



HAL
open science

A survey on addressing privacy together with quality of context for context management in the Internet of Things

Sophie Chabridon, Romain Laborde, Thierry Desprats, Arnaud Oglaza,
Pierrick Marie, Samer Machara Marquez

► To cite this version:

Sophie Chabridon, Romain Laborde, Thierry Desprats, Arnaud Oglaza, Pierrick Marie, et al.. A survey on addressing privacy together with quality of context for context management in the Internet of Things. *Annals of Telecommunications - annales des télécommunications*, 2013, 69 (1-2), pp.47-62. 10.1007/s12243-013-0387-2 . hal-01285786

HAL Id: hal-01285786

<https://hal.science/hal-01285786v1>

Submitted on 6 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 12723

To link to this article : DOI :10.1007/s12243-013-0387-2
URL : <http://dx.doi.org/10.1007/s12243-013-0387-2>

To cite this version : Chabridon, Sophie and Laborde, Romain and Desprats, Thierry and Oglaza, Arnaud and Marie, Pierrick and Machara Marquez, Samer *[A survey on addressing privacy together with quality of context for context management in the Internet of Things](#)*. (2014) Annals of Telecommunications, vol. 69 (n° 1-2). pp. 47-62. ISSN 0003-4347

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

A survey on addressing privacy together with quality of context for context management in the Internet of Things

Sophie Chabridon · Romain Laborde ·
Thierry Desprats · Arnaud Oglaza · Pierrick Marie ·
Samer Machara Marquez

Abstract Making the Internet of Things (IoT) a reality will contribute to extend the context-aware ability of numerous sensitive applications. We can foresee that the context of users will include not only their own spatio-temporal conditions but also those of the things situated in their ambient environment and at the same time, thanks to the IoT, those that are located in other remote spaces. Consequently, next-generation context managers have to interact with the IoT underlying technologies and must, even more than before, address both privacy and quality of context (QoC) requirements. In this article, we show that the notions of privacy and QoC are intimately related and sometimes contradictory and survey the recent works addressing them. Current solutions usually consider only one notion, and very few of them started to bridge privacy and QoC. We identify some of the remaining challenges that next-generation context managers have to deal with to favour users' acceptability by providing both the optimal QoC level and the appropriate privacy protection.

Keywords Internet of Things · Context management ·
Privacy · Quality of context

S. Chabridon (✉) · S. M. Marquez
Institut Mines-Télécom/Télécom SudParis, CNRS UMR 5157
SAMOVAR, 9 rue Charles Fourier, 91011, Evry, France
e-mail: Sophie.Chabridon@telecom-sudparis.eu

R. Laborde · T. Desprats · A. Oglaza · P. Marie
Université de Toulouse, IRIT UMR Toulouse, 5505, France

R. Laborde
e-mail: Romain.Laborde@irit.fr

1 Introduction

The next big step in our electronic society will be the realization of an Internet of Things (IoT) [6, 50]. The Internet will connect not only people but also machines and smart objects or things, thanks to wireless connectivity. The “Anywhere, anyhow and anytime” communication paradigm of mobile and ubiquitous computing gets extended to “Anything, anyone and any service” with the IoT [92].

Thanks to a variety of sensing and potentially mobile devices, it is becoming possible to perceive events and changes within the ambient space surrounding users. This implies to sense precisely the current context of users in order to determine what is their situation and what should be the behaviour of a context-aware system. The importance of this notion of context has been identified by Coutaz et al. [27] in demonstrating that context information can enrich human activities with new services able to adapt to the circumstances in which they are used. As context information is central to the decisions that context-aware systems must take, the quality of context (QoC) has to be known and evaluated carefully. Without this knowledge, the service provider would be unsure about the quality of the context information received and inadequate service adaptations could be triggered by data of unknown quality. Besides, as context data may reveal sensitive information about persons, like their location or their activity at any time of the day or night, the success of a context-aware service is highly dependent on how this service enables users to feel in control of their privacy.

Context-aware services under the paradigm of ambient intelligence (AmI) were so far limited to closed environments (e.g. a room, a house or a building) where a number of specific functions known at design time can be supported. The IoT expands the AmI paradigm to open scenarios where

new functions or services need to be accommodated at run time without them having been necessarily considered at design time [70]. This paradigm shift is being recognized by the research community [86] under various terms like emerging pervasive environments [11] or cyber-physical systems [26]. The IoT thus brings new opportunities by enabling enriched context-aware services, but it also raises new challenges for managing the QoC of the tremendous amount of collected information while at the same time preserving the privacy of the users. Privacy is one of the major ethical concerns of users with respect to the IoT [50] and is a crucial open issue that may limit the realization of the IoT vision [70]. We consider that new research directions studying privacy jointly with QoC are the most promising to take up the challenge for a context management in the IoT that will be technically effective and well accepted by the users.

In this article, we first outline in Section 2 the new challenges faced by context management in the IoT resulting from the complex interdependencies between QoC and privacy. In Section 3, we precisely define what is QoC and discuss the recent solutions and open issues for QoC management. Similarly in Section 4, after defining privacy, we survey current privacy protection techniques and analyse their applicability for context management in the IoT. We finally discuss in Section 5 the works that started to bridge privacy and QoC and identify challenging issues related to the IoT before concluding this survey in Section 6.

2 Challenges for context management in the IoT

The IoT is the concept that aims to extend the regular Internet to the real-world physical objects [6]. Consequently, a large number of things can be, at any time during their life cycle, either temporarily or permanently, connected to the global network infrastructure. These things have to be identified and accessed ubiquitously. Some of them can be directly connected by natively embedding communication capabilities, while others are classical raw things or physical infrastructures to which one or more additive communication devices are associated. Moreover, some things are situated in fixed locations while others can be mobile because they move or are moved from some place to another one. Things, by using technologies such as RFID, wireless sensors networks, smart objects networks, etc., can provide some data related to their own identity, location, state, behaviour and/or to those of the environmental conditions they can perceive. The control of this data dissemination is left either to the owner of the object or to a trustworthy mandated operator.

Making the Internet of Things a reality will contribute to extend the “context aware” ability of numerous sensitive applications. We can envisage from now on that the context of users will include not only their own spatio-temporal conditions but also those of the “things” that are situated in their ambient environment and at the same time, thanks to the IoT, those that are located in other remote spaces. Consequently, next-generation context managers, like the ones envisioned in the INCOME project [4], have to interact with the IoT underlying technologies. More precisely, they have to (1) collect data that are originally produced by the devices spread around the IoT, (2) iteratively process and propagate it as context data within a network of intermediate consumer/producer entities and (3) provide computed high-level context information to context-aware applications assisting and helping users.

