


Internet of Things (IoT)

Societal Challenges & Scientific Research Fields for IoT



*This document was written and coordinated
by Emmanuel Baccelli.*

Contributions – *Thanks to the following individuals
for their textual suggestions and input: N. Anciaux,
K. Bhargavan, G. Casiez, F. Gandon, N. Georgantas, J.M. Gorce,
S. Huot, E. Lank, A. Lebre, W. Mackay, K. Marquet, N. Mitton,
JP Richard, E. Rutten, M. Tommasi, P. Vicat-Blanc, O. Sentieys,
B. Salvy, M. Serrano, B. Smith, M. Vucinic, T. Watteyne.*

Acknowledgements – *Thanks to the following individuals
for their useful feedback: C. Adjih, C. Bormann, I. Chrisment,
F. Cuny, F. Desprez, MA Enard, JF Gerbeau, P. Guitton,
V. Issarny, P. Jacquet, L. Mé, V. Roca, H. Tschofenig, A.Viana.*

Publication date: November 2021



Contents

Executive Summary	03
PRELUDE: Alice's Cocoon	06
Alice's Dream	06
Alice's Nightmare	07
Enabling IoT Optimism & Reducing Pessimism	08
PART 1 - The Internet of Things (IoT)	09
1. IoT: Past, Present & Perspectives	10
1.1 A Brief Prehistory of IoT	11
1.2 Innovation Supporting the Emergence of IoT	12
1.3 IoT Today: the Tipping Point	14
1.4 Perspectives with IoT: Billions of New Ways to Sense & Take Action	15
2. IoT Societal Challenges	17
2.1 Legal Framework: Balancing Permissionless Innovation vs Ethical Concerns	18
2.2 Public Trust: Gaining & Retaining it	18
2.3 Sovereignty	20
2.4 Standardization	21
2.5 Education	21
2.6 Fighting Climate Change & Resource Depletion	22
2.7 Inria Contributions Addressing IoT Societal Challenges	23
3. IoT Scientific & Technical Challenges	25
3.1 How to Preserve Privacy with Pervasive IoT?	26
3.2 How to Boost Resilience, Safety & Security in IoT?	26
3.3 How to Ally Machine Learning with IoT?	27
3.4 How to Expand Last Hop Connectivity for IoT?	28
3.5 How to Stretch End-to-End Networking Principles Further for IoT?	28
3.6 What Human-Machine Interfaces does IoT Require and Enable?	28





3.7	How to Bridge the Gaps between IoT, Control and Robotics?	29
3.8	How to Enable Millimeter-Scale IoT Devices?	29
3.9	How to Tend Towards Net-Zero Resource Footprint with IoT?	29
3.10	Inria Contributions Addressing IoT Scientific & Technical Challenges	30

PART 2 - Fields of Research for IoT **31**

Communication Networks for IoT	32
Data Representation for IoT	38
Distributed Systems for the Cloud-Edge-Thing Continuum	40
Cryptology for Low-end IoT	48
Data Processing & Privacy with IoT	54
Safety, Reliability & Certification for IoT	59
Human-Machine Interaction with IoT	61
Control with IoT in the Loop	65
Security for IoT	68
Low-Power Hardware Architecture, Programming and Compiling	77
Global Resource Footprint Optimization	83

CONCLUSION **88**





© Julien Canavezes

Executive summary

Just as the Internet radically reshaped society, the Internet of Things (IoT) will have an impact on all areas of human life: from our homes, vehicles, workplaces and factories, to our cities and towns, agriculture and healthcare systems. It will also affect all levels of society (individuals, companies and state-level), from urban to rural and the natural world beyond. This makes it essential to have a proper understanding of IoT and the challenges which relate to it. The primary aims of this document are to:

- determine the scope of IoT, its origins, current developments and perspectives;
- identify the main societal, technical and scientific challenges linked to IoT.

It seems inevitable that IoT will become increasingly omnipresent. Indeed, it is set to penetrate every aspect of all of our lives, connecting everything (billions of new heterogeneous machines communicating with each other) and measuring everything: from the collective action we take at a global level, right down to our smallest individual physiological signals, in real-time. This is a double-edged sword, in that it simultaneously gives people cause for hope (automation, optimisation, innovative new functionalities etc.) and cause for fear (surveillance, dependency, cyberattacks, etc.). Given the ever-evolving nature of the IoT, new challenges

linked to privacy, transparency, security appear, while new civil and industrial responsibilities are starting to emerge.

IoT is centred around an increasingly complex set of interlinked concepts and embedded technologies. At an industrial level, this growing complexity is making the idea of having full control over all components of IoT increasingly difficult, or even infeasible. However, as a society, we must get to grips with the technological foundations of IoT. One challenge for education will therefore be to gradually increase awareness of IoT, both in order to protect individuals' sovereignty and free will, and to initiate the training of our future scientists and technicians. A public research institute such as Inria can contribute towards understanding and explaining the technological foundations of IoT, in addition to preserving sovereignty in Europe.

IoT will inevitably increase dependency on certain types of embedded technology. It is hence necessary to identify the new risks that entail, and to devise new strategies in order to take full advantage of IoT, while minimising these risks. Similarly to the situation in other domains where one must continually seek to preserve ethics without hindering innovation, creating a legal framework for IoT is both necessary and challenging. It nevertheless seems clear already that the best way of facing up to industrial giants or superpowers is to take action at the EU level, as shown by recent examples such as GDPR. Furthermore, given the growing influence of technological standards on society, playing an active role in the process of standardising IoT technology is essential. Open standards and open source – conceived as a common public good – will be pivotal for IoT, just as they have been for the Internet. Last but not least, massive use of IoT can help better capture and understand the environmental challenges we are currently facing – it is also expected IoT will help to mitigate these challenges. The goals in this context are not only to reduce the quantities of natural resources consumed by IoT (for production, deployment, maintenance and recycling). We must also aim to more accurately evaluate the overall net benefit of IoT on the environment, at a global level. This requires determining and subtracting IoT's environmental costs from its (measured) benefits, which is currently a challenge. The growing impact of IoT underscores the importance of remaining at the cutting edge when it comes to scientific research and technological development. This document therefore aims to:

- highlight the wide range of research fields which are fundamental to IoT;
- take stock of current and future research problems in each of these fields.

A number of links are made throughout the document to contributions made by Inria. These contributions are, by their nature, diverse (basic and applied research, open source software, startup incubation) and concern the majority of research fields on which IoT is based.

PRELUDE: *Alice's Cocoon*

Alice's Dream

Alice appreciates her freedom. If Alice wants it, her smart wearable devices, garments and implants (combining sensors, actuators and local wireless communications) collaborate with one another. Alice can also easily interconnect her devices with other smart devices nearby, or with more remote computing resources of her choice, reachable via the network when needed. As a whole, the system provides Alice with a personal cyberphysical cocoon which 'cushions' her experience wherever she may be: at home, on the road, or at work.

On the way to her workplace at the factory, Alice's self-driving vehicle interacts with smart city infrastructure to automatically take the less polluted route, according to Alice's preferences, and to locate a parking spot near the factory where she works. Predictive-maintenance and advanced, real-time environment monitoring used in the factory ensure a safe and productive work-place, while optimizing energy consumption. Back home, Alice's cocoon continues to customize her cyberphysical experience by orchestrating interaction with her appliances. Most importantly: Alice remains in control and can trust the system, which operates securely and safely, while preserving her privacy.

For instance, Alice uses her cocoon to provide her with advanced predictive healthcare. If Alice wants it, she can deactivate the system and delete the data at any time through a simple but meaningful interface, and she can easily swap devices, or swap remote computing resources. She can grasp the key aspects of the system at work, and self-assess her health status e.g. be informed of potential health warning signs. If Alice wants it, she can choose to share some of her health data with her doctor, temporarily or permanently. On demand, Alice's e-cocoon can also actively participate in advanced medical treatments, or in coordinated prevention of specific virus spread.

Alice can thus benefit from the best health, on a voluntary basis, and at an optimized cost for her, for her employer, and for society as a whole. And all this, even as Alice just moved to a remote place in the countryside!

Alice's Nightmare

Alice's self-driving taxi drove her into a tree. The taxi company later explained that their fleet of vehicles was disrupted by spoofed GPS signals. No serious harm done for Alice, luckily, but one of her precious smart implant was broken in the process. Managing to get her e-cocoon working again (after swapping for a replacement device) took much longer than Alice anticipated, however. On top that, to make things worse, Alice must urgently pay a hefty ransom after her other smart implant was hacked – pirates had remotely exploited vulnerabilities in its software.

Alice must now decrease her living costs. Forced by her insurance company which aims to optimize risk and profit, Alice must agree to use additional devices tracking her vitals, and must consent to her data being sold to 3rd parties. Alice's privacy decreases drastically. Surpassing Orwellian fiction, Alice's behavior is tracked in excruciating detail in real-time, and various actuators inconspicuously distort her reality.

At work, Alice is constantly spied upon by her superiors, who abuse the fine-grained tracking and actuating capabilities of the factory's IoT deployments – which themselves are increasingly targeted by cyberattacks, endangering both productivity and safety at her workplace. In her “private” life, Alice is prey to advanced surveillance capitalism; Alice's behavior is often influenced via her cyberphysical cocoon by various profit-driven entities.

Following a pilot study, the government proposes a new policy generalizing the mandatory use of the e-cocoon health tracking – aiming to decrease the country's public debt. This reform is discussed amid mounting suspicion that the recent elections were swayed through advanced profiling using real-time health tracking data, leveraged for voter manipulations at massive scale...

Ubiquitous connectivity and engineered technological dependency makes it impossible for Alice to escape the grip of the smart city. Alice feels lost, imprisoned in her e-cocoon, pondering what is left of her privacy, and indeed of her own free-will, in this mad society.

Enabling IoT Optimism & Reducing Pessimism

Alice's reality leverages the Internet of Things (IoT), which is the source of both optimism and pessimism at various levels, simultaneously creating novel opportunities and raising new issues. While such issues do not all have a technical solution, science and technology can certainly contribute to decreasing potential negative effects. In order to enable IoT optimism and to reduce causes for pessimism, technology must combine progress in a variety of scientific fields including:

- **computer networks** (so Alice's devices can communicate and interoperate),
- **miniature energy-efficient hardware** (so Alice's devices are long-lived and convenient to wear),
- **low-power embedded software** (so devices can cooperate, durably, on a small battery),
- **distributed computing** (so Alice remains flexible as to where and how her data might be processed),
- **privacy-preserving data processing** (so Alice can keep under control personal or sensitive data, and manage its use),
- **control and robotics** (to efficiently pilot Alice's sensors and actuators),
- **human-computer interfacing** (to offer simple but powerful control of the system),
- **system safety** (to ensure the actuators are not dangerous for Alice or others), and
- **system security** (to defend Alice against potential hackers).

Inria, with its 200+ project-teams in eight research centers, is active in all of these scientific areas. This document presents Inria's views on the main trends and challenges in IoT, and how its teams are actively conducting scientific research, software development and technology transfer around these challenges.

Extending into other domains, this document also identifies key societal challenges in a world depending on IoT, ranging from ethical concerns to transparency, sovereignty and education.

A number of technically-oriented whitepapers have previously covered aspects of IoT. Some have primarily focused on a subset of IoT, such as on telecommunication and regulation aspects (see [ARCEP](#) or [AFNIC](#) whitepapers), or on open source software for IoT (see [Systematic](#)'s whitepaper). Yet others have overviewed IoT from an industrial hardware vendor point of view (see [NXP](#)'s whitepaper), or from an IT service provider point of view (see [Atos](#)'s whitepaper). In this document, we provide instead a fundamentally holistic coverage of the Internet of Things, anchored around the scientific research challenges that pertain to IoT.

PART I


The Internet of Things (IoT)

In this first part, we start apprehending IoT, its genesis, its current state, and its perspectives. Then, we identify the main societal, technical and scientific challenges related to IoT.





IoT: Past, Present & Perspectives



Beyond the buzzword, **defining IoT is difficult**. Indeed, aspects of IoT range from hardware to software, from network technologies to data science, from services to infrastructure deployment, from wireless sensor networks to cloud computing. Is IoT a vision still? Or is it already happening? What is IoT? Answers to these questions are multiple and debatable – not unlike answers to the question “what is the Internet?”.

In this document, we consider IoT as the embodiment of an important part of the next-generation Internet. From this point of view, IoT consists of a set of general-purpose technologies which:

- bridge the gap between the digital world and the physical world;
- bridge the gap between Internet technologies and an increasing variety of embedded systems.

IoT terminology is not entirely settled. In this document we consider IoT as roughly equivalent to what is referred to as the *Internet of Everything* (Cisco/W3C terminology), the *Physical Web* (Google terminology), *Physical Computing* (Arduino terminology), *Machine-to-Machine* (M2M), *Cyber-Physical Systems* (Control theory terminology) or the *World-Sized Web* (a term coined by B. Schneier).

A Brief Prehistory of IoT

Even before the Internet, domotics products such as [X10](#) had already started to appear on the market. Then, in the early 1990s, futuristic visions such as the [Digital Desk](#) imagined “augmented objects” cooperating over the network and interfaces for tangible interactions blurring the frontier between the digital and the physical world. Around the same time, IoT was also anticipated by visions such as Mark Weiser’s [ubiquitous computing and embodied virtuality](#), which is gradually becoming reality.

At the end of the 1990s, the Auto-ID Center pioneered the rise of RFID tags, foreshadowing a world where virtually every object could be uniquely identified and addressed via the network. In this preliminary system, each tag featured a simplistic microchip storing only a basic serial number (to keep price down) which could be advertized in the vicinity via local wireless communication. Data associated with the serial number on the tag was stored separately in a database accessible online.

In the 2000s, new concepts and techniques enabled wireless sensor/actuator networks (WSAN, or WSN): tiny battery-powered computers first collaborate to establish (multi-hop) wireless networks, then use such networks to transport their sensors' data, or to distribute actuator commands.

Generalizing such concepts, the term “pervasive computing” (somewhat analogous to ubiquitous computing) was introduced to capture the trend of embedding some computational and communication capabilities into everyday objects. Quite synonymous, the term “Internet of Things” itself started to be widely used in the 2010s.

Since then, over the last decade, a variety of augmented objects have appeared, which offer different levels of computing power and cooperation over the network.

Innovation Supporting the Emergence of IoT

Recently, the emergence of IoT has been accelerated by innovation in the domains of *low-power embedded hardware*, *low-power networking*, *embedded system software*, and *edge computing*. An extremely wide variety of industrial players (SMEs and Big Tech) contribute to innovate in these domains, at different levels. Both new standards developing organizations (SDOs) and new IoT technological standards have emerged. For interested readers, we provide concrete references: this list is far from exhaustive, since our goal is not to establish relevance but to illustrate the diversity of this innovation.



*FIT/IoT Lab sensor network installed in the Inria Grenoble Rhône-Alpes research center.
© Inria/Photo H. Raguet.*

Embedded hardware innovation

On the one hand, cheaper single-board computers have become mainstream (e.g. RaspberryPi, NVIDIA Jetson Nano...). On the other hand, vendors such as STMicroelectronics, Microchip Technology Inc, Espressif, or SiFive develop new ultra-low-power IoT device hardware using novel micro-controller architectures designed by companies such as ARM Ltd., or open source hardware standards such as RISC-V. Complementarily, new energy-efficient security and crypto co-processors are provided by companies such as NXP Semiconductors or Nordic Semiconductors for example.

Low-power networking innovation

Smaller devices, as small as sensors/actuators, are networked using new low-power radios (e.g., LoRa, 802.15.4, BLE...), tiny general-purpose network protocol stacks (e.g., 6LoWPAN), and a quasi-infinite pool of unique network addresses (with IPv6). Battery-less, energy-harvesting communication hardware is emerging, spear-headed by companies such as EnOcean, Onio. Consortiums and SDOs such as the LoRa Alliance, IEEE, IETF, W3C or 3GPP produce new open specifications for network communication protocol accommodating low-power devices. Companies such Semtech and Texas Instruments provide new low-power radio chips used by hardware vendors and by new IoT-focused operators such as Sigfox or Actility – or by larger traditional operators such as Orange, SFR or Bouygues.

Embedded software innovation

Compact embedded distributions of Linux have emerged as the most prominent software platforms for single-board computers (high-end IoT devices based on microprocessors), to which a large pool of developers from smaller and larger companies contribute open source. On smaller devices (low-end IoT devices based on microcontrollers) new embedded open source system software platforms aggregate low-power software development, such as FreeRTOS (Amazon), Zephyr (Intel), Arduino, or RIOT, for instance.

