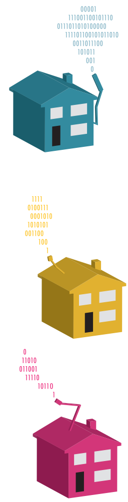

In our first Little Book[1] in the PETRAS series we explained the term Internet of Things (IoT) as follows:

> *"... the term [is used] to describe objects or things that can be interconnected via the Internet. This allows them to be readable, recognizable, locatable, addressable, and/or controllable by computers. The things themselves can be literally anything. Later in the book we use examples such as a kettle, a door lock, an electricity meter, a toy doll and a television but it's important to remember that there is no limit on what is or is not an IoT thing. Anything that is connected to the Internet is arguably part of the IoT including us."*

In this book, we focus on IoT products and services targeting the consumer market, in particular, those for use in our homes. These connected products are often referred to as 'smart' and our IoT-enabled homes are often called, 'smart homes'. The promise of smart homes filled with connected products is frequently promoted as a way of making our lives easier and more convenient. For example, the *Roomba* robotic vacuum cleaner claims to allow you to *"Forget about vacuuming for weeks at a time"* and that it [the robot] is smart enough to know if your cat has tracked its litter through the house.[2]

[1] Coulton, P., Lindley, J. G., & Cooper, R. 2018. The Little Book of Design Fiction for the Internet of Things. Lancaster: Lancaster University.

[2] https://www.irobot.co.uk/home-robots/vacuuming

The promise of the smart home envisages a future where our lives are easier, giving us more time to do things whilst consuming less energy and saving money. However, what is frequently absent from these discussions is the tsunami of data which will be generated and collected as we add millions of IoT products and services to our networks.

For example, many consumers are increasingly conscious of the data they generate through social media use, but they are less conscious of the data they generate elsewhere, such as while browsing the internet or watching streaming media. IoT products and services add to this largely unconscious data production and rarely make it clear to their owners the extent to which data is collected, where and how it is stored, and what it is being used for and by whom. While awareness of these Human-Data relationships may not be of immediate concern to most users, when this activity is unexpectedly brought to the fore they often challenge our existing expectations for personal privacy in our homes.[3] For example, returning to the *Roomba* vacuum cleaner we discussed above, many owners were shocked to learn that the latest versions of the device produced detailed maps of their homes. These were then relayed to the manufacturer who could potentially have shared these with 3rd parties. While an automatic vacuum cleaner seems attractive, a digital device which maps the interior of your home in order to, potentially, sell that map to the highest bidder, is clearly a more complicated proposition. Following this example, we can see the tension that home data collection places on our expectations of privacy. This could negatively affect the adoption and acceptability of IoT products and services, and in this Little Book we will explore how this tension can be softened through new design strategies which we discuss later in the book.

[3] Akmal, H & Coulton, P. 2018. Using Heterotopias to Characterise Interactions in Physical/Digital Spaces. in C Storni, K Leahy, M McMahon, P Lloyd & E Boehmia (eds), *Proceedings of the Design Research Society Conference 2018*. vol. 1, Limerick, pp. 269-278.

Technology Innovation in the Home

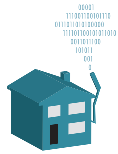

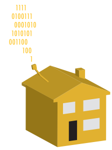

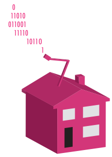

Whilst many of us may complain about the chores we do around our homes, we have largely forgotten that no more than a century ago, people's lives (primarily women's lives), were almost entirely consumed with housework. It wasn't until the widespread provision of electricity (1926 in the United Kingdom), and subsequent development of home electrical appliances, that this fundamentally started to change. Electricity originally provided lighting as a replacement for candles and gas lamps in homes – immediately removing a significantly time-consuming task – but it also paved the way for the development and subsequent adoption of appliances such as irons, vacuum cleaners, toasters, refrigerators, washing machines, dishwashers, clothes dryers, hair dryers, freezers, and so-on.

The first appliances were aimed at reducing the amount of time and physical effort required to perform a task although they still required oversight by a human. Later developments focussed on making their operation increasingly automated. The increasing automation of home appliances coincided with the rise of Home Economics being taught in schools. In 1969, Honeywell proposed a Kitchen Computer with the tagline, *"if she can only cook as well as Honeywell can compute"*.[4] Although they never sold a single device, computing technology *was already being embedded* in home appliances in the form of Integrated Circuits (ICs).

[4] http://www.thecatalogblog.com/2017/04/02/can-cook-well-honeywell-can-compute/

These ICs provided the means to automate various processes so that they could be started by simple button or dial configurations. A good example of such behaviour can be commonly seen on automatic washing machines where most of the controls are driven by ICs hidden behind a panel of buttons and dials. Full home automation (sometimes referred to as 'domotics') began to be seriously proposed in the 1970s, but as is frequently the case with emerging technologies, those ideas were often based on earlier visions. For example, themes from General Motors' Designing for Dreaming (a short film produced for their Motorama in 1956, which features Frigidaire's 'Kitchen of the Future'[5]) and the famous 1960s Jetsons cartoon series, were both influential on home automation concepts. Aspirational visions of futuristic automated devices in domestic settings have resurfaced numerous times in the 20th and 21st centuries. In the late 1970s, the X10 standard for communication via powerline signalling became an industry standard, however, the widespread adoption of Wi-Fi in the home has provided a ubiquitous, flexible and affordable communications infrastructure. This is where the story of the IoT meets the story of technology in the home.

The increasing number of IoT products being released (Figure 1) highlights that the home is seen as a significant market segment. These products span many types of application—baby monitors, doorbells, toys, ovens, entertainment systems, door locks—which tend to look very familiar as they make up part of our everyday lives. However, if we consider the devices listed in figure 1, it is clear they are based on current understandings of how we live. As the modern world changes the way we live (e.g. co-housing, intergenerational occupancy, home working, etc), and the technologies we have around us become network-connected, new ways of looking at these 'socio-technical' spaces is increasingly necessary.

[5] https://www.youtube.com/watch?v=GjoLaV2y-tA

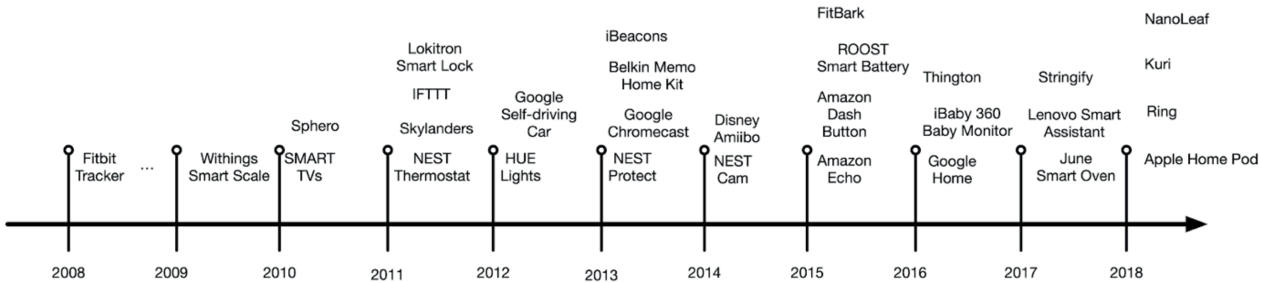

Figure 1. IoT Product Release Timeline

An Internet of Superfluous Things!

Under a façade of innovation, a significant majority of the current IoT products and services being created are evident of a design culture displaying a penchant for creating *superfluous things*. These products tend towards what Evgeny Morozov describes as 'solutionism', as frequently these devices appear to be *"solving problems that do not really exist."*[6] For example, self-driving baby strollers, connected underwear, smart dental floss, and connected showerheads, are just a few of the hundreds of examples of IoT products and services that are largely superfluous to our lives.[7]

These devices are similar to what the science fiction author and technology critic Bruce Sterling describes as 'Gizmos' in his 2005 non-fiction book *Shaping Things*. Sterling characterises classes of objects into their varying human-object relationships. These are:

- Artefact: made by hand, used by hand, and powered by muscle.
- Machine: complex artefacts with integral moving parts and with a non-human/non-animal power source.
- Product: non-artisanal, uniformly mass-produced artefacts,

[6] Morozov, E. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. Public Affairs.