A context manager is a software entity computing high-level information from various sources of data. Its functionalities include context data acquisition, context data processing (fusion, aggregation, interpretation, inference) and context data presentation to context-aware applications. These applications are classically named as final context data consumers, while entities providing raw data are considered as producers. A component within a context manager that operates context data transformation is considered as both an intermediate context data consumer and producer.

Consequently, such a context manager can, from the huge amount of data it collects from various connected things that may belong to multiple owners, build high-level context information that is then delivered to any interested context-aware application. A single, direct and strong coupling between one original context data provider and one final context data consumer does not exist anymore: they do not have to know or be aware of each other.

Nevertheless, even more than in the case of ambient systems, the same requirements towards a context manager remain for both original producers and final consumers. The final consumers have to deal with the knowledge about the quality of the context information provided by the context manager. According to the QoC level, they can adjust their own context sensitive reaction: the more the QoC is explicit and precise, the more the algorithms leading to relevant decision taking will be sophisticated. Similarly, context data owners and indirectly original context data providers need to be able to express privacy requirements about the data they accept to provide to a context manager. As illustrated in Fig. 1, next-generation context managers have to support the measurement of the QoC at each step of the context data life cycle: from its acquisition from the IoT to its delivery to context-aware applications. At the same time, these context managers also have to respect privacy requirements.

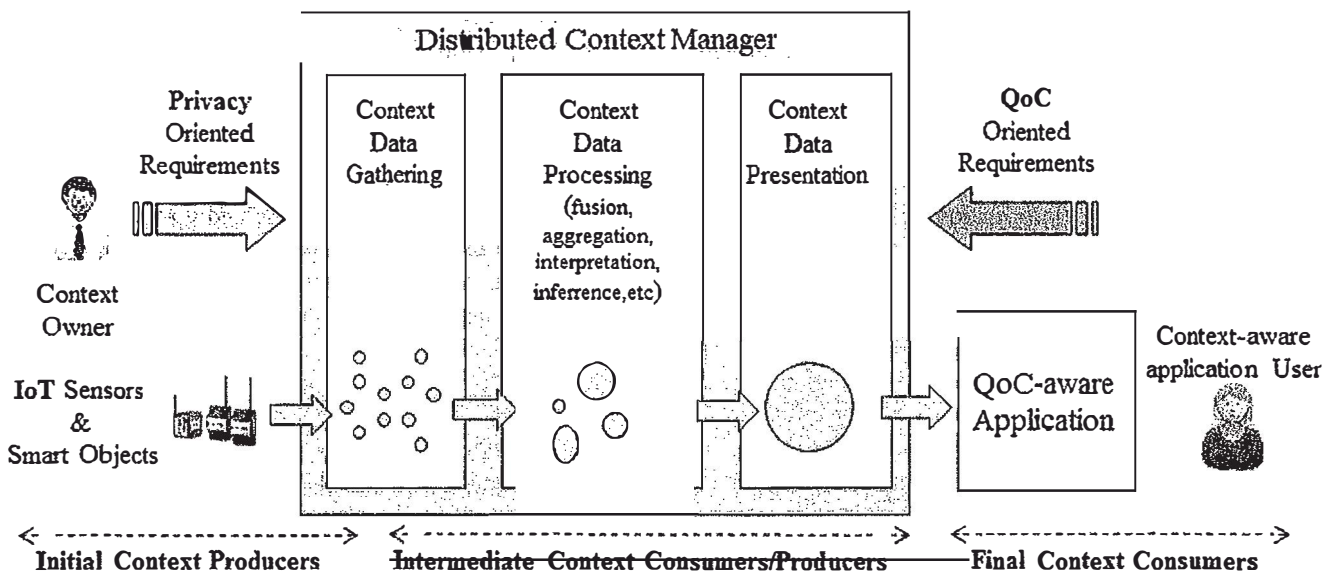


Fig. 1 Logical view of a privacy- and QoC-aware context manager

We identify three new challenges brought by the IoT with respect to the management of privacy and QoC in context-aware computing:

1. *Context data production/consumption decoupling*: To provide effective and efficient context management at the scale of an IoT, context producers and consumers should be decoupled both in time, working at different times and speeds, and space, not having to know each other. This calls for new solutions to protect the privacy of the users who are the owners of context information that could be exchanged and exploited without them being aware of it. The very high amount of data provided by context producers in a decoupled and asynchronous way calls for an efficient and effective context data distribution solution. As mentioned by Bellavista et al. [9], this includes powerful context aggregation and filtering techniques in order to reduce the final management overhead. Filtering data as close as possible to the node that generated them avoids useless transmissions. Moreover, filtering context data on their QoC according to the QoC requirements of context consumers allows to further improve the filtering efficiency.
2. *QoC-aware privacy*: As noticed in [9], context data security including all the mechanisms to grant privacy, integrity and availability of data is still a neglected issue. This can be explained by the existence of many efficient solutions for addressing security problems, for instance by exploiting access control and encryption mechanisms. However, these solutions usually fail

to consider the privacy loss that can result from indirect inferences by fusing different pieces of information coming from various sources as in the case of context management in the IoT. Also, the high computational overhead that can result from the use of security mechanisms such as cryptography requires to limit their use to specific cases. Context data may have different levels of privacy, and only some of them require proper security mechanisms. One promising research direction is to consider the specificities of context information and their associated QoC metadata. These metadata bring additional knowledge that could be exploited for enhancing privacy. This would enable to define QoC-aware privacy policies that may evolve dynamically according to the changes in the context of the user and to the quality of this context information. For instance, users might allow the access to some of their private information but only under some specific conditions: during a meeting on a particular topic, with a timeliness constraint to activate the policy at most 30 s after the beginning of the meeting and with a particular group of persons.

3. *Dynamic interdependency of QoC and privacy*: Even though the interdependency of QoC and privacy has already been tackled in ambient intelligence and closed context-aware systems, this issue gains importance in the case of an open IoT when dealing with an abundance of information generated by a very large number of highly diversified data sources distributed over large geographical areas calling for dynamic solutions [26]. QoC may impact privacy in several ways. In order to

have high quality context information, fine-grained and precise data are necessary. This may be in contradiction with the privacy requirements of users who would prefer to deliver only coarse-grain data. Moreover, a large amount of context information may need to be collected and then aggregated or fused, allowing to infer new context information unforeseen in the first place. This is not compliant with the data minimization principle of privacy laws requiring to limit data collection as much as possible and this may allow to derive user's new information that was unknown or hidden so far, increasing the risks of privacy violation. Privacy also has consequences on QoC. It may influence the QoC level of the context information that is to be delivered to context consumers by setting both its upper and lower bounds. On the one hand, more precise information may be considered as more intrusive. On the other hand, low quality data may allow to infer false context information with potentially bad repercussions on the user's privacy.

