Building Accountability into the Internet of Things

Andy Crabtree, Tom Lodge, James Colley and Chris Greenhalgh
School of Computer Science
The University of Nottingham
Email: {First name. Last name}@nottingham.ac.uk

Richard Mortier
Computer Laboratory
UNIVERSITY OF CAMBRIDGE
Email: richard.mortier@cl.cam.ac.uk

Abstract. Privacy is a matter of widespread societal concern, which will be further exacerbated by the ongoing effort to connect billions of data-generating devices to the Internet. The promise of large-scale social and economic benefits motivates the development of the Internet of Things. As a primary engine of personal data production in the near future it also carries with it real threats to privacy. Proposed legislation seeks to tackle these threats head-on. This paper seeks to explore principled ways in which systems design might engage with and respond to legal requirements to enable effective data protection and the projected utility of personal data in the digital economy. Key to this is the need to build *accountability* into the IoT to ensure compliance.

Keywords. Internet of Things, privacy, data protection, accountability, IoT-Databox model.

1. Introduction

> " … observers have made a number of predictions for the IoT. We are told that, in 2015, the world will have 25 billion connected devices; the number of smart home devices will reach nearly 25 million; and IoT software platforms will 'become the rage' … the IoT … has the potential to provide enormous benefits for consumers, but it also has significant privacy … implications … In the not too distant future, many, if not most, aspects of our everyday lives will leave a digital trail … Connected devices are effectively allowing companies to digitally monitor our otherwise private activities … companies are investing billions of dollars in this growing industry; they should also make appropriate investments in privacy …" Edith Ramirez, Federal Trade Commission, 2015 [19]

Edith Ramirez's comments underscore the projected benefits of the Internet of Things (IoT) and the risks that accompany it, particularly the threat to privacy. Widespread recognition of this threat is driving new approaches to data protection in Europe and the US [6, 25]. Key to proposed legislation is the notion of *accountability*, which refers to the demonstrable implementation of enforceable privacy policies and procedures. Accountability is also a key concept in HCI and the sub-field of Computer Supported Cooperative Work (CSCW), where it refers to the 'observable and reportable' character of human and system actions. Clearly the two concepts are not the same, but they are not mutually exclusive either and the possibility exists to align the two to address the privacy threat occasioned by the IoT. This paper explores the potential alignment of legislative and CSCW concepts of accountability and elaborates a distinctive interactional model and associated interaction mechanisms enabling the demonstrable implementation of enforceable privacy policies and procedures. We first explore legislative conceptions of accountability, followed by the concept of accountability in CSCW, which for the sake of convenience we label 'computational' accountability in a manner consonant with Dourish and Button's notion of accountability in systems design [11]. We then turn to consider the potential alignment of the two conceptions, focusing on practical measures that policy-makers envision will enable the 'external' accountability required by legislation; measures that are intended to allow users to exercise control over the flow of data in the digital economy and thereby manage their privacy. The external accountability requirements of proposed legislation motivate the design of the IoT-Databox model - a computational model which makes the actors, activities and 'things' implicated in data harvesting and use accountable to end-users, thus enabling them to exercise control over their personal data, and in doing so allows 'data controllers' or parties who would process an individual's data to demonstrate compliance with proposed data protection regulation.

2. Accountability in Proposed Data Protection Legislation

Two fundamental categories of accountability are to be found in proposed EU and US legislation: 'internal' and 'external'. We briefly review each of these in turn below.

2.1. Internal Accountability

Internal accountability refers to the policies and procedures that a 'data processing' organisation – e.g., government department, commercial entity or charity – puts in place to *demonstrate to itself* that its data processing operations comply with the requirements of proposed data protection legislation. The EU proposal [6] lays down a raft of legally binding requirements that 'data controllers' (i.e., those who commission and are responsible for the processing of personal data) must comply with as detailed in Articles 22 to 43. These requirements stipulate detailed documentation of data processing operations, security provisions, and the risks that attach to the categories of data being processed. The latter is called a 'privacy impact assessment' or PIA, and authorisation is required from a supervisory authority to process high risk data (e.g., genetic and biometric data). Proposed EU legislation also requires that any enterprise that routinely engages in data processing and employs more than 250 people, or any enterprise whose core operations rely on the regular and systematic monitoring of 'data subjects' (i.e., individuals), or any public body that processes personal data, must appoint a data protection officer who must oversee data processing operations, operate independently of the controller, and report directly to management. No such requirement is mandated by proposed US legislation, which seeks to implement a voluntary code of conduct enforced through Federal acts prohibiting unfair or deceptive practices, though it does recommend the use of PIAs as an effective means of conducting 'structured assessments' of the potential privacy issues arising from new IT systems [25]. While the US proposal is very different from the EU proposal, both agree on the need to put rigorous internal accountability policies and procedures in place.

2.2. External Accountability

More relevant to this paper is the notion of 'external' accountability. External accountability refers to the policies and procedures a data processing organisation puts in place to *demonstrate to others outside the organisation* that its operations are in compliance with the requirements of proposed data protection legislation. Those 'others' fall into two essential categories: regulatory authorities and the individuals whose data an organisation processes. Being accountable to regulatory authorities means that an organisation can demonstrate that it has put internal data protection policies and procedures in place, including policies and procedures that make novel re-uses of data and third-party transfers accountable to regulatory authorities and individuals alike, and that it has also implemented external policies and procedures, *particularly with regards to the individual*. In proposed EU legislation such policies and procedures fall under the auspices of 'consent' as laid out in Articles 6 to 20. These articles specify that data processing must be lawful and is only lawful if the data subject has given consent to the processing of their personal data for one or more specific purposes [6].