[7] See the twitter account The Internet of Shit (@internetofshit) for a running commentary on superfluous IoT devices.

supported by large transport, finance, and information infrastructures.

- Gizmo: user alterable and programmable multi-functional objects commonly linked to network service providers.
- Spime: networked objects with extensive and rich informational support that are designed on screens, fabricated on screens, and tracked in space and time throughout their lifespan.
- Biot: is an entity that is both object and person that provides data to the network.

Spimes are crucial to delivering an optimistic, sustainability-centric future, which is enabled by the IoT, however, it is the Gizmo that currently dominates.[8] Arguably, many of the IoT Gizmos being developed for the home user are little more than disposable novelty devices. In addition to novelty, however, they are often part of surveillance capitalism,[9] facilitating mass data collection for corporate profit. Social platform Facebook is notorious for its highly-effective targeted advertising, which tracks usage, not-only on Facebook's platform, but across vast swathes of the web as well.[10] This type of data collection is based on the assumption that if something is described by a sufficient volume of data, then the certainty of a decision made based upon the data, will be increased. Because of this assumption, and because of the huge financial wealth generated by corporates like Facebook and Google through collecting vast amounts of data, IoT designers are often tempted to turn their Gizmos into data-collection devices. Whether this data turns out to be valuable or not remains to be seen (it is likely that the complexity of human experience and social context will render much of it useless). It is also important to note how data collected now may have unintended consequences in the future. For example, Fitbit collects your health data on a daily basis directly from your wrist; in a not-so-distant dystopic future, your health data might be used by insurance companies to increase your premiums based on your lack of physical activity. While insurance is a common

[8] Stead, M.R, Coulton, P, Lindley, J.G & Coulton, C. 2019. The Little Book of Sustainability for the Internet of Things. Lancaster: Lancaster University.

[9] We discuss surveillance capitalism in more detail on pages 18-19.

[10] https://www.wired.com/story/facebooks-targeted-ads-are-more-complex-than-it-lets-on/

example of this problem, it may be relevant to what products you are offered, where you are allowed to live, what romantic opportunities are presented to you, or the price retailers offer their products to you. In this sense, while the products themselves are often superfluous, their potential consequence on everyday life could be enormous.

While superfluous Gizmos are often laughable and can be amusing, there are more serious concerns related to them. For example, Morozov also critiques the very notion of smart products and services seeing the drive to make everything efficient and convenient as blocking reasonable consideration of alternatives. If this is true, he warns, it leads us down a path towards an algorithm-driven world where technology companies, rather than elected governments, are determining the shape of our futures.[11] The demonstrable impact of data on recent elections is, perhaps, an early sign of this happening. In recent elections, data was collated about individuals (some of which was generated by IoT devices), while algorithms and models were then used to analyse that data and determine what kind of message would affect and influence a specific individual. Finally, using digital platforms such as Facebook, these individualised election-related messages were delivered to their targets. At each stage of this process, what Morozov warned seems to be evident and is helping fuel the increased concern about whether the potential benefits of mass data collection are worthwhile.

The IoT's role in these concerns is significant. When users (and perhaps regulators or policymakers) do not understand the context of data collection, it is hard to control it. A lack of clarity about data collection, and how it might be used, exacerbates this issue. Moreover, when there is a hidden market of data aggregators who trade in personal data, often not knowing its provenance, the scope for potential scandal and privacy infringement is large. Together, these factors pose real problems for the adoption and acceptability of IoT devices. If the more utopian vision afforded by Spimes is to be realised, then these risks must be reduced. In the following section, we consider how 'better' design strategies may offer some of these reductions, with particular consideration for IoT in the home.

[11] Morozov, E. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. Public Affairs.

Strategies for 'better' IoT design

The notion of what makes something 'better' in design terms is always going to be a tricky subject. 'Better' design is subjective and a number of different opinions about any design are inevitable. However, in this book we suggest that 'better' IoT designs are those that reflect the complex and dynamic nature of the home environment and take into account the impact of IoT products and services on their environment. To support this kind of betterness, we need a design approach that reflects both product design and software engineering perspectives being considered together as, usually, they are viewed separately. In the following sections, we will discuss two new design stances, Human Data Interaction (HDI) and More-Than Human Centred Design, which can be useful tools in addressing the issues that we have highlighted so far and are relevant to design for the IoT both in the home, and beyond.

Human-Data Interaction (HDI)

The Human-Computer Interaction (HCI) field has evolved over the last 30 years. Initially, it mainly considered the interactions between humans and computers as physical artefacts, but now in the early 21^{st} century, it is concerned with how people interact and live with the services enabled by networked computing devices. The role that data plays, as people increasingly engage with these networks, is considered to be the next step. The term Human-Data Interaction (HDI) was coined to describe this new field:

"HDI places the human at the centre of these data flows, and HDI provides mechanisms which can help the individual and groups of people to interact explicitly with these systems and data."[12]

While HDI is still a new field, three core design principles for data-enabled products and services have been identified: *legibility*, *agency* and *negotiability*.

- *Legibility*: This recognises that the full extent of our interactions with data flows and data processes are generally opaque. So, legibility is concerned with ensuring that data and associated algorithms are made clear and understandable to users. For example, owners of *Vizio* smart televisions were unaware that 100 billion data points related to their viewing habits were being collected every day until it was made public in 2016.
- *Agency*: This relates to how users of data-enabled systems are able to manage their data and who has access to it. Aside from the basic ability to opt-in or opt-out of data collection, agency also relates to how data is stored and used, including the ability to modify data and the inferences that may be ascribed from it. Consider the 'smart' meters that are being rolled out currently in the UK. Users have little agency to optimise their tariffs or control who has access to the data which reveals a great deal about the users' lives.
- *Negotiability*: This acknowledges the transactional nature of data collection as we are often trading functionality or access to our data. This is to facilitate an ongoing engagement by users so that they can withdraw completely or in part and derive value from data collection themselves. For example, if you choose not to connect your *Roomba* to your Wi-Fi you lose some of the features offered through the mobile app such as: remotely scheduled cleanings, customised cleaning features, and any voice control functionality provided by *Amazon's Alexa* or *Google Assistant*. In this instance, the trade-off for losing this functionality is increased certainty that your data is secure (as it is not leaving your house), however, the negotiation is very one-sided. In the Roomba's case (as is frequently true) the terms are 'give us your data or we do not provide functionality'.

[12] Mortier, R, Haddadi, H, Henderson, T, McAuley, D, Crowcroft, J, and Crabtree, A. 2016. Human-Data Interaction: *The Encyclopedia of Human-Computer Interaction*, 2nd Ed. Interaction Design Foundation.

As an example of HDI in practice, we highlight the *Databox*[13, 14] platform, which is an edge computing device. This means that instead of distributing personal data to remote cloud servers for processing, processing takes place on-the-box, which means no personal data needs to leave the home. The *Databox* itself can be considered as a special form of home router which acts as a bridge between the things generating the data within the home context and outside organisations wishing to access this data. The owner of the databox is provided with direct control of the bridge, thus they are given agency over any distribution of data they may wish to provide. For example, it might enable a media provider to access data about an individual's viewing habits, and of others in the room, and offer up bespoke content of mutual interest without disclosing personal data to the provider. *Databox* is not limited to privacy-preserving functionality, but also enables actuation of IoT devices and became the basis for the *Living Room of the Future*, which we discuss later in the book.

More-Than Human Centred Design

Unlike HDI, which seeks to maintain the perspective on the human being as the central consideration, More-Than Human Centred Design approaches see the human as just another thing within the hyper-connected and data-mediated assemblages that make up the IoT. For example, this approach sees the 'things' within such networks are much more than their physical forms and extend to include algorithms, humans, data, business models, etc. Each of these aspects brings with it independent-but-interdependent motivations and perspectives.[15]

This understanding of the IoT takes inspiration from Object Oriented Ontology (OOO)[16] and has been encapsulated within the design metaphor

[13] Mortier, R., Zhao, J., Crowcroft, J., Wang, L., Li, Q., Haddadi, H., Amar, Y., Crabtree, A., Colley, J., Lodge, T. and Brown, T. 2016. December. Personal Data Management with the Databox: What's Inside the Box? In Proceedings of the 2016 ACM Workshop on Cloud-Assisted Networking (pp. 49-54). ACM.