Edge computing innovation

This trend aims to bring computation and data storage physically closer to the sources of data. New ecosystems accelerate edge computing deployment, operation and maintenance. For instance, embedded machine learning software development is facilitated by platforms such TFLite led by Google. Novel hardware co-processors and software provided by companies such as ARM, or Greenwaves

Technologies, improve the operational capabilities using Neural Networks on IoT devices located at the edge of the network. Edge computing is also facilitated by new open source management tools dedicated to IoT such as those developed by the Eclipse IoT Foundation, or by fleet management tools such as Kubernetes. Cloud services cater for IoT and edge computing with adapted services such as those provided by Microsoft Azure IoT Hub or Amazon IoT GreenGrass.

As such innovation integrates more easily with generic Internet technologies, and with common cloud infrastructure, IoT deployment is booming.

Beyond its research activity, Inria also developed support frameworks for Deep Tech projects which innovate by applying research results. In effect, **Inria incubates and spins off SMEs** with its Startup Studio program. Recent examples of such spin-offs related to the IoT domain include for instance Falco, Stackeo, Statinf or CryptoNext.

FALCO is a company providing IoT hardware, software and services helping marina ports environmental management: real-time monitoring of available slips, optimized control of resources, fight against pollution and awareness of good practices. Falco products leverage among others low-power wireless IoT technology developed by Inria researchers.

STACKEO is a software company which spun off Inria, which helps enterprises industrialize their connectivity and IoT solutions at scale. Based on a Software-as-a-Service business model, Stackeo provides a suite of tools to articulate an IoT strategy, aligning IT/OT teams and piloting sustainable value chains. Stackeo develops the concept of IoT-Architecture-as-Code, based on a dedicated modeling language, and on patented systemic methodology.

STATINF is a company providing tools enabling statistical analysis of embedded software's temporal behavior, for real-time systems involving multi-core processors. With such tooling, highly critical embedded systems from industries such as avionics, automotive, drones, aerospace can interpret the possible run-time variations of the software that uses their embedded IoT hardware.

CRYPTONEXT is a cybersecurity company which spun out of Inria and Sorbonne University, which provides software libraries implementing cryptographic algorithms designed and optimized for post-quantum security. CryptoNext also provides consulting in the realm of quantum-resistant cybersecurity.

IoT Today: the Tipping Point

Let's consider IoT devices as being any type of consumer and B2B connected devices, excluding all phones, tablets, laptops/desktop PCs. Then, between 2010 to 2020, [the number of connected IoT devices has grown a staggering 1000%](#). Approximately 10 billion IoT devices were deployed and interconnected during the last decade. To give perspective about the gradient of this evolution: in 2010, the proportion of IoT amongst connected devices was 10%. In 2020, this proportion jumped to more than 50% : in other words we are now over the tipping point. Onwards, IoT devices officially outnumber non-IoT devices. Global [market analysis reports](#) that in 2019, IoT weighed over 300 billion dollars in revenue globally – a fast progressing number.

Many IoT devices rely on low-power microcontrollers. More than [28 billion microcontrollers were shipped in 2018](#), and it is estimated that there were over [250 billion microcontrollers in use worldwide in 2020](#). Not all microcontrollers are networked, but more and more deployments rely both on microcontrollers running increasingly complex code and on them being connected either directly or indirectly to a network. The bulk connects via wireless personal-, local-, wide-area networking, or cellular. Such deployments proliferate in an extremely large variety of segments including (in decreasing [market share](#) order): automotive applications, industrial automation, personal/home consumer devices, smart meters in smart grid systems, healthcare applications, aerospace and defense domains. In the industrial segment, for instance, the network now taps into a multitude of processes which traditionally relied on local sensor/actuator control loops. Data aggregated in real-time and “digital twins” aim to better assess as a whole complex chains of supply, production and sales – and to control these, globally.

Perspectives with IoT: Billions of New Ways to Sense & Take Action

IoT deployments are planned in virtually all markets and verticals: smart home, net-zero energy buildings, e-healthcare, Industry 4.0, precision agriculture, wildlife and environmental real-time monitoring, smart cities, goods transportation supply chain, car/bike/scooter sharing... Rather than attempting to list them exhaustively, interested readers can consult these [references](#). [Projections](#) predict that [tens of billions of additional IoT devices](#) will be deployed everywhere on the planet, interacting over the network, dwarfing non-IoT connections in the near future.

In essence, IoT provides the Internet with limbs. IoT deploys new ways to communicate, to sense (gathering data), to reason (using this data), and to then act physically via actuators. New varieties of control loops via the network can leverage large remote computing resources, if necessary. By realizing and exploiting these new control loops, IoT is radically changing the landscape:

Sensors are used to observe processes more closely and more precisely. For instance, sensors can capture inefficiencies in complex systems, with a view to drastically decreasing operational expenditures. Industrial IoT leaders project that industrial processes can be optimized by combining industrial hardware, IoT wireless connectivity and a sophisticated data pipeline: issues and breakdowns can be detected much faster (in minutes or hours, instead of days or weeks with current practices), thus cutting waste by hundreds of thousands of dollars per year per plant – and generating double-digit growth in output and gross profit through increased steady-state line speed.


Actuators enable quick, adaptive, automated (re)configuration of complex cyber-physical systems. For instance, autonomic energy-saving strategies can be installed at various levels (e.g., a building complex, or an entire city) and can interact dynamically to achieve dramatic decreases in energy consumption. There is also the expectation that the reduction of our environmental impact enabled by IoT-enhanced pollution-detection and energy-consumption awareness etc. will globally outweigh the environmental impact of producing, deploying and maintaining IoT.

Cyber-physical control loops based on IoT data processing at various timescales enable new levels of control and prevention. For instance, by coupling advanced monitoring with machine learning, predictive maintenance can be performed well before issues and breakdowns actually occur on industrial hardware. Beyond merely improving existing processes, there is also the expectation that cyber-physical control loops will create entirely new processes and new services, with a high impact on society.

Robotic components both result from, and contribute to, IoT deployments. On one hand, robot-like systems emerge by combining sensors and actuators, and interacting with cyber-physical control loops. On the other hand, fleets of robots, drones, and other autonomous devices are also deployed to complement IoT infrastructure where necessary; these devices can take part in collecting data, and in dynamically providing resources.



IoT Societal Challenges



IoT has already started transforming our society, yielding a number of fundamental challenges of a societal nature. This section offers an overview of prominent challenges in this category, and a glimpse at Inria's contributions towards addressing such challenges.

Legal Framework: Balancing Permissionless Innovation vs Ethical Concerns

Inheriting from the Internet, IoT is subject to extremely rapid growth and driven by partly conflicting interests.

On the one hand, Internet market dynamics and technologies are fundamentally designed to foster fast, permissionless innovation. On the other hand, ever-growing ethical concerns emerge about the use of these technologies – for instance, with respect to privacy or the environment.

Average IoT users cannot be expected to fully grasp the implications of their usage of IoT products. Thus, evolving appropriate legal frameworks are crucial to guide IoT.

However, the speed of adoption typically outpaces the setup of legal frameworks and regulations, creating grey zones and gaps. With the expected “penetration” of IoT – in terms of scale, nature, and the granularity of data collected – such gaps could have dire consequences.

Careful action is thus required to minimize such negative impact without hindering innovation. Here, capacity and responsibility are primarily in the hands of governments and regulation bodies. In this domain, the benefits of acting at EU-level have been demonstrated by the impact of GDPR, for instance.

Public Trust: Gaining & Retaining it

In order for IoT adoption to thrive, more transparency is needed, as well as end-user empowerment with respect to governments and industry.

IoT deployments tying together multiple stakeholders are often driven by interests that do not remain aligned over time. Without the necessary technical or legal tools, such misalignment cannot be handled in a fair and appropriate way.



Capture of ambient noise via cell phones to establish a collaborative map of noise pollution.
© Inria/Photo C. Morel

In particular, compared to a B2B context, consumers in a B2C context are more exposed to such issues, and may require extra care.

The public has become increasingly aware of how Internet technologies are being used for state-driven mass surveillance (e.g. Snowden scandal), for profit-driven [online piracy](#) or for [surveillance capitalism](#) using advanced advertisement personalisation.

Push-back and criticism is taking shape in the public sphere against plans for more [ambitious IoT pilot deployments](#) which could foster opaqueness, or favor de facto industrial monopolies. Meanwhile, user trust erodes as controversies flare when more security flaws (often quite [basic!](#)) are exposed on countless IoT devices, and more privacy-threatening [processes](#) or [hidden functionalities](#) are uncovered. Further controversies are to be expected: for example, potential trade-offs of privacy versus governments needs for surveillance, in a national security context.

A major characteristic of the Internet – and arguably, one of its strengths so far – has been reliance on approaches fostering end-user empowerment to [make choices that are not dictated by a specific vendor or government](#). However, recent political and technical trends raise fears that end-user empowerment and transparency is losing ground. A crucial challenge for IoT in this context is

to contribute to regaining and retaining these characteristics – and thereby, public trust. If IoT is obviously devoid of such potential, then IoT adoption may not turn out as expected.

Sovereignty

If not kept in check, IoT could increase dependence on technology which may be neither geopolitically neutral, nor privacy-compatible.

As technology plays an ever more central role in people's lives, aspects of sovereignty become prominent, both at the individual level, and at the state level.

- For individuals, the challenge is to maximize and maintain one's ability to break free from dependence on providers and technology which could threaten privacy. For example, in order to interact with their IoT devices, a sovereign user might prefer a setup via an intermediate machine under his full control, compared to setups offering more "seamless" integration with a cloud service provider. Large "walled-garden" companies might discourage some approaches, considering that [the customer is the attacker who must be contained](#).
- For states, the challenge is to minimize dependence on technical solutions which can be weaponized. To give a concrete example: dependence on Linux can be considered rather neutral from a geopolitical standpoint, whereas deciding to rely on Android is a significantly less neutral decision – even though Android is technically based on Linux. Making such choices can be tough in areas which have long been neglected (or outsourced).

In any case: there is no free lunch. Sovereignty comes at the cost of additional effort. Typically, for individuals, this translates into reduced convenience. For states, the tradeoffs mimic those typical of corporate environments: sovereignty generally requires significantly more investment in research, development and maintenance. Furthermore, if the desire for sovereignty leads to the development and use of more separate, concurrent systems, then other issues kick in – interoperability, for example. The crucial challenge is therefore to optimize such overhead, to obtain the "biggest bang for the buck".

Standardization

As IoT technology becomes more tightly woven into society and our individual lives, the design of IoT standards becomes ever more critical.

The essence of IoT consists of enabling heterogeneous embedded systems to connect and interoperate over the network, potentially at large scale. To no one's surprise, standardization is thus prominent in the IoT space, and in particular IoT network communication standardization. For instance, standards development organizations (SDO) such as the World Wide Web Consortium (W3C), the Internet Engineering Task Force (IETF), or the Institute of Electrical and Electronics Engineers (IEEE) are prominent venues where key technical decisions are made which have endless ramifications – including at societal levels.

In practice, [no standards development organization is value-neutral](#) and this applies to IoT standardization as well. In particular, the concentrated power of Big Tech over society, combined with their decisive influence in shaping the next standard (or de facto standard) creates situations which must be addressed. Further, there is a need to expose the relationship between IoT [technology standards and human rights](#), for example. The challenge here is thus to increase the general awareness of both how standard technologies fundamentally work, and how their characteristics might have impact on societal aspects.

Education

The more IoT harnesses “invisible” computers – such as unconsciously used sensors and actuators – the more it becomes necessary for education curricula to include the basics of cyber-physical systems, and to point out their potential pitfalls.

Education for computer science in general is a challenge, even more so for IoT. This is in part due to the fact that IoT technology is still extremely fragmented – especially for low-end IoT. There are several aspects to this challenge:

- ***Developing the required hard skills:*** on a purely technical level, education issues are obviously exacerbated in regions where generic technical computer science education is already lagging behind. Some technical skills required for IoT, like deep embedded programming, are skills that are too rare.
- ***Developing the necessary soft skills:*** at a less technical depth, it is still desirable to have some level of education to “raise awareness”, and develop common sense and critical thinking with respect to what IoT products can/should or cannot/should not do. Such critical thinking may be on the verge of becoming uncommon – a potential shift for our society.

From a philosophical standpoint, the very way we will apprehend reality in a fully IoT-enabled world is in itself a subject of study. Through sensors and actuators, not only virtual reality, but also physical reality may be customized and experienced differently. On one hand, efforts such as [postphenomenology](#) aim to characterize the interplay between humans, the natural world and modern technologies – which are increasingly becoming non-neutral mediators. On the other hand, some designers already work on the design of [socio-technical systems](#): that is, a design that not only considers digital and technological material, but also human users themselves as material.

Fighting Climate Change & Resource Depletion

IoT can be an instrument to fight climate change, but the proliferation of IoT gadgets is also a vector for resource depletion.

Deploying billions of IoT devices will globally consume vast amounts of energy and resources (including plastics, metal, batteries) to produce, ship and operate these devices. In the context of the current ecological crisis, the relevance of these devices should be scrutinized. In principle, IoT provides key components and tooling that are necessary to precisely track both climate change and the



Predicting frost events in peach orchards. © Inria / Photo G. Scagnelli.

effect of policies put in place to curb it, in real-time. Furthermore, IoT can provide the means to automate dynamic adjustments and to implement optimizations in a wide variety of complex systems that are resource and energy-hungry – for example industrial workflows, or smart home/building energy management systems. There is thus no doubt that IoT can be useful to fight climate change. However, the energy and resource-efficiency of IoT should be continuously investigated, and evolving legal constraints should be crafted to guide this efficiency.

Inria Contributions Addressing IoT Societal Challenges

As a governmental agency combining broad scientific expertise and substantial applied research output in IoT-related domains, with impact on education, sovereignty, standards and legal frameworks, Inria contributes to gaining public trust with regards to using new IoT technologies.

Inria contributes in practice to enabling **sovereignty** in IoT, in multiple ways. By conducting scientific research peers consider excellent, and by systematically publishing results in open access venues, Inria contributes to enabling **sovereignty**. But beyond such involvement from a scientific point of view, Inria contributes on other levels too. Inria teams engage in international, cross-disciplinary technical collaborations, and oversee the deployment of new IoT technologies in various real-world applications.


Inria produces, publishes and maintains open source software. Large communities of users have snowballed around several high-impact open source software projects spearheaded by Inria (RIOT, Scikit-learn...). Such software building blocks are platforms which can be depended upon to both deliver top performance and ensure geopolitical neutrality. Another example of Inria's contributions to **sovereignty** in IoT is how Inria applies its scientific expertise to open **standardization** efforts. As a result, Inria regularly co-authors the new **open standards** specifications for IoT, published by high-impact standards development organisations e.g., IETF, and W3C.

Inria also uses its expertise in diverse scientific domains related to IoT to contribute to debates regarding **legal frameworks and safe-guards** for IoT. For instance, research activities at Inria include evaluations of practical compliance issues with the GDPR European regulation, for consumer IoT products in the smart home segment.

Inria actively contributes to **education** in the field of IoT, on several levels. Many Inria researchers also teach, in the context of a university professorship or equivalents. But Inria also contributes to **education** beyond this traditional involvement of researchers in teaching. For instance, Inria provides massive open online courses (MOOC) tailored for wider audiences in the field of IoT, followed by thousands of participants around the world. Inria also shaped educational programmes aiming at younger age categories, equipping 15-year-olds with knowledge to better grasp digital technology. Another prominent example of Inria's involvements in educational activities is Inria's series of white papers and white book publications, such as this present document.



IoT Scientific & Technical Challenges



What we expect from IoT ranges from empowering individuals with revolutionary services, to having major impact on society and industry, and possibly saving the planet by globally fighting climate change.

What scientific challenges lay ahead, in order to realize such high expectations? We summarize below a shortlist of questions that we consider key for IoT.

How to Preserve Privacy with Pervasive IoT?

Permanent tension exists between IoT data exploitability and IoT user privacy. At one extreme, IoT user data should be massively exploitable when it's a matter of saving lives. At the other extreme: pervasive IoT deployment could contribute key building blocks for an electronic "[panopticon](#)". Such a system could potentially annihilate individual privacy entirely. A crucial question is thus: how, and to what extent, can we guarantee strong privacy protection while retaining IoT data usefulness?