Proposed EU legislation specifies a range of conditions that attach to the giving of consent, including providing information concerning the processing of personal data to the data subject in an 'intelligible form, using clear and plain language' (Article 11). The information provided must minimally consist of the identity and the contact details of the controller and data protection officer; the purposes of data processing, including the contract terms and general conditions; whether the provision of personal data is obligatory or voluntary, as well as the possible consequences of failure to provide such data; the recipients of the personal data; the period for which the personal data will be stored; if applicable, that the controller intends to transfer personal data to a third country or international organisation and the level of protection afforded in such circumstances by reference to an adequacy decision by the Commission; any further information necessary to guarantee fair processing. The controller shall provide this information at the time when the personal data are obtained from the data subject, or at the time data is recorded, or within a reasonable period after data collection (Article 14). The controller shall also inform the data subject of their rights, including the right of access (Article 15) the right to rectification (Article 16), the right to be forgotten and to erasure (Article 17), the right to data portability (Article 18), the right not to be subject to measures which produces legal effects based solely on automated processing of data pertaining to performance at work, economic situation, location, health, personal preferences, reliability or behaviour (Article 20), and the right to lodge a complaint to the supervisory authority and the contact details of the supervisory authority. Consent is a legal requirement of all data processing operations in Europe, though reuses of data that are 'compatible' with the original purposes for which it was gathered are allowed without further consent [4]. Failure to comply with the requirements of consent may result in a fine of up to €1,000,000 or up to 2% of an organisation's annual (worldwide) turnover.

While the US proposal is voluntary and punitive sanctions are therefore absent, the emphasis on external accountability to the data subject is just as pronounced, this time under the auspices of 'notice and choice' enshrined in 6 of the 7 principles that constitute the 'consumer privacy bill of rights' (the other principle focuses on security and the implementation of adequate internal accountability measures to ensure risks are identified and controlled) [25]. Thus, the Individual Control principle stipulates that consumers should be offered clear and simple choices about personal data collection, use, and disclosure, and means to withdraw or limit consent that are as easily used as the methods for granting consent in the first place. The Transparency principle stipulates that consumers should be provided with clear descriptions of the personal data to be collected, why the data is needed, how it will be used, when it will be deleted or de-identified, and whether and for what purposes it may be shared with third parties. The Respect for Context principle stipulates that data controllers should limit their use and disclosure of personal data to those purposes that are consistent with both the relationship that they have with consumers and the context in which the data was originally disclosed. The Access and Accuracy principle stipulates that consumers should be provided with reasonable access to personal data, and appropriate means and opportunity to correct inaccuracy, limit data (re)use, or request its deletion. The Focused Collection principle stipulates that only as much personal data as is needed to accomplish purposes specified under the Respect for Context principle is gathered. And the Accountability principle stipulates that data controllers should be accountable to enforcement authorities and consumers for adhering to these principles.

2.3 Accountable Data Processing

While proposed legislation is clearly different in the EU and US, one being extremely detailed and mandatory the other outlining key principles and voluntary, the emphasis on external accountability to the data subject is pronounced in both and, sanctions aside, both are very similar. Thus, and for example, external data subject accountability turns in either case upon processing data for specific purposes or purposes that are consistent/compatible with those for which it was originally collected, and conveying information related to data processing in ways that clearly describe the purposes of data processing, data retention policies, data transfer policies, and the data subject's access, rectification and data deletion rights. There are differences, the EU legislation clearly attaches more complex rights to the processing of personal data, but both recognise and stipulate the need to implement external data subject accountabilities in the processing of personal data:

> "Privacy protection depends on companies being accountable to consumers as well as to agencies that enforce consumer data privacy protections." [25]

Accountability 'frames' the processing of personal data. It does not specify *how* accountability is to be achieved, only *what* it must consist of, and what it must consist of and provide is a demonstration that adequate privacy-preserving measures have been put in place:

> "Accountability refers to a company's capacity to demonstrate the implementation of enforceable policies and procedures relating to privacy (whether adopted voluntarily or as a result of legal obligations)." [25]

The demonstration is key and applies equally in Europe and America, not only to companies but to public bodies, not-for-profit organisations or any entity that processes personal data as part of its daily business. The demonstration is two-sided, consisting of 'internal' and 'external' accountability. Internal accountability consists of implementing policies and procedures that enable a data processing organisation to demonstrate to *itself* that it has put adequate data protection measures in place, e.g., privacy impact assessments and data protection officers. External accountability, conversely, consists of demonstrating to *others* that adequate data protection measures are in place and it is here that making data processing accountable to the individual is key.

3. Accountability in HCI

Accountability is an important concept in HCI where it has been used to describe the challenges involved in making the behaviour of computational systems 'observable and reportable' in order to better enable human-computer interaction [11]. By way of explaining the notion of accountability at work here Dourish and Button first note that,

> "The act of systems design is the creation and manipulation of abstractions … Abstractions help us manage complexity by allowing us to selectively hide it. In systems design, abstractions typically function as 'black boxes'. They are defined by the nature of their interactions with the outside world (human users or other pieces of code - the 'clients' of the abstraction), which are typically defined in terms of the available functionality, procedure call conventions and return values - what we typically refer to as the 'interfaces' to the abstraction. The system's internal mechanisms, which describe and control how it goes about doing the work it does, are intentionally not available to inspection … Human users interact with abstract interfaces to the system's functionality (such as a print dialogue, or a direct manipulation view of a filesystem) which provide simple, consistent interaction by hiding the complex realities of the system mechanisms (creating a Postscript file and sending it to the printer, or copying files from a local disk to a server across a long-distance network connection)." [11]

Dourish and Button go on to elaborate the benefits of making computational abstractions accountable to end-users in presenting a familiar 'file-copying scenario'. The scenario involves dragging and dropping files into a folder over a network, the progress of which is reflected in a file transfer percentage bar. At 40% the copy operation fails. The authors ask, what resources are available to the user to understand what has happened, and to understand what options are now available? Have 40% of the files been read from the local drive or been written onto the network drive, or have 40 % of all the files been read or written? There is no way of telling. The computational abstraction ordering file copying hides the details upon which such understandings could be based. The alternative is to make the underlying file copying mechanism 'translucent', which involves surfacing the 'structural properties' of system hidden by the file copying abstraction. For example,

> "Between the file source and destination are arrayed a number of staging posts (data buckets). File data flows from the start-point to the end-point by moving from one bucket to another along a data path. As data flows from one bucket to the next, the buckets are related to each other by flow strategies, by which the movement of data from one to the next is regulated ... The flow of data through these, and the activation of the flow strategies, provides a framework for the relationship between the action in which the system engages and the reading and writing of data files." [11]

Thus, surfacing underlying behaviour of computational systems, e.g., making data buckets and data flow observable and reportable, provides resources that users can draw on to answer the sorts of questions raised above – i.e., to monitor the flow of data and determine just where within it copying has reached when it stalls, for example, thereby allowing the user to make some sense of what the percentage-done bar is actually reporting.

