A roadmap for interdisciplinary research on the Internet of Things: Technology

24 September 2012

Lead author: Professor Hamid Aghvami
King's College London

Contents

1 Rationale

The Internet of Things (IoT) is regarded as the evolution of today's internet, attracting huge interest from governments, businesses, industry and academia in the USA, Europe, China and Japan. There is no doubt it will have a huge impact on every aspect of the personal, professional and social lives of citizens. The IoT is changing everything about the way we live, work and entertain ourselves, and about our environment. The IoT aims to support a diverse set of applications including smart grid, e-health, intelligent transportation and logistics, to name just a few. In this regard, it could have huge benefits for UK businesses and citizens in health care, energy saving, transport, etc.

The IoT is envisioned as global interconnections of a huge number of devices and objects (RFID and NFC tags, sensors, actuators, nano-devices, embedded devices, etc.) with different capabilities (in terms of energy, computing and memory), interacting with each other and connecting to large databases and networks, with unique addressing to support a diverse set of applications. IoT applications include:

- intelligent transportation and logistics
- healthcare
- smart grid
- personal and social interactions
- smart environment (smart city, smart home, smart office and smart manufacturing)
- future applications.

The devices and objects are being equipped with intelligence and connectivity, transforming today's internet into a global Internet of Things. This enables the objects to sense, collect and transmit a vast amount of data from their proximate surroundings. By processing, analysing and distributing the data to extract useful information, and transform the information into knowledge and economic value, it gives businesses and citizens a competitive edge (see Figure 1).

Figure 1 Humans turn data into wisdom (Cisco IBSG, April 2011)

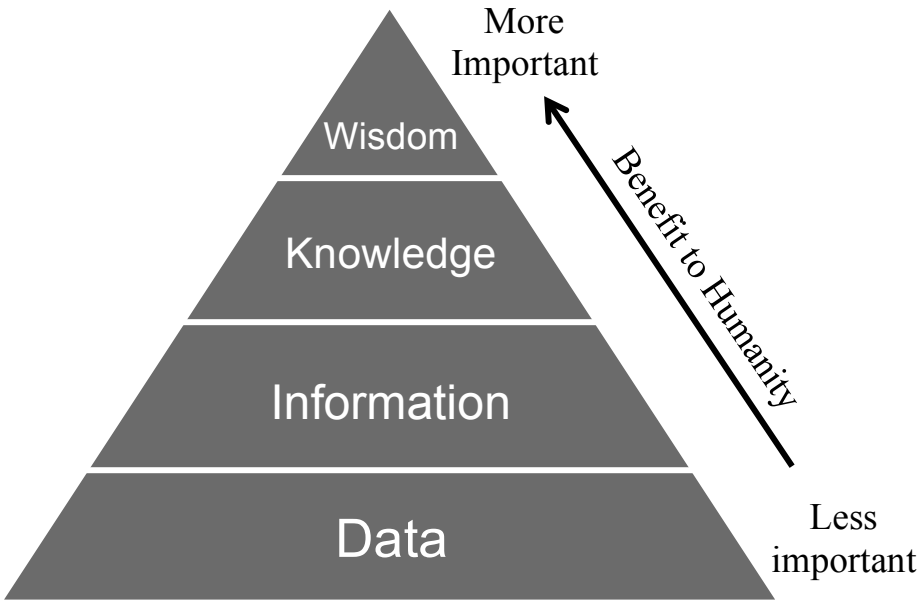

The IoT is regarded as a network of networks, as shown in Figure 2. Important enablers of the IoT are radio frequency identification (RFID) for identification of things, sensors for sensing physical changes around things and collecting data, and wireless short links and communication networks for connecting things. The infrastructure of the IoT is an integration of several networks, as shown in Figure 2.