3 Quality of context

In this section, we first define the concept of QoC. As a single largely accepted definition of QoC is still missing, we present the most representative definitions. We then present the main criteria used in research works to evaluate QoC. We finally compare and analyse these works.

3.1 Quality of context definitions

In the domain of context-aware computing, context information is known to be inherently uncertain [10, 46]. Henricksen and Indulska [46] identified four types of imperfection of context information. A context attribute may be *unknown* when there is no information about it, leading to incomplete context information. It may also be *ambiguous* as there is a risk of having contradictory information from different context sources. An attribute is *imprecise* when the reported information is correct but not provided with a sufficient degree or precision. As context data are by nature dynamic and very heterogeneous, they also tend to be *erroneous* and not exactly reflecting the real state of the modelled entity. Therefore, one solution that has been used for a decade is to attach metadata to context information representing its quality. Historically, the importance of QoC as a first-class concept for the design of context-aware systems has first been identified by Buchholz et al. [15]. The authors define QoC as “any information that describes the quality of information that is used as context information. Thus, QoC refers to information and not to the process nor the hardware component that possibly provide

the information”. Note that they consider that QoC is intrinsic to the information and different from the quality of the computing service (QoS) and from the quality of the hardware device (QoD). In this seminal work, the authors have identified five criteria from their experience as the most important ones, namely precision, probability of correctness, resolution, trustworthiness and up-to-dateness. A description of these criteria is presented in Section 3.2.

The notion of worth has then been added by Krause and Hochstatter [54] to introduce the point of view of the targeted application. Manzoor [62] has pushed this notion further by differentiating the objective and the subjective views of QoC. The *objective view* is independent of the situation in which context information is used and of the consumer requirements for that context information. The objective view of QoC is determined by the characteristics of the sensors that have collected the context information and how the measurement of the context value took place. The *subjective view* of QoC illustrates how much a piece of context conforms to the requirements of a particular consumer application. Considering the objective and subjective nature of QoC, Manzoor [62] then proposes to define QoC as an indication of the degree of conformity of the context collected by sensors to the prevailing situation in the environment and the requirements of a particular context consumer. This means that QoC may vary with different context consumers and that it cannot be measured independently of a context consumer and of the intended purpose.

Bellavista et al. [9] revisit the initial definition of Buchholz et al. [15] for the case of context data distribution for mobile ubiquitous systems. In addition to the traditional and extremely data-focused notion of QoC, they consider the quality of the context data distribution process (e.g. data delivery time, reliability, etc.). Whereas Buchholz et al. [15] handle separately the three different quality dimensions that they have identified as QoC, QoS and QoD, Bellavista et al. [9] argue that it is not always possible to clearly separate these three quality dimensions. They therefore prefer a broader QoC definition, dealing both with the quality of the context data and the quality of the context data distribution.

3.2 Main QoC criteria

We present in this section the main criteria, i.e. quantifiable properties or parameters, that have been proposed in the literature to measure QoC. The same term is sometimes associated with different definitions, and also several terms may represent the same concept. Indicators, i.e. targeted values of the criteria, are usually given, but little is said on the metrics that can be used, i.e. the mechanisms or algorithms allowing to measure the criteria. This shows that there is no

consensus on a standard framework for QoC evaluation and that more work is still needed.

We separate QoC criteria into two categories: simple criteria that do not depend on other criteria to be evaluated and usually provided by sensors and composite criteria requiring the knowledge of some other criteria for their evaluation.

Simple criteria

Precision In metrology [49], precision indicates the degree of dispersion of a set of measures [62, 72]. Buchholz et al. [15] use this term in the meaning of the accuracy (see below). Filho [37] defines precision as the level of details in which the context information is describing an entity of the real world and evaluates it as the ratio of the current precision level to the maximum precision level (by configuration).

Resolution Resolution denotes the granularity of information [15] which corresponds to the degree of detail with which a sensor can collect data [62], and is part of the sensor characteristics. Filho [37] computes it as the ratio of the current granularity level with the maximum granularity level (obtained from a configuration file).

Accuracy In metrology [49], accuracy represents the closeness of agreement between a measure and the true value as for [37], [63] and [89]. Buchholz et al. [15], under the term of precision, use bounds to represent the deviation from the exact value. Kim and Lee [52] estimate accuracy using confidence intervals obtained by a classical statistical method.

Timeliness For Buchholz et al. [15], up-to-dateness or, for Kim and Lee [52] and Sheikh et al. [98], freshness represents the age of context information corresponding to the time elapsed between the determination of context information and its delivery to a requester. However, Filho [37] and Manzoor et al. [64] redefined it as the degree of rationalism to use a context object for a specific application at a given time. They compute it as the ratio of the age of the context information with the time period and normalize it on [0..1].

Composite criteria

A number of composite QoC criteria can be derived from the simple criteria presented previously and also from the combination of sensor characteristics, measurement conditions and consumer requirements.

Completeness Completeness indicates the amount of information available in a context observation [52]. Manzoor [62] associates a weight to each context attribute

to represent its importance and computes completeness as the ratio of the sum of the weights of the available attributes to the total of the weights of all the attributes of a context object. Han et al. [44] define a reliability criterion as the minimum number of sensor data that should be collected within some time units. Filho [37] redefines completeness to integrate timeliness so as to avoid unnecessary computations and manipulate only current context information.

Significance This criterion was proposed by Manzoor [62] to represent the worth of some context information for a specific application. It is also considered by Filho [37]. It is computed as the ratio of the critical value of the context information to the maximum critical value that a context object of that type can have.

Usability Usability depicts how much a piece of context information is suitable for use with the intended purpose by the context consumer application [62, 64]. It is equal to 1 when the provided granularity level of collected context information is larger than the required granularity level and 0 otherwise.

Probability of correctness The probability of correctness estimates how often some context information is unintentionally wrong due to internal errors [15, 88, 98]. Brgulja et al. [13] propose to combine different QoC criteria like timeliness, context source trustworthiness and precision, to calculate the probability of correctness of context information. Filho and Agoulmine [38] extend this work by taking into account context dependencies among context data that may affirm or contradict simultaneously the truth of the characterization of a given situation; therefore, a high probability of correctness indicates that the chances are small that the associated context is in contradiction with other context information.

Trustworthiness Filho [37] identified two approaches for measuring trustworthiness: (1) measuring how trustworthy is the entity that provided the context information like in [15] and (2) measuring the belief one can have directly on the context information. In the first approach, Huebscher and McCann [47] propose a learning model that calculates trustworthiness based on binary positive/negative feedback from the users. Manzoor follows the second approach in [64] and later renamed this criterion *reliability* [62] and proposed to evaluate it as the inverse of the distance between the sensor and the entity about which context information is collected. Neisse [72] argues that trustworthiness is related to the capability of the context provider to reliably describe the QoC levels it can guarantee and, as such, should not be part of the QoC criteria.