Accountability in HCI – computational accountability as we shall call it – thus refers to the surfacing or making visible of computational behaviours or actions to better enable human-computer interaction. These behaviours are traditionally masked by abstractions in systems design to hide complexity. The IoT is a massively complex computational environment, distributed across an increasingly dense and differentiated layering of people, activities and 'things' – e.g., sensors in people's homes, distributing data to multiple cloud servers, and processing data in manifold locations. How are we to get a handle on this and make the IoT computationally (and legally) accountable? One potential starting is provided by Robertson and Wagner [20], who construe of the IoT as a social-technical network – i.e., a distributed network of 'things that socialise' with one another and with human beings. IoT devices may therefore be seen as 'things' situated in *cooperative arrangements* of data processing. Robertson and Wagner thus suggest that key concepts in CSCW (a sub-field of HCI) might enable design to negotiate the boundaries between (networks of) objects and people, making them transparent, understandable and adaptable.

The most fundamental concept in CSCW is that of cooperative work, at the core of which is the notion of *interdependence in work*, which is to say that the parties to some job of work (e.g., of data processing) are mutually dependent in their work and are therefore required to cooperate in order to get the work done. Interdependence in work means that any cooperative effort thus involves *secondary activities* to mediate cooperative relationships. Cooperative work is not, then, simple a matter of A doing this and B doing that. Individual tasks must be 'articulated' if they are to be effectively meshed together. As Schmidt and Bannon put it,

> " … by entering into cooperative work relations, the participants must engage in activities that are, in a sense, extraneous to the activities that contribute directly to fashioning the product or service … That is, compared with individual work, cooperative work implies an overhead cost …" [22]

The idea that cooperative arrangements involve a necessary overhead, which involves the 'articulation' of cooperative relationships, is key to the design of cooperative systems and rooted in the sociological studies of Anselm Strauss [24]. Strauss studied the conduct of work in organisations, which he conceived of as taking place within a division of labour. He was not the first to conceive of the division of labour, it had long been recognised in sociology, economics and other fields before it, but he did adopt a unique orientation to and understanding of it. Thus Strauss sought to conceptualise the division of labour in terms of work itself, that is as something brought about through the very *doing* of work, rather than as an external structure which contained work and shaped it. This 'interactionist' orientation conceives of work as consisting of a distinct 'project', e.g., building a house, or a computer, or processing personal data. The accomplishment of the project consists of an 'arc of work', i.e., the totality of the tasks involved in accomplishing the project. Strauss was interested in how discrete tasks within complex projects and arcs of work are 'sequentially and simultaneously' carried out, such that *one gears in with the other*, and discrete 'clusters of tasks' and then larger 'segments' are *meshed together* to form the whole. Strauss called this gearing in and meshing together or integration of tasks 'articulation work', which he described as a 'supra type of work', an unavoidable add-on to the doing of individual tasks.

There is no time out from articulation work within the division of labour. Individual tasks *must* be articulated if they are to be geared into one another, and this minimally involves the 'allocation of articulation tasks', including what is to be articulated, by whom, when, where, how, and so on thereby making every participant in the work 'accountable'. Accountability here means that every worker is responsible for doing their assigned portion of the arc of work and rendering its accomplishment *reportable* to others according to certain criteria; e.g., when it was done, where it was done, how it was done, by whom, how much of it was or has been done, what is left to do, etc. It is tempting to see parallels here with the legal sense of accountability – i.e., the demonstration that proper procedure has been implemented and complied with – and this is underscored to some extent by Strauss's notion of an 'accountability ladder' along with the idea that accountability can be 'lateral', 'downwards', or upwards to 'higher authority' and is 'enforced by law, custom, organisational rule, or mandate'. We should be careful about reading Strauss's notion of accountability as being equivalent to legal conceptions, however. For Strauss-the-interactionist accountability is a key ingredient in the very *doing* of work itself:

> " … accountability requires the work of reporting accountability"; … reporting is a crucial condition for further interaction … in the total division of labour; … [thus] built into any division of labour is a system of accountability which has direct bearing on the carrying out of types of work and their implicated tasks." [24]

On this view then accountability is an integral feature of the ongoing accomplishment of collaborative action. It does not *frame* action, as per the legal conception, but is an *inherent feature of its conduct* essential to the gearing in and meshing together of cooperative work. It consists of making action *reportable* to others so that they can *see* what's been done, what's going on, what needs to be done next, etc., and can integrate their actions accordingly. Accountability in CSCW thus refers, as it does in HCI more generally, to the 'observable-reportable' character of action, but here the emphasis is on articulating individual actions within cooperative arrangements rather than surfacing computational behaviours masked by abstractions. The concept of accountability in CSCW thus extends the computational concept of accountability from a matter of making computational behaviours observable and reportable to enabling arrangements of computer supported cooperative work as well.

4. Reconciling Legal and Computational Accountabilities

The legal notion of accountability is not the same as the computational notion. Legally construed, accountability requires that data processing organisations demonstrate that they have implemented adequate privacy-preserving measures, with the adequacy of the matter being defined by legislation. Computationally construed, accountability requires that the behaviour of computational systems be surfaced and articulation work between actors be supported. How can we reconcile these different notions? Indeed, why should we? It would not be unreasonable to suggest that we can take the privacy problem in contemporary life as read and simply busy ourselves with working out computational solutions. This might not be an unreasonable response, but it is not sufficient to simply acknowledge the problem. If we are to address it *effectively* we need to be able to respond *appropriately* to the requirements of proposed data protection legislation generally, and for the purposes of this paper, what they might amount to for the IoT as a primary engine of personal data production in particular. So how might we reconcile legal and computational accountabilities? The external data subject accountability requirements of proposed legislation provide a prima facie point of connection. External accountabilities in legislative proposals make the data controller – the party or body that would consume personal data – accountable not only to supervisory authorities and enforcement agencies, but to the data subject or individual as well, and proposed legislation anticipates systems design to intervene and provide computational tools to enable accountability in this respect.