Figure 2 Network of networks

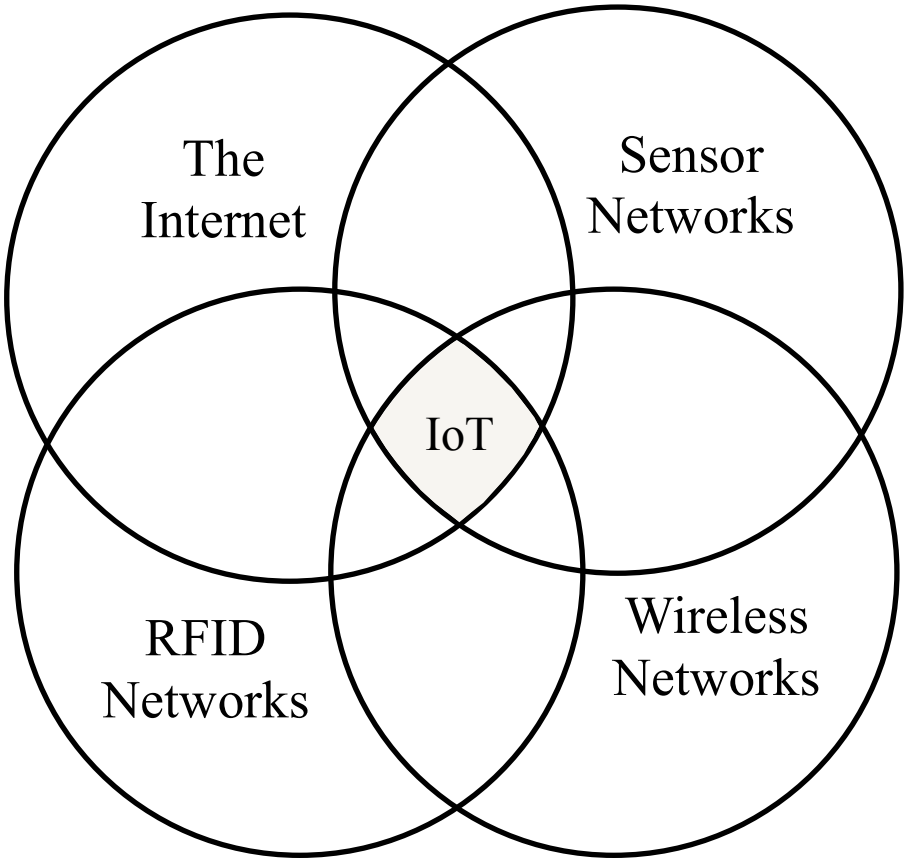

The IoT in its simple form already exists. It is a collection of fragmented and disparate networks connecting several billions of devices to the internet. The integration of these networks is in its infancy but is increasing gradually towards full integration. According to a Cisco IBSG prediction, there will be 25 billion devices connected to the internet by 2015 and 50 billion by 2020. In such a scenario, connecting and identifying billions of objects and devices in an integrated manner will present great social and technological challenges.

The technological challenges are mainly related to addressing and networking aspects of the integration of several technologies and communication networks, as stated above. The integrated network of networks should be scalable, high capacity, with low energy consumption levels, and low cost. To realise such an infrastructure, that enables billions of devices and objects to access it with a common identification scheme, and supporting a huge amount of traffic (several zettabytes in 2020) in a global scale, requires a huge coordinated research effort by British businesses, industry, academia and standards bodies.

Our task was to consider technical challenges worthy of further investigation, relating specifically to the IoT technology. We considered two timescales: less than three years for Technology Strategy Board (TSB) projects, and more than five years for the Engineering and Physical Sciences Research Council (EPSRC), and identified a number of major challenges in this domain.

Although our focus is entirely on challenges faced by the IoT domain, we hope that many of these challenges will be addressable using approaches developed for other domains, in which case the challenge becomes to identify, adapt and adopt those approaches for the

IoT. However, we believe that there are significant nuggets within these challenges that probably do require novel thinking and development. Sometimes the sheer difference of scale at which the IoT operates requires qualitatively different solutions from the way things are done today. Underlying all our thinking is the view that today's IoT applications are highly vertically integrated, with often just one provider engineering the entire stack from top to bottom, but the IoT applications of tomorrow will be heterogeneous, multi-vendor and emergent.

2 Challenges

Technological challenges can be classified in four categories: data processing and applications; architecture (networking and addressing); things (objects, devices, etc.) and radio links; and security and privacy.

2.1 Data processing and applications

2.1.1 Short-term challenges

Semantic connectivity in addition to data connectivity

Currently 1:1 closed IoT implementations are concerned mainly with data connectivity, and the actual meaning of the data is often implicit (held only in the minds of the engineers who designed the complete system). For example, the same vendor that designed a river sensor also designed the cloud service to analyse the data from it. As such, there is implicitly a common understanding of what each of the data streams that come from the sensor and get stored and manipulated in the cloud service actually 'mean' and therefore how they should be treated. In many cases, processing the data requires this background knowledge and descriptive information. For example, an application might be required to know that a 'temperature' data stream is in degrees centigrade, accurate to within a tenth of a degree, updated every minute, measured the centre of a river, at a particular location, etc.

So in the emergent, open IoT of the future, where different parties design the different parts, a significant challenge is to enable these parts to establish a common interpretation of the data being exchanged. To address this issue and provide interoperability for information exchange, common semantic description frameworks are needed. This will require the design of light-weight ontology models and the creation of efficient semantic description frameworks – query languages that are suitable for volatile, distributed and dynamic IoT environments and can be adapted in the resources that have limited memory, power and computing and communication resources (i.e. directly or via gateway nodes).