[14] https://www.databoxproject.uk.

[15] Coulton, P & Lindley, JG. 2019. 'More-Than Human Centred Design: Considering Other Things' The Design Journal. https://doi.org/10.1080/14606925.2019.1614320

[16] Lindley, JG, Coulton, P, Akmal, H, Hay, D, Van Kleek, M, Cannizzaro, S & Binns, R. 2019. The Little Book of Philosophy for the Internet of Things. Lancaster: Lancaster University.

of *constellations*. The metaphor is intended to take this rather mind-boggling aspect of philosophy and put it in a form that is easy to grasp. The metaphor originated from Walter Benjamin's description of how the meaning of any collection of things varies wildly and is dependent on the observer's perspective, noting "*ideas are to objects as constellations are to stars*."[17] This idea naturally resonates with the More-Than Human perspective of the IoT.

In short, the metaphor suggests that, as with the constellations of stars in the night sky, IoT things are simultaneously 'stars' in their own right as well as being part of groups or constellations. For example, the physical Amazon Echo device is a thing in its own right, but it is also part of a constellation which involves the Echo and the cloud service which makes it work. An interesting property of constellations of stars is that depending on what perspective an observer takes, how they appear varies wildly. So, if you look at the stars from the Southern Hemisphere they appear 'upside-down' compared to the Northern Hemisphere perspective. Similarly, while IoT devices exist individually, their meaning and significance is wildly different depending on what angle one views the constellation from, so the user speaking to their Amazon Echo has a different perspective compared to the software engineer designing the cloud service which processes the user's voice.

Take a look at Figure 2 below. This shows some of the possible constellations formed for a voice assistant system such as Alexa or Google Home. These interdependent relationships serve to highlight different parts of the constellation, and how they have independent-but-interdependent perspectives. For example, viewpoint 1 represents the user perspective as might be considered with 'Human-Centred' approaches. The user is focussed only on the task of interacting with the voice assistant and the remaining system is basically invisible to them. If we shift focus to viewpoint 2 and consider the system from the service provider's perspective (which also includes the user and their device), we can see that the viewpoint is heavily influenced by their business model. This implies the user is seen primarily as a means of providing data both to improve the Natural Language Processing part of the service, but also to infer behavioural patterns of the user which can be used to better target advertising. This data, in its own right, may then also be traded on the open market.

[17] Benjamin, W. 1999. The Arcades Project, trans. Howard Eiland and Kevin McLaughlin.

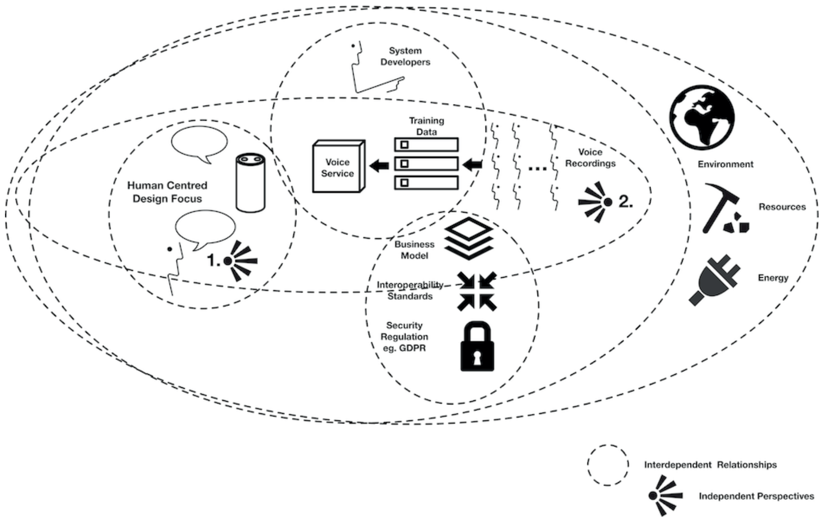

Figure 2. Potential Voice Assistant Constellation Diagram

So far, we have explored the rationale for HDI and More-Than Human Centred Design. However, it is important to explore how technologists and designers might put these approaches into practice. Using the HDI principles of Legibility, Agency, and Negotiability as our guide, the following sections will discuss a number of PETRAS research projects which have addressed these principles using the practical scenario of the home alongside HDI and More-Than Human Centred Design.

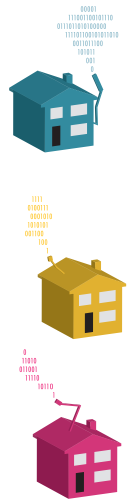

Legibility

It has often been said that *"a person's home is their castle"*. The castle's sense of security is based on its ability to deal with external threats; moats, draw bridges, and high, thick walls. They are not, however, as well designed against internal threats, such as sneaky invaders or treacherous members of the community.

Like the castle, our homes have walls that give us a sense of security and an expectation that the things or people brought inside are trusted and worthy of being invited. So, when we bring IoT devices into our homes, we compromise some of our security. By inviting these sophisticated, connected devices capable of automatically doing things into the very private and personal space of our home, we open ourselves up to a large problem: how do we trust that these devices are doing what they say they are, and how can we be sure they will continue to do what we want them to do?

Many IoT devices do things in the home that many people would object to. For example, eavesdropping on and recording everything said and done in the living room, bathroom, or even bedroom, then sharing this information with various others, including the police, governments, or advertisers. Despite our dislike of being watched like this, some devices today are already doing this, which highlights a crisis of legibility: *people cannot easily "read" their devices, in order to see what they're doing*. This includes whether the devices are turned on, when they're listening, and when they're sharing information. Smart home devices today don't even show users to whom such data are disclosed, for what purposes it will be used, and how long such data will last.

At this point, many of you may be thinking, *"What? But I didn't agree to any of this!"* And this brings us to the current dominant business model of the digital economy, which is known as surveillance capitalism. In this business model, data about individuals is seen and treated as a raw material for the profit-orientated manipulation of end-users.[18] In *surveillance capitalism*,

[18] Zuboff, Shoshana. "Surveillance capitalism." Esprit 5 (2019): 63-77.

a lack of legibility is seen as a benefit because a greater awareness of data collection practices are more likely to inspire people to object, find ways to avoid or stymie them, and thereby undermine their efficiency.

Although surveillance capitalism is widespread, it is the view of many that this sort of data collection is both unethical and depending on the interpretation of regulations, in some cases illegal (e.g. meaningful consent provisions within the European Union's General Data Protection Regulation). Thus, the path forward will very likely soon be towards making devices for the home not only legible but also *accountable* for their actions. Accountability extends the notion of legibility to make it possible for people to easily understand the reasons that a device did something, such as why they are collecting and transmitting data to particular entities or companies.

Together, legibility and accountability could also bring massive benefits beyond privacy. As systems become increasingly sophisticated, they become able to do advanced tasks that previously required a human assistant. Voice assistants, for example, can already help plan a weekend away, including providing personalised recommendations on hotels and sightseeing based on interests, schedules, and available discounts. But how can users tell whether the virtual assistant has their best interests in mind, or whether they are acting as an unscrupulous travel agent might? Many virtual assistants are created and run entirely by big platform companies, such as *Amazon*, *Apple*, and *Google*; how are users to know these platforms aren't manipulating (and surveilling) us for their own gain? Whose interests do IoT platforms prioritise?

Such conflicts of interest result in what is called a *moral hazard* in economics, when a lack of accountability protects a system's true actions or intentions from being discovered, those systems are likely to cheat and manipulate (due to the likelihood that they will get away with it). By the same argument, increased legibility and accountability incentivises systems to act honestly and respectfully towards their users, treating their needs genuinely. This is likely to, in turn, build trust, and accelerate adoption, by decreasing barriers to adoption and increasing the willingness to use advanced automation.