One type of issue is the design of novel paradigms and pre-processing techniques applicable directly on the IoT devices, that obfuscates parts of the data which do not correspond to a specific "legitimate interest". Another type of issue is the design of novel cryptographic primitives, applicable even on low-power IoT devices, potent even in face of post-quantum attackers.

How to Boost Resilience, Safety & Security in IoT?

As we depend more on new services built on top of IoT, its resilience in face of partial infrastructure outage or subsystem malfunctioning becomes crucial. IoT deployments involve increasingly complex distributed system architectures, whereby resilience-by-design remains a major challenge. The level of dependability also impacts safety and security in IoT.

Cybersafety used to primarily concern cyberspace: keeping digital information secure, using good 'netiquette' (internet etiquette) etc. With IoT, however, cybersafety extends from virtual to physical space: cybersafety is now also about securing your physical integrity, and about protecting one's environment in the real world.



Map of Things: data collection and information for users of connected objects. © Inria/Photo C. Morel.

What are the new security or safety risks with IoT? For which benefits? Mere “IoT gadgets” may not be worth the risks. New models are needed to capture both attackers and safety aspects, in complex IoT contexts. Then, based on these models, novel mechanisms are required to provide guarantees on IoT software, hardware and communication for IoT devices – throughout their life-time, which can be decades. This challenge is exacerbated for low-power IoT devices, which are the new “weakest link”.

How to Ally Machine Learning with IoT?

Applications using artificial intelligence and machine learning (ML) are being developed in a fast-growing variety of domains. From this perspective, IoT is simultaneously a huge provider of precious data necessary to train ML models and a consumer of inference capabilities based on these models.

A key question is how to harness much more IoT data, with much less strain on privacy and on network load. One challenge is to design alternatives to centralized ML model training, which distribute training, for example in a peer-to-peer, federated fashion. How robust and efficient can distributed learning be? How far down can we push resource requirements for a learning peer? Conversely, another challenge is how to fit inference capabilities on smaller IoT devices, with less performance loss, while retaining the flexibility of modifying the models used on these devices down the line.

How to Expand Last Hop Connectivity for IoT?

IoT crucially depends on the connectivity of billions of new devices located at the edge of the network. To accommodate the resulting surge in traffic, next-generation network protocols are needed to cope with communication medium scarcity, and enhance the capacity of the segments device-to-infrastructure, and device-to-device.

With so many devices including ultra low-power devices, next-generation network technologies for the last hop must be extremely affordable both in terms of capital expenditure, and operating expenses. How can we foster an Internet-like dynamic, whereby “connectivity is its own reward”? At the same time, better penetration indoors and lower power requirements are expected, as well as longer range communication outdoors to reach more remote locations. The notions of “last hop” and “IoT devices” will have to expand to extraterrestrial domains, as shown by the current orbital-and space-races.

How to Stretch End-to-End Networking Principles Further for IoT?

The Internet has scaled quickly based on the principle that intelligence is placed at end-points, rather than hidden inside the network. IoT challenges this principle. To which extent can we expand the fundamentals of end-to-end networking?

Over the last decade, end-to-end principled mechanisms such as the open standard 6LoWPAN IPv6 network protocols and preliminary work on the Web of Things (WoT) hint at what end-to-end IoT might be in the future. However, challenges remain to complete an architecture adequate for IoT. Delegating parts of the intelligence to proxies seems unavoidable for very low-power IoT devices; but how, and to which extent? Furthermore, new protocols and semantics are required to automate additional levels of Machine-to-Machine communication.

What Human-Machine Interfaces does IoT Require and Enable?

On the one hand, IoT adds billions of new interfaces to the physical world, from cyberspace. Conversely, how can human IoT users better interact with cyberspace, via such interfaces?

Improving both the accessibility and the control of such interfaces is crucial. One aspect is identifying the adequate levels of control with which to empower humans. Humans are a heterogeneous crowd, with varying capacities and requirements to grasp what happens under-the-hood of IoT. Another aspect is to design novel ergonomics that can physically embody some parts of cyberspace. The challenge is to help users learn, understand and appropriate IoT technology, even as it grows and changes over time.

How to Bridge the Gaps between IoT, Control and Robotics?

As IoT interconnects sensors, actuators, with computing capacity, available locally or remotely, over the network, new varieties of control loops appear.

Which loops can be exploited? How should the control be tuned? On the one hand, research in the field of Industrial IoT aims to establish and optimize control over extremely complex chains of supply and production. Distributed hybrid systems mixing continuous states and discrete events are very challenging to model and to control. On the other hand, micro-robots and swarm robotics can be seen as a next frontier, as they require addressing many open research questions simultaneously, including latency predictability, mobility, localization and resource parsimony matching low-power requirements.

How to Enable Millimeter-Scale IoT Devices?

Recent developments in micro-electronics push the limits of extreme miniaturization: prototypes of chips the size of a grain of rice (or less) can sense, compute and communicate via wireless, with no additional components required. This *Smart Dust* has the potential to revolutionize micro-wearables and swarm robotics, but presents unique challenges with a view to interoperate with standard low-power wireless communication, in terms of embedded programming, and calibration.

How to Tend Towards Net-Zero Resource Footprint with IoT?

IoT can be a tool to execute environmental policies, or to measure their effect. But what about the resource footprint of IoT itself?

Minimizing the environmental cost of producing and operating individual devices is key to diminishing the footprint of IoT, and to make the necessary massive deployment affordable. Producing electronics using significantly fewer non-renewables such as plastic and metal is one aspect. Another concerns designing new embedded hardware, software and networking paradigms that can exploit intermittent ambient energy, and offer new tradeoffs in terms of performance vs energy.

The evaluation of the global impact of IoT remains challenging. Globally, what are the net benefits, and IoT's footprint? A complex interdisciplinary effort is required to capture the full picture of direct and indirect impacts, complete life-cycles, and induced effects.

Inria Contributions Addressing IoT Scientific & Technical Challenges

In order to address the scientific questions identified above, we need computer science in diverse domains, including:

- *communication networks,*
- *data representation,*
- *distributed systems,*
- *cryptology,*
- *data processing & privacy,*
- *safety, reliability & certification,*
- *human-machine interaction,*
- *control,*
- *security,*
- *low-power hardware architecture, programming & compiling,*
- *global resource footprint optimization.*

Inria researchers conduct scientific activities contributing to advance state-of-the-art in each of the above fields of IoT research. In the second part of this document, we take a closer look at the array of IoT research challenges within each of these domains.

PART II

Fields of Research for IoT

We now dive further into different fields of computer science research which must be harnessed, and which require advances, in order to address the key scientific questions we identified, concerning IoT.

> The below chapters (capturing broad research fields) intend to be rather self-contained, readable in any order. For example, the reader may choose to read the chapter on Cryptology for Low-end IoT before reading the chapter on Communication Networks for IoT. Within these chapters, the document does not aim for exhaustive coverage of related research topics, but instead aims at highlighting the diversity of problems being worked at.

> Speed readers and those least interested in further science and technology dives may briefly skim through the next chapters/section titles, and skip to the end of the document



2.1 Communication Networks for IoT

Both the edge of the network and the core of the network expect significant impact, as IoT and Machine-to-Machine (M2M) communication ramp up.

IoT Access Network Protocols Optimization

At the edge of the network, the multitude of IoT devices relies on efficient and available connectivity to attach to the network. Studies already show that we are approaching a state where [75% of devices and connections on the Internet come from the consumer segment](#). A particular challenge concerns wireless access protocols.

High-end IoT devices, increasingly throughput-hungry, will inevitably challenge state-of-the-art wireless. Announced high-power (Watt-range) wireless access technology such as Wi-Fi 6 or cellular 5G access, will be challenged by upcoming network-heavy applications such as fully immersive, IoT-enhanced virtual reality, augmented reality or autonomous vehicles. The design of optimizations and new protocols in this space is thus an ever-promising avenue for research. The scientific challenge is to approach the transport capacity limit of the wireless medium (imposed by laws of physics and information theory), to make the best of the available spectrum of radio frequencies. An associated technical challenge is the design and production of novel chips dedicated to optimized radio communication, which require enormous R&D investments.

Techniques such as massive MIMO are promising to tackle the density challenge, but require deeper investigation to reach full potential. It is also worth mentioning that new algorithmic solutions become appealing based on machine learning techniques (e.g., reinforcement, deep learning). We thus expect that the design of optimizations and new access protocols for high-end IoT devices will be an active and challenging area of research in the future.

➤ At Inria, the **MARACAS** project-team combines information theory with statistical signal processing, control theory, game theory and machine learning to explore new technologies optimizing communication in wireless physical (PHY) layers. On top of wireless PHY layers, projects-teams **TRIBE** and **EVA** work on wireless multiple access (e.g. modern random access) optimizations, also exploiting machine learning.

Exploring complementary and radically different communication techniques is necessary to alleviate pressure on the radio frequency medium. We thus expect work and challenges stemming from investigations on alternative communication paradigms such as nanocommunications or visible light communications (VLC). Such emerging technologies, at the frontier with physics research, deserve attention and require further pioneering work.

➤ At Inria, project-teams **AGORA** and **FUN** investigate novel alternative access network communication technologies using visible light communications (VLC). RFID communications, and hybrid infrastructures, together with new services they could provide.

Enlarging IoT Access Network Coverage

Compared to high-end IoT devices, low-power IoT devices rely on different technologies and protocols for network access. Instead of targeting maximum throughput assuming Watt-range devices, low-power protocols target minimum energy consumption (milliWatt range or less) for low-to medium-throughput IoT devices. The field of low-power Personal/Local Area Network (PAN/LAN) **radio technologies tackles lower and lower energy consumption** for 1m-100m distances, and aims at being frugal enough to only require ambient energy harvesting. On the other hand, the field of low-power Wide Area Network (lpWAN) radio technologies aims to trade off somewhat lower throughput for much longer distances (10km or more) for the same low-energy budget.

➤ At Inria, the **DIONYSOS** project-team works on the design of wireless access control mechanisms for large NB-IoT networks.

An overarching challenge is to **extend network coverage** to a myriad of low-cost, low-power IoT devices. Beyond the issue of efficiently managing crowded access networks, this extension is challenging both geographically and economically. Geographically, coverage must reach both deeper indoors, and farther outdoors in more remote areas. Economically, the **network access technology must be kept extremely affordable for low-cost devices**.

Nevertheless, the kilometer range coverage is insufficient to cover either very remote devices (such as in smart agriculture scenarios where sensors can be deployed in very wide fields far away from urban infrastructure) or to provide enough throughput to comply with application requirements. For instance, in this domain, it remains to be determined what is the achievable performance for IoT access connectivity via deployed fleets of low-orbit satellites.

Multi-hop wireless communications is another alternative – possibly leveraging multiple radio technologies – but requires the design of new routing protocols tailored for this environment applications and environments.

↗ At Inria, project-teams including **AGORA**, **FUN**, **MARACAS** and **TRIBE** investigate performance optimisations for low-power wireless sensor deployments and multi-hop routing and data dissemination protocols in this context, using methodologies combining theoretical modeling, simulation, and experiments.

Another research avenue in this domain aims to complement the fixed infrastructure with a fleet of temporarily dispatched robots, deploying swarms of unmanned aerial vehicles (UAVs, drones) or ground robots. Such additional resources are to be deployed quickly, to restore connectivity after a failure, to monitor a one-off event, to explore or supervise an unknown/hostile area or simply to periodically offload data from remote devices out of reach otherwise. The perspective of deploying mobile unmanned entities, either rolling or flying, is becoming attractive as more manoeuvrable devices become commercially available. However, challenges remain to enable full autonomy for such devices, which much both abide to very constrained budgets in energy, computing, and storage capacity, while ensuring great coordination, robust communication and task sharing.

↗ At Inria, project-teams including **ACENTAURI**, **FUN** and **DANTE** design self-deployment algorithms for UAV and ground robots, aiming to optimize robot placement and/or to increase the autonomy of robots.

Protocol Optimization for Core & Edge Networks

Currently, the core of the network has to carry data measured in exabytes (billions of billions of Bytes) monthly. In the future, driven by a steep growth in IoT traffic, Machine-to-Machine, **the core of the network will have to carry through orders of magnitude more data.**

A significant challenge is thus to limit the speed of growth for IoT traffic flowing through the core of the network. For that, an active areas of research **explore alternative network protocol architectures which aim to better leverage in-network data storing and processing**, as close as possible to the origin of the data, such as Edge computing and Information-Centric Network.

An associated avenue for research concerns latency. Average latency between sensing and acting with network in the loop can be a few milliseconds. However, some real-time IoT applications (e.g. tactile Internet, tele-surgery...) strictly require latency below a few tens of milliseconds. **Achieving combined requirements of ultra-low latency, cost and complexity for real-time IoT applications** is thus a major challenge.

↗ At Inria, the **DIANA** project-team designs, implements analyses new network architectures, services and protocols in the context of hundreds of billions of wireless devices, that target service transparency and better control on user data.

Communication Standardization for IoT

The essence of IoT consists in enabling heterogeneous embedded systems, to connect and interoperate over the network, potentially at large scale. To no surprise, network communication standardization bodies are thus prominent in the IoT space, including:

- **IEEE**, working on physical and link-layer communication protocols, in particular the IEEE 802.15 working group, developing Bluetooth and IEEE 802.15.4 wireless standards;
- **IETF**, working on network, transport and application layer communication protocols, including many working groups developing 6LoWPAN network protocol and IPv6 communication security standardization for low-power IoT;
- **W3C**, and in particular the Web of Things working group, developing web resource identifier schemes for semantic interoperability between IoT service suppliers and consumers.
- **3GPP** is the prominent standardization organization which develops protocols for mobile cellular telecommunications (NB-IoT, 5G etc.).

↗ At Inria, project-teams including **EVA**, **PRIVATICS**, **TRIBE** and **WIMMICS** contribute actively to the design and standardisation of protocols and data models for low-power networks, within standardisation bodies such as the IETF and W3C.

However, beyond IEEE, IETF, W3C and 3GPP, a maze of standardization bodies is flourishing – and constantly re-arranged. As a result, **the moving landscape of partly overlapping, partly competing standards is challenging to navigate.**

Network access via radio varies between LoRa, Sigfox, NB-IoT, Bluetooth LE, ZigBee, Dash7, EnOcean, WirelessHART, DECT ULE, UWB... Network access via wire varies between PLC, KNX, BACnet, CAN... A wide variety of higher-level communication standards and semantic data models are issued by a large variety of bodies and alliances including, but not limited to: OMA SpecWorks (developing the LWM2M standards), OPC Unified Architecture (OPC UA), OASIS (developing MQTT standards), DotDot (Zigbee Alliance), One Data Model (OneDM)...

Unfortunately, not only the landscape, but also IoT protocol specifications themselves often turn out complex and challenging in view of guaranteeing actual interoperability and security between devices from different vendors. An ever-growing set of **alliances developing interoperability and certification frameworks emerge in various verticals** including (but not limited to) the Thread Group, Connected Home over IP (recently rebranded as Matter), Zigbee Alliance, Open Connectivity Foundation (AllJoyn), WiSun...

A significant challenge ahead is IoT communication technology standards landscape consolidation. With the advent of IoT, technical people from very different domains and cultures must suddenly collaborate: for instance embedded hardware engineers vs the community developing Internet protocols and software. These communities often experience trouble talking to one another, hinting at a basic lack of common terminology.

Even within each technical domain, excessive heterogeneity is a harsh reality which must be coped with in IoT, at different levels including hardware, software, network protocols, technology standards. Technologies in these fields have to converge towards a handful of standards (some de facto) based upon which a “sweetpot” will develop, fostering quicker progress and large-scale interoperability, while avoiding the pitfalls of “monoculture”. In a nutshell: market consolidation is yet to happen concerning IoT technologies.

End-to-End Low-power Network Protocols

Low-power, general-purpose internetworking protocol stacks have emerged. The best example is probably the IPv6 protocol stack based on 6LoWPAN and CoAP, standardized by the IETF. However, in practice, nagging protocol mismatch persists at different levels. A typical mismatch is IPv6 which is often not supported natively at the edge of the network, which supports only IPv4. Another typical mismatch often appears between cloud providers and device vendors: the latter speaks CoAP over UDP above the transport layer, while the former speaks only HTTP over TCP. Yet another mismatch is the data encoding/model that is used: the semantics of data aggregated from different IoT devices will typically not correspond.