Directories, discovery, data federation

Within the machine–human or human–human world of the Web, services such as Google and Twitter enable discovery – a search phrase submitted and a list of matching links containing that phrase received. Today, most IoT applications are hard-wired to a particular publishing mechanism (for example, Cosm) but in the future it will be necessary to be able to use generic publishing mechanisms, and therefore it will not be possible to know a priori where to 'go' to find the data. Building tomorrow's machine-to-machine (M2M)/IoT applications out of diverse parts requires the ability to publish datasets, and then discover

these data sets using generic search-and-discovery mechanisms. The volume and diversity of IoT directories and databases will mean that they have to be machine-maintained. For example, an application might search for 'live temperature data for inner London updated at one minute intervals'. Other related issues are access rights, provenance, revocation, defining commercial models, and dealing with privacy issues (as cross-correlated datasets can often reverse-engineer anonymisation process).

2.1.2 Long-term challenges

Management of device/resource/service lifecycle

People already struggle to manage their networks in homes, offices, etc. and the IoT will hugely increase the number and variety of devices on these networks, potentially making them unmanageable. It is not just a problem of scale – today's devices are typically human-facing, enabling people to interact with them during installation, to check that they are working, and to diagnose problems. In contrast, most IoT devices are not human-facing (indeed have no user interface), and may have such obscure functionality that typical end users, as distinct from the engineers who installed the device, will not understand their function or even know of their existence – and nor should they need to. Existing IoT service providers have implicit knowledge about how to install and maintain their devices, but in the multi-service heterogeneous future, a service provider must be able to manage a device from various vendors and providers.

A key challenge arises from the need to define standards for the 'care and feeding' of IoT devices and gateways, making it possible for third-party IoT service providers to install, upgrade and configure devices as well as maintain devices and diagnose any malfunctions of devices from other vendors and providers (which its developers have not been able to anticipate). This challenge probably also includes the need to define a way to hold repositories of device-management information (not just their data) which can be accessed by multiple heterogeneous services for discovery and lifecycle management and maintenance purposes. (Note: this challenge does not include data management per se.) Ideally, this would result in a plug-and-play experience for the end devices, and service integrators, application developers and end users. Devices can automatically find gateways that can manage them, or connect through gateways to services that can manage them.

Deriving quality knowledge from incomplete and untrustworthy data

Conventional 'enterprise' IT systems such as databases and internet services often run centrally, or if distributed are at least connected by fast, reliable networks. In contrast, IoT applications are generally much more spread out physically, running to a variety of locations, through different gateways, connected over slow and unreliable links, with many more parties involved. It is therefore much more challenging to know and ensure the quality of the data flowing through the system. There will be inaccuracy, incompleteness and flaws in data where, for example, a device setting is not correct, a communication link has temporarily gone down, or some parts of the data are lost over an unreliable means of communication. There will be data-collection quality problems caused by damaged or poorly designed edge-devices (for example, the uncalibrated Pachube/Cosm radiation

monitors used around Fukushima produced differing measurements for similar locations[1]). There can also be malicious actors trying to disrupt the system by purposely injecting false data. Therefore datasets/streams may be mutually inconsistent.

Thus a notable challenge is to define standards for declaring data quality and develop methods and solutions that can process the data and semantic information, and extract knowledge from the underlying data is another challenge. This will also require methods for testing and proving data quality, including stress-testing. Another important issue is to define methods for deriving high-quality, high-confidence knowledge from multiple low-quality inputs. Establishing an IoT framework for trust is not about just encryption; it also considers issues such as the trust and reliability definition for the sources of the data.

Deploying and running the IoT application in distributed environments

If we define an IoT application as 'a piece of value-added functionality running within an IoT service' then frequently the application is actually distributed as computer programs running on multiple IoT entities. For example, some parts of the processing may happen within an end device, some within a gateway, and some within the cloud. Reasons for partitioning the functionality include scalability, data aggregation, reliability and latency. In some of today's IoT services, developers make fixed decisions about how to partition their applications – where to run different parts of the software – based on characteristics of end devices, gateways, etc. This means that they are coding for a special case, and probably even writing different parts of the applications in different languages, which makes development slow and complex. Such an approach makes IoT implementation dependent on the specifics of a particular IoT configuration, not portable between IoT frameworks, and brittle in the face of change.

Thus another challenge is to explore whether it is possible to design IoT applications in a 'write once, run anywhere' way. Is it possible for an IoT fabric to automatically partition code and flow parts of it to run in the appropriate parts of the infrastructure, possibly even dynamically as the situation changes? If so, then developers will no longer have to hard-wire the partitioning. Examining the feasibility of this may include exploring formal high-level languages for IoT applications development.