Draft US legislation is explicitly concerned to enable users to 'control how personal data flows in the digital economy' and provide them with 'better tools to do so' [25]. Accordingly, the bill anticipates 'innovative technologies' that provide detailed privacy settings, do not track, and opt out mechanisms, with the caveat that while such mechanisms are beginning to 'show promise' they require further development. While the EU proposal does not explicitly encode an Individual Control principle and is technologically 'neutral', the issue of individual control prefaces proposed legislation (see the explanatory memorandum). The issue of individual control provides HCI and systems design with a common *focal point* for building accountability into the IoT, one that becomes sharper when we turn to the agencies tasked with translating legislation into 'best practice recommendations' - notably the Article 29 Data Protection Working Party (WP29) in Europe and the Federal Trade Commission (FTC) in the US.

Both of these agencies have recently issued data protection guidance for IoT development. Of particular note here are those recommendations that speak to the general principle of individual control. The FTC proposes a number of practical measures to put the individual in control of personal data generated by IoT devices [12]. These include 'general privacy menus' enabling the application of user-defined privacy levels (e.g., low, medium, high) across all their IoT devices by default. The use of icons on IoT devices to 'quickly convey' important settings and attributes, such as when a device is connected to the Internet and to enable users to quickly ‘toggle the connection on or off’. The use of out of band communications to relay important privacy and security settings to the user via other channels, e.g., via email or SMS, and the use of management portals or ‘dashboards’ that enable users to configure IoT devices and accompanying privacy settings.

> “Properly implemented, such ‘dashboard’ approaches can allow consumers clear ways to determine what information they agree to share.” [12]

WP29 also proposes a number of practical measures to facilitate the application of EU legal requirements to the IoT. These include providing users with ‘granular choices’ over data collection, including the time and frequency at which data are captured, and scheduling options to ‘quickly disable’ data capture. Users should also be in a position to administrate IoT devices ‘irrespective of the existence of any contractual relationship’ and easily export their data from IoT devices ‘in a structured and commonly-used format’. Furthermore, settings should be provided that enable users to distinguish between different individuals using shared devices so that they cannot learn about each other’s activities. Data portability aside, these recommendations complement the dashboard approach towards putting the principle of individual control into practice, insofar as they are concerned to put computational mechanisms in place that allow end-users to *specify privacy settings*. Thus accountability essentially becomes a matter of making data collection *transparent* through design and enabling user *control* over data capture. However, the WP29 recommendations go a step further:

> “Device manufacturers should enable local controlling and processing entities allowing users to have a clear picture of data collected by their devices and facilitating local storage and processing without having to transmit the data to the device manufacturer.” [5]

The *local control recommendation* is radical, undermining the current approach to privacy that is being widely adopted by industry – i.e., encryption – which puts personal data online for cloud processing *before* making it available to the user. This is not to dismiss the need for security only to recognise, as Winstein [28] puts it, that it does not prevent abuses of personal data by industry, it just prevents ‘eavesdropping’ by others.

The security-based model is necessary to put the principle of individual control into practice but *not sufficient* to enable privacy protection as data is not simply a technical object consisting of bits and bytes, but a ‘relational object’ embedded in cooperative arrangements of data sharing, processing and use [8]. Thus, the *flow* of personal data in the IoT is not simply a matter of bits and bytes being transferred between machines. Rather these machines, and the data they generate and compute, are embedded in a *division of labour* that minimally implicates data subjects (who the data is about in some way), data controllers (who consume their data), data processors (who process it, and may themselves be machines); and, while not directly implicated in the flow of personal data, data officers (who monitor data processing) and enforcement agencies (who regulate it). Thus, building external accountabilities into the IoT to enable individual control requires that we make the arrangements of cooperative work implicating both the actors and the machines involved in the processing of personal data observable and reportable, i.e., accountable, so that the individual can reason about the demands being placed on their data and ensure its protection as needs be. We need, in other words, to make the IoT computationally accountable to make it legally accountable and thus enable it to be demonstrated that adequate data protection procedures have been implemented in compliance with data protection regulation with respect to the external data subject accountabilities required by proposed legislation. The question is, how might we make the IoT accountable to the data subject and meet the overarching purpose of legislation, thus enabling the individual to control the flow of personal data in the digital economy?

5. Building Accountability into the Internet of Things

We are of course not the first researchers in HCI to consider the challenges posed by the IoT and how the privacy threat might be handled [21, 23, 26, 29]. Prior research has explored the development of privacy aware systems [e.g., 13, 14] and personal data management services [e.g., 1, 9, 17], with the latter being taken up on a commercial scale. The World Economic Forum (WEF) reports that more than one new personal data management service was launched per week between January 2013 and January 2014 [27]. The WEF sees these kinds of services as key to 'empowering' the individual and rebalancing the current 'asymmetry in power' in the digital economy. However, despite growing commercial interest, public uptake of personal data management services has been problematic. A recent report suggests that this is due, somewhat ironically, to public perceptions of 'privacy and security risks' [15]. The situation is compounded by the fact that personal data are distributed across a great many cloud-based silos having no standard data formats, no standard APIs for access, and no easy way of obtaining a *holistic* overview. Needless to say, this makes it very difficult if not impossible for individuals to control the flow of personal data in the digital economy.

An alternative approach advocated by Chaudry et al. [7] might be more apposite for addressing the privacy threat occasioned by the IoT and posits a physical device as a gateway to a distributed platform. This 'Databox' approach is based on prior research that sought to develop a new wave of personal digital services and applications for individuals and resulted in the 'Dataware model' [16]. The Dataware model posits a socio-technical division of labour which implicates *the user* (by or about whom data is created), *data sources* (e.g., IoT devices, which generate data about the user), a *personal container* (which collates the data produced by data sources and can be accessed via APIs), a *catalogue* (which allows the user to manage access to the personal container), and *data processors* (external machines exploited by data controllers who which wish to make use of the user's data in some way). The Dataware model is a logical entity formed as a distributed computing system. The arc of work implicated in its use involves human readable data processing requests being sent to the catalogue, which are approved or rejected by the user. If approved, the catalogue issues a processing token to the data processor for permitted requests. The processor presents the token to the personal container, which accepts the token, runs the processing request on the relevant data sources, and then returns the results to the data processor.