Legibility and accountability have an additional final benefit - helping people figure out what happened when things went wrong, and to help them understand whether the system might get it wrong again in the future. For instance, if a smart lock accidentally let a stranger in, a homeowner might be very interested in ascertaining whether the system would do so again in the future, in deciding whether to trust the system going forward.

In the following sections, we present *Data X-Ray* and *ARETHA* to illustrate how we might increase the legibility and accountability of IoT devices to their owners even when it is not being provided by existing smart home device makers.

Data X-Ray

The Data X-Ray project uses a combination of analysis methods to find out how smart devices are capturing, handling and disclosing data. It gathers intelligence from a number of sources:

1. Publicly available information about devices.
2. Applying what is called program analysis methods on available program code wherever it is available (such as Android apps).
3. Most important of the three, network analysis methods by watching all the network traffic passing out of devices to the Internet (see Figure 3).

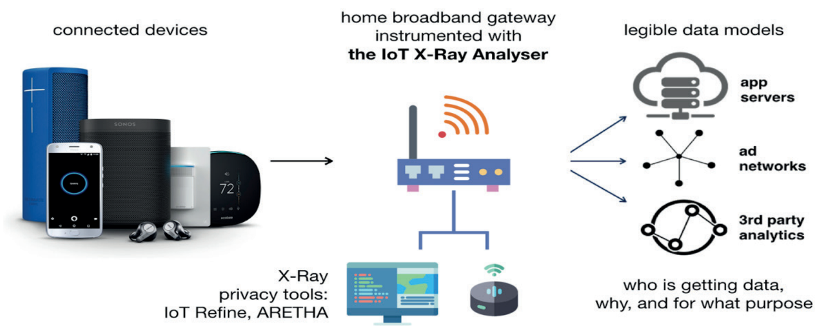

Figure 3. Data X-Ray System

The end result is that X-Ray makes legible all of the data-sharing activities, including the organisations (typically companies) that receive data from each device individually, and, the kinds of data that are shared and the purposes for which the data are likely used.

These data can be presented in several ways, including the IoT Refine interface (Figure 4), which is a *privacy disaggregator* that helps people improve their privacy by continuously monitoring and showing detailed information about data disclosures. It shows who receives data, what devices generate data, and who is responsible for data protection in the physical locations where data are stored.

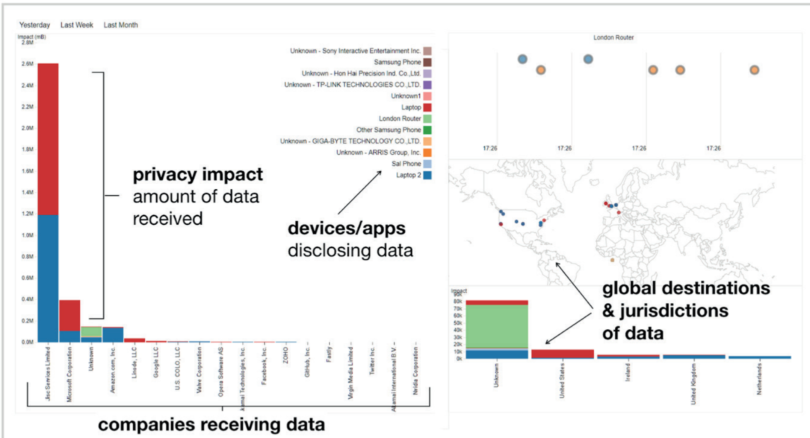

Figure 4. Data X-Ray Refine Interface

Respectful Devices and Project ARETHA

While the X-Ray approach increased the legibility of the data handling aspects of devices, it did not make other potentially important aspects of their operation legible or more accountable. A second important limitation was that while X-Ray's approach to displaying the data disclosure and use practices had the *potential* to help people take control of their privacy, realising such potential, in practice, required traversing a significant chasm: that of being able to interpret these data practices sufficiently to make reasonable judgements about their acceptability.

This led researchers to pursue legibility and accountability more intensively with a project called ARETHA (Artificial Respect Enabled Trust enHancing Agents). This project explores ways to increase end-to-end legibility and accountability in the design of systems from the bottom up, to think about how such devices might be re-designed around gaining users' trust. The first ARETHA system is a virtual assistant, which, like *Amazon's Alexa* or *Google Assistant*, will eventually have a large range of behavioural capabilities. But the ARETHA assistant differs from these platforms in several important ways:

1. Respectful by Design - ARETHA is designed to always act in the service of users' needs and wishes, following the philosophical (Kantian) notion of *respect*: to treat people genuinely as ends in and of themselves, above all others.[19] This is managed in three ways:
    - ARETHA will keep all sensitive data locally on the device, which will enable it to be one of the first privacy-respecting virtual assistants (most others send data as soon as it's captured to a cloud computing service to carry out voice recognition). In contrast, ARETHA will only use entirely local processing, so that what people say to ARETHA and what ARETHA hears stays local to the device.
    - All software that is part of the ARETHA will be created, owned, and maintained openly by a community tasked with ensuring the principles of respect are not violated.[20] Most virtual assistants are instead owned and run by large corporations who are incentivised to use virtual assistants to further their own commercial gain, and, in turn, keep the code secret and private.

[19] Van Kleek,M., Seymour, W., Binns, R., & Shadbolt, N. 2018. "Respectful things: Adding social intelligence to 'smart' devices," Living in the Internet of Things: Cybersecurity of the IoT - 2018, London, 2018, pp. 1-6.

[20] Whilst the formation of such a community is beyond the scope of the research presented we envision it being similar to the open source movement which has produced models and licenses which creators are expected to adhere to.

- ARETHA will be made legible and accountable by making the outcomes of its behaviours actually by the users themselves.

2. Explanations By Default - Whenever ARETHA uses machine learning or data-driven AI systems, ARETHA will ensure legibility through *machine explanations* and i*nterpretable models*, which provide succinct summaries of *how* data is used to arrive at specific results.
3. Tutoring - Since understanding is a key barrier to making informed decisions, ARETHA tools have intelligent, computer-based tutoring built-in which aims to help people become better informed about the complexities of the digital world through periodically updated digital skills curricula. ARETHA currently focuses such tutorials on privacy related concepts, but the educational role could be expanded to other activities that explore areas such as risk or security.

Whilst it is clear both X-Ray and ARETHA both demonstrate the HDI principle of agency, they do so in a way that draws upon More-Than Human Centred Design in that they effectively reveal the wider constellations in which they operate. Whilst legibility provides useful information to device owners they must be able to act upon this information, thus, in the next section, we discuss another principle of HDI - agency.

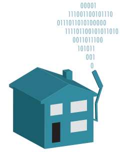

Agency

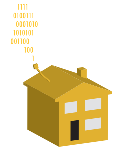

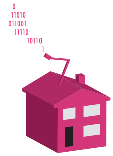

Previously we defined agency in HDI terms as providing users with control mechanisms over how their data is stored and utilised and we highlighted the current lack of agency provided by smart meters. In this section, we explore another project that focuses on agency by using the example of energy management in the home.

The promise of smart systems enhancing our lives has now reached the infrastructural level. In Distributed Energy Systems, for instance, batteries and algorithms balance the supply and demand of energy production and consumption. This is often referred to as a key challenge in energy futures. Distributed energy systems rely on technologies such as blockchains, smart contracts and high-capacity batteries. Using these innovations, batteries can connect directly to the power grid and carry out transactions. These both 'pull' and 'push' data as well as buying and selling energy.

The most widely tested technology in this area is called Vehicle-to-Grid (V2G), which uses decentralised storage to optimise price and convenience to final consumers and also to redistribute power to the grid. This provides an alternative way of balancing supply and demand. Used in this way electric vehicles play a dual role as consumers of energy but also as distributors of energy. Similar technologies have been tested in domestic appliances such as vacuum cleaners, which buy electricity and recharge themselves at optimal times through autonomous algorithms. Standards for digital 'tokens' facilitate these transactions, while also helping to quantify the energy and translate it into money.