2.2 Architecture and networking

2.2.1 Short-term challenges

Decoupling the identification and addressing of things allowing mobility and tracking of things

In order to facilitate mobility and the generality of applications, the decoupling of identification and addressing is needed. There must be an agreement on the unique and global identifier to be used, and then on the design and implementation of efficient ways to make the mapping between the identity of a thing and its location. This can be supplied by a service that gets the required information of a thing based only on its supplied

[1] http://tinyurl.com/c5dlloy

identification; the same service can map the identification of a thing and its current location. This service is called the object naming system (ONS) and is a naming resolver. The ONS should be dynamic and lightweight when and where necessary, allowing good efficiency in fixed and mobile scenarios. For example, a gateway could probably provide an ONS for an RFID cloud.

The need for faster establishment of interoperable standards has been recognised as an important element for IoT applications deployment. Current IoT solutions are fragmented; as such, there is a need to unify solutions through standards. However, standards are normally too complex; they need to be simpler to make them more adaptable to evolution. The simplification of standards is another challenge. There is a need for more structure in the standards that allows building applications in the short term, as well as allowing IoT evolution in the long term.

Standard and solutions for IoT middleware, routing and transport (lightweight and lower power consumption)

There is a clear need for a lightweight and open middleware based on interacting components that abstract resources and network functions. So far, the middleware architectures proposed in many IoT projects have been based on the service oriented architecture (SOA). SOA allows the decomposition of a complex application into a simpler system composed of small, well-defined components that interact with each other. SOA requires standard interfaces and protocols. However, the current solutions are not lightweight and also fail to provide the loose coupling and proper separation between types and instances that are needed in domains that involve 'things'. This is a challenge to be addressed in the short term (preferably) in order to allow interoperable IoT applications with the transparency needed to users.

In the layers below, routing and transport are also a challenge for the characteristics of the IoT. There are current proposals for more hierarchical routing that should be looked at (an example of routing proposition is RPL from M2M standardisation). However, if the emphasis is on 'things', new ways to perform routing must be investigated (e.g. content- or context-based (information-based) instead of addressing-based).

The current transport protocols require excessive buffering to be implemented in objects; so they may be too heavy for resource-constrained devices. Also, the connection setup and congestion control mechanisms may be useless for certain IoT applications. Protocols that address these characteristics must be developed. Another challenge is that IoT applications may generate data traffic with patterns significantly different from the ones observed in the internet today, therefore it will be necessary to define new quality of service (QoS) requirements specific for IoT applications.

2.2.2 Long-term challenges

Things with automatic access and selection of networks

Mobility is still a huge challenge for resource-constrained devices if it is to be done autonomously. A major problem is the completely automated access across networks. The state of the art currently falls short of support for autonomous discovery of, selection of and access to different networks. This is a longer-term innovation challenge for the IoT (i.e. self-

organisation, self-selection and self-configuration). This challenge is not only technical, but also demands improvements in the business models currently used by operators.

Self-managed, self-configurable, self-healing, etc.

In the definitions of the IoT it is expected that 'things' are networked in systems that are reliable, intelligent, self-managed, context aware and adaptable to network technology, network discovery and network management. This means that IoT applications must be self-organising, easy to use, energy-efficient (able to harvest energy if necessary), able to operate in any circumstances, even in disasters or emergencies, by being adaptable and responsive to different modes of communication. The research community is therefore expected to provide novel solutions to guarantee the needs of the IoT to be self-capable of 'everything' (i.e. have self* properties).

2.3 'Things' and radio links

2.3.1 Short-term challenges

Low cost, low-power repeater nodes powered by harvesting

The range of feasible IoT applications becomes much wider as the devices that provide them become cheaper and use less energy. Hence, the challenge is to minimise both cost and energy consumption, to the point where nodes can be powered by energy harvesting without requiring an external power supply or a battery. It is proposed to develop a node that can function as a repeater, partly because networks for many IoT applications seem likely to require multi-hop networks, and also as a demonstrator of the main functions required by any IoT node. Note that it is assumed that IoT networks will include a higher functionality 'gateway' node, which interfaces with the rest of the internet and acts as a concentrator for the very low-cost, low-energy sensor/repeater nodes. Components of the challenge include:

- low-power radio frequency (RF) design, including implementation of functions in the RF domain to reduce the energy used in signal processing, and including low-power reception (low-power low noise amplifiers (LNAs)), as well as power-efficient high power amplifiers (HPAs)
- minimising the energy usage of signal processing, including low-power digital signal processing (DSP), and system design to move energy-consuming functions to gateway nodes
- selection and development of energy-harvesting techniques for specific applications (e.g. exploiting thermoelectric effects or body motion for body-mounted devices, etc.)
- system and network design for minimum energy, taking into account radio propagation
- intellectual property licensing costs for standardised physical layers that may be a significant component, unless either appropriate licensing deals are struck with IP owners or generic technology can be developed.