The Dataware model represents a distinctive approach to personal data processing, that not only seeks to enable user control but also data minimisation. Thus, the Dataware model not only disrupts direct machine-to-machine processing, it also limits processing to specific requests and limits data sharing to the *results of processing*; the raw data remains 'in the box' under the users control. The Dataware model is currently being reconfigured around the Databox concept [7]. The Databox concept embeds the Dataware model in a physical object situated in the physical environment (e.g., a networked mini-computer in the home) under the direct control of the individual. It allows the individual to collate a diverse array of personal data sources – digital (i.e., online) and physical (i.e., IoT devices) – in a single place. The data sources remain distributed across digital and physical locations but Databox indexes them and allows the individual to control access to them. From an IoT perspective Databox holds data from local IoT devices in stores that can be written to by data sources but are isolated from reading by processors until appropriate permissions are presented. Data from IoT devices that process data online can also be collected by the Databox insofar as APIs are available, though this raises the possibility of privacy breaches insofar as the service provider has access to the data. Fortunately, not all IoT devices are locked into the cloud by default.

Databox proposes that data stores are implemented as 'unikernels', i.e., application-specific virtual machines that eschew use of a general purpose operating system with the attack surface and management problems it entails, for a library operating system approach where only the specific system-level code required by the data store is linked into the resulting unikernel. The Databox approach enables raw data to be retained by the individual and supports both local processing of requests and local hosting of computation, including running algorithms on the box to deliver local services and/or to drive actuation. Raw data can be released to specific data processors should the individual wish to do so, though processing can be restricted to only those operations that have been explicitly permitted by the individual. In each case data is encrypted and tagged with sticky policies [18] in a bid to prevent data processors using the data for any but the specified purposes. Below we consider extensions to the Databox concept that are needed to make the IoT accountable to users and meet the external data subject requirements of proposed legislation.

5.1. Extending the Databox

Computational accountability requires that the Databox make IoT devices and the actors implicated in cooperative arrangements of data processing 'observable and reportable'. The extended Databox model – which we call the IoT-Databox model - assumes that a number of distinct actors, of which there may be many of each, are directly implicated in data processing:

- The individual or 'data subject'.
- The 'data controller' or party who wants to consume an individual's data for some (lawful) purpose.
- A 'data processor' or party who carries out data processing on the controller's behalf, which we assume will be a machine.
- An intermediary, which enables data controllers to discover data subjects and vice versa.

This model of data sharing goes beyond 'walled in' data transfers between, for example, IoT device manufacturers and individuals to enable the broader use and reuse of personal data in the digital economy.

The IoT-Databox model puts in place a set of interactional arrangements and supporting system architecture that enables data subjects and data controllers to exploit an individual's data for mutual benefit and, at the same time, enables *demonstrable* compliance with proposed legislation, particularly external data subject accountability requirements it requires including transparency and consent, granular choice, data portability, and access. Interaction between the parties to personal data processing (i.e., the actors) is provided for in the following ways.

The data subject first *configures* IoT data sources. This entails associating physical (local) and online (remote) data sources with the Databox. Data sources may then be assigned ownership (e.g., collective, shared by specific individuals, or a single individual) and be annotated (e.g., fridge smart plug, kitchen humidity sensor, etc.). Individual accounts may also be created to enable individuals to manage the data sources they own (including shared data sources). The data subject may then register with an intermediary *discovery service*. This entails establishing a secure association with the service, e.g., setting up an account and lodging an authenticated public encryption key. The data subject may then post metadata about the data sources they own and wish to make available to data controllers and their processors; only meta-data about the data sources is published and gross location (e.g., the first part of a zip or post code), thus preserving the user's privacy.

Before a data controller can access a data subject's data sources they must also register with the discovery service and create an account. This also entails establishing a secure association with the service, as well as declaring the legitimate purposes for which the controller seeks to process personal data, the kinds of data sources it wishes to exploit, and registering any data processor APIs. The latter enable individual *access* to processed data and allow data subjects to inspect data uses, retention, third-party sharing, etc. The data controller can also post containerised [10] *apps* on the discovery service, which can be downloaded by data subjects and enable data processing or the local hosting of computation on the Databox. Apps may also be provided by other parties (e.g., open source developers) who wish to support data processing (e.g., providing data visualization, analytic, risk evaluation apps). Apps are made automatically available to data subjects based on the data sources they have registered with the discovery service.

The discovery service reviews a data controller's application, *rates* the controller based on the information provided (e.g., no processor API might result in a poor rating) and issues a revocable machine-readable token that will allow the data processors acting on the controller's behalf to search the data source registry for the required data sources. The discovery service also enables individuals to post *reviews* about data consumers and rate them. Reviews and ratings are lodged with the controller's account. Ratings are displayed alongside apps on the Databox, from where reviews can also be accessed, and the data subject can actively search the service via the IoT-Databox to find reviews and ratings for other data controllers should they wish.

Interaction between data subjects and data controllers is mediated by a 'multi-layered notice' [3] providing an easy to read description that identifies the controller, the purposes of processing, and the other mandatory information that is required to be provided to the data subject prior to data processing by proposed legislation. The multi-layered notice also defines the *benefits* of data processing, and the *risks* that attach to particular categories of data (e.g., that occupancy can be inferred from $C0_2$ data). The data subject may use data visualisation apps to preview the data that is requested by the controller, and also exercise *granular choice* over data collection, configuring which data sources may be used and setting data sampling frequencies. This may reduce the service options that are available to the individual, which is dynamically reflected in the notice itself.

Multi-layered notices – or 'manifests' – are attached (like terms and conditions) to apps. An app cannot be used without a manifest being in place and data cannot be transferred to a controller's processors without a manifest being completed by the data subject. Manifests are machine configurable, though it is assumed that they will initially be drafted by human actors (i.e., the controller's representatives). Once a manifest has been configured by the data subject and accepted it assumes the status of a service level agreement or *SLA* (Figure 1), which the IoT-Databox transforms into a set of machine readable policies that govern a data processor's access to the particular data sources agreed upon by the data subject and regulate subsequent data processing operations.

Data subject interactions are provided for through the IoT-Databox Catalogue. In addition to the interactions outlined above (data source configuration, discovery service registration, metadata publication, app discovery, ratings and reviews, and SLA configuration), the Catalogue enables data processing *auditing*. Auditing enables the data subject to inspect all data processing operations, historical and live, that have been permitted on the IoT-Databox and the SLAs that attach to them. The audit log may be lodged and updated dynamically with a trusted third party or distributed ledger (as may the processor's). The Catalogue also provides a messaging service that enables data subjects and data consumers to communicate with one another (e.g., to identify faulty data sources, such as a sensor, or faulty appliances that may, for example, need replacing).