Work-arounds do exist (e.g. for the above mismatch, IPv6-in-IPv4 tunneling, or CoAP-HTTP proxying, data encoding translation). However such mismatch is a barrier of entry which is still substantial (if not fatal) for most non-specialists. Eliminating these barriers thus remains a challenge. Eliminating these barriers will not only significantly lower IoT's cost of entry, but it will also pave the way for advanced in-network storing and computation paradigms standardization for IoT, that are needed to realize edge computing, and to relieve core network load.

Beyond workarounds (proxies, gateways...) the challenge for new low-power network protocol standards is wide adoption end-to-end, throughout the cloud-edge-thing continuum. In particular, since low-power IoT devices are extremely constrained in terms of resources, end-to-end solutions cannot simply resort to "adding another protocol layer on top of legacy". Low-power network protocol standard integration is a crucial issue.

A key question is thus: how far can the end-to-end networking principle stretch towards reaching low-power IoT devices themselves?


2.2 Data Representation for IoT

As we connect more and more objects to the Internet, they become visible to various applications running throughout the network, but as a whole, they produce raw IoT data that is typically not straightforward to exploit en masse. Heterogeneous IoT devices (‘Things’) often use different data representation schemes, and semantics vary. This reality fragments the IoT further, at the application layer. New approaches are needed to allow developers to build applications that span a disparate variety of objects and technologies — to link Things with other parts of the system, a unified framework is required.

A prominent proposal in this space is the [Web of Things \(WoT\)](#) which consists in relying on the Web as universal application platform for connected objects — literally [envisoning](#) a **Web of everything, and running on anything**. This opportunity comes with the challenge of evolving classical Web techniques to address the size, heterogeneity and specificities of IoT devices and networks.

For instance, key concepts brought by the Web include URIs to identify devices, services and players, providing the default dereferencing mechanism to obtain rich descriptions for newly discovered identifiers. To realize the Web of Things, new levels of flexibility are required from Web-based standard models that cater for:

- the vendors, to describe their products and services’ features;
- the platform and application providers, to expose these features;
- the users, to express their goals and requests, etc.



Distributed machine learning for IoT applications to drive the creation and evolution of complex networks. © Inria / Photo L. Jacq.

The flexibility of the description is also vital to allow manufacturers to differentiate their products from their competitors and offer a range of models with different features, while maintaining a Web-wide interoperability. These generic models must enforce shared standard when appropriate (e.g. core data types, physical units) and provide programming language neutral representations that support cross-platform and cross-domain interoperability. Yet they must also support extensions for application and domain-specific models such as more advanced typing and definitions for more complex data structures dedicated to specific usages and scenarios.

Languages of the Web, and in particular **Linked Data and Semantic Web**, can provide interoperability at higher levels and standard formats for descriptions of objects, operations, inputs, outputs, etc. There is a need for description in the WoT – a special case of an open question in the field of Semantic Web: how to formalize cross-platform and cross domain vocabularies? The challenge is to provide languages and vocabularies with the adequate expressivity to formally represent the web of “digital twins” of things. In the WoT, Things are viewed as software resources identified on the Web with characteristics, operations and events to be described and linked in order to support discovery, interoperability and composition. The models to be designed must not only allow us to describe very heterogeneous Things (their needs, capabilities and characteristics) but also the platform and the cyberphysical context.

As we need to represent, publish, query, validate and infer from these metadata and the data exchanged on the WoT, we have to design and standardize abstract models, concrete syntaxes and efficient serializations, processing languages applicable in a resource-limited and dynamic context, etc. The classical approaches of linked data on semantic Web can bring solutions (e.g. ontology languages) but also face specific challenges stemming from the Web of Things. A significant challenge for semantic Web models is the ability to cope with the dynamicity of the WoT including flows of data (e.g. sensor output) and the rapid reconfiguring as things are connected or removed. Another challenge is the need to capture and adapt to the application context in terms of scope, distribution, limitations, privacy, user profiles, etc. Standard data models must also support the provision of Web-based scripting languages for platform-independent Thing-to-Thing interaction, applications and management. Linked data query and validation languages can be used in accessing and validating the descriptions against constraints. Ontology-based representations and reasoning can support composability and interoperability with pivot languages and transformations. The pivot languages and transformation mechanisms provided by the Semantic Web can also support the capture and mitigation of the differences between available protocols (e.g.

HTTP, WebSockets, CoAP, MQTT) to provide uniform means to send and receive messages to things and services such as declarative bindings. Security schemes can also be expressed and exchanged among heterogeneous systems.

↗ At Inria, the **WIMMICS** project-team works on the integration of autonomous agents to the Web of Things (WoT), relying on Semantic Web languages and linked data principles to bridge the Web architecture and the Multi-agent architecture style. The goal is to provide an open standard environment to deploy intelligent agents and behaviors in IoT systems such as in factory automation scenarios.

2.3 Distributed Systems for the Cloud-Edge-Thing Continuum

Designing, developing and running applications in the IoT requires mastering and managing its complexity in terms of distribution, heterogeneity, dynamicity and scale.

Middleware for IoT

Middleware traditionally undertakes the handling of such issues, transparently, for distributed systems, either in the shape of a software layer operating on each device, and/or a software entity somewhere in the network acting as a broker. However, middleware design faces new unique challenges when it comes to supporting smart spaces and applications in the IoT.

The main challenge is to cope with the high level of uncertainty characterizing the IoT execution environment, which contrasts with the typical software engineering process, whereby a system is finalized during its design phase. The IoT context is ever-changing, and the complexity of changes (that the IoT system must adapt to) is such that it cannot be tackled at system design-time. Due to their automated, dynamic, environment-dependent composition and execution, IoT systems emerge in non-anticipated ways. Both systems and their properties take their complete form only at runtime and typically evolve down the line, requiring unforeseen levels of interoperability.

A prominent application of IoT middleware is massive participatory crowdsensing e.g., for urban environmental monitoring. A related challenge for IoT middleware is the design of novel algorithms and protocols able to efficiently integrate the massive adoption of smartphones and other user-controlled IoT devices, and to manage large scale, dynamic user involvement in IoT. Besides physical sensing – where a device's sensor passively reports the sensed phenomena – social sensing comes into play, where the user is aware of and participates in the sensing of the environment.

Yet another associated challenge yields from the low accuracy of sensors and the uncontrolled conditions in participatory sensing. It is particularly difficult to raise opportunistic crowdsensing to a reliable means of observing environmental phenomena. One area of research is the design of spontaneous, decentralized coordination schemes among mobile sensors.

➤ At Inria, project-teams including **MIMOVE** and **SPIRALS** work on middleware solutions supporting smart spaces and applications in the IoT. For example, DeXMS is a middleware developed by Inria aiming to provide dynamic system interoperability, composition and scheduling for emergent IoT systems, while relying on computational resources at the network edge. Other middleware platforms developed by Inria, such as APISENSE, or SenseTogether target mobile crowdsensing. Such middleware aims to enhance IoT data quality, to leverage context-awareness, as well as edge resources, including the mobile crowdsensors themselves.

Test Platforms for the Cloud-Edge-Thing Continuum

IoT contributes billions of connected devices with small computing capacity. These devices can collaborate over the network. Collaboration can take place either amongst IoT devices, or with some nearby machines providing medium computing capacities (a.k.a edge or fog-computing) or with more remote machines providing large computing capacity (cloud-computing). Thus a Cloud-Edge/Fog-Things continuum emerges, all along which computation can be distributed adequately, depending on high-level requirements (which can evolve over time).

↗ Inria leads the development of large test infrastructures for experimental research on the Cloud-Edge-Thing continuum. For instance, Inria develops **SILECS**, an open-access test platform enabling researchers to generate and deploy the full stack, software and protocols, end-to-end, from small connected objects to the large data centers. The goal is holistic, fine-grained capture of events, from sensors/actuators, to data processing & storage, radio transmissions & dynamic deployment of edge computing services.

Orchestrating Cloud/Edge/Fog/Things Resources

Cyberphysical system designers are tasked with programming the continuum Cloud-Edge/Fog-Things. In particular, vastly different techniques and paradigms are used so far to program and manage these different categories of machines (cloud, edge, thing). As a result, it is difficult to grasp and assess the system as a whole. Dynamic orchestration of Cloud-Edge/Fog-Things resources thus remains a challenge. Other stemming challenges include holistic assessments of a cyberphysical system's global security and privacy characteristics. A research avenue in this field is the design of a unifying syntax to program all cyberphysical system components, which would enable global cyberphysical characterization. Such an approach aims to tackle and enforce security across the continuum, (instead of tackling security component by component) and to simplify the programming IoT's non-trivial temporal behavior mixing synchronous and asynchronous activities.

↗ At Inria, the **INDES** project-team works on designing secure, multi-tier programming languages for IoT, syntax capturing the continuum, from microcontrollers to the cloud. For instance, **INDES** develops Hop.js and HipHop.js. With these JavaScript dialects, servers, clients, and IoT devices are all programmed in the same code, simplifying IoT system design, and enabling global security enforcement.

DevOps for Cyberphysical Systems

Toolchains are required to reduce the time between committing a change to a (large, distributed) system, and that change being deployed in production, while ensuring high quality. DevOps is both a technical and research field which uses and develops such toolchains, for software, encompassing all stages of coding,

building, testing, packaging, releasing, configuring, monitoring during the lifetime of the system. In practice, DevOps is necessary to allow the agile development and maintenance expected from modern, Internet-age software. Traditional toolchains are typically not applicable on smaller IoT devices, due to constraints stemming from low-power networking and on-board resources. With cyberphysical systems spanning the full continuum (Cloud-Edge/Fog-Things) a challenge is thus the design and implementation of novel, comprehensive toolchains, which extend support to smaller, low-power IoT devices. A naturally associated challenge is assessing and guaranteeing the security properties of these extended toolchains.



Fixing bugs in remotely operated IoT objects with Pharo. © Inria / Photo Raphaël de Bengy.

↗ At Inria, the **TRIBE** project-team works on designing secure DevOps approaches which can stretch down to ultra low-power IoT devices, applicable not only on machines using microprocessors, but also on machines using smaller connected microcontrollers.

Optimal Computation Placement in post-Cloud IoT

IoT data crunching in the cloud can be complemented (or entirely bypassed) by using computing power nearer the origin of raw IoT data. Not only can intermediate edge/fog devices contribute computing power, but also to some degree the low-end IoT devices themselves. Post-cloud IoT potentially offers data (pre-processing

capabilities anywhere in the network, end-to-end. With such a possibility, a challenge which emerges is to identify and implement strategies for the optimal placement of IoT services, data computation or pre-processing, along the continuum (cloud/edge/fog/things). An associated challenge is the automation of service migration along the continuum, to dynamically adapt the cyber-physical system to evolving high-level requirements. For instance, research communities such as [COIN](#) are exploring alternative software-defined, data-centric network architectures and network function virtualization (NFV) which could contribute to addressing this challenge.

↗ At Inria, the **STACK** project-team works on designing system mechanisms as well as software abstractions to manage and use next generations of Utility Computing infrastructures, combining Cloud, Fog, Edge, and beyond. Such techniques aim to efficiently manage the life cycle of applications running on the continuum Cloud-IoT, and take into consideration the costs such as energy consumption, applications delay requirement, bandwidth constraints, and security as crosscutting dimensions.

New wireless access architectures emerge which push towards more softwarization of network infrastructure (a prominent example is cellular 5G and beyond). The trend is swap expensive, dedicated (proprietary) access-point hardware, for “plain” antennas paired with backend software running on cheap generic servers in the cloud (or edge). Such software-defined networks (SDN) significantly increase the flexibility of the wireless access infrastructure – and may disrupt vendor/operator business models. Such architectures can accommodate not only more advanced “slicing” of network access resources, but can also provide new user-specific computing capacities, dynamically. New trade-offs emerge in terms of cost, latency, reliability etc.

A related aspect is IoT user mobility, which creates both challenges and opportunities. By exploiting spatio-temporal patterns of users' mobility, the system could provide a base for more precise and agile resource allocation strategies. Challenges in this domain include on one hand characterizing (and predicting) IoT user mobility, and on the other hand designing dynamic computation offload and placement schemes which can leverage such patterns.

↗ At Inria, the **TRIBE** project-team works on designing new computing offloading and placement strategies for IoT exploiting user mobility, and on evaluating the impact of task offloading to edge computing, on energy consumption and latency in mobile IoT contexts.

Maintainability of Distributed IoT Software

While IoT deployments multiply in a wide variety of verticals, too many IoT devices lack a built-in secure software update mechanism. Others may have built-in software update mechanisms, but these are not put to use, for non-technical reasons, e.g. lack of incentives for the vendor (or the vendor has gone bankrupt). However, without the availability (and the use) of such mechanisms, critical security vulnerabilities are not fixed, and IoT devices become a permanent liability, as demonstrated by recent large-scale attacks. In effect, a massive number of deployed IoT devices which have been operating autonomously for years (or decades) run unmaintained software. For instance, many Mirai IoT botnet elements (as well as other potential targets of this botnet) are still running unpatched to this day, despite the fact that Mirai was uncovered several years ago and patches have since been developed, and though these machines are relatively much less resource-constrained compared to low-power IoT devices!

On the legal side, in terms of security updates, the associated questions concern the duty of care: what does the *duty of care* mandate for to-be-deployed IoT? What does it mandate for already-deployed (legacy) IoT?

On the technical side, a substantial challenge is enabling and automating legitimate security software updates for IoT devices. On the one hand, enforcing legitimacy for software updates can also lead to so-called [treacherous computing](#) which can paradoxically prevent necessary software updates. On the other hand, on low-power IoT devices, the challenge is exacerbated by the stringent resource constraints in terms of network throughput, energy and memory budget in the first place.

Prior research in this field has mostly focused on simple use-cases where updates target single-binary, single-stakeholder software. However, as IoT software evolves and complexifies, this simplification no longer holds: more and more, IoT software mimicks Internet software in that it becomes a patchwork of components developed, maintained and updated by different stakeholders. In an average company today, less than 5% of the software used is home-grown, while more than 95% is from 3rd parties and/or open-source.

An open challenge is thus how to efficiently secure multi-stakeholder software on low-end IoT devices, whereby stake-holders have limited mutual trust. There are several aspects that comes into play, which must combine the output of research domains including:

- novel deep embedded system mechanisms to host and sandbox different software components,
- enabling and securing the supply-chain of IoT software,
- efficient low-power IoT network mechanisms to transport modular updates.

In practice, maintainability also entails capacity to monitor and manage IoT devices at run time, remotely, over the network. An active area of research is thus the development of low-power protocols and management data models adequate for IoT device connected over low-power radio. An associated area of research is the efficient instrumentation of IoT devices with debug/monitoring code snippets at low-level, inserted and removed on-demand at run-time, remotely, over the low-power network. Further research is also to be carried out in the domain of adaptive software, to transition from merely adaptive software to self-adaptive software. Indeed, advanced automation could enable software to be self-healing.

➤ At Inria, the **SPIRALS** project-team works on self-adaptive systems, aiming to introduce more automation in the adaptation mechanisms of software systems, targeting primarily self-healing and self-optimization properties.

Last but not least, the maintainability of IoT devices heavily depends on the co-evolution of low-level hardware and software. For instance, from a maintenance perspective, a software implementation of a cryptographic function is preferable (easily modified down the line to fix bugs, or implement new regulation). Therefore, research challenges also associated with maintainability are next-level performance for software implementations of critical low-level embedded system functionalities, and the design corresponding more future-proof (generic) hardware accelerators.

A NOTE ON BUSINESS MODELS.

Most business models so far focus on making money by deploying and/or exploiting IoT. Less is known about making profits by maintaining IoT – which is nevertheless crucial.

After IoT products have been rolled out and put to use, it is necessary to enable less [feudal](#) relationships between IoT end-users and providers. In particular, software updates for IoT products must become easier in general, including in cases where the original manufacturer does not deliver them (because it was acquired, went bankrupt, or attempted a [forced phase out](#) of the product, for example). Beyond technical barriers, legal barriers (such as breach of contract/guarantee) often make software updates difficult, if not impossible. This is mimicking a “walled garden” model, which allow users to add only authorized hardware components, or to only buy repair services from authorized dealers.

For instance, IoT vendors required to certify their IoT products for safety infrequently update software on these products. Users, on the other hand, may wish to update the software more often, e.g. to gain functionalities. [Reports](#) have shown how such tensions have already produced unfortunate situations where users resort to pirated IoT software (!) which only complicates matters.

2.4 Cryptology for Low-end IoT

Cryptography provides the fundamental protocols and basic algorithms (“primitives”) for authentication, identification, and encryption on which all secure systems are built.