Antenna design for small devices

Antennas are difficult to design for small devices: there is an inherent relationship between size and both gain/directionality and bandwidth. IoT networks are likely to require self-configuration, which will require antennas that can be reconfigured under system control to optimise their performance. This may also be most efficiently carried out in the RF domain.

2.3.2 Long-term challenges

Efficient use of the radio frequency spectrum

Probably the greatest challenge in the medium to long term for the IoT is the availability of radio spectrum to provide connectivity for millions of devices. Hence, it is crucial to use the available spectrum most efficiently and avoid congestion. Moreover, IoT applications have a wide range of requirements in respect of data throughput, latency and reliability: it is essential to consider how to meet these without causing congestion by over-provisioning for applications that do not require it.

Some options that could constitute components of the challenge include:

- cognitive radio for low-power devices to allow the secondary usage of those parts of the spectrum that are under-used by the primary user. The current 'TV white space' model of cognitive radio (using a spectrum availability database) is certainly not the complete solution for large numbers of low cost devices, and hence power-efficient spectrum sensing technology is required, with frequency-agile RF technology
- spectrum-/power-efficient multiple access protocols/etiquettes
- ultra-wideband signalling, which allows multiple users to share the spectrum without any coordination
- non-radio media, e.g. infra-red/visible light, acoustic, power-line communication.

Control of signalling overhead

IoT networks are likely to exchange very large numbers of small data packets: existing standards and the whole protocol stack on which they are based are likely to be highly inefficient for this purpose, since they require large overheads per packet, and generate large numbers of additional acknowledgement packets. Moreover, common multiple access techniques operating at the media access control (MAC) layer respond very inefficiently to even moderate congestion, since colliding packets are lost. A radical redesign of the protocol stack will therefore be required for the IoT, going beyond cross-layer design to full integration of layers, allowing more advanced signal processing to be brought to bear on functions traditionally performed at MAC, logic link control (LLC) or network layer.

In addition, for many IoT applications propagation conditions mean that multi-hop architectures are essential, and hence there are routing functions to be performed within the wireless network. Protocols that work on a per-link basis will result in a further multiplication of signalling overhead. It has already been realised that this effect severely limits capacity in multi-hop ad-hoc networks, but it will be more severe in multi-hop IoT networks.

The application data processing needs also to be carefully considered to avoid unnecessary alarms giving rise to congestion. Appropriate data processing at nodes should ensure that only meaningful information is transmitted.

End of life issues

Efficiency of resource use includes also reuse/recycling of devices, as well as avoiding polluting the environment – this is especially important for IoT applications that involve very large numbers of nodes. In some cases, these may be embedded in buildings or city infrastructure that has a much longer life than the likely useful life of the nodes: the nodes should ideally be designed in a sufficiently flexible way so that they can be reused for new applications, or at least recycled in order to reclaim materials. In other applications (e.g. environmental sensing), large numbers of cheap, disposable nodes may be used, and it is important that if these cannot be recovered and recycled that they do not cause pollution. One approach would be to develop biodegradable nodes, even nodes that can be remotely triggered to begin to degrade.

2.4 Security and privacy

The IoT implies multi-layered, heterogeneous, mass-scale interoperation of things (including of people and services). As such, the distributed network-centric architecture that is to support the IoT will have to satisfy many constraints of which the highest priority has to be usability, security efficiency and security cost-effectiveness, which implies matching the end-to-end security and privacy protection goals to an acceptable user experience and business model for viable and scalable IoT services.

2.4.1 Short-term challenges

Identity, AAA, security–privacy–trust (IASPT) and non-repudiation management matched to the deployment context of the respective IoT business model, users' needs and network and devices resource limitations

The IoT involves a deeper integration of human life-style and work-style support services; increasingly underpinned by a convergence of the Web-of-Things (future Internet of services), mobile, always-on wireless sensor networks (WSNs) and heterogeneous network topologies and communication protocols. It is thus important to note that in order to fulfil the expectations of transparency, security (including safety), privacy and accountability (including non-repudiation) in service delivery end-to-end, we require either:

- a unified security-and-privacy regime applied to all the layers of the IoT architecture

or

- semantic security interoperability between security-and-privacy regimes as selected and optimised to best suit the requirements and constraints that prevail within each IoT layer.