Architecturally the IoT-Databox model consists of three key components: the Databox, a controller's processors, and the discovery service (Figure 2). The Databox is a small form factor computer consisting of a `web server` and webapp containing the `catalogue UI`, which supports user interaction with apps and data. `Apps`, like data stores, run within isolated containers (e.g., using Docker) and interact with APIs provided by the Databox to perform a task. For example, apps may use the Databox's `datastore API` to query data sources for processing, or the `comms API` to send data to external machines. The `comms API` is responsible for recording transactions which are encrypted and signed/countersigned by the Databox and recipient of data and stored in the `transaction log`. `Accounts`, `raw data and indexes`, and `metadata` are also stored on the box. Restrictions on the use of the APIs is determined by an app's SLA. Apps are installed/removed/updated using the Databox's `app manager API`.

The discovery service is a cloud-based service, which is interacted with using standard internet protocols (principally HTTPS). It consists of a `web server` and webapp containing the `discovery UI` providing for human interaction, and a `query API` providing for programmatic (machine-based) interaction. It contains a key and security association manager or `key server`, which is utilised by Databoxes and a data controller's processors for signing data transactions. The `discovery API` allows data subjects to upload data source metadata

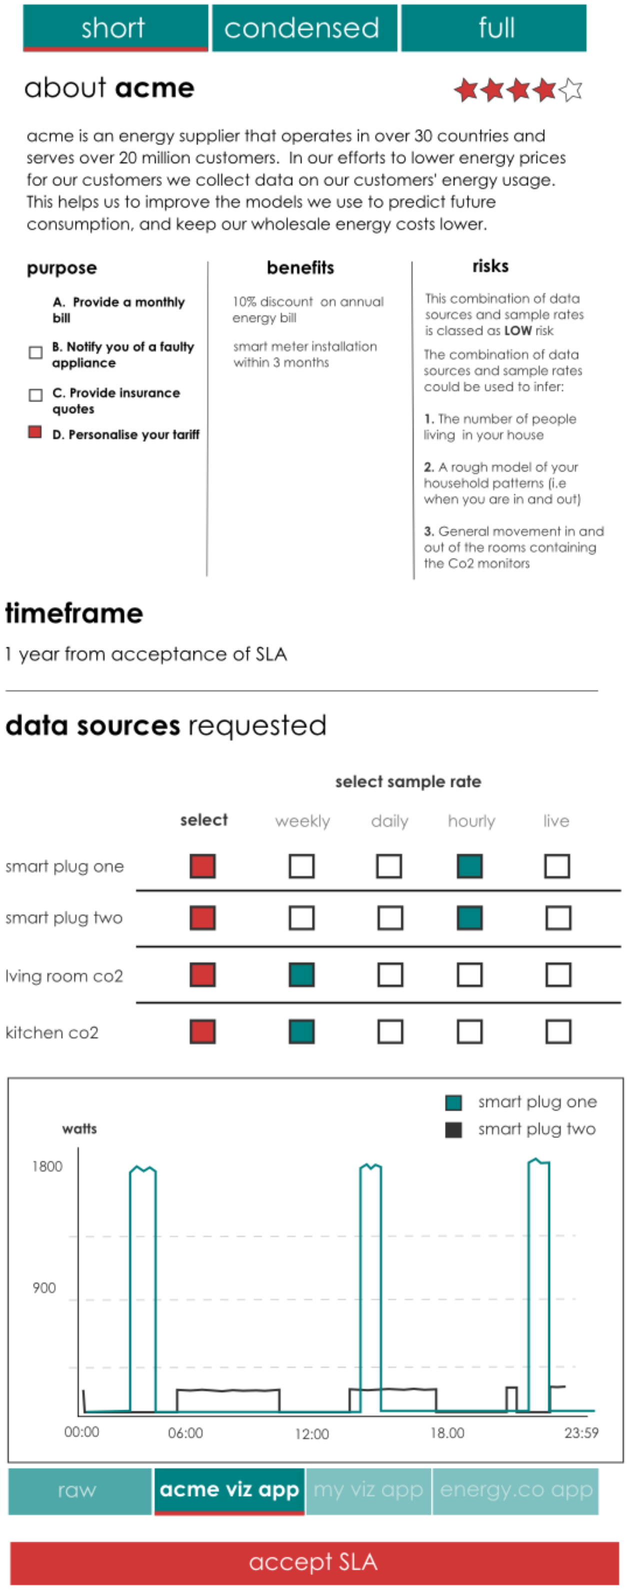

about **acme** ★★★★☆

acme is an energy supplier that operates in over 30 countries and serves over 20 million customers. In our efforts to lower energy prices for our customers we collect data on our customers' energy usage. This helps us to improve the models we use to predict future consumption, and keep our wholesale energy costs lower.

purpose

A. Provide a monthly bill

□ **B. Notify you of a faulty appliance**

□ **C. Provide insurance quotes**

☑ **D. Personalise your tariff**

benefits

10% discount on annual energy bill

smart meter installation within 3 months

risks

This combination of data sources and sample rates is classed as **LOW** risk

The combination of data sources and sample rates could be used to infer:

1. The number of people living in your house

2. A rough model of your household patterns (i.e when you are in and out)

3. General movement in and out of the rooms containing the Co2 monitors

timeframe

1 year from acceptance of SLA

data sources requested

		select sample rate			
	select	weekly	daily	hourly	live
smart plug one	☑	□	□	☑	□
smart plug two	☑	□	□	☑	□
lving room co2	☑	☑	□	□	□
kitchen co2	☑	☑	□	□	□

Figure 1. The IoT-Databox manifest/SLA.

(via the catalogue UI) and is stored as `Databox metadata`, which is made available to humans via the `discovery UI` or machines via the `discovery API`. The discovery service manages an `app repository` of all apps, which are uploaded via the `app API` and indexed by the metadata they operate on. It provides tools to help data controllers publish apps, one of which will be a set of skeleton `SLA templates` that can be specialised according to the particular aims of an app. And it manages `accounts` for data subjects and consumers, along with `rating/reviews`.