Confidence Different research works define the confidence attached to some context information. However, the way confidence is evaluated was initially limited to a single criterion. It can derive from the generation time of context information [14, 96]. Korpipää et al. [53] compute the probability of correctness of context information, while Ranganathan et al. [91] rely on the precision of sensor measurement to indicate the confidence in context information.

More recently, some works were proposed to combine several QoC criteria to measure confidence. McKeever [68] quantifies the imperfections of vague context (fuzzy membership), erroneous or conflicting context (reliability), imprecision (precise membership) and out-of-dateness (freshness). These values are then combined to evaluate context event confidence. Manzoor [62] proposes a confidence inference system that uses fuzzy logic to infer the value of confidence. It combines different QoC criteria such as reliability, timeliness, completeness, significance and usability and takes into account the application QoC requirements to provide the value of confidence on context.

Yasar et al. [111] intend to improve the efficiency of communication in large-scale vehicular networks based on *aggregated quality*. It is close to the notion of confidence and is derived from two other criteria: QoC and peer reputation. QoC here is computed using four criteria: temporal relevance, completeness, significance and spatial relevance. Spatial relevance is specific to vehicular networks; it determines the distance and the direction of the destination node relatively to the source node, providing that contextual information is more relevant when it is received from a neighbouring node moving in the same direction.

3.3 Synthesis of works on QoC

Table 1 presents a synthesis of the main works on QoC criteria measurement. For simple criteria, a confusion exists in the terms used and we recommend to rely on standard definitions from the metrology domain or from the ISO where available. With regard to temporal aspects, the notion of timeliness brings additional knowledge and should be favoured. Concerning composite criteria, the two criteria of trustworthiness and confidence present commonalities which require to be further studied. Besides, the most recent works of Filho [37] and Manzoor [62] proposed new criteria with their associated indicators. However, their relevance should still be validated on concrete applications.

Recently, Perera et al. [86] have presented the results of a study of 50 research projects on context-aware computing covering a decade and analysed their readiness to address the issues raised by the IoT. A little less than one-third of these research projects (15 over 50) consider the feature of QoC although, like we do, Perera et al. identified this feature as a requirement of context-aware frameworks for tackling

the IoT. More work on this topic is thus clearly needed. Context data coming from billions of sensors may be collected in the IoT. Reasoning on all these context data at once is not feasible due to the processing power and time or storage capacity that would be necessary and a selection of the appropriate input context data must take place. QoC therefore appears as a way to filter out this abundance of context data and can help to identify what sensors should be used by ranking sensors on QoC criteria. For addressing the new challenges of high amount of collected data, openness and dynamicity brought by the IoT with regard to context-aware computing, we consider that QoC management frameworks should be flexible and efficient and limit overhead by computing only the relevant QoC criteria as proposed in [1, 21]. They should also be extensible by enabling the definition of new QoC criteria including their associated computation algorithm as investigated in [65].

4 Privacy

As emphasized by the ITU in its report on the IoT, privacy is crucial for the control of this new complex and moving environment: “Invisible and constant data exchange between things and people, and between things and other things, will occur unknown to the owners and originators of such data. The sheer scale and capacity of the new technologies will magnify this problem. Who will ultimately control the data collected by all the eyes and ears embedded in the environment surrounding us?” [50]. In this section, we first discuss privacy definitions and then review the state of the art of privacy technology from the perspective of the IoT.

4.1 Privacy definitions

Langheinrich presents a thorough analysis of privacy issues in the domain of ubiquitous computing [57]. The author revisits old references to privacy in law texts, but notice that it is still unclear exactly what privacy means today, especially with the new usages provided by communication and computing technologies. Warren and Brandeis [109] described privacy as “the right to be let alone”. However, Langheinrich [57] states that preserving privacy through isolation is no longer an option in today’s information and communication world. Privacy is now usually perceived by users as an expectation of being in a state of protection without having to actively pursue it. Users actually feel concerned when their privacy gets violated. Marx [66] identifies four personal border crossings that are perceived as privacy violations:

A natural border prevents your presence (or feelings or emotion) from being perceived through one of the

Table 1 Synthesis of the QoC criteria used in surveyed works

QoC criterion	Research works	Definition
Simple criteria		
Precision	[62, 72] [15] [37]	ISO definition: dispersion of a set of measures Confidence interval. Use in place of accuracy Level of precision. Ratio: current precision level/maximum precision level
Resolution	[15, 62] [37]	Degree of detail or granularity Level of detail. Ratio: current granularity level/maximum granularity level
Accuracy	[37, 89] [15] (precision), [52] [62]	ISO definition: closeness with true value. Combines precision and trueness Confidence interval around the true value Sensor characteristic
Up-to-dateness or freshness	[15, 52, 98]	Age of context information
Up-to-dateness (new definition) or timeliness	[37, 64] [62]	Rationalism of using a context object for a specific application at a given time
Composite criteria		
Completeness	[52] [62] [44] [37]	Amount of available context information Ratio of weight of available attributes to weight of total number of attributes Min. number of sensor data to be collected within some time For available and timely context information
Significance	[37, 62, 64]	Ratio of critical value level to max. critical value
Usability	[62, 64]	1 if granularity level > requested granularity
Probability of correctness	[15, 88, 98] [13, 38]	Chances there are no unintentional errors Combination of QoC criteria. Bayesian approach
Trustworthiness with respect to the context provider	[15] [47]	No evaluation method explicitly given based on feedback from users
Trustworthiness with respect to context information or reliability	[64] [62]	Inverse of distance between entity and sensor
Confidence	[14, 53, 91, 96] [62, 68, 111]	Single dimension Aggregation of QoC criteria

human senses. Walls, doors, clothes, darkness, sealed letters, telephone and email messages represent natural borders to observation.

A social border involves expectations that persons with certain social roles (doctors, clergy members, lawyers) will not disclose confidential information.

A spatial or temporal border separates information from various periods or aspects of one's life.

Borders due to ephemeral or transitory effects: Such borders assume that interaction and communication are ephemeral and transitory like actions that one hopes to get forgotten soon or old pictures and letters that one puts out in the trash.

Solove argued that no single privacy definition can be workable, but rather that there are multiple forms of privacy [101]. He proposed a privacy taxonomy with an overview of the activities that might lead to privacy violations:

Information collection: Although information collection is usually done with the consent of a person (data subject), hidden and forced collections lead to surveillance or interrogation activities that violate the data subject's privacy.

Information processing: Information gets stored, combined and searched, threatening the data subject's privacy.