Cryptographers have decades of experience in the design and analysis of efficient cryptosystems and protocols for relatively powerful devices (such as PCs, servers or smartphones) on one hand, and for more constrained devices such as smart cards on the other. The rise of IoT, with ubiquitous interconnected low-power devices, brings a fascinating new challenge for cryptographers, as it mixes the application requirements of the PC paradigm with the hard physical constraints of low-end devices. Put simply, we know how to provide some security for microcontrollers on smart cards, but smart cards were never meant to be connected to the Internet; and we know how to provide Internet security for powerful processors, but not on a stringent low-energy budget. The challenge for cryptographers is to develop full-strength primitives operating within the special constraints and requirements of the IoT paradigm.

Cryptographic Primitives for Secure IoT Communications

High-performance, high-security cryptographic primitives are now standardized, and widely deployed in protocol suites such as TLS (Transport Layer Security) for secure Internet communication. However, these algorithms have traditionally been developed and optimized for higher-powered platforms: from servers and PCs down to smartphones. When we move to more limited low-end IoT devices, the resource constraints are tightened to the point where conventional primitives are often too costly for the device in question. A critical challenge is thus the development, optimisation and adoption of alternative cryptographic primitives providing adequate building blocks for secure, low-power IoT communications.

Symmetric & Asymmetric Cryptographic Primitives

Cryptographic primitives are divided into two fundamental classes, symmetric and asymmetric, according to the function and application of the primitive. Data encryption and data authentication, for example, are symmetric primitives;

key exchange and signatures are asymmetric primitives. In general, symmetric primitives have much higher throughput and lower resource consumption. On the other hand, asymmetric primitives offer essential functionalities (such as digital signatures) that symmetric cryptography is literally incapable of providing. But these asymmetric primitives come at the inevitable cost of comparatively bigger keys, larger internal states, and more intensive computations, with the time, memory, and battery requirements that they entail. Optimized symmetric and asymmetric cryptography are both needed for secure IoT communication in practice.

Pre-quantum & Post-quantum Cryptographic Primitives

With the rise of [quantum computing](#), there is a second important distinction to be made, between pre and post-quantum primitives. This distinction represents a change of attack model: if a cryptosystem is designed to resist attacks by adversaries equipped with quantum as well as conventional computers, then it is labelled post-quantum. The construction of quantum computers with sufficient power to attack modern cryptosystems is a major challenge for physicists and engineers. While we can only speculate as to when, or whether, they will succeed, we must prepare IoT for a quantum future regardless—after all, we cannot predicate future security on the inadequacy of science.

Symmetric Cryptography Optimized for IoT

Symmetric primitives require a secret key shared between the communicating parties. Important examples include algorithms for data encryption, such as ChaCha20 and AES (the NIST, and de facto international standard), and message authentication, such as HMAC and Poly1305. Hash functions (such as SHA-3), while typically keyless, are also included in the symmetric family.

Looking towards low-end IoT devices, we enter the world of lightweight cryptography. Lightweight cryptography aims to provide efficient symmetric primitives with extremely small resource footprints, though often at substantially lower security levels. Lightweight cryptography is interesting for IoT for two reasons: first because it enables cryptographic operations in extremely limited devices, and second because it is required for the low-end devices that will communicate with lightweight ones. NIST (the influential US standardisation body) is currently running a [competitive standardization process](#) for candidate lightweight primitives – the result of this competition will have an important impact on symmetric cryptography in the IoT space.

When designing cryptography for low-end IoT devices, a challenge is to steer away from the “double penalty” experienced on microcontrollers compared to the PC/smartphone space: not only are CPUs weaker and slower (performance penalty 1), but also hardware accelerators may be absent, forcing a software-only approaches (performance penalty 2). For example, some microcontrollers lack the hardware support for AES (the NIST standard) that is taken for granted in the PC world. Instead of implementing AES in software to mimick the PC world, alternative symmetric crypto primitives should be designed and used on low-end IoT devices. Experience has shown for instance that switching from AES to a software-focused encryption scheme such as ChaCha20 can yield a 30% performance improvement for some applications.

➤ At Inria, the **COSMIQ** project-team works on the design and analysis of lightweight symmetric primitives. **COSMIQ** researchers are involved in submissions to the NIST Lightweight Cryptography standardization process.



Scuba: a tool chain for the security of connected objects. © Inria / Photo D. Betzinger.

Symmetric Cryptography for Post-quantum IoT

Post-quantum symmetric cryptography has mainly been studied in response to threats such as Grover’s algorithm, which (very) roughly speaking allows us to search keyspaces in a time proportional to the square root of the number of keys (while conventional computers require time linearly proportional to number of keys). Conventional wisdom suggests that for many elementary symmetric crypto primitives, ensuring post-quantum security is a matter of doubling keylengths. The question is, of course, much more subtle, even more so when more complicated symmetric systems and operations are considered. Determining the true post-quantum security of existing symmetric primitives is the subject of active research. But even if the simple fix of doubling key lengths suffices, this will have an important impact on IoT security: besides doubling the space required for keys, the throughput and resource consumption of the algorithms will degrade by variable factors. For example, passing from the crypto primitive SHAKE128 to the corresponding primitive with doubled key length (SHAKE256) will have no impact on memory requirements, but could impose at least a 20% decrease on throughput.

➤ At Inria, the **COSMIQ** project-team works on the security of post-quantum symmetric cryptosystems.

Asymmetric Cryptography Optimized for IoT

In contrast to symmetric primitives, asymmetric primitives depend on each party maintaining a private secret that is never disclosed, and publishing a matching “public key” to other parties. For example: Alice signs a message with her private key; later, having received the message, Bob can verify Alice’s signature using her public key. This asymmetry, reflected in the distinction between private and public keys, allows many new primitives that cannot exist in the symmetric paradigm, including not only signatures but also Diffie-Hellman key exchange, which is essential for establishing shared secret keys to allow symmetrically encrypted secure communication.

The public and private keys are tightly related: in essence, the public key presents an instance of a mathematical problem (such as an elliptic-curve discrete logarithm), and the private key represents the solution to that problem. The problem is chosen such that solving it is computationally infeasible—or at least: far

from worthy of the effort. In general, then, working with asymmetric cryptosystems means working with intensive computations in mathematical structures—and this comes at a heavy cost in terms of memory and energy, a cost that is often simply too heavy for low-end IoT applications. Improving the performance and applicability of asymmetric cryptosystems in the IoT space is an important area of research.

In the pre-quantum setting, public-key cryptography for IoT is dominated by Elliptic Curve Cryptography (ECC). ECC's unique selling point is its particularly small keys: 32 bytes is enough to store a high-security ECC key, and a high-security ECC signature fits in 64 bytes. However, using ECC means carrying out a large number of computations modulo 32-byte (not 32-bit!) integers. In the realm of PCs this is no problem, and ECC is now ubiquitous for key exchange and signatures on the internet. In the realm of low-power IoT, however, these number-theoretic computations represent a nontrivial memory footprint, a serious drain on energy reserves, and a painful runtime with consequent latency issues. One active area of contemporary research is thus adapting ECC protocols, and developing new ECC algorithms, to do more with less: maintaining high security for low-end IoT devices, while substantially reducing the run-time costs.

↗ At Inria, the **GRACE** project-team works on side-channel security of microcontrollers, and on efficient public-key (asymmetric) cryptographic primitives, including primitives targeting IoT devices such as the qDSA signature scheme.

Asymmetric Cryptography for Post-quantum IoT

Looking towards the post-quantum future, asymmetric cryptography faces a major problem: the existence of sufficiently large quantum computers running Shor's algorithm would destroy the security of virtually all widely-deployed asymmetric primitives. Current research aims to develop and study further asymmetric cryptography systems that are quantum-safe, based on a wide variety of approaches, including for instance lattice-based, hash-based, code-based and isogeny-based systems.

NIST has initiated a multi-year international process to select candidate algorithms for post-quantum signatures, and key encapsulation (essentially replacing pre-quantum Diffie–Hellman). Several finalist candidate algorithms under consideration already push down the required size of public keys, signatures and

computational cost. However, these algorithms have not been designed for IoT applications, and their potency in this domain is yet to be explored. Developing efficient, practical, and proven post-quantum key establishment and signature schemes targeting the IoT space is a fascinating and very exciting source of problems for researchers in cryptography.

↗ At Inria, the **ARIC**, **COSMIQ**, and **GRACE** project-teams work on the design, analysis, and efficient implementation of post-quantum asymmetric cryptosystems. Inria researchers are also involved in submissions to the NIST post-quantum standardization process.

2.5 Data Processing & Privacy with IoT

IoT enables pervasive data collection, tracking various physical systems, capturing environmental observations, industrial processes or other human activities, often in real time. Thus, while gathering and exploiting such data does fuel advances in many domains (e.g. health, and sustainable development, to name a few) preserving privacy becomes a major issue.

This issue yields challenges both at the political level and at the scientific and technical levels.

At the political and legal levels, new regulations can partly solve privacy issues. Notable examples are frameworks such as the EU's General Data Protection Regulation (GDPR), WP29 recommendations, and similar guidelines for ethics and trustworthiness.

At the scientific and technical levels, specific challenges become prominent to design IoT data exploitation techniques preserving privacy inherently.

↗ At Inria, the **PRIVATICS** project-team works on tracking and characterizing the exposure of personal data in IoT use cases, and designs mechanisms enhancing transparency for IoT users and enabling adequate IoT user consent.

Privacy-preserving Paradigms

Gathering data and data science are at the heart of modern information systems. IoT devices play a key role in collecting such data.

However, basic centralized data gathering leads to a privacy dead-end – as well as excessive network load when too much data must be transferred. New paradigms are thus needed, and current research explores alternatives.

For instance, instead of sharing its raw data, an entity can share pre-processed data. This approach can work in various organisational models, for example client-server or fully decentralized (peer-to-peer). The motivation for providers of

data is to fully control the distribution of information derived from their data. This requires some computation, storage or communication costs to be shared among more diverse machines in different parts of the network, potentially at its edge.

In order to enable pre-processing, one must first be able to *program* the IoT devices. There is thus the preliminary need for *adequate embedded software platforms* offering an adequate base, both in terms of openness and performance. Data usage sharing approaches can then build upon techniques such as:

- **transforming the data** (e.g. adding some noise) to gain some guarantees on anonymity (differential privacy, K anonymity, L-sensitivity...),
- **specific cryptographic primitives** such as homomorphic encryption to protect data in multi-party computations etc.
- **lossy data compression to partly obfuscate information** e.g. communicate only data aggregates, (partial) prediction models or intermediate computations like gradients or statistics, or combinations thereof.
- **embedding data management techniques**, to externalize results selectively (e.g. issue an alert based on the occurrence of a conjunction of events, rather than all events collected) and to aggregate the information collected (e.g. issue a computed statistic, rather than extensive raw datasets).
- **trustworthy computing techniques** embedded in (and distributed on) secure IoT hardware devices, to attest that the expected processing code was used, along with the appropriate input data, to produce a given result.

The typical privacy challenge is to design mechanisms maximizing the utility of data while preserving privacy. An additional challenge here is to minimize the ratio cost vs privacy benefits for data providers – which may be very resource-constrained IoT devices.

↗ At Inria, the **PRIVATICS** project-team works on the design and implementation of algorithms for privacy-preserving information sharing, applicable to wearable IoT devices ranging from smartphones to smart tokens.

↗ At Inria, the **ARIC** project-team works on the design and analysis of full homomorphic encryption schemes, and on its use for privacy-preserving computations.

↗ At Inria, the **COMETE** project-team works on the design and analysis of differential privacy schemes used for privacy-preserving computations.



Smart container, a solution for tracking and monitoring containers. © Photo Raphaël de Bengy.

Decentralized Machine Learning with IoT

Sharing data usage instead of data itself is a principle that can also mitigate privacy issues for Machine Learning (ML) exploiting potentially privacy-sensitive IoT data. In this context, Federated Learning (FL) is an active field of research¹, whereby a number of clients collaborate via a central server to train a model while each keeping their own training data. The principle is again to minimize data collection, aiming to both eliminate privacy concerns and network bottlenecks (when too much data must be transferred²). One key challenge which arises here is the lack of central authority which controls and distributes data among peers. Because data stays where it was collected, ML algorithms at the heart of decision systems have to deal with non identically and distributed (iid) data. Indeed, the 'identically' assumption is the major assumption for ML to assert that learned model will behave as expected for future values and the 'independently' property strongly simplifies the complexity of the class of possible models.

1. Peter Kairouz et al. Advances and Open Problems in Federated Learning. Technical report, arXiv:1912.04977, <https://arxiv.org/abs/1912.04977>, 2019.

2. For instance, autonomous cars equipped with camera and sensors collect huge amount of data, and many locations experience intermittent network connection.

Further research explores fully decentralized federated learning, without a central server, in a fashion resembling peer-to-peer systems. An opportunity which arises in this context is to take advantage of the massive number of peers to improve privacy. One key challenge is however to learn with whom to collaborate and how to optimize communication costs. Here, typical P2P security and trust issues pop up, and must be revisited and mitigated in this context (e.g. detecting malicious peers, collusion between peers). Specific issues also arise depending on privacy-preserving techniques: for instance, if noise is added to partly obfuscate data that is shared, utility can dramatically drop when only little data is available. Last but not least, computation and state (model and training data) must be minimized in order to fit the tinier resources available on low-end IoT peers.

↗ At Inria, the **MAGNET** project-team works on designing methods for privacy-aware Machine Learning (ML) using data anonymization techniques to feed ML, and using fully decentralized peer-to-peer algorithms, relaxing the core assumptions of Federated Learning.

Trustworthy Decentralized Database Computation in the IoT

Database processing, and in particular Big Data, is of paramount importance in the context of IoT. Sharing the “use” of data rather than the data itself leads to a new paradigm, where data processing operations are pushed to the edge of the network and, at the extreme, within the IoT devices themselves. This approach fosters data confidentiality and privacy (notably by locally enforcing access control and confidentiality rules) as well as energy savings (by avoiding the transmission of seldom used data, and by producing aggregated results instead of all the raw data collected).

From a database point of view, this leads to consider large sets of IoT objects (endowed with storage and computing resources) as a distributed or federated database, on which global database processing can be launched. The research challenges posed by such a vision are the following:

- (1) making database techniques (storage and indexation, query evaluation algorithms) compatible with the severe hardware constraints of smart objects, and especially, scarce RAM vs comparatively large amounts of Flash memory;
- (2) designing new secure distributed query evaluation techniques on IoT devices, at a very large scale without resorting to a trusted central server.

The first challenge stems from the conflicting hardware constraints of the IoT devices with respect to data management techniques : scarce RAM calls for massively indexing the data (because intermediate results cannot be constructed in RAM at runtime), but the specificity of NAND Flash memory penalize small random writes (which are needed to maintain the indexes). Pioneer works are conducted at Inria to solve such constraints enabling storing and processing millions of stored database entries in a smart object thanks to new design principles. The next research challenge is to generalize such results to the case of data series, the main data type present in IoT devices.

Regarding the second challenge, a parallel can be drawn with private data federations, an active research area in the database field, where the goal for a set of data owners is to contribute with their own data for answering a global query, without disclosing their (potentially sensitive) data to one another. Several approach are currently investigated, resorting to various techniques like cryptography (secure multiparty computation), noise addition (differential privacy) or trusted computing (hardware-based security). The transposition of these techniques in the IoT context is difficult, as it requires considering huge sets of "data owners" (potentially millions of smart objects involved) with resource and energy constraints. Some preliminary studies rely on the secure hardware present in some IoT devices (e.g. secure chips or TPMs) to handle some Big Data computations (MapReduce) with confidentiality and integrity guarantees, but the issue remains today a broad research perspective.

More generally, solutions to these challenges allow to restore the agency of individuals on their personal data, and to help citizens to collectively contribute to database computation of any kind (SQL, Big Data, AI, etc.) in a trustworthy manner – without necessarily resorting to any trusted third-party.

➤ At Inria, the **PETRUS** project-team works on designing trustworthy database architecture for Personal Data Management Systems (PDMS), with security and privacy guarantees, enabling distributed queries involving very large sets of PDMSs. **PETRUS** has developed PlugDB, a PDMS hardware/software platform using trusted hardware and microcontroller technologies.

2.6 Safety, Reliability & Certification for IoT

Full-fledged IoT will increase our vital dependency on embedded systems, and networked sensors and actuators. In many applications, the term “vital” is here used literally – because with IoT, digital devices can trigger direct physical effects.