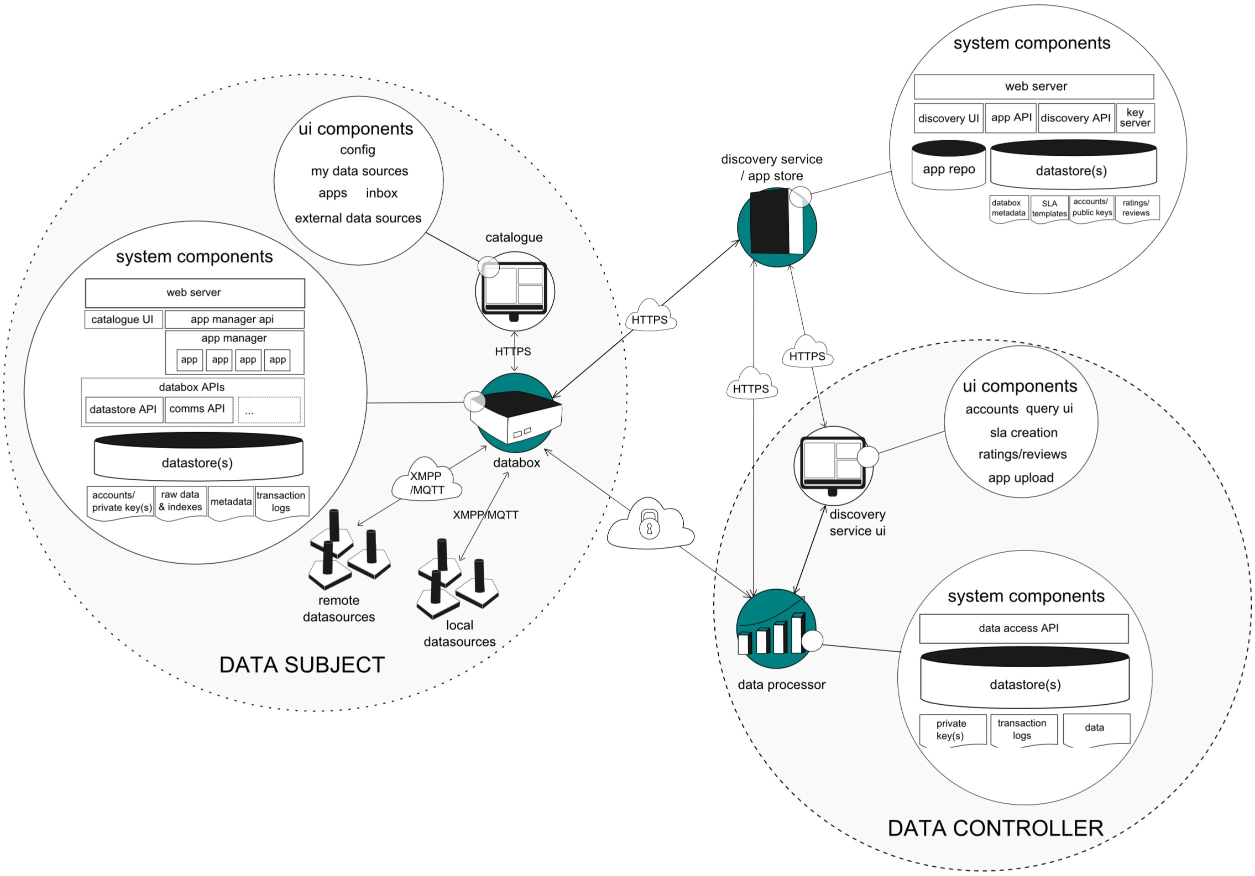

Figure 2. The IoT-Databox Model.

The discovery service provides most of the resources a data controller requires to exploit the IoT-Databox model. However, the controller will need to put in place sufficient resources to support their own operation. While the specific components needed to process data will vary from case to case, all controllers will need to deposit an app in the `app repository` to support their operations and we anticipate that they will want to keep a record of data transactions and thus require a `transaction log` to meet their accountability requirements to supervisory authorities. They will also need to store their `private keys`. Minimally, perhaps optimally, a controller's app might perform local processing and/or data visualisation but entail no data export. A manifest/SLA will still be required but no further components are needed here. Where an app exports data however, the data controller is responsible for providing a `secure data endpoint` and an encrypted connection for data transfer. Controllers are also encouraged to provide a `data access API`. This is not mandatory but it is desirable as it enables a controller to meet the data subject accountability requirements of proposed legislation and thus allows individuals to gain further assurances on how their data is used.

6. Conclusion

The IoT-Databox model is currently under construction. Nevertheless, the interaction mechanisms provided by this computational accountability system make data processing 'observable and reportable'. This enables the actors involved in the arc of data processing work to articulate not only their own actions but also, in surfacing relevant behaviours of computational systems (IoT devices, data processors, data processing requests, etc.), the actions of underlying machines. The IoT-Databox model thus allows the actors involved the complex socio-technical arrangements of cooperation involved in the processing of personal data to mesh their actions together. In articulating the actions of actors and machines the IoT-Databox model reflexively provides the transparency and consent, granular choice, access and data portability required to make data processing accountable to the data subject and enable individuals to control the flow of data from IoT devices into the digital economy. In doing so it also enables data controllers to demonstrate compliance with the external data subject accountabilities required by proposed legislation. The IoT-Databox is, then, of potential benefit to both data subjects *and* data controllers who would make use of the IoT and the personal data it generates, enabling them to benefit from the use of personal data while at the same time managing the widespread threat to privacy occasioned by the IoT.

> "Data protection must move from 'theory to practice' … In the discussions on the future of the European and global data protection framework, accountability based mechanisms have been suggested as a way of encouraging data controllers to implement practical tools for effective data protection." [2]

Acknowledgements

The authors acknowledge the support of the EPSRC, grants EP/M001636/1, EP/M02315X/1, EP/N028260/1, and EU FP7 grant 611001.