Information dissemination: When information is disseminated, confidentiality may be breached in multiple ways. Disclosure may happen with the publication of truthful facts that might affect the person's reputation. Exposure of private details may occur. The accessibility of already public information may increase by disseminating it, like for telephone numbers.

Invasion: Invasion into personal data may occur through intrusion into one's life and through decisional interference.

As pointed out by Langheinrich [57], even though this taxonomy intends to be used for legal protections, it may also be useful for technologies. Technology providers should systematically analyse whether some software or technology might increase the chances of such problem to occur, and how to mitigate it. This taxonomy matches the context management activities and is identified as the most comprehensive by Shen and Pearson [99]. It will therefore guide our analysis of current privacy technology.

4.2 Privacy technology

We present in this section a synthesis of recent reviews on privacy-enhancing technologies (PETs), namely [30, 84, 99, 104, 108]. The concept of PET can be traced back to the early 1980s with the work of Chaum [25], Pfitzmann [87] and others focusing on confidentiality in communications as an application of the recent invention of public key cryptography. This concept later came out of the academic and research corner in the mid-1990s and was promoted by the Dutch and Canadian Data Protection Authorities [16, 20] for IT product development [106].

For a clear presentation of this large variety of solutions, we refine the categorization proposed initially by Danezis and Gürses [30] and organize this section in three parts with privacy as confidentiality, privacy as control and privacy as transparency. We renamed the third category, entitled privacy as practiced in [30], to insist on the recognized necessity to provide users with transparency tools for them to easily and efficiently control their privacy [57].

4.2.1 Privacy as confidentiality

Confidentiality is usually present in some form in existing privacy technologies as the first objective of privacy is to protect personal context data from being accessed by unauthorized persons. If personal data become public, confidentiality and hence privacy are lost. Privacy as confidentiality represents the solutions for anonymizing the collected data, anonymizing communications and minimizing the collection of data.

Anonymity of data It relies on cryptographic solutions in order to achieve properties like unlinkability (two information items or two actions of the same user cannot be related), undetectability (an attacker cannot distinguish whether an information item exists), unobservability (it is not possible to detect whether a system is being visited by a given user) and communications content confidentiality. The k -anonymity [103] approach claims that an individual cannot be identified within a set of k users. Several variants have been proposed like l -diversity [60] where a block of data is l -diverse if it contains at least l well-represented values for the sensitive attribute S . Other anonymity metrics have been proposed [51]. However, what degree of anonymity is sufficient for a particular use case is dependent on legal and social consequences of a data breach and is still an open question [34]. Differential privacy aims to provide means to maximize the accuracy of queries from statistical databases while minimizing the chances of identifying its records [36]. Borcea-Pfitzmann et al. [12] underline that data minimization as the prime property of privacy has reached its limits and is now balanced by the users' demands of more and more functionality.

Anonymity in communication Anonymizing communications aims at protecting traffic data from concealing who talks to whom. Even if the content of a communication is kept confidential, sensitive information may be leaked by traffic data which include locations and identities of the communicating parties, time, frequency and volume of the communication. Providing anonymous communication is challenging since many communication protocols use unique identifiers [67]. Approaches like the pioneering Mix-Net protocol [25] also known as onion routing and Tor (The Onion Router) [35] where encrypted messages are routed in an unpredictable path provide a solution. Mobility protocols such as MIPv6 and HIP need to resort to global identifiers (resp. Home Address and Host Identity) that can be used to track users and their location. Some new architectures/mechanisms partially address this location privacy issue such as IP² [81], Turfnet [95] or Blind [112].

Data minimization It aims to limit the collection and processing of personal data. It can be enforced by encrypted aggregation techniques like those described in [17, 69]. Other approaches include perturbation and obfuscation. Perturbation means that data get systematically altered using a perturbation function (e.g. adding random numbers [2]). Obfuscation means that a certain percentage of data get replaced by random values (e.g. replace with the mean). Borcea-Pfitzmann et al. [12] underline that data minimization as the prime property of privacy has reached its limits and is now balanced by the users' demands of more and more functionality.

4.2.2 Privacy as control

Privacy as control refers to the ability to control what happens with personal data and to prevent abuses. This encompasses technologies for specifying and enforcing privacy policies. We first review in this section the legal status of privacy protection and then discuss the notions of access control and usage control policies.

Privacy as a fundamental right International guidelines have been defined to protect privacy in [77, 79] and more recently in [78] from which Wang and Kobsa [108] identify a set of 11 fundamental privacy principles:

1. Notice/awareness: Make policy statements clear and explicit.
2. Data minimization: Carefully evaluate the necessity, effectiveness and proportionality of a new technology before deployment. Prefer the least privacy-invasive solutions.
3. Purpose specification: Specify the purpose of data collection at the collection time.
4. Collection limitation: Set limits to the collection of data.
5. Use limitation: Personal data should not be used or disclosed for purposes other than those specified.
6. Onward transfer: Do not transfer data to a 3D party if it does not ensure adequate protection.
7. Choice/consent: Individuals should be provided with mechanisms, such as opt-in and opt-out mechanisms, to decide on the collection, use and disclosure of their personal data.
8. Access/participation: Individuals can access and inspect their stored data.
9. Integrity/accuracy: A data controller should ensure that the collected personal data are sufficiently accurate and up-to-date to the intended purpose.
10. Security: Protect data against risks such as loss, unauthorized access, destruction, use, modification or disclosure.
11. Enforcement: Include mechanisms to enforce privacy principles.

Following the work of Cavoukian [19], one of the Ontario privacy commissioners, Privacy by Design (PbD) has been recently accepted in 2010 [90] as a concrete way to control and protect personal data, ensuring the process of compliance with law from the projecting phase and not after, when everything has already been carried out and the system is in operation. PbD is the key to the future of privacy and can be seen as an evolution of PETs with a focus on accountability and law application.

The legal directives also include the Privacy Impact Assessment (PIA) that aims to evaluate the risks threatening

privacy and data protection. PIA requires that organizations or industries consider in advance what kind of risks on privacy an envisioned project may generate. Several data protection agencies recommend PIAs to ensure legal compliance to national and international privacy laws.