Here from an identity and trust management viewpoint, issues include IoT identity-key setup, group setup and membership management, mobility support, privacy-aware identification and interoperation, the management of unique identities for physical objects and devices, and multiple identifiers and cross-referencing for people, locations and

identifiers (e.g. encrypted privacy-preserving pseudonyms and respective authentication credentials), whilst also supporting non-repudiation. Clearly, IoT identity management has to provide for interoperability with legacy identifier schemes (e.g. URL) and faster and less-energy-consuming encryption algorithms, as well as efficient key distribution schemes.

Accordingly, the IoT needs to deliver authentication, trust-binding and privacy-preserving technologies that can operate with smaller more resource-constrained devices. For this, different ID schemes to suit the deployment context of various applications, and convergence of IDs and addressing schemes have to be supported as part of the IoT security-privacy architecture.

Multi-perspective, context-aware IoT security management mix (variable grade security)

Context-awareness can range from a single device-ID-linked reference at the WSN layer to more extensive descriptors at the middleware and application layer to represent the situation of an entity (e.g. person/place/object) and other information essential to service provisioning and interaction with the user at the application layer (e.g. who, where, what, when, how, …); plus occasionally including some other relevant entity information such as profile elements to express the involved Thing's role(s) and responsibilities and, possibly, also some expression of their goal-stack. Thus a security context descriptor set would be expected essentially to include information elements such as the Thing's identity (including data-type, event, location and device capabilities etc.). Device capability descriptors could be, for example, statelessness, state-fullness, memory and energy storage capability, safety criticality (e.g. the security context for a pacemaker will have several descriptor fields which will be different from those of a hallway thermometer). Similarly, trust management can simply depend on a lightweight pairwise key (and shared secret) at the edge layer whereas at the middleware, core network and application layers more extensive models of trust may be feasible (including trust transitivity resolution) if deemed essential for the particular business model being served by the IoT.

It is thus evident that we need to be cognisant of the distinct security contexts that exist in the IoT end-to-end as may be distinguished according to the requirements and constraints that prevail within each of the various IoT layers and domains, including the needs of specific business models (deployment contexts) to be sustainable into the future.

Security sensitivity context analysis addressing the prevailing constraints and stakeholders' requirements from the viewpoint of various security contexts

This includes scoping the essential context-sensitive interoperable security–privacy for each IoT domain/layer (e.g. edge layer, access link layer, core network layer, application layer) including to cover the WSNs' lifecycles (sensor introduction and security bootstrapping, operational life, obsolescence/death or dis/re-use); core network access control, and application layer secure services matchmaking, including static/dynamic service composition and orchestration (application security contracting and cloud services) for various business eco-systems, for example involving smart clothing, smart home/transport/city/factory etc. to support sectoral business models; supporting open scalable *prosumer* chains and value constellations for competitive co-creative innovation.

Thus identity, security and non-repudiation management has to be best matched to the needs and limitations of any given network device and layer, for example for WSNs, due to the lower capabilities of their devices, lightweight pairwise approaches can be deployed provided the storage/computational overhead can be sufficiently low for viable and scalable operation:

- Scalable and distributed pairwise security protocols are thus needed, including mechanisms for pre-distribution of keys, establishment of logical paths and probabilistic key and secret sharing, scoping of the number of shared secrets (e.g. for sensor peers and the local hubs in a WSN).
- Cross-layer unified public key authentication and mutual trust schemes need to be developed such that the memory requirement for operating them remains sufficiently low for the more resource-constrained devices in any layer across the communication chain (e.g. WSN hub-to-core network).

Furthermore, static rule-based and/or dynamic model-based personalised security–privacy management are needed to cover core network access and the application layer trust-binding [(device-Id + key material + security policy)] as well as application layer security, to allow only trusted instances of an application to communicate with the IoT.

Security–privacy impact assessment (static/dynamic)

For the IoT, dynamic forms of security–privacy policy evaluation and impact assessment would be necessary so as to be able to respond to changing environments and unexpected situations. At a policy execution level, this may involve identifying and selecting which set of access control rules to apply in situations where access control policies are in conflict or apply only partially and may have to be overridden due to unforeseen contextual changes in the environment.

- Security–privacy impact-level sensitivity analysis for prototypical IoT deployment scenarios

  There is a need for systematic deployment-context-specific modelling from the viewpoint of the stakeholders' response to particular security–privacy regimes of given IoT value propositions and the stakeholders' potential acceptance and uptake of IoT innovations based on their perception of their security–privacy needs and the level of fulfilment of such needs by the oT in particular contexts; for example, relating to data integrity, ownership and application integrity in particular sectoral IoT deployment scenarios.