References

[1] Allard, T., Anciaux, N., Bouganim, L., Guo, Y., Le Folgoc, L., Nguyen, B., Pucheral, P., Ray, I., and Yin, S. (2010) "Secure personal data vaults: a vision paper", *Proceedings of 6th International Conference on Very Large Data Bases*, pp. 25-35, Singapore, VLDB Endowment. http://www.vldb.org/pvldb/vldb2010/pvldb_vol3/R02.pdf

[2] Article 29 Data Protection Working Party (2010) *Opinion 3/2010 on The Principle of Accountability*, WP173. http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2010/wp173_en.pdf

[3] Article 29 Data Protection Working Party (2013) *Opinion 02/2013 on Apps and Smart Devices*, WP202. http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2013/wp202_en.pdf

[4] Article 29 Data Protection Working Party (2013) *Opinion 03/2013 on Purpose Limitation*, WP203. http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf

[5] Article 29 Data Protection Working Party (2014) *Opinion 8/2014 on Recent Developments on the Internet of Things*, WP233. http://ec.europa.eu/justice/data-protection/article29/documentation/opinion-recommendation/files/2014/wp223_en.pdf

[6] COM (2012) 11 Final - Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2012:0011:FIN:en:PDF

[7] Chaudry A., Crowcroft J., Howard H., Madhavapeddy A., Mortier R., Haddadi H. and McAuley, D. (2015) "Personal data: thinking inside the box", *Proceedings of Critical Alternatives*, pp. 29-32, Aarhus, ACM. http://dx.doi.org/10.7146/aahcc.v1i1.21312

[8] Crabtree, A. and Mortier, R. (2015) "Human data interaction: historical lessons from social studies and CSCW", *Proceedings of ECSCW*, pp. 1-20, Oslo, Springer. http://dx.doi.org/10.1007/978-3-319-20499-4_1

[9] de Montjoye, Y., Wang, S. and Pentland, A. (2012) "On the trusted use of large-scale personal data", *Bulletin of the IEEE Technical Committee on Data Engineering*, vol. 35 (4), pp. 5-8. http://sites.computer.org/debull/A12dec/p5.pdf

[10] Docker, https://www.docker.com/what-docker

[11] Dourish, P. and Button, G. (1998) "On 'tehnomethodology' – foundational relationships between ethnomethodology and system design", *Human Computer Interaction*, vol. 13 (4), pp. 395-432. http://dx.doi.org/10.1207/s15327051hci1304_2

[12] FTC (2015) *Internet of Things: Privacy and Security in a Connected World*. www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-thingsprivacy/150127iotrpt.pdf

[13] Hong, J. and Landay, J. (2004) "An architecture for privacy-sensitive ubiquitous computing", *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services*, pp. 177-189, Boston [MA], ACM. http://dx.doi.org/10.1145/990064.990087

[14] Langheinrich, M. (2002) "A privacy awareness system for ubiquitous computing environments", *Proceedings of UbiComp*, pp.237-245, Gothenburg, Springer. http://dx.doi.org/10.1007/3-540-45809-3_19

[15] Larsen, R., Brochot, G., Lewis, D., Eisma, F. and Brunini, J. (2015) *Personal Data Stores*. https://ec.europa.eu/digital-agenda/en/news/study-personal-data-stores-conducted-cambridge-university-jud

ge-business-school

[16] McAuley, D., Mortier, R. and Goulding, J. (2011) "The dataware manifesto", *Proceedings of the 3rd International Conference on Communication Systems and Networks*, pp. 1-6, Bangalore, IEEE. http://dx.doi.org/10.1109/COMSNETS.2011.5716491

[17] Mun, M., Hao, S., Mishra, N., Shilton, K., Burke, J., Estrin, D., Hansen, M. and Govindan, R. (2010) "Personal data vaults: a locus of control for personal data", *Proceedings of CoNEXT*, Article No. 17, Philadelphia, ACM. http://dx.doi.org/10.1145/1921168.1921191

[18] Pearson, S. and Casassa, M. (2011) "Sticky policies: an approach for managing privacy across multiple parties", *IEEE Computer*, vol. 44 (9), pp. 60-68. http://dx.doi.org/10.1109/MC.2011.225

[19] Ramirez E. (2015) "Privacy and the IoT: navigating policy issues", *International Consumer Electronics Show*, Las Vegas. www.ftc.gov/system/files/documents/public_ statements/617191/150106cesspeech.pdf

[20] Robertson, T. and Wagner, I. (2015) "CSCW and the Internet of Things", Proceedings of the 14th European Conference on Computer Supported Cooperative Work, pp. 285-294, Oslo, Springer. http://dx.doi.org/10.1007/978-3-319-20499-4_15

[21] Roman, R., Zhou, J. and Lopez, J. (2013) "On the features and challenges of security and privacy in distributed internet of things", *Computer Networks*, vol. 57, pp. 2266-2279. http://dx.doi.org/10.1016/j.comnet.2012.12.018

[22] Schmidt K. and Bannon L. (1992) "Taking CSCW seriously: supporting articulation work", *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, vol. 1 (1), pp. 7-40. http://dx.doi.org/10.1007/BF00752449

[23] Sicari, S., Rizzardi, A., Grieco, L. and Coen-Porisini, A. (2015) "Security, privacy and trust in Internet of Things: the road ahead", *Computer Networks*, vol. 76, pp. 146-164. http://dx.doi.org/10.1016/j.comnet.2014.11.008

[24] Strauss A. (1985) "Work and the division of labor", *The Sociological Quarterly*, vol. 26 (1), pp. 1-19. http://dx.doi.org/10.1111/j.1533-8525.1985.tb00212.x

[25] US Consumer Privacy Bill of Rights (2012) *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy*. https://www.whitehouse.gov/sites/default/files/privacy-final.pdf

[26] Weber, R. (2010) "Internet of Things - new security and privacy challenges", *Computer Law & Security Review*, vol. 26 (1), pp. 23-30. http://dx.doi.org/10.1016/j.clsr.2009.11.008

[27] World Economic Forum (2014) *Rethinking Personal Data: A New Lens for Strengthening Trust*. http://www3.weforum.org/docs/WEF_RethinkingPersonalData_ANewLens_Report_2014.pdf

[28] Winstein, K. (2015) "Introducing the right to eavesdrop on your things", *The Agenda Magazine*, July edition. http://www.politico.com/agenda/story/2015/06/internet-of-things-privacy-concerns-000107

[29] Ziegeldorf, J., Morchon, O. and Wehrle, K. (2014) "Privacy in the Internet of Things: threats and challenges", *Security and Communication Networks*, vol. 7 (12), pp. 2728-2742. http://dx.doi.org/10.1002/sec.795