From static to contextual access control policy models

Access control models have been designed to formalize how to write authorization rules. From mandatory and discretionary access control to modern-attribute-based models that include the concept of context, we will see that the more their power of expression increases, the more the difficulty to write policies for non-technical users increases too. Access control models consider three main entities: subjects, objects and permissions. Subjects are users or applications who can perform actions in the system. Objects are resources or services that subjects want to control access to. Permissions determine how subjects can access resources. Even if these three entities appear in all models, their representation has evolved over time to adapt to the requirements of modern systems. First access control models from the 1970s to the 1990s such as identity-based access control models, mandatory access control models or role-based access control (RBAC) models were designed for stable computing environments involving few mobility. As a consequence, they did not consider the notion of context. Some of these models, especially RBAC, were extended to cope with this issue in the 2000s such as [28, 55] or [56]. However, roles were initially defined as a job function in an organization and they are not efficient in a dynamic, open and context-aware system. A more flexible approach called attribute-based access control has emerged. Authorization rules are specified based on any security characteristic of the subject, resource, action or environment. This approach has been reused for integrating the context of a situation like that in [7] and [29]. Enterprise Privacy Authorization Language (EPAL) [5] is a formal language to define internal policy. An enterprise first defines an EPAL vocabulary and can then specify its own EPAL-customized policies. This offers to define rich policies. However, Wang and Kobsa [108] mention that EPAL lacks a finer granularity in the writing of rules. Extensible Access Control Markup Language (XACML) [76] is a popular general-purpose policy language which outweighs EPAL in expressing not only access control policies but also privacy control policies [108]. It is being investigated in [80] to describe context-aware policies.

Usage control policies In addition to authorizations, privacy also includes obligations describing how context information is handled after access is granted. Obligation policies specify some actions that must be achieved to control the usage of context information. These actions

are required to be performed by a consumer before, during or after the usage of the context information. The main well-known policy-based languages and/or related frameworks that support obligations expression are usage control (UCON) (see [58] for a survey on UCON-based approaches), XACML [76], Ponder2 [105] and the OSL-based framework of Neisse et al. [73]. Because the process chain of context data is complex and usage control policies must be respected at each step, one of the usage policy distribution approaches usually consists in bounding machine-readable policies to personal data. This approach is called sticky policies. Obligations are kept travelling with data along the context processing chain. Access to data can be as fine grained as necessary. Encryption mechanisms supporting the stickiness of policies to data and a related key management allow data attributes to be encrypted based on the policy. Access to data is mediated by a Trust authority that checks for compliance to policies in order to release decryption key [85]. An example of using sticky policies to distribute privacy policies was proposed in the EU FP7 PrimeLife project [33].

Identity management Modern identity management systems separate (1) the entity that provides a personalized service to users by using information of users from (2) the entity that authenticates users and stores data about users. These entities are respectively called service provider (SP) and identity provider (IdP). Selective disclosure of personal data sent by IdPs to SPs has been a core issue in the field of identity management. For example, Shibboleth [100] allows IdP administrators to be defined whose attributes are sent to specific SPs and OpenID [82] allows users to control their personal information sharing by defining personas. U-Prove [83] is currently one of the most advanced solutions for user-centered identity management focusing on privacy providing unlinkability and users' control. The PrimeLife project has also proposed an interesting approach that couples usage policies with an advanced selective disclosure mechanism [33].

4.2.3 Privacy as transparency

Transparency tools intend to improve the users' understanding and control of their data profile. Castellucia et al. [18] identify four characteristics that such tools should possess:

1. Provide information about the intended collection, storage and/or data processing
2. Provide an overview of what personal data have been disclosed to what data controller under which policies.
3. Provide online access to the personal data and how they have been processed

4. Provide counter profiling capabilities helping the user to guess how the data match relevant group profiles, which may affect future opportunities or risks

Privacy as transparency is an important issue because most PETs are useless if people cannot use them efficiently. Privacy as transparency is even more critical for the next IoT-based distributed systems than it is in the existing web-based ubiquitous applications. The users (i.e. context data owners) will not only have to control the personal data which can be propagated from the terminals with which they directly interact (smartphone, laptop) but they will also have to handle the control of the data automatically produced by the connected things they own, which surround them or which are located in their life environments (home, office, etc.). These IoT data could be scattered across a large distributed system while facing issues like heterogeneity, scalability, etc. Despite the importance of this issue, very few research works have studied it. For example, Castellucia et al. [18] indicate that there is yet no tool supporting characteristic 4 although this is highly desirable as shown by the FIDIS project [45].

Several research works have been conducted to simplify the users' interaction with their electronic security. For example, the P3P Project ("Platform for Privacy Preferences" [107]) has defined a standard to simplify users' data confidentiality policies of websites to allow people to understand how websites manage their data. These policies are then evaluated according to users' preferences by ad hoc mechanisms. Reaching the same goal, Inglese et al. [48] proposed a constrained natural language for the specification of authorization policies. Stepien et al. [102] worked on a non-technical notation for XACML policies.

Some approaches try to involve users in privacy management. Lederer et al. [59] proposed to improve users' understanding of the privacy implications by providing them feedback. Privacy mirrors [75] allow users to set their privacy controls and then check how their private data are seen from the point of view of the other people. Oglaza et al. [80] presented an approach based on multi-criteria decision support techniques for facilitating the process of writing authorization policies. The PrimeLife project published an interesting analysis on how to build human-computer interaction interfaces for privacy purpose [43].

4.3 Synthesis of works on privacy

We analyse in this section the state-of-the-art privacy-enhancing technologies with respect to context management activities as depicted in Fig. 1. These activities correspond to the first three types of Solove's taxonomy [101] which helps to deduce which PETs should be used [99].

During the collection of context data, anonymization techniques like k -anonymity can be used to protect data from being linked to the user. However, de-anonymization attacks, like those in [71], have demonstrated that such techniques are still not safe enough. Differential privacy [36] has therefore proposed to rather use perturbation techniques to make datasets indistinguishable, which is a promising research direction.

At the time of processing context data, anonymization can also be used together with identity management techniques to control data disclosure. As discussed in [30], traditional monolithic identity management solutions will slow down the informational self-emancipation required by privacy as transparency. On this matter, OpenID [82] and the friend-of-a-friend social principle [31] are promising approaches and should be further investigated. In addition, data minimization techniques like perturbation [2] and obfuscation can limit privacy violations caused by data aggregation and inference. One issue of current PETs for context data processing is to cope with the variety of anonymization and data protection mechanisms that can be used along the processing chain. This calls for new policy languages.

In the last step of context data presentation and dissemination towards context consumer applications, higher risks of privacy harms exist as stated by Solove [101]. Several techniques should be associated, encompassing confidentiality techniques, identity management and access and usage control policies.

Finally, solutions for transparent privacy are valuable in all the phases of context management. Users should have the choice of the data being collected and processed and should be informed of how they are used and for what purpose. There is a clear trend in the solutions for privacy as transparency to offer more control to users and further work is still urgently needed as learned from the PrimeLife project [43]. As shown in the recent comparison of 50 research projects on context-aware computing by Perera et al. [86], only 11 projects over 50 (about 20 %) did provide some security and privacy solutions. Privacy solutions are key to the success of future commercial context-aware services and applications in the IoT, and more work is required on this topic.