  There is also the need to build on the so-called framework for privacy and data protection impact assessments (PIA) mechanism for data protection compliance, privacy sensitivity and risk assessment monitoring of a given IoT application, supported by a trusted certification and audit framework.

- Security–privacy by co-design

  For security to be assured at design and execution time, trust in the security–privacy–governance integrity of the IoT has to be built in at the design stage (including at the initial business model design stage). This has to include 'privacy by co-design', actively involving all stakeholders in co-design including for security–privacy impact in novel IoT-enabled e-government applications (e.g. border control, passports, medical cards,

hospital security, etc.), and in particular for IoT-enabled behavioural monitoring of mobile customers for marketing purposes.

The social impact of IoT requirements has to be considered, for example according to the 'fair information practice principles'. This includes informed consent, transparency and involving the users/citizens in the management of their data.

Privacy by co-design, privacy impact assessments, and privacy-enhancing technologies should all be considered as a means to promote societal trust and confidence in the IoT.

Additionally, the impact of the security provision on other operational aspects of the IoT needs to be taken into account, such as interoperability, scalability, discovery of things, accountability, manageability, cost effectiveness, user's adoption tipping points, business model viability, cross-elasticity of supply and demand, competitive strategy, etc.

Certification and audit management

To enable authentication and secure interoperability, nodes may have both static certificates through public key infrastructures (PKIs) and/or dynamic cryptographically generated identifiers and trust binding to a peer group. Scalable, expressive and efficient assertion mechanisms are required to allow authentication and semantic interoperability between the end points in the IoT chain of network layers.

2.4.2 Long-term challenges

- Zero knowledge leakage (ZKL) assurance mechanisms including semantic security resolution: this is to provide the maximum possible semantic control of security and privacy management. A user may require such a level of protection for certain highly security–privacy sensitive applications of the IoT in, for example, the medical domain. The objective of a desired ZKL regime can be realised through de-linking of the *authentication* context from the *transaction* context such that, even if the data from either of the above two contexts happens to be exposed to a security risk, the fact that the two types of user data (authentication and transaction) are de-linked will ensure that no knowledge of the activity of the user will be known to anyone who may gain unauthorised access to the network.
- By building on the semantic security resolution framework that has been developed through recent advances in IoT-middleware (e.g. SmartLink), it should be possible to develop middleware that can be applicable across all possible smart environments, such as smart homes, smart vehicles and smart city, including trust management exploiting semantic trust expression as part of the resource description. Trust management procedures can be realised as a set of specific business scenarios in the form of agent configuration plans.
- Dynamic negotiation and enforcement of quality of service of security in IoT environments should be provided, including in relation to resource constraints and/or intermittent connectivity.
- More efficient security solutions need to be implemented pairwise between 'last mile' connections of devices to more capable middleware endpoints so that the small devices can be securely maintained and suitably assessed for continued integrity by the middleware end point. This can exploit any self-state-reporting capabilities of a device

i.e. self-state-awareness (self*) properties as may be feasible and if required, as referred to earlier (e.g. simply, in relation to battery level).

- Dynamic allocation of security mechanisms at runtime in order to fulfil general security protection goals, including confidentiality, integrity and availability, application-specific requirements and the ability, for security reasons, to effectively refuse contact with communication counterparts in IoT environments, including for instance protection from communication flooding, will remain an important security issue. This will particularly be the case where the integrity of IoT-enabled devices is of critical importance to the safety of the individual and/or environment and a breach of device security may have a physical impact (consider for example an IoT-enabled pacemaker). Critically important devices must be able to protect themselves from all types of potential attacks and must be able to, for instance, disconnect from networks in order to retain operational integrity. Such security self-management behaviour will be essential for critically important devices that operate in constantly evolving and potentially hostile operational environments such as public spaces.
- Proactive identification and protection of the IoT from arbitrary security attacks (e.g. denial of service (DoS) and distributed denial of service (DDoS) attacks), abuse and malicious software. There are also security issues arising from the long-term mass-deployment of IoT devices, mostly un-updated and some even 'abandoned' by operators but remaining operational. The security status of such devices needs to be determined and potentially compromised devices need to be identified and isolated. Furthermore, intelligent means for identifying and mitigating the effects of potentially compromised devices will need to be put in place in order to ensure the continued operational integrity of an IoT environment.
- Protection is needed against possible breaches of privacy through the availability of unprecedented numbers of connected sensors in environments such as smart cities. To address this, potential trade-offs between the utility of an environment and the willingness of an individual to share information with that environment need to be made explicit, made controllable by the individual, and need to be enforced when an individual interacts with that environment. As far as possible, the identity of an individual should remain protected in such scenarios while maintaining the integrity of actions taken by the same individual, for instance by enforcing non-repudiation of transactions.
- Context-aware IoT access control policies need to operate, including an enforcement mechanism such that the successful fulfilment of specified obligations prior to, in parallel with or forensically after handling an access request can be enforceable.
- Building on recent advances in combining lightweight and regular cryptographic protocols and off-line key management schemes is recommended.
- Technologies have to be developed to support a unified global ID and object description as the ultimate target. This is to accommodate: identity management, identity encoding/encryption, pseudonimity (revocable), anonymity, authentication of parties, repository management using cost-effective identification, authentication and addressing schemes, and the creation of global directory lookup services and discovery services for IoT applications. It is important to ensure consistent integration of the existing IPv6-focused architecture, Web of Things and Restful with such a global