A trend of smart implants is emerging, involving consumer embedded hardware, a wireless network loop and open source code. For instance, an [Artificial Pancreas System](#) (APS) is designed to automatically adjust an insulin pump’s basal insulin delivery to keep blood glucose (BG) in a safe range overnight and between meals. And if the APS malfunctions, the life of the patient is immediately on the line.

More broadly, a series of recent IoT safety and reliability incidents involve a variety of machines, including vehicles such as [connected cars](#) and even [airplanes](#). Due to IoT shortcomings, the lives of users were either shown to be in grave danger, or lost in a real accident. Collectively, the string of recent incidents hint at the following:

- even extreme safety-oriented code can contain fatal and/or exploitable bugs;
- even highly sensitive, closed source code can get leaked.

As IoT use-cases extend beyond funny gadgets, safety aspects of IoT become ultimately more prominent. Issues arise because [IoT devices are likely to be used in unforeseen ways](#), or in potentially hostile contexts, prone to cyberattacks. Combining both the safety requirements of IoT and the security requirements of IoT is a daunting challenge. A speaking example is certification: often, an IoT product is a moving target (due to software updates) and the context in which it is used unrestricted (think: consumer electronics). Then *what exactly should be certified, and how?*

To address challenges in this domain, the safety and the security research communities—traditionally rather distinct— must collaborate more tightly and revisit fundamental concepts together, from the ground up.

On one hand, guaranteeing safety and security is no longer a “single-handed” task. It involves a complex set of stakeholders: software vendors, operators, regulators, individual users... In such contexts, legal frameworks must be capable

of determining who is responsible for what. On the other hand, increased interdependence between systems mandates reconsidering what we consider “critical”. For instance, studies have shown that a mid-sized botnet of [IoT enabled air conditioners and heaters can be weaponized to disrupt a national electricity powergrid](#). A challenge is thus the safety and security of mixed system involving not only traditional critical components, but also consumer off-the shelf low-cost IoT components. Certification needed for such components is challenging not only because it is difficult to formalize in this context (what is good-enough cybersafety?), but also because it needs to remain very cost-effective, as these are low-cost devices.



Mapping of the IoT room at the Lyon Lab. © Photo C. Morel.

2.7 Human-Machine Interaction with IoT

With IoT, computers “disappear” more and more, and increasingly complex machine-to-machine (M2M) interaction happens under-the-hood. Through sensors and actuators, and as interaction with the system takes new forms (e.g. gesture-based), physical reality itself may be customized and experienced differently. Postphenomenology studies the interplay between humans, the world and modern technologies – the latter being non-neutral mediators. Paradoxically, with M2M, the human factor becomes even more crucial. As humans are more systematically bypassed to achieve more benefits, humans may be more impacted by a system malfunction, or by a lack of understanding of how the system works. For example, workers increasingly risk being deskilled – or even replaced.

As IoT grows mainstream, IoT risks creating a new digital divide, where some users are severely disadvantaged by the technology, others gain more detailed control, and still others gain power through savvy control of APIs. In this context, proper design of novel human-machine interaction is critical.

Empowering humans with the adequate levels of control on IoT (what is the right level?)

Let's look at an example: thermostats. Thermostats were initially designed for controlling temperature, typically with minimalistic interfaces such as a rotary knob, and do nothing more. With IoT, such devices may have many more capabilities and functionalities: they can be programmed, used for controlling some other systems, linked to other devices or sensors, or to some online service, etc. However, although they are IoT-augmented, many of these devices retain the original minimalistic interface. This retro-fitting approach raises a number of issues. On the one hand, such interfaces are easy and familiar. On the other hand, basic substitute interfaces (e.g. through a smartphone screen) fall short of the full potential of cyberphysical interactions. While the technological aspects of embedding technology in the real world is well advanced, we have yet to achieve a seamless transition between the two.

Understanding what happens under the hood of IoT

Achieving an adequate level of transparency is linked to how users understand IoT. IoT infrastructures are inherently complex, made of interconnected heterogeneous nodes and interactive devices, with underlying “ambient intelligence”. This complexity raises challenges in terms of interaction, for the users to control such systems. The general trend over the last decades in interactive technologies has been to oversimplify interfaces to make interaction simpler. This has certainly helped to democratize the use of technology for newbies, but at the cost of decreased expressiveness. In the context of IoT, this approach is likely not going to scale due to the complexity of the infrastructures and the particular means of interaction they offer (e.g. multiple devices, reduced input/output, distant and distributed interaction).

One challenge is thus to better accommodate complexity by designing appropriate interfaces and interactions so that the users can progressively build an appropriate mental model of the system, by which we mean:

- understanding and anticipating how the system will react to their actions;
- having a clear and correct view of the system’s current states and errors; and
- progressively acquiring skills to finely control the system.

Another challenge is that of shared control. IoT systems often have a certain degree of autonomy and can take the initiative in order to perform tasks or propose actions to users. This aspect can become critical. Users need to understand the current state the system, since it may have changed autonomously; users need to know how they can take the control back when needed; and when the system does act on its own, users need to feel in control. In particular, users must be able to

- identify appropriate actions and communicate their intentions;
- trust the system in terms of its consistency given similar situations; and
- monitor the system such that errors can be identified and corrected.

Examples of shared control gone horribly wrong are the [Boeing 737MAX crashes](#) which were in part due to hiding an undocumented behavior “under the hood”. Research in the field of HCI has started to analyse [incidents of this kind](#), aiming to characterize users’ understanding and visibility of the states of the system.

One pitfall here is a situation where the system controls the user, more than the user controls the system. Nevertheless, an approach based on computational intelligence through tracking and task inference could be leveraged to anticipate user actions, thus removing the need for an explicit and always-enabled interface. One promise of ubiquitous technology, first explored with Welner's [Digital Desk](#) research, has been environments that anticipate the user. Essentially, through models of user and task, the environment can automatically reconfigure itself whilst letting the user maintain control. Ultimately, this approaches could also support adaptation of the interfaces to several kind of audiences, by providing users with appropriate levels of controls according to their needs, skills and contexts of use.

↗ At Inria, project-teams including **AVIZ**, **EXSITU**, **ILDA** and **LOKI** study and design new interaction methods and interactive systems that empower users by better accounting for their capabilities and expertise.

Delivering on the potential for physically-augmented interaction

Research on tactile and tangible interaction shows considerable potential to expand the (so far) reduced input capabilities that are typical of IoT devices. Indeed, research in HCI (Human Computer Interaction) constantly explores advanced sensing techniques that can detect and differentiate between the ways we touch a surface (e.g., fingers identification, precise contact point detection, applied pressure, finger inclination). Similarly, the way a physical object is grasped and manipulated can strongly inform on the user's intention. For instance, imagine a pencil case: one would grab it differently whether the intention is to open it, to store it, or to hand it over to someone else. This principle applies to other objects, and IoT devices are no exception. In conjunction with studies on users' ability to leverage tactile and tangible sensing technologies, research on physically-augmented approaches aim to improve the interaction bandwidth between users and IoT devices without any additional/external interface: a single button could for example trigger different actions according to the way it was touched or grasped.

However, these approaches only partially address the issues of limited visibility for the different actions possible with the system, and lack of feedforward (what to do) or feedback (what was done) from the system. On the one hand, making the interaction "physical" could help the user to transfer knowledge from other contexts. On the other hand, augmenting the interaction capabilities (e.g., by adding touch sensing to a physical button) could contribute to improving the

visibility and discoverability of actions and functionalities. This is where Mixed-Reality (MR) offers a promising way to making these “invisible interfaces” reappear, e.g. with smartphones or wearables devices such as glasses, superimpose subtle clues about interaction over the physical devices, or even by overlay complete tutorials, upon user demand.

MR could also help provide users with better feedback from the system, which could also be complemented with haptic solutions such as vibro-tactile feedback. A very active avenue of research in HCI related to interaction with IoT devices focuses on finding new ways to enable such feedback on any surface (e.g. actuators, electrovibration), and on understanding how users perceive them and what kind and quantity of information they can convey. Another area is detection of human movement to support a wide range of rehabilitation and creative applications.

➤ At Inria, project-teams including **AVIZ**, **EXSITU**, **ILDA** and **LOKI** explore novel interactive materials and devices to create novel forms of tangible interfaces for a wide variety of home, work and creative applications.

2.8 Control with IoT in the Loop

A primary purpose of IoT is to enable advanced monitoring and control for distributed systems deployed in a variety of environments (from smart home & buildings, to smart cities, and industry 4.0). In particular, IoT has the potential to bring advanced monitoring and performance through optimized control to small-scale systems that cannot afford dedicated control systems.

The use of multi-purpose shared networks to control spatially distributed elements results in very flexible architectures. The drawback is that asynchronous dynamics are added in the loops, which can strongly degrade the performance and, even, the stability. Estimating the effects of the network and designing robust networked control systems (NCS) is a motivating challenge, involving hybrid systems mixing continuous states and discrete events, delay systems, etc. Algorithms must also manage intricate hierarchies of subsystems involving those asynchronous, multi-scale dynamics with extremely varied time scales (from months to microseconds) and extremely varied geographic scope (from on-chip to planetary scale). Designing closed-loop control over such nondeterministic networks imposes ultra-resilience in face of (inevitable) variations in terms of latency, jitter, throughput.

↗ At Inria, the **VALSE** team works on modeling and analyzing highly distributed, uncertain dynamical systems found in IoT and cyberphysical systems. **VALSE** designs robust estimation and decentralized control algorithms using the concepts of finite-time/fixed-time/hyperexponential convergence and stability.

A related **challenge in this field is autonomic feedback loop management in IoT middleware**. Typically, IoT control and monitoring entails the use of middleware for supervision and management of the infrastructure. IoT middleware aims to enable centralized or distributed management of complex, distributed logic components, on very heterogeneous infrastructures (for example small devices with limited computation power, home gateways, local nodes of a cellular network, or data-center in the Cloud). IoT middleware must thus provide usable abstractions for the high variability in operating systems and communication protocols.

A crucial issue in this domain is automating cyber-physical feedback loops, e.g. with Autonomous Computing. Such control loops aim to enable continuous

self-adaptation of the cyber-physical system, by reacting to monitored information with decisions, taken on the basis of a representation of the system, and implemented through actions, in order to enforce a high-level strategy or policy. This is done in face of (potentially high) dynamics which happen either in the physical environment being monitored and controlled (which is the classical object of Control Theory), or in the computation and communication system infrastructure itself (e.g. varying load, fault-tolerance, self-protection).

↗ At Inria, the **ACENTAURI** project-team works on new paradigms increasing the autonomy of robotic systems, enabling task oriented behavior, and exploiting multi-sensory perception and control.

Challenges include designing and optimizing automatic reconfiguration schemes and software architectures for application-level functional aspects, as well as for infrastructure-level computational aspects (e.g. service migration, self-scaling), with requirements for separation of concerns between these different levels.

↗ At Inria, the **CTRL-A** project-team works on designing methods for Autonomic Computing controllers, leveraging Control Theory to enable application-aware management of reconfigurable computing architectures, in the IoT as well as in HPC.

Bridging the Gap between Robotics and Industrial IoT

Micro-robots are emerging as low-power, low-cost, and tiny tools whose size enables new applications in robotics. Coordinated swarms of these tiny robots are of particular interest. Swarms have the potential to outperform monolithic robots in applications where spatial diversity has advantages, such as distributed sensing. Swarm robotics can be seen as a next frontier for industrial IoT research as it requires addressing many of the open research questions, simultaneously, to enable control and interaction with large numbers of micro-robots.

A first challenge is mobility. While low-power IoT protocols have now been largely standardized and are being rolled out, they were mostly designed for interconnecting devices that are statically deployed in an area. Having some, or all of these devices move is not well supported in today's standard industrial IoT protocols (such as 6TiSCH).



FIT (Future Internet of Things) experimental platform. © Inria / Photo C. Morel.

A second challenge is low, and predictable latency. Industrial IoT networks can currently ensure generated data is delivered, for example to a gateway. But while latency can never be guaranteed (wireless being unreliable) scheduling approached in the network allows the latency to be predictable. This determinism opens up the possibility of running control loops through IoT networks, as highlighted previously.

A third challenge is accurate, parsimonious localization. Foundational localization techniques such as a UWB or BLE angle-of-arrival have been developed. An open research challenge is to co-design the localization solution with the communication protocols, resulting in on-demand localization which is compatible with the low-power requirements of most IoT applications.

↗ At Inria, the **EVA** project-team works on experimental swarm robotics. For example, **EVA** has developed the Atlas swarm robotic simulator, and is building DotBot, a large testbed for centimeter-scale robot swarms of up to 1,000 units. The **AVIZ** project-team works on the design of swarms of tiny robots able to perform physical visualizations and a variety of other tasks. For example, **AVIZ** has designed Zooids.

2.9 Security for IoT

Cyber attacks involving entities across national borders are now the new normal. Profit-driven and state-driven online piracy is happening at unprecedented levels: World War 3 is online. There is also a current trend towards most crimes involving some cyberphysical components.

In this context, security challenges arise and cross-cut all aspects of IoT. As security and resilience are only as strong as the weakest links, securing low-power IoT becomes all the more essential.

Beyond the most basic cybersecurity attacks (phishing, social engineering etc.) a growing variety of attacks, require mitigation and new security mechanisms at all levels of the system.

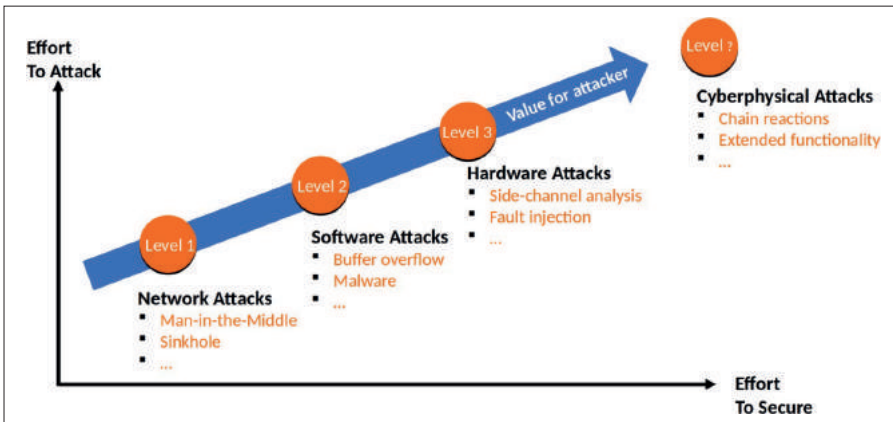


Figure 1: IoT Attack Surface.

Defining attacker model(s) with IoT

Traditional cyberattack modes of operation remain effective with IoT: sequence of tentatives to exploit different vulnerabilities, incremental priviledge escalation etc. However, *risk assessment changes significantly with IoT*. If more and more IoT-controlled actuators surround you, physically influencing your environment (or even your biological state, e.g., a smart implant) the risk levels you tolerate is significantly lower. If IoT-enhanced sensors gather more and more intimate and fine-grained data (think: your heart and perspiration-rate, in real-time) the impact of privacy breaches is significantly larger.

As per traditional IT security alone, [a huge attack surface needs to be addressed](#), spanning from network attacks (man-in-the-middle, sinkhole...), to software attacks (malware, buffer overflow) and hardware attacks (fault injection, side-channel...), not to mention the human vector and social engineering attacks. But *with IoT, new attack vectors appear*.

It is now possible to remotely trigger catastrophic [cyberphysical chain reactions](#), Domino effects massively exploiting botnets and the increased inter-dependence between systems which were formerly isolated from one another, such as the electricity power grid and the Internet. On the other hand, [extended functionality attacks](#) can weaponize an IoT controlled device by transforming its use, in a completely unexpected way. Such attacks further expand the traditional attack surface of networked systems.

Additionally, as novel user cyber-physical interfaces emerge, they also offer new attack vectors. For instance, the emergence of voice commands enable [new attacks on Voice Personal Assistants](#) whereby authentication of the voice assistance by the user (or vice versa) is difficult and can easily be abused. In the near future, it is expected that novel cyber-physical user interfaces will include advanced smart implants similar to the [brain cyber-physical interface prototype](#) Neuralink, which will drive upwards the stakes in terms of security and safety requirements.

A crucial challenge is thus defining novel attacker models which capture this context.

Securing IoT network protocols

The benefits brought by IoT rely on looping in new devices via the network, that were previously either absent or operating stand-alone, isolated. On the flip-side, these benefits come at the price of opening new avenues for cyber-attacks via the network. Securing IoT network communications is thus crucial.