5 Privacy and QoC in the Internet of Things

We discuss in this section the research works that started to consider both privacy and QoC aspects at the same time. This opens the way for new enriched privacy solutions and for a better control of the QoC of context data. However, it appears that it also brings an additional complexity calling for efficient and effective middleware solutions for

implementing next-generation context managers able to tackle the new challenges raised by the IoT.

5.1 Works considering both privacy and QoC

Wishart et al. [110] propose privacy protection mechanisms able to disclose context information at different granularities. The disclosure requirements are expressed following a preference model, where the privacy preference states whether access is given to a certain type of context and the granularity preference indicates what maximum level of granularity is accepted. This work is promising in the use of obfuscation techniques; however, it only considers one QoC characteristic which is the granularity of context information.

Freytag [40] reviews research works which intend to preserve privacy in location-aware systems and shows that fuzzy location information does not automatically imply a loss of quality in the result. For Neisse [72] and Neisse et al. [74], QoC can be seen as a means to protect the user's privacy through the use of obfuscation techniques associated to the QoC level of the context information provided.

Chakraborty et al. [22–24] address the concept of behavioural privacy as opposed to traditional identity privacy and propose solutions to prevent the disclosure of some of the unintended inferences on user information. Their work relates to social networks, and they use a trust graph to identify the possible collusion possibilities between receivers and use it to determine how to adjust the quality of the data shared. The authors propose the criterion of resolution to designate the obfuscation level necessary to preserve privacy. It is a combination of several QoC criteria such as accuracy, precision, timeliness and completeness. With this work being dedicated to social networks with known relationships between context producers and consumers, the static approach it proposes for defining the trust graph makes it unapplicable in practice to the case of the IoT. The very high number and the variety of both context producers and context consumers in the IoT call for more dynamic solutions.

In addition, QoC raises new issues of confidentiality that are not yet addressed by current research. We illustrate them through three scenarios [61]:

Choosing the proper QoC level is not easy: Let us consider that John has a mobile phone equipped with positioning technology. He wants to share his location that is encoded using the Google address component types format [41]. John defines a policy to provide only the region where he is, which consists in sharing only the attribute of the *administrative_area_level_2* type according to the Google format. By using this mechanism, John believes

that nobody will be able to track him. However, when he is at the border of three regions R_1 , R_2 and R_3 , within a short period of time (10 min for instance), the location data history will include the three different values R_1 , R_2 and R_3 . A third-party system may then deduce that John is at the crossing border of the three regions. As a consequence, the actual level of detail is then much more precise than the one expected by John. In addition to this simple example, a recent study conducted by de Montjoye et al. [32] has proven that in an anonymized dataset “where location of an individual is recorded hourly and with a spatial resolution equal to that given by carrier antenna, four spatio-temporal points are enough to uniquely identify 95 % of the individuals”. The spatial resolution of antenna is from 0.15 to 15 km². They highlight the fact that “a point on the MIT campus at 3AM is more likely to make a trace unique than a point in downtown Boston on Friday evening”. These examples show that even low quality context data are useful for data mining algorithms and reliable QoC information will improve the efficiency of such algorithms.

QoC is sensitive information: The first step in preparing a security attack on networked systems is the “Reconnaissance phase” whose objective is to collect information about the target system in order to detect possible known vulnerabilities. For example, TCP/IP stack fingerprinting consists in collecting configuration values (e.g. the initial TTL and window size fields) from a remote device during standard network communications. The combination of parameters values may then be used to infer what is the remote machine’s operating system because different operating systems, and different versions of the same operating system, set different default values for these parameters. QoC might ease system fingerprinting if different systems set different default QoC values. As a consequence, QoC is also sensitive information that must be protected too.

QoC change is sensitive information: Context-aware computing in the IoT does not allow people to have the power to switch off the system or to easily disconnect from it if wanted. Some researchers have proposed to use the concept of *white lie* to provide people with this capability [3]. However, QoC will make white lying much more complex to perform. Let us consider the case of Mary who is a teenager who provides her location to her parents with a high degree of detail. This is Friday evening and Mary tells her parents that she is going to visit her grandmother. Actually, she is lying and wants to see her friends who live near her grandmother’s house. Thus, she uses the obfuscation mechanism that changes the granularity level of her location information. However, when her parents notice that the granularity level has changed, they can deduce that their daughter lied

to them by using algorithms for detecting changes [8]. Consequently, people cannot use white lies if QoC information is reliable. Thus, obfuscation mechanisms must consider that changing QoC level carries information or at least allow people to provide false QoC information.

5.2 Privacy, QoC and the IoT

While humans are an essential part of current privacy solutions, the need to broaden the scope of these solutions to the IoT has recently been acknowledged by the research community [11, 86, 97]. The IoT enables various kinds of communication patterns such as human to human, human to thing or thing to thing. It therefore offers opportunities for new emerging applications such as smart grid management, road traffic management, supply chain monitoring, crowd and participatory sensing applications or environmental control. These applications have in common relied on the fast-paced analysis of a huge amount of streaming data gathered from heterogeneous collections of sensory sources, in a loosely coupled unpredictable manner and possibly across multiple administrative domains [11]. Realizing such applications implies to revisit traditional privacy and QoC solutions.

In Section 4.2.3, we explained that privacy as transparency is far from being achieved due to the complexity for understanding and controlling the context management system. Associating QoC metadata to context makes the understanding and the control of the system even more complex since additional information has to be considered. One illustration of that point comes with the format of data. The granularity of location information using the Google address format (like “1600”, “Amphitheatre Parkway”, “Mountain View”, “CA”, “US”) is different from the granularity of geographic coordinates (like latitude 37.423021 and longitude -122.083739). Obfuscation mechanisms manipulating the QoC level of context data thus depend on the data format. With the address format, it is possible to remove an address attribute (ex: `keep type=administrative_area_level_1` and country would retain only “CA”, “US”). For geographic coordinates, a fake point should be calculated adding computing overhead. Because data format can be diverse especially with the heterogeneity of context sources in the IoT, there should be some standard QoC level with predefined values associated to context data and they should be easily interpretable by users. Solutions to master this complexity could benefit from mechanisms for improving users’ feedback interfaces as proposed by the EU FP7 uTRUSTit project [42].

A distributed IoT architecture, as compared by Roman et al. [93] to a centralised architecture, brings multiple benefits in terms of scalability, autonomy and applicability to the

real world. With such an architecture, all entities connected to the IoT have the ability to retrieve, process, combine and provide information and services to other entities. However, this requires that these entities have sufficient processing and storage capabilities to take part to the IoT. Calculating QoC, filtering on QoC and enabling obfuscation based on QoC require computational power. Additionally, most of the models for QoC-based access control are derived from RBAC which requires to store the list of the users associated to roles on the device. This calls for more studies to evaluate the feasibility and performance of IoT deployments where all the resources accessible from some device may be used in a cooperative manner like in the cloudlet approach [94] or with the cyber-foraging paradigm [39].