Distributed systems, in contrast to centralised grids, require a lower initial investment. This allows smaller enterprises to produce and sell energy in a free market economy. However, these systems also present challenges of data management, and how to guarantee security and tracing of energy transactions. They also include some degree of autonomy to mediate transactions and balance energy availability and pricing; which raises questions of ownership and control. For example, who defines algorithms that support autonomy, for whose benefit do these algorithms operate, and how is this communicated to end consumers? Here agency is distributed between consumers, companies, and the system itself. What end-consumers see in this context is only the 'tip of the iceberg'. This hinders discussion on the impact of such systems in a wider sense. Ultimately, autonomous systems are designed in a way that balances the needs of different stakeholders. They process and share information, making decisions in the background, which inevitably may prioritise one of these stakeholders at particular times. When seen in terms of traditional user-centered design, a significant question remains; who has the power to make these decisions?

GigBliss Hairdryers

To discuss issues of control, autonomy and ownership in autonomous systems we designed three hairdryers as part of the GigBliss series; *GigBliss Plus*, *GigBliss Balance* and *GigBliss Auto*.[21] Similarly, to V2G systems, the hairdryers contain a battery that can connect to the grid and carry out energy transactions, buying and selling energy according to implicit rules embedded in the different models.

GigBliss Plus

In the first model, GigBliss Plus, end-users have higher levels of agency, as they are put in control of energy transactions. In this case, users can follow fluctuations in the market through a small display integrated into the device and can buy, store and trade energy through three buttons above and below the display to buy (pull) and sell (push) energy. The device can be activated through a third "on" button. In our imagined scenario, GigBliss Plus would be acquired at a high price but would allow end-users to make a profit, which could also cover the initial outlay. This concept alludes to a growing number of online platforms that support flexible and ad-hoc forms of employment, and where workers use their own tools. This form of work contrasts to "traditional" forms of work as it includes no security and little guarantee of long-term contracts. In this way, individuals would use their hairdryers and potentially other appliances to store energy and make money (perhaps as part of other

[21] Pschetz, L., Pothong, K. and Speed, C., 2019, April. Autonomous Distributed Energy Systems: Problematising the Invisible through Design, Drama and Deliberation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (p. 387). ACM.

"gig-economy" activities such as driving their Uber taxis or renting out spare bedrooms through Airbnb).

Figure 5. GigBliss Plus

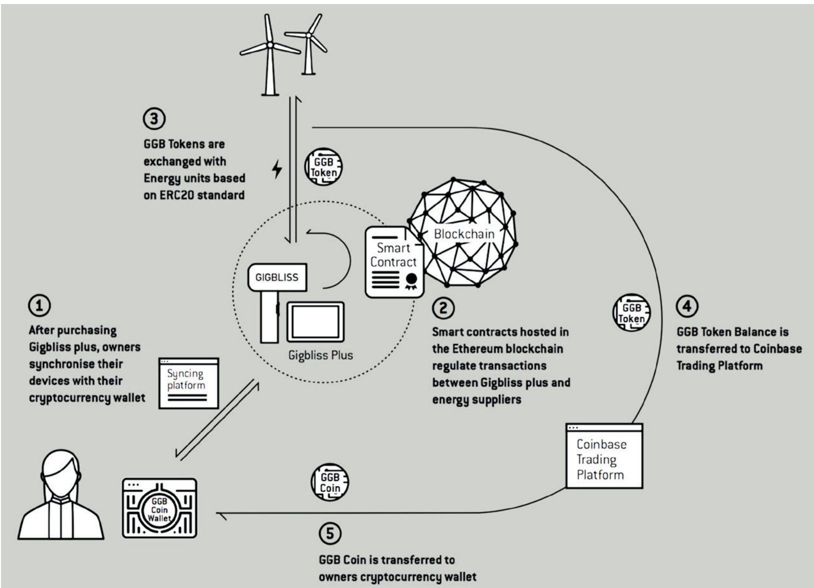

Figure 6. GigBliss Plus Transactional Model

GigBliss Balance

The second model - the GigBliss Balance - explores a scenario in which the GigBliss Corporation would have higher levels of agency in the system. The Balance is very cheap to buy (it could even be loaned, with no initial outlay for the customer) and would be returned to the corporation when it is no longer needed. Because they give the device out so cheaply the corporation recoups the costs by using the device to carry out energy transactions through predefined smart contracts hosted on a blockchain. While users can operate the device through an “on/off” button on the interface, its background operations are beyond their control. Such operations are indicated through an LED light that changes colour (green, yellow or red) to indicate that the device is available, available with a waiting time or unavailable, respectively. Waiting time and time of usage are indicated through a numeric display. While the user gives up agency, in return they have a very cost-efficient hair drying experience.

Figure 7. GigBliss Balance

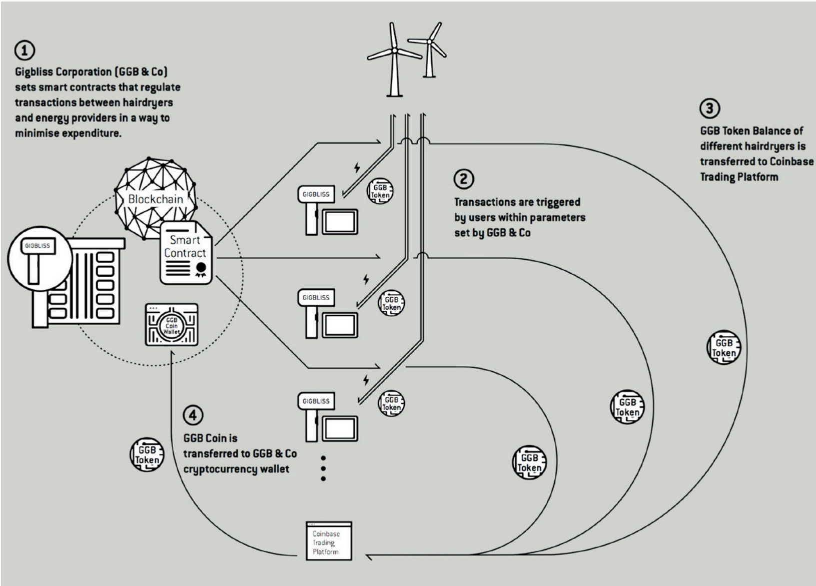

Figure 8. GigBliss Balance Transactional Model

GigBliss Auto

The third concept, GigBliss Auto, represents a model where a third party would subsidise costs of both devices and electricity supply for the device for a particular group. In this example, a local Council, community service, or charity establishes an agreement with an energy provider and the hairdryer company, paying a fixed price and setting up conditions for energy access. Here we invite reflection on the interests of stakeholders and what happens if this agreement attempts to reduce costs in order to maximise the number of households served by this scheme or if, for example, energy provision is set to occur during very specific or off-peak times in a way that regulates people's actions rather than vice-versa. Illustrating a rather dystopian scenario, GigBliss Auto has no buttons on its panel and offers users no control. An LED light and bar display indicates whether the device is about to turn on and for how long. The result is a hairdryer that turns on, on its own, at the times which most suit the needs of the energy agreement.

From these examples, it is clear that agency, as with legibility, is a complex notion to apply in practice and requires technologists and designers to address the complexities of the relationships between users and suppliers. In particular, it highlights that these new design approaches can help redress some power imbalance in favour of the user.

Whilst legibility and agency will undoubtedly improve our interaction with IoT devices we also need to recognise that in many cases there may be a cost. Users will likely need to negotiate the level of data they are prepared to make available against the functionality (and potentially the cost) of the services they wish to use. Thus the principle of negotiability is the subject of the next section.