The network communication stack is traditionally divided into abstract layers. Each layer provides services to the layer above, and uses the services of the layer directly below. The dominant model is that of today's Internet, consisting of the application, transport, network, link and physical layers. Specific security mechanisms are needed at each layer of the network protocol stack.

To comprehend challenges in this context, one must first grasp some peculiarities of IoT devices and networks, compared to average machines connected at the edge of the Internet.

TINY DATA TRANSFER RATES

IoT local networking (or IoT's last-mile) often relies on low-power radios, which exhibit unusual constraints in terms of physical data rates: from 250 kilobit/s announced for short-range technologies (providing connectivity within 10s of meters indoors, 100s of meters outdoors), down to approximately 10 kilobit/s for long-range technologies (providing up to 10 km outdoors connectivity range). Roughly, this is somewhere between 0,01% and 0,1% of announced data rates for modern WiFi, or 4G cellular.

SPECIFIC DATA TRAFFIC PATTERNS

Data originated at IoT devices is often stored at intermediaries before reaching its final consumer. It is therefore no longer sufficient to put trust in the data based on the identity of the other communicating peer, as often done in the traditional Internet. Instead, a producer-consumer model is needed for security purposes, to provide security guarantees at the application layer.

MICROSCOPIC BUDGETS FOR ON-BOARD RESOURCES

IoT networks devices which a very small budget in terms of power, processing or memory. For example, in terms of memory, this is 0,001% of the resource budget available on a traditional machine connected to the Internet.

DIFFERENT HUMAN-IN-THE-LOOP FACTOR

Low-power IoT devices often lack common user interfaces such as a display or a keyboard. With push buttons and LEDs being the only means to interact with a device, the configuration step in the field (for commissioning or debugging purposes) becomes significantly more difficult. Furthermore, IoT devices tend to have an inherently 1:N ratio with humans (think: all your sensors/actuators/implants & smart gadgets), whereas other machines connected at the edge of the network tend more to 1:1 (think: your smartphone, your laptop).

Bootstrapping security without a user interface

A common assumption for the communication security solutions defined by the standardization bodies is that the trust relationship between the different entities involved in the communication has already been established through common keying material. At manufacturing time, the trust relationship is typically established between the IoT device and the manufacturer. However, the domain where the IoT device will be installed is not known at the manufacturing time. Before the IoT device can join a given domain, it needs to be provisioned with domain-specific credentials. Bootstrapping this trust relationship between the IoT device and the domain owner is typically considered out of scope for the

standards bodies, yet it is a non-trivial task as most IoT devices lack common user interfaces (display, keyboard...). Prompting a low-end IoT device for a password is simply not an option and the use of automated authentication mechanisms becomes favored. Companies typically resort to out-of-band channels (e.g Near Field Communication, ad-hoc wireless network, pre-shared keys printed on the back of a device, serial port). First, this approach opens up various vulnerabilities as the “bootstrapping” protocol ends up being designed in-house, without a thorough review of the community and security experts. Second, this approach does not scale (think: bootstrapping secure communication for dozens of sensors at once...). One challenge is thus the definition of appropriate bootstrapping protocols taking into account on the one hand the device and network constraints, and on the other hand operational constraints and device lifecycle of IoT products.

➤ At Inria, the **EVA** project-team works on the design of zero-touch security protocols and performance evaluations of communication security standard candidates in IoT use cases.

Securing IoT's data-oriented paradigm

Internet communication have been designed to be endpoint-oriented: inter-connecting machines that are simultaneously responsive, and establishing the security of relatively long-lasting communication channels carrying streams of data between such communication endpoints. The traditional approach (both in the standardization and the research communities) has been to reduce communication overhead through a more efficient encoding, but to not compromise on the security level. For instance, researchers proposed lightweight versions of IPsec and (D)TLS protocols, reducing the communication overhead through compression of the protocol fields that are not critical for security.

In large parts, however, IoT communication incurs a different, data-oriented traffic pattern: one-shot communication (think: a periodic sensor measurement, or a firmware update) involving machines that can be in power-saving (sleep) mode most of the time, such that, somewhere along its transit over the network, IoT data will be stored and will rest temporarily in some repository.

The requirements imposed by this data-oriented paradigm have only recently been addressed through an effort on defining new mechanisms based on the “object security” primitives, which apply the protection mechanisms at the application layer. New lightweight protocols and mechanisms must be defined and standardized. One example is the research in the field of information-centric

network (ICN) protocol design and its security extensions. Another example is the on-going activity around the standardization of OSCORE for object security, i.e. the protection of web transfer using CoAP, and the EDHOC protocol for key exchange at the application layer. These categories of solutions promise appealingly low communication overhead, and native support for the security of IoT data while at rest somewhere along its transit over the network – a major improvement over traditional solutions.

Nonetheless, providing a fair, common ground for protocol comparisons is often a tricky task. An on-going challenge in the academic community remains to establish and conduct relevant comparisons amongst new IoT protocols and traditional protocols.

Security of unseasoned IoT protocol specifications

Many IoT solutions (the specifications or their implementations) have only been appeared recently. The relative novelty of these protocols means that compared to seasoned protocols (such as TLS for example), they have been subject to less analysis, not just on the level of security proofs against protocol specifications vulnerabilities, but also in terms of efficient algorithmic implementation and optimisations.

A major challenge in the academic community is therefore the formal analysis of these novel, lightweight solutions defined by the standardization bodies such as the IETF, formally proving the advertised security guarantees in conjunction with the targeted performance in terms of energy efficiency.

↗ At Inria, the **PROSECCO** project-team works on the design of formal verification methods applicable to low-power IoT protocol specification.

In the end, it is the role of the network administrator to decide which protocol combination to use, which implementations to run, and how to configure those, to mitigate the threats of a given deployment. The skills set and background knowledge of network administrators vary, especially with newer protocols. Experience shows that default parameters are often left untouched – leading to well-known vulnerabilities in the system. Usable IoT security thus mandates not only solid specifications, but also the possibility for (and the availability of) implementations with adequate default security guarantees on *all* devices. A challenge is thus to drive the co-evolution of IoT protocol specifications and their implementations, in order to obtain usable security.

IoT Security Incident Detection & Handling

An essential aspect of securing a distributed system is the capabilities related to network and service management. Global system audits can uncover potential vulnerabilities before they are actually exploited. Distributed system monitoring can detect security incidents when they happen. When the system includes cyberphysical IoT components, the complexity and the heterogeneity of the system explodes, from which stem specific challenges.

One challenge is collecting relevant information (through passive or active scanning) which effectively tracks IoT device characteristics and activity in vivo, which is especially difficult to achieve within low-power resource budgets. Here, new protocols need to be designed and/or instrumented, to provide adequate vantage points in IoT deployments.

Another challenge is to automate cross-checking of the specific information collected in a specific deployment, with globally maintained security information bases (CPE, CVE, CAPEC, CWE, suspicious information flows...). Here, a promising approach under investigation exploits machine learning techniques to automate and optimize the performance of complex IoT security audits, and improve the speed and accuracy of IoT security incident detection. IoT security incident handling also yield challenges. One example is the design of novel attack confinement strategies and mechanisms mimicking fault-isolation in safety domain.

➤ At Inria, the **RESIST** project-team works on the design of network management platforms for IoT, and novel techniques facilitating auditing and monitoring functions, in order to automate the security assessment in IoT environments. For instance, SCUBA is a platform developed to accelerate security audits on heterogeneous connected objects.

Securing IoT Software

Implementations targeting IoT use cases are flourishing. Until recently, software on low-end IoT devices has been proprietary, closed-source and sometimes, worse: it would rely on security-by-obscurity, weak by design. There is an on-going trend towards more open source implementations for low-power IoT, a side-effect of which is, that security by obscurity is not an option anymore. With this evolution, we can expect that security will thus necessarily improve.

Since they are recent, however, critical parts of these implementations have not been formally verified for software level vulnerabilities. A challenge lying ahead for the formal verification community is to get involved in the study of such IoT implementations. Writing secure software is hard, and there is no one-size-fits-all verification approach. The main challenge is to produce verified IoT software for low-end IoT devices, without incurring significant performance penalties, and without sacrificing versatility (low-level IoT software tends to target a wide variety of hardware and use-cases).

↗ At Inria, the project-teams **TEA** and **PROSECCO** work on automating the proofs of software building blocks embeddable in low power IoT devices. A particular focus is put on proving key security components, such as cryptography primitives, and on verifying the functional correctness and memory safety of minimal bootloaders. By developing new workflows designed around the formal language Fstar, producing verified and efficient embedded software modules targeting low-power hardware becomes practical. A prominent example of such an IoT software module is the crypto library HACL.

A priori, it is impractical (both technically and economically) to formally verify all the software that is shipped and deployed. Furthermore, even if some piece of software is formally verified before deployment in the field, IoT software can still have bugs and vulnerabilities which can be exploited³. The reason for this is that code is proven against a security model (assumptions on the attacker etc.). Prior verification offers no guarantee if the model does not hold in practice – because an IoT device is used in an unexpected way, or in an unexpected context... and chances are: it will.

Therefore, it is necessary to complement a priori formal verification with means to periodically update software on IoT devices, to fix bugs and vulnerabilities uncovered a posteriori, after the software has been deployed. Although it is a security feature, software updating is also an attack vector. For instance, a software update might lace legitimate software with malware. Conversely, a functional and necessary software update could be blocked because no authorized party provides a digital signature. A crucial challenge for IoT is thus the design an appropriate and secure supply chain for IoT software, which should remain in operation throughout the life-time of low-power devices.

3. Donald Knuth 1977: "Beware of bugs in the above code; I have only proved it correct, not tried it." <http://www-cs-faculty.stanford.edu/~knuth/faq.html>

Associated challenges combine research on low-power cryptography, reproducible software, remote attestation, and deeply embedded system software design. Beyond academic research, new standards are also needed and expected in this domain, as demonstrated by on-going work on the [SUIT specifications](#) for instance.

➤ At Inria the **TRIBES** project-team works on the design of a secure IoT firmware update supply-chain, tailored for cheap, low-power IoT devices but without compromises on security.

➤ RIOT-fp is a cyber-security project initiated by Inria, which targets resource-constrained, microcontroller-based IoT devices. RIOT-fp contributes practical building blocks for an open source IoT solution improving both software durability and the functionality versus risk tradeoff, for end-users. These building blocks combine high-speed, high-security, low memory IoT cryptographic primitives, frameworks offering guarantees for software execution on low-end IoT device, and secure supply-chain for IoT software updates over low-power networks.

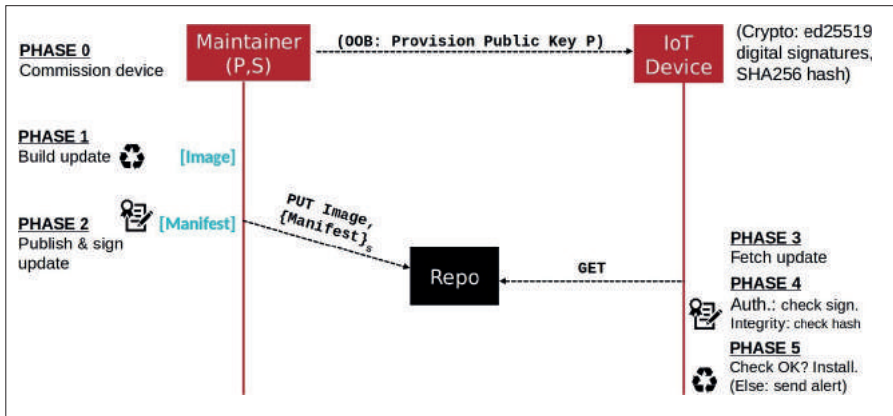


Figure 2: Secure IoT software update workflow (work-in-progress SUIT specification).

Securing IoT Hardware

On the one hand the inherent flexibility of IoT applications operating in a multi-standard environment mandates the interoperability, ease of use and product update facilities typically provided only by high-level software development. On the other hand, hardware (HW) acceleration can increase energy efficiency by several orders of magnitude, while pure software (SW) approaches are often incompatible with resource constraints on-board IoT devices.

Clever hybrid design – mixed HW/SW architecture – is therefore a promising avenue which must be further explored. Examples of critical functionalities for which hybrid design is necessary include (but are not limited to) IoT cryptography.

On the flip side, hardware also offers an attack surface which must be mitigated by specific mechanisms. In particular, with low-power IoT, the physical access and security of devices needs to be re-assessed. Compared to physically capturing your smartphone or laptop, it may be easier for attackers to capture some IoT device (think: one of the dozens of sensors/actuators scattered in the vicinity) and to submit this hardware to elaborate side-channel attacks. In this context, challenges include:

- novel HW accelerators for security functions (symmetric and asymmetric cryptography, hashing, authentication, signature, random number generations, etc.) with a specific focus on energy efficiency and ultra-low-power,
- specialized crypto-processor including protection against attacks such as randomization,
- optimizing compilers targeting resource-constrained crypto-processors,
- hardware accelerated dynamic binary translation (DBT) as a mean to enhance protection of software, and
- new techniques for efficient hardware protections against side-channel attacks and fault injection, both in SW and HW.

↗ At Inria, the **CAIRN** project-team works on hardware acceleration for low-power cryptography primitives, as well as energy efficient crypto-processor architectures with hardware countermeasures. The **PACAP** team explores security mechanisms inserted by the compiler to improve programmer productivity, and application robustness against side-channel attacks.

2.10 Low-Power Hardware Architecture, Programming and Compiling

As already mentioned in this book, one of the most prominent transversal challenges is energy-efficiency – and more generally, resource-efficiency. Many IoT devices are required to operate for years on a small battery which is neither envisioned to be changed nor recharged. The challenge is on one hand to better provision IoT devices in energy and on the other hand to reduce their energy consumption both from a hardware and a software perspective. Also, from a more global perspective, IoT devices are expected to be billions. Even a small decrease in the individual energy consumption of billions of devices leads to saving significant amounts of energy, to substantial cost reductions and to reduced environmental impact. An array of complementary research directions must thus be investigated. We detail some of them next.

Pushing ultra-low power towards net-zero energy

The use of wires or batteries to power embedded systems is inconvenient because of form factor, cost or maintenance considerations. Research efforts are to be continued in designing self-powered communicating devices, powered neither from battery nor via wires: what we here call *net-zero* IoT devices.

Passive RFID is an example in which the tags are battery-less. They are only composed of a chip and an antenna and are powered upon data reading by a reader. This technology is thus very scalable from the energy perspective since a single reader can be used for an infinite number of tags. Nevertheless, the applications enabled by passive RFID are limited since tags cannot communicate when there is no reader around – and the reader generally consumes more energy than a simple (active) IoT device.

Some research focuses on developing novel battery-less IoT hardware harvesting ambient energy, e.g. light, heat, vibrations/movements or radio waves. However, such process provides very low current levels. Therefore, net-zero IoT devices must be designed to consume as little energy as possible.

For example, techniques of power gating allow to shut off the current to blocks of the circuit that are not in use. Also, the use of non-volatile memory (NVM such as NVRAM) makes it theoretically possible for a device to suffer power shortages without losing data, allowing it to continue its task rather than restarting them. Current NVM technologies still suffer from slow write times, high write energy and limited write endurance. Research challenges in this domain thus include thinking up new hardware which better captures ambient energy and offers better levels of operation with the resulting ultra-low current.

A prominent related challenge is the design of efficient and robust software and network protocols running on this type of hardware. On the toolchain side, one challenge is to design compilers which can better assist programmers, through program analysis, in the context of intermittent-powered embedded device programming. Such program analysis is required to identify efficient check-pointing strategies: what program state to store and when.

↗ At Inria, the **PACAP** project-team works on compilers and program analysis designed to facilitate check-pointing on intermittently powered systems.

Embedded software architecture aspects offer additional challenges.

For instance, naively replacing traditional RAM with NVRAM has undesirable side-effects on the embedded system. Because power losses are frequent, they can occur in the middle of the modification of a non-volatile data structure. When the platform reboots, the program restarts with inconsistent data. This issue is sometimes referred to as the “broken time machine” problem. In fact, unless all bits of all pieces of memory of a device (CPU and memory, but also peripheral devices!) are made non-volatile, this problem can occur. All software layers are thus impacted by such architectural choices.