identity scheme, as well as the integration of targeted security-privacy solutions with the future internet (FI) architecture.

- Unique as well as multiple and group IDs need to be developed where needed to interconnect every object (or object group as may be required for multi-cast) for transparent and intelligent event (context)-triggered security control and invocation of services (including from cloud computing).
- Support is required for real-time controlled and/or stateless IoT devices (e.g. car brakes, temperature sensor).
- Support is required for security, privacy, and trust for M2M social networks.
- Power-efficient approaches to security–privacy authentication encryption and access control need to be investigated.
- Support is required for context detection and for semantic modelling and managing of data.
- Assessment of the impact of the realisable levels of trans-border, cross-domain privacy protection resolution with the IoT and cloud security, for example consistent security resolution when cloud services may also enforce their own security policies for load-balancing and providing shared caches.
- Impact of (linked) open and big data analytics on privacy need to be evaluated.
- A 'security testbed cloud' needs to be established for the IoT industry.
- IoT context-specific cybercrime execution stack modelling would be helpful in informing the managed mix of security approaches and mitigation solutions as may be deployed in each layer of the IoT to support the goal of rightsizing the cost-effective IoT security for each IoT deployment context, and at each respective IoT layer.
- The provision of an end-to-end fair information practice principles compliance management regime for the IoT would help to enhance the potential uptake.
- M2M privacy should also be protected within IoT environments, as IoT devices often can be directly associated with their owner, or even with other individuals nearby.

3 Recommendation

A considerable amount of research on IoT technology is being carried out worldwide. However, the research efforts are disparate and carried out in isolation from each other. Also, many network infrastructures such as RFID, sensor networks and future internet have been developed, but no attempt has been made to integrate them. Furthermore, existing solutions and standards are fragmented.

Different design approaches, algorithms, protocols and techniques have been proposed and investigated for each element of IoT in isolation without considering their inherent interactions. There is no single network infrastructure platform (NIP), physical or virtual, that has integrated all disparate platforms and elements of IoT. There is a need for a unified and coherent approach.

An NIP for the IoT is not only essential for the technology aspect of the IoT, it is also an important step towards the development and validation of emerging and new applications, and their commercial and social aspects. The development of such an NIP is a huge

challenge that requires a huge coordinated effort by Research Councils, British businesses, industry, academia and standards bodies.

Workshop participants

This is a list of all participants in the technology group at the roadmapping workshop held in July 2012 in Loughborough. This report summarises the output of the technology group at the workshop, and does not necessarily represent the views of those listed here.

Main contributors

Professor Hamid Aghvami	King's College London (lead author)
Professor Atta Badii	University of Reading
Dr Payam Barnagh	CCSR/University of Surrey
Pilgrim Beart	AlertMe.com Ltd
Dr Eliane Bodanese	Queen Mary, University of London
Professor Alister Burr	University of York

Other contributors

Dr Peter Bull	Loughborough University
Dr Ruzanna Chitchyan	University of Leicester
Joshua Ciioer	Hildebrand
Lipika Deka	Indian Institute of Technology, Guwahati
Dr Robert Foster	Queen Mary, University of London
Matt Gallop	BBC Learning
Paul Green	Arkessa
Dr Lin Guan	Loughborough University
Melissa Jenkins	Temeletry Ltd
Costis Kompis	Vodera Ltd
Dr Jonathan Loo	Middlesex University
Dr Miranda Mowbray	HP Labs
Dr Amyas Phillips	ARM Ltd
Stuart Revell	ICT KTN
Professor Tom Rodden	Nottingham University
Dr Monika Solanki	Birmingham City University
Alastair Somerville	Acuity Design
George Spanoudakis	City University London
John Stenlake	Living PlanIT
Paul Tanner	Virtual Technologies
Hugo Vincent	ARM Ltd
Peter Ward	WMG, University of Warwick
Dave Whydall	Dan & Adam Ltd
Dr Dave Wisely	BT