Figure 9. GigBliss Auto

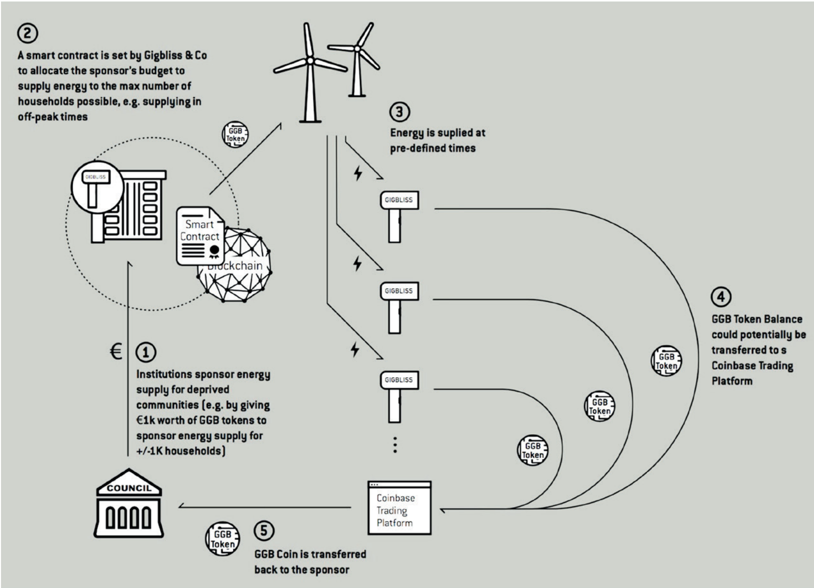

Figure 10. GigBliss Auto Transactional Model

Negotiation

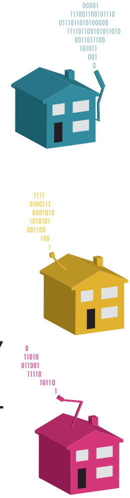

In this section, we acknowledge the transactional nature of data collection in that users are often trading functionality for access to their data. Currently, the vast majority of services offer little negotiation, often adopting an agree to all or get nothing approach. In order to address this topic, we created the Living Room of the Future (LRofTF), which is an Experiential Design Fiction.[22] The LRofTF mixes both real and fictional elements to allow us to situate the audience in a near-future world in which the negotiability of data access is brought to the fore. The LRofTF explores this by looking at how media broadcasters may utilise the potential of a technology called Object-Based Media (OBM) to deliver more immersive experiences to audiences in home environments. OBM allows the customisation of media, like radio and television, based on audience data. For example, it may customise a soundtrack based on your music preferences. OBM delivers personalised viewing experiences by breaking the programme into smaller parts (these are called media objects) and describing how they need to semantically relate to each other in order for the programme to make sense. This allows them to be dynamically reassembled into many possible personalised versions of the same programme. In addition to using OBM, the LRofTF uses IoT devices and external data sources to personalise the media even further, for example, by using the current weather to alter how the screen media is displayed, or use personal data to create customised sound-tracks for programmes. Finally, the IoT objects provide a physical means to contribute to an immersive media experience; for instance, the smart lights may adjust their colour and brightness automatically to match the overall look and feel of the content being shown on the screen.

We chose to build this prototype around a living room since this is a universally understood space within our own homes which most people are familiar with. Whilst the initial version of the LRofTF was designed as part of

[22] Coulton, P, Lindley, JG, Gradinar, AI, Colley, J, Sailaja, N, Crabtree, A, Forrester, I & Kerlin, L 2019, Experiencing the Future Mundane. in Proceedings of RTD 2019. Research through Design 2019 , Delft, Netherlands, 19/03/19.

a public exhibition for a specific installation in the FACT gallery in Liverpool (UK), it has been redesigned with new interactions and a new narrative for events at the Victoria and Albert Museum and the Tate Modern in London before going on permanent display in a 'Future Home' at the Building Research Establishment in Watford. You can see what the LRofTF looks like in the images shown in Figure 11, and it is this version which we are describing below. This takes into consideration the design decisions that went into the experience created by looking at the relationships of the three types of objects previously described; physical, media and data.

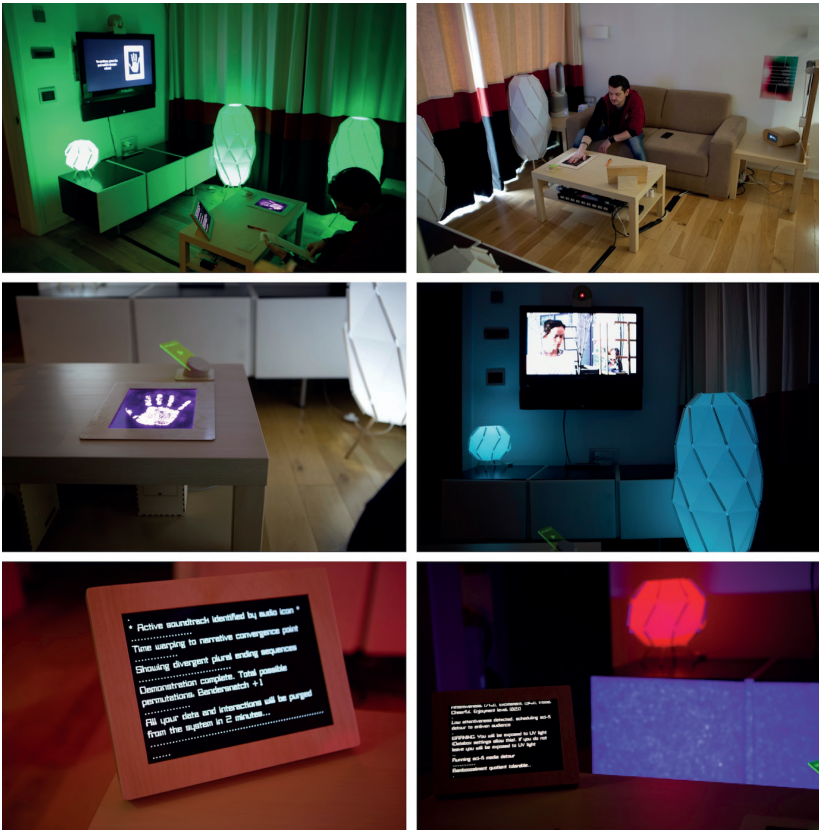

Figure 11. LFofTF System at Building Research Establishment, Watford (UK).

Physical Objects

The LRofTF was designed to represent a potential near future, and we identified a selection of off-the-shelf IoT products to put in it, which included programmable lights, a heating/cooling fan, window blinds, and smart plugs. In addition to these commercially-available products, we included a clock-radio whose speaker provides ambient sounds as part of the media; a series of sensors to detect audience interactions with objects in the room (including a drink coaster and a remote-control device); a coffee table with in-built hand sensor, display, and thermal printer; and a voice-activated LED 'eye' which provides a personality for the living room and acts as the camera for its face-scanning technology. Whilst the commercial products and the printer may be considered as 'outputs' of the LRofTF, the sensors should be seen as 'inputs' that generate data the system then uses to personalise the experience.

Media Objects

The expanded version of the LRofTF uses a short drama called The Break-Up, specially commissioned by the BBC's Research and Development department to highlight the potential of dynamic programming for television. The Break Up is a special type of programme which utilises an innovation known as 'Perceptive Media' which had previously been demonstrated for radio.[23] The Break Up was written and filmed in such a way that the entire narrative can be shifted to accommodate the viewer. For example, there are two contrasting endings (one positive, one negative) and two paths through the story (one emphasising the male character, and one highlighting the female). There is even an alternate version of the story where the gender roles of the characters are swapped, and rather than existing in the present day it takes place in a Science Fiction alternate universe. Further, it allows for dynamic changes of the soundtrack to better reflect the experience. By breaking the media into its constituent 'objects' the LRofTF can utilise the OBM system to provide a personalised and unique experience to each viewer.

[23] Gradinar, A, Burnett, D, Coulton, P, Forrester, I, Watkins, M, Scutt, T & Murphy, E 2015, Perceptive media: adaptive storytelling for digital broadcast. in Proceedings of INTERACT 2015 The 15th IFIP International Conference on Human-Computer Interaction, Bamberg, Germany, pp. 586-589, 14/09/15. https://doi.org/10.1007/978-3-319-22723-8_67

Data Objects

The media objects of the film can be reconfigured based on data objects. This means that the system will attempt to construct a new version of the programme based on data which provides contextual information about the viewer and the environmental conditions. For example, data points such as the current weather, the location, time, number of people in the room, music preferences, political leanings, can all be used as data inputs. These data points are combined with further information from the IoT sensors in the room (e.g. whether the audience is smiling, if they are looking away from the screen, or whether they are moving around in their seats). Combined together—and processed by the OBM system—these data are used to deliver a completely personalised version of the film which is optimised to maximise immersion.