The main challenges include :

- ensuring data consistency, avoiding unreasonable performances loss – involving both runtime and compile-time techniques;
- designing efficient multi-tasks system in a context where power outages occur often;
- designing network protocols exploiting ambient energy while avoiding unreasonable performance loss when nodes in the network reboot very frequently;
- providing uninterrupted service in face of intermittent IoT connectivity and/or intermittent power.



Scheme of an experimental board including a microcontroller with NVRAM. © Inria / Photo C. Morel

➤ At Inria, the **SOCRATE** project-team works on designing robust embedded software architectures to support energy-harvesting, and intermittent power on net-zero IoT devices.

➤ ZEP is an interdisciplinary research project initiated by Inria which designs tiny wireless, battery-less IoT devices, harvesting energy in the environment, based on a novel architecture embedding non-volatile random-access memory (NVRAM). In order to benefit from the hardware innovations related to energy harvesting and NVRAM, and to optimize energy usage, ZEP designs novel software mechanisms, active at compilation-time on one hand, and at run-time on the other hand, combining architectural, compilation, and operating systems aspects.

Designing faster, smaller, cheaper IoT hardware

CPU hardware technology scaling is reaching a limit. Onwards, the most relevant technique to increase energy efficiency (the number of computations per time unit and per Watt consumed) is hardware specialization. Domain-specific hardware accelerators can come with a 100x gain (or more) in energy efficiency when compared to general-purpose computers. This gain comes mainly from

moving data closer to the computation and from removing the energy cost of full programmability (instruction fetch, cache, speculation, etc.). Small embedded devices also need specialized hardware to operate under stringent power/energy constraints. In the next ten years, we expect specializations to become even more common to meet increasing demands for performance.

The strong demand for pushing/keeping more intelligence at the edge of the network is a driver for more energy efficiency on IoT devices. For instance, machine learning and inference engines run today in a remote datacenter. Instead, next-generation of neural networks will be deployed at the edge, to take advantage of real-time sensors collecting training data and to limit the energy cost of moving raw data over the network. Designing appropriate hardware accelerators for neural networks on low-power IoT devices is thus desirable.

However, the more specialized hardware is, the harder it is to “program”. Hardware accelerator designers for IoT devices operating in the milliWatt range (or less) face a number of challenges. Designers need to explore a massive design space encompassing both hardware and software. Associated research challenges include:

- identifying and defining the IoT software stack functionalities amenable to acceleration – while the rest is kept on low-power programmable processors;
- designing domain-relevant programming interfaces for the accelerators, leading to a seamless frontier between SW and HW;
- the run-time reconfigurability of the accelerators while maintaining efficiency;
- offering a sufficient level of programmability in the accelerator (think about defining the equivalent of the GPU of IoT) to adapt to evolving standards or usage.

↗ At Inria, the **CAIRN** project-team works on designing ultra-low-power hardware computing platforms for the IoT, specialized-yet programmable, and new abstraction levels for domain-specific hardware accelerators. The **CAIRN** team also works on enhancing embedded CPU hardware architectures with mechanisms articulating non-volatile RAM and intermittent power.

Enabling Millimeter-Scale IoT Devices (Smart Dust)

Recent developments in micro-electronics have led to the first [prototypes micro-motes of a size smaller than a grain of rice](#), which can sense, compute and communicate with no additional components – in particular: no need for a printed circuit board (PCB), and no multi-die wirebonding. This extreme miniaturization

is made possible by removing the external crystal oscillators, relying only on internal RC-based oscillating circuits inside the chip. Such a micro-mote can be very small and very cheap, and this is an exciting technological breakthrough. However, significant challenges still remain, which pertain to accurately keeping track of time on such devices.

In particular, the downside of using an RC-based oscillator is clock drift, which is in the order of 16,000 ppm for a crystal-free micro-mote, compared to 40 ppm for a crystal oscillator. Drift is also very sensitive to temperature. A crystal-free micro-mote requires careful manual calibration of its clocks, to successfully communicate with other off-the-shelf devices. The absence of crystals in micro-mote impacts the very foundation of low-power wireless research. Virtually all low-power wireless platforms are equipped with a radio that communicates reliably and is able to measure time accurately. The new availability of mm-scale IoT devices devoid of crystals thus opens up a research domain and has the potential of deeply changing the field of low power-wireless research.

↗ At Inria, the **EVA** project-team develops calibration algorithms and protocols that allows mm-scale IoT devices to communicate with off-the-shelf devices, and to form coordinated networks. In partnership with UC Berkeley, **EVA** developed SCuM, the world's first micro-mote complying with wireless communication standards.

Taming Low-power Hardware Polymorphism

Elsewhere on the Internet, common computer hardware has mostly converged to a quasi-ubiquitous configuration combining 64-bit processors (x86 or ARM).

Comparatively, **low-power IoT hardware diversity is extreme**. Encountered processor architectures vary wildly, from an extensive variety of vendors, from 8-bit to 16-bit, 32-bit and 64-bit.

This extreme hardware diversity is a technical challenge in itself: **choosing the adequate hardware is difficult, and IoT software developing too often requires exotic skills**, while interoperability issues are exacerbated.

Low-power hardware innovation continues to appear, at a rate which does not decrease. A recent example is the extremely polymorphic family of CPU architectures such as RISC-V, which will shake up the status-quo, and rival increasing domination of ARM Cortex-M CPUs. On the radio side, new categories of self-powered (battery-free) chips emerge, which will disrupt the very notion of low-power. All the while, extreme miniaturization promises next-generation System-on Chip solutions which will essentially amount to “Smart Dust”, regarding its size.

Here, the challenge is thus initially to harness this extreme diversity, and subsequently to **drive an evolution towards less than a handful of standard, generic low-power hardware platforms.**

Standard Embedded Software Platforms for Low-power IoT Hardware

Requirements including low-power, cybersecurity, interoperability, IoT device management functionalities significantly increase the complexity of embedded IoT software – also on low-power devices based on microcontrollers. In the past, software embedded on such devices has been single-purposed, mostly immutable, proprietary, hardware and/or vendor-specific. These characteristics are evolving as IoT software complexity grows. A larger part of this software is now expected to mimick typical Internet-age software dynamics: more general-purpose, open-source, reusable across heterogeneous hardware and vendors, implementing a set of common standards and APIs. It has become necessary to foster generic IoT software across industrial sectors (e.g. same control algorithm same implementation, applied in different industries). This evolution has driven the emergence of a plethora of embedded operating systems which aim to provide an adequate software platform. Many vendors and Big Tech players push their own platforms, highlighting that such platforms are crucial, and that market consolidation is to be expected. The challenge for these deeply embedded software platforms is to balance performance (ultra-low energy and latency, tiny memory footprint...), with safety and security guarantees, while facilitating deeply embedded code development/portability across extremely diverse low-power IoT hardware.

↗ At Inria, project-teams including **TRIBE** and **EVA** work on designing compact, low-power embedded IoT software platforms. One example is the operating system RIOT. Another example is OpenWSN, the reference open source 6TiSCH network stack.

↗ RIOT is a general-purpose, vendor-independent operating system, for small IoT devices that cannot use Linux due hardware resource constraints. RIOT offers a free, open source platform, developed by a large grassroots community gathering companies, academia, and hobbyists, distributed all around the world, co-founded by Inria. The goal of this platform is to implement and to bundle the building blocks necessary for a more durably up-to-date, secure, transparent and privacy-friendly Internet of Things.

2.11 Global Resource Footprint Optimization

On the one hand, performance is multi-faceted e.g. speed, accuracy, security guarantees, fairness etc. On the other hand, performance must not only be evaluated locally, but also in the larger global system context. Facing stringent resource frugality constraints, locally on IoT devices, new trade-offs must be explored. More globally, facing a global ecological crisis, a comprehensive assessment of the footprint / benefits ratio of IoT is needed.

Exploiting new performance vs energy tradeoffs

Most computing today is performed with significant over-provisioning in terms of output quality (for example, in terms of precision). However, in many cases, acceptable results can be produced based on inexact or approximate computations. Both traditional applications (signal, image, vision, wireless communications, etc.), and emerging applications (machine learning, data mining etc.) exhibit inherent resilience to errors. Less performance for less power consumption is thus a traditional trade-off, which must be revisited in IoT.

For instance, exploiting the tradeoff between energy and accuracy (while keeping functionality within acceptable bounds) is a promising approach for improving energy efficiency, complementary to hardware acceleration. For example, more than 50x gain in energy efficiency can be achieved by swapping a low-precision 8-bit operation suitable for vision for a 64-bit double-precision floating point operation necessary for high-precision scientific computations (considering storage, transport and computing of the data). Optimizations have so far focused primarily on low-level representations of arithmetic computation, which do not scale to large IoT applications. Similarly, enormous gains can be obtained by severely quantizing the weights of a Machine Learning model, so as to fit the tiny memory and the small CPU capacities available on a low-power IoT device.

Optimizations have so far focused primarily on low-level representations of arithmetic computation, which do not scale to large IoT applications.

The challenge is now to design higher levels of abstraction to improve scalability and to identify high-level transformations that affect accuracy. Approximation acceptability is based on domain-specific knowledge, which must be up-to-date (may evolve) and be exploited efficiently. The degree of approximation can be tuned by the programmer at design-time or at run-time. A challenge here is the integration of compiler analysis and transformations (e.g., identifying promising regions, hierarchical decomposition of large programs, and algorithmic transformations) into accuracy tuning.

Based on similar principles, related research challenges include exploring other trade-offs, such as for instance speed vs power consumption, or increased security vs lower power consumption.

↗ At Inria, the **CAIRN** project-team explores accuracy versus energy trade-offs, designing methods to optimize low-precision domain-specific computing architectures for IoT, and inference/training for deep neural networks at the edge of the network. The **TRIBE** team studies trade-offs in terms of communication and computation costs as well as accuracy and privacy, with hierarchical machine learning models approaches, which aim to split and distribute inference along the IoT continuum – from the constrained device to the cloud, via the edge.

IoT devices consume energy not only for sensing, computing, processing data but also for communicating over the network. Therefore, an ever-promising approach to reduce energy consumption is to reduce the frequency and the size of data transmissions. However, sending less data may decrease the accuracy of the data available remotely. There is thus a tradeoff between data accuracy and communication pruning.

An interesting area of research in this domain is the design of dual-prediction mechanisms, whereby machine learning techniques are used to infer the next data transmission based on previous data transmissions. The new data is transmitted only if the prediction differs too much from the new data. The challenge here is to adapt machine learning engines and models to the extremely small resource (memory/computation) budgets available on IoT devices.

Uninterrupted service with intermittent connectivity / intermittent power

In order to reduce the overall consumption of IoT devices, research is lead on communication stack in order to set duty cycles such that every node can switch its communication interface(s) off regularly. The node is then said to sleep. Indeed, communication is the more energy consuming task for an IoT node (compared to sensing and processing). But when the communication interface is off, the node is disconnected and is not able to receive any message. If a message is sent to a sleeping node, the message will be lost and the energy used to send it is wasted. In a traditional IoT system, several mechanisms can apply to advertise a node when it wakes up that it should remain awoken. The system could be completely synchronized and thus nodes know when to wake up and listen and when they can sleep. (And acquiring accurate synchronization in highly distributed networks is a challenge as is!) But in asynchronous networks (as most of IoT networks), the sender has to generally send either a short beacon or the full data regularly until the receiver wakes up, receives the beacon or the data and acknowledges the sender.

IoT device hardware typically provides aggressive power-saving modes (sleep modes) which consume negligible power, but require temporarily inactivating CPU and network interfaces. However, successful network communication requires a sender and a receiver that are simultaneously active. A tradeoff appears to allow each device to sleep as much as possible, while still activating it at proper times to ensure global functionality.

Synchronous IoT networks solve this problem by scheduling in advance when devices wake up and listen, and when they sleep. A difficult challenge in this context is to design smart scheduling and accurate synchronization mechanisms, with less overhead, in highly distributed networks.

Asynchronous IoT networks (the bulk of IoT so far) yield different challenges to be addressed. In this domain, active area of research is the design of “wake-up radio”, whereby devices are equipped with dual radio interfaces. The main interface, used for data transfer, is switched off by default. A secondary, ultra-low power interface is used to receive wake-up signals. Passive wake-up receivers are being investigated, the challenge being to increase their sensitivity, without increasing transmit power.

A complementary approach is caching data, on behalf of sleeping nodes, somewhere in the Cloud-Edge-Thing continuum. The challenge is to determine and assess novel strategies of caching and cache replacement which:

- optimize where to store IoT data : the cost is reduced when it is closer to the requester but this could be far from the data source;
- optimize how many duplicates to store : multiple locations increase the cost of saving data but can reduce the cost of retrieving it since it is more likely to be closer to the requester;
- optimize the frequency of updates: more updates improve accuracy but increase cost.

Evaluating & minimizing the global footprint of IoT

The current ecological crisis urges researchers from all fields to evaluate the environmental impact of different technologies, currently in use or upcoming. Among others, IoT allows a better waste sorting and recycling, a better environment friendly street lightning or a better road traffic management. IoT can thus help reduce our impact on the environment in many ways, as assessed by several studies^{4 5 6}.

Still, research often focuses too much on some potential optimization an IoT technology could provide, and misses a comprehensive analysis evaluating the conditions under which net gains may in fact occur, and whether or not these conditions are likely to be met.

Direct environmental impact must be evaluated, taking into account the whole lifecycle of devices, from their production (e.g. extraction of mineral resources), to their operational expenditure (e.g. maintenance and energy consumption), to their end-of-life (e.g. potential recycling). A huge research challenge in this domain concerns improving the recyclability of small IoT electronic components.

All IoT devices are made of electronic components. A fundamental challenge is thus to build and design hardware using less resources. For instance, some research focuses on miniaturizing the printed circuit board (PCB), e.g. to use less metal and plastics, or on enabling the use of new promising material such as graphene.

4. 5 ways the IoT is helping the Environment.

5. Where IoT Meets The Environment: Building a Greener Future.

6. IoT for Environmental Sustainability.

Research challenges also have to be addressed in designing antennas (that can sometimes be printed with biodegradable ink⁷) and resource-friendlier batteries.

Furthermore, indirect impact must be evaluated, encompassing induced effects, rebound effects etc. as new uses are likely to counterbalance optimizations allowed by IoT. For example, trillions of microscopic batteryless wireless devices powered by energy harvesting, though leveraged for advanced services in the next-generation Internet, may still lead, globally, to more greenhouse gas emissions.

Typically research focuses on improving potential direct impact, but stops short of answering more global questions: improvements to what extent in practice? Until when? For what cost? And above all, for what net benefits, globally? In this field, complexity rises because extremely interdisciplinary skills are required, combining not only technological knowledge but also social, economical, political knowledge.

A challenge in this domain is thus the design of adequate conceptual frameworks which can better capture and evaluate the full spectrum of IoT's environmental impact. Existing frameworks typically do not capture indirect impact, although indirect impact might dwarf direct impact. Although direct impact can theoretically be captured by existing frameworks, these are challenged because (i) data is difficult to gather, and (ii) technology and uses change at a very fast pace.

7. D. Iba, et al. "Development of smart gear system by conductive-ink print," Proc. SPIE, 2019.


Conclusion

The Internet of Things has gained fundamental importance in the landscape of technologies which will shape tomorrow. In the world that is coming, entities (countries, organizations, companies, individuals) aiming to preserve their sovereignty must raise their awareness, and must devote the means necessary to lead substantial research activities and deep tech development, in several domains which underlie IoT.


These domains are diverse, spanning from next-generation communication networks to pervasively distributed computing, from embedded system software to low-power hardware, from Human-Machine interaction to cyber-physical system control and resilience, from cyber-security and safety, to privacy-preserving data processing.

Furthermore, as IoT technology becomes more tightly woven into society and our individual lives, the design of IoT technology standards becomes ever more critical. In this context, in order to be in a position to preserve the geopolitical neutrality of IoT technology, active participation in relevant standards development organizations is crucial.

Last but not least, as the environmental crisis mounts, there is the hope that our impact on nature can be reduced thanks to more IoT-enhanced mechanisms, to be deployed and used massively. Substantial complementary efforts are nevertheless required, to assess that this reduction will indeed globally outweigh the environmental impact of producing, deploying and maintaining these IoT mechanisms, throughout the entirety of their life-cycles.



Anechoic room of the FIT (Future Internet of Things) experimental platform. © Inria / Photo C. Morel.



Inria

Domaine de Voluceau, Rocquencourt BP105
78153 Le Chesnay Cedex, France
Tel.: +33 (0)1 39 63 55 11
www.inria.fr