Technical Challenges

Designing and building the LRofTF was a significant challenge and required a number of bespoke technical layers to function properly. Given that each IoT device uses different protocols and shares data differently, we had to create a bespoke system for translating this data into a format that the LRofTF could understand. In order to address the privacy challenges of the experience (which acted as if it accessed various kinds of highly personal data) all data moving through the system is managed through a Databox providing a unique ecology for exploiting personal data in privacy-preserving ways. For example, Databox can enable a media provider to utilise algorithms that process data referring to an individual's viewing habits in order to offer bespoke content, but can do so *without* disclosing personal data directly to the provider. Instead of distributing personal data to remote cloud servers for processing, processing takes place on-the-box (this is called 'edge computing'[24]), preserving privacy by ensuring no personal data needs to leave the home or be accessed remotely.

[24] Mortier, R., Zhao, J., Crowcroft, J., Wang, L., Li, Q., Haddadi, H., Amar, Y., Crabtree, A., Colley, J., Lodge, T. and Brown, T., 2016, December. Personal data management with the databox: What's inside the box?. In Proceedings of the 2016 ACM Workshop on Cloud-Assisted Networking (pp. 49-54). ACM.

LRofTF Experience

The experience of the LRofTF is split into three parts. To begin, the participants seat themselves on the sofa in front of the television screen. The experience is then introduced using a voice user interface which seeks to gain consent from users to collect, process and store their data (the LRofTF prints out a permission slip using the thermal printer embedded in the coffee table, which the audience must sign to proceed). In the second part of the experience, the film is played based on a profile generated by the system. During this phase, various IoT objects in the room begin to contribute to the immersion. For example, at the start of the film, the blinds come down and the room's lighting adapts to each scene (the system 'knows' the outside weather and picks up a relevant colour gradient). When the lead character in the film is outdoors, the fan switches on, matching the wind blowing her hair. The music within the film is chosen dependent on the profile generated by the system, as is the chosen ending (which depicts the character either leaving or staying with her abusive partner). By the middle of the story, if the audience appears bored (based on sensory inputs) an IoT smart plug is triggered by OBM to turn on an Ultra Violet (UV) light during which a short section of the science fiction film is shown before returning to the main drama. The impact of particular data interactions which affect the drama does not immediately affect the media objects, which means that while each experience was uniquely tailored to the audience, they would not necessarily be able to see why or how. Therefore, the tablet on the coffee table highlights when data is being collected and subsequently used. Finally, as the audience has only experienced one of the many possible variations of the drama, an explainer video shows the variations which could have occurred, and why, as shown in Figure 12. At the end, the system generates a receipt for the audience for the data generated. We do this as part of the experience to further highlight the transactional nature of the majority of personal data collection, in that we are trading system functionality for the participants' data.

Although partly a means of including a fictional consent procedure within the experience, the living room's voice actually highlighted and prototyped the HDI aspect of negotiability. The consent involved introducing each sensor in turn, starting with the fictional face recognition system and at each point, the audience was asked to indicate their willingness to have the data

Figure 12. Explainer Video for The Break-Up

collected. Whilst this provided legibility, we purposefully did not provide a choice other than 'Yes' or 'No'. This was intended to highlight the lack of real choice that many consent systems actually provide. When the audience in the LRofTF said 'no' the system would say that this was a shame as they would miss out on the video but they could exit through the gift shop. However, during the hundreds of times the experience was run, only one person declined to consent, and only once, and then quickly changed their mind. This perhaps indicates how engaging voice can be as an interaction, which was further illustrated when we asked the audience to hold up their receipt at the end of the consent section and smile at the camera which they all dutifully did. The beguiling nature of voice is perhaps a problem for future IoT systems in that, if their security is compromised, they may present nefarious hackers a highly effective means of phishing.

The need to 'design in' negotiability is also made clear when considering 'More-Than Human Centred' theory.[25] Most designed things, and the components that make them up operate familiarly and there is no need to negotiate consent around their use. For example, things such as taps, doorknobs, light switches, and cars have, through a process sometimes referred to as 'mediation'[26] or 'domestication',[27] become so very familiar that virtually anyone knows without thinking what to expect from them. Occasionally, technological innovation upsets our familiar relationships with things, and we need help in renegotiating them. For example, car wing mirrors that increase the field of view but make objects appear smaller highlight this to drivers. In some countries these wing mirrors must carry a disclaimer, this begins a kind of dialogue with the user: because the technology has changed, it must increase its negotiability. In the case of connected products in the home, rather like the wing mirror, although outward appearance remains largely similar, the inner workings are often very different. For this reason, our relationships must be renegotiated.

[25] Coulton, P & Lindley, JG 2019, 'More-Than Human Centred Design: Considering Other Things' The Design Journal. https://doi.org/10.1080/14606925.2019.1614320

[26] Verbeek, P., 2015. Beyond interaction. interactions, 22(3), pp.26–31. Available at: http://dl.acm.org/citation.cfm?doid=2767137.2751314.

[27] Silverstone, R., 2006. Domesticating domestication. Reflecting on the life of a concept. In T. Berker et al., eds. Domestication Of Media And Technology. Open University Press, pp. 229–247.

Conclusion

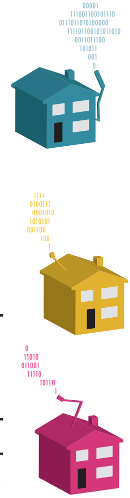

Through the projects presented in this Little Book, we have considered alternative design approaches inspired by HDI and More-Than Human Centred design. In particular, each project highlights the potential benefits of these approaches by presenting experiences through which users directly experience potential near futures. In doing so, we allow users to challenge the design practices imposed by technology and open up the debate on what may be the consequences introduced by our increasingly data-enabled homes.

The colonisation of the home by IoT has already begun and with it comes the creation of a 'data exhaust', which largely goes unnoticed by the people residing in these homes. In the same way that motorcars offer many tangible benefits, the exhaust they produce is a necessary—but largely undesirable—by-product. The data exhaust from IoT products isn't noxious like petrol fumes, and in fact it has the potential to produce many positive aspects for future home residents—greater energy efficiency, better health care in the home, more immersive media experiences—but it also has the potential to increase our exposure to highly sophisticated personalised cyber-attacks.

In this book, and throughout the PETRAS project, we have sought to re-imagine the way IoT products and services are currently designed so that users are not simply treated as data-point providers but willing collaborators with the technology. This is achieved by addressing issues such as legibility, agency, and negotiability using approaches such as HDI and More-Than Human Centred design. Whilst new challenges will emerge along the way, it is important we address the tricky elements of IoT data collection before they become highly problematic and detrimentally affect the future adoption of IoT products and services in the home.

References

Akmal, H & Coulton, P 2018, Using Heterotopias to Characterise Interactions in Physical/Digital Spaces. in C Storni, K Leahy, M McMahon, P Lloyd & E Boehmia (eds), Proceedings of the Design Research Society Conference 2018. vol. 1, Limerick, pp. 269-278.

Benjamin, W. 1999. The Arcades Project, trans. Howard Eiland and Kevin McLaughlin.

Coulton, P., Lindley, J. G., & Cooper, R. 2018. The Little Book of Design Fiction for the Internet of Things. Lancaster: Lancaster University.

Coulton, P, Lindley, JG, Gradinar, AI, Colley, J, Sailaja, N, Crabtree, A, Forrester, I & Kerlin, L 2019, Experiencing the Future Mundane. in Proceedings of RTD 2019. Research through Design 2019 , Delft, Netherlands, 19/03/19.

Coulton, P & Lindley, JG 2019, 'More-Than Human Centred Design: Considering Other Things' The Design Journal. https://doi.org/10.1080/14606925.2019.1614320

Gradinar, A, Burnett, D, Coulton, P, Forrester, I, Watkins, M, Scutt, T & Murphy, E 2015, Perceptive media: adaptive storytelling for digital broadcast. in Proceedings of INTERACT 2015 The 15th IFIP International Conference on Human-Computer Interaction, Bamberg, Germany, pp. 586-589, 14/09/15. https://doi.org/10.1007/978-3-319-22723-8_67

Lindley, JG, Coulton, P, Akmal, H, Hay, D, Van Kleek, M, Cannizzaro, S & Binns, R 2019, The Little Book of Philosophy for the Internet of Things. Lancaster University.

Morozov, E., 2013. To save everything, click here: The folly of technological solutionism. Public Affairs.

Mortier, R, Haddadi, H, Henderson, T, McAuley, D, Crowcroft, J, and Crabtree, A. 2016 Human-Data Interaction: The Encyclopedia of Human-Computer Interaction, 2nd Ed. | Interaction Design Foundation.

Mortier, R., Zhao, J., Crowcroft, J., Wang, L., Li, Q., Haddadi, H., Amar, Y., Crabtree, A., Colley, J., Lodge, T. and Brown, T., 2016, December. Personal data management with the databox: What's inside the box?. In Proceedings of the 2016 ACM Workshop on Cloud-Assisted Networking (pp. 49-54). ACM.

Pschetz, L., Pothong, K. and Speed, C., 2019, April. Autonomous Distributed Energy Systems: Problematising the Invisible through Design, Drama and Deliberation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (p. 387). ACM.

Silverstone, R., 2006. Domesticating domestication. Reflecting on the life of a concept. In T. Berker et al., eds. Domestication Of Media And Technology. Open University Press, pp. 229–247.

Stead, MR, Coulton, P, Lindley, JG & Coulton, C 2019, The Little Book of Sustainability for the Internet of Things. Imagination Lancaster, Lancaster.

Van Kleek,M., Seymour, W., Binns, R., & Shadbolt,N. 2018 "Respectful things: Adding social intelligence to 'smart' devices," Living in the Internet of Things: Cybersecurity of the IoT - 2018, London, 2018, pp. 1-6.

Verbeek, P., 2015. Beyond interaction. interactions, 22(3), pp.26–31. Available at: http://dl.acm.org/citation.cfm?doid=2767137.2751314.

Zuboff, Shoshana. "Surveillance capitalism." Esprit 5 (2019): 63-77.

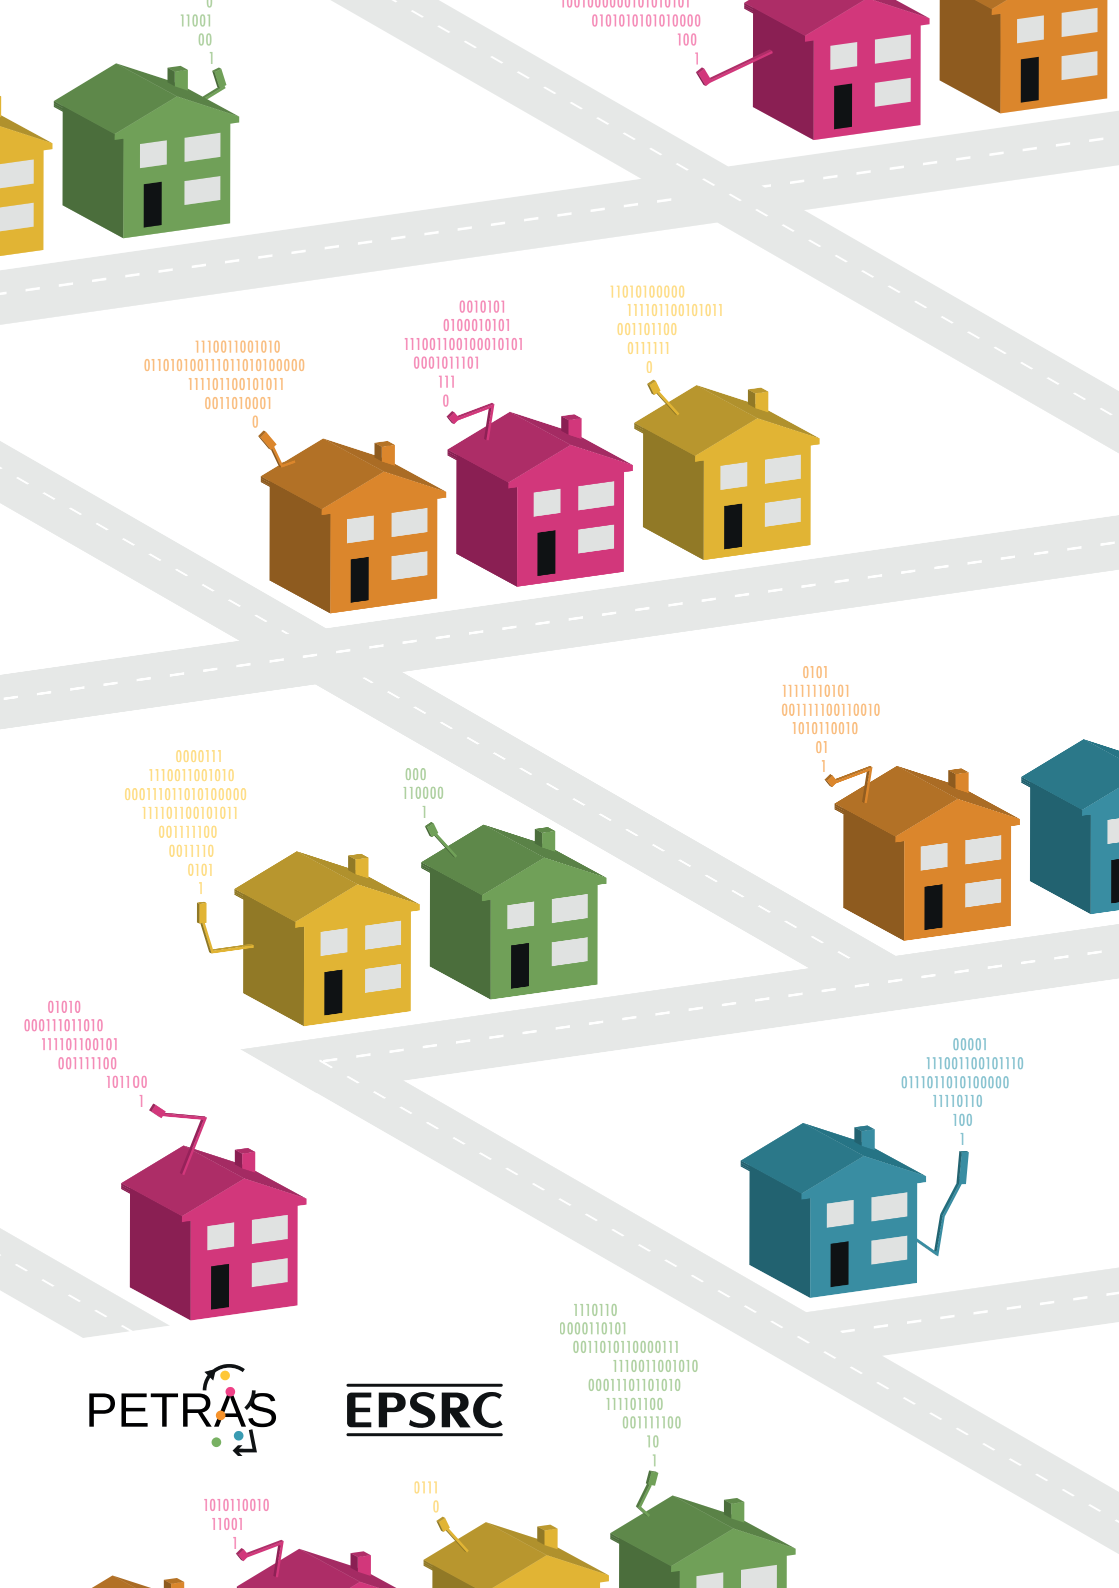
PETRAS
EPSRC