6 Conclusion

In this article, we analyse the maturity of privacy protection techniques to prepare context management for the IoT and envision that privacy solutions should take into account the quality of the personal context data manipulated. The IoT paradigm brings not only new opportunities by enabling enriched context-aware services but also new challenges faced by next-generation context management. We identify three main challenges, namely (1) the decoupling of the production and consumption of context data, (2) the need to enable dynamic QoC-aware privacy policies and (3) the complex interdependency of QoC and privacy.

QoC evaluation has been a growing research field over the last decade, and we establish a list of the main QoC criteria proposed in the literature. However, the relevance of some of these criteria must still be validated on concrete applications. Additionally, we underline that next-generation QoC management frameworks should be flexible for preserving performance and extensible for allowing to define new QoC criteria if needed. The different families of PETs are valuable to protect personal context data in the context management activities of context data collection, context data processing and context data presentation and dissemination. We, however, show that there are still remaining issues like the definition of new policy languages allowing to deal with the variety of anonymization and data protection mechanisms that can be used along the context data processing chain. Also, transparent privacy solutions are urgently needed for users to feel in control of their privacy and adhere to the IoT vision.

The few recent research initiatives exploring how to bridge privacy and QoC have started to identify some early solutions, but they are not yet sufficient to cater for the dynamicity and the various spatio-temporal scales of next-generation context management. New models, new

languages and new frameworks are required and imply to gather the various research communities of model engineering, knowledge and context management, security and privacy as targeted by the ongoing INCOME (<http://anr-income.fr>) project.

Acknowledgments This work is part of the French National Research Agency (ANR) project INCOME (ANR-11-INFR-009, 2012-2015).

References

1. Abid Z, Chabridon S, Conan D (2009) A framework for quality of context management. In: First international workshop on quality of context. Lecture notes in computer science, vol 5786. Springer, Berlin
2. Agrawal R, Srikant R (2000) Privacy-preserving data mining. In: ACM SIGMOD conference
3. Alcalde Bagüés S, Zeidler A, Fernández-Vakdivielso C, Matias I (2007) Disappearing for a while—using white lies in pervasive computing. In: Proceedings of the ACM workshop on privacy in electronic society, ACM, pp 80–83
4. Arcangeli J-P et al (2012) INCOME—multi-scale context management for the Internet of Things. In: International conference on ambient intelligence (AmI). Lecture notes in computer science, vol 7683. Springer, Berlin
5. Ashley P, Hada S et al (2003) Enterprise privacy authorization language (EPAL 1.2)
6. Atzori L, Iera A, Morabito G (2010) The Internet of Things: a survey. *Comput Netw* 54(15):2787–2805
7. Bai G, Gu L, Feng T, Guo Y, Chen X (2010) Context-aware usage control for android. In: Security and privacy in communication networks, Springer, New York, pp 326–343
8. Basseville M, Nikiforov I et al (1993) Detection of abrupt changes: theory and application, vol 104. Prentice-Hall, Englewood Cliffs
9. Bellavista P, Corradi A, Fanelli M, Foschini L (2012) A survey of context data distribution for mobile ubiquitous systems. *ACM Comput Surv* 44(24):24:1–24:45
10. Bettini C, Brdiczka O, Henriksen K, Indulska J et al (2010) A survey of context modelling and reasoning techniques. *Pervasive Mob Comp* 6(2):161–180
11. Bisdikian C, Sensoy M, Norman TJ, Srivastava MB (2012) Trust and obfuscation principles for quality of information in emerging pervasive environments. In: IEEE international conference on pervasive computing and communications, PerCom 2012, 19–23 March 2012, Lugano, workshop proceedings, pp 44–49
12. Borcea-Pfutzmann K, Pfutzmann A, Berg M (2011) Privacy 3.0 := data minimization + user control + contextual integrity. *Inf Technol* 53(1):34–40
13. Brgulja N, Kusber R, David K, Baumgarten M (2009) Measuring the probability of correctness of contextual information in context aware systems. In: 8th IEEE international conference on dependable, autonomic and secure computing, Washington
14. Bu Y, Gu T, Tao X, Li J, Chen S, Lu J (2006) Managing quality of context in pervasive computing. In: Sixth international conference on quality software, QSIC 2006
15. Buchholz T, Kupper A, Schiffers M (2003) Quality of context information: what it is and why we need it. In: 10th international workshop HPOVUA, Geneva

16. Canadian and Dutch Protection Authorities (1995) Privacy-enhancing technologies: the path to anonymity. <http://www.onlta.on.ca/library/repository/mon/10000/184530.pdf> Accessed 15 Feb 2013
17. Canny J (2002) Collaborative filtering with privacy via factor analysis. In: 25th ACM SIGIR
18. Castellucia C, Druschel P, Fischer Hübner S et al. (2011) Privacy, accountability and trust—challenges and opportunities. Technical report MSU-CSE-00-2, ENISA
19. Cavoukian A, Chibba M (2009) Advancing privacy and security in computing, networking and systems innovations through privacy by design. In: Proceedings conference of the Centre for Advanced Studies on Collaborative Research, Toronto pp 358–360
20. Cavoukian A, Tapscott D (1996) Who knows: safeguarding your privacy in a networked world. McGraw-Hill, New York
21. Chabridon S, Conan D, Abid Z, Taconet C (2012) Building ubiquitous QoC-aware applications through model-driven software engineering. *Sci Comput Program* 78:1912–1929. doi:10.1016/j.scico.2012.07.019
22. Chakraborty S, Charbiwala Z, Choi H, Raghavan KR, Srivastava MB (2012) Balancing behavioral privacy and information utility in sensory data flows. *Pervasive Mob Comput* 8(3): 331–345
23. Chakraborty S, Choi H, Srivastava MB (2011) Demystifying privacy in sensory data: a QoI based approach In: Percom workshops
24. Chakraborty S, Raghavan KR, Srivastava MB, Bisdikian C, Kaplan LM (2012) An obfuscation framework for controlling value of information during sharing. In: IEEE statistical signal processing workshop
25. Chaum D (1981) Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun ACM* 24(2):84–88
26. Conti M, Das SK, Bisdikian C, Kumar M et al (2012) Looking ahead in pervasive computing: challenges and opportunities in the era of cyber-physical convergence. *Pervasive Mob Comput* 8(1):2–21
27. Coutaz J, Crowley JL, Dobson S, Garlan D (2005) Context is key. *Commun ACM* 48(3):49–53
28. Covington M, Long W, Srinivasan S, Dey A, Ahamad M, Abowd G (2001) Securing context-aware applications using environment roles. In: 6th ACM symposium on access control models and technologies
29. Covington M, Sastry M (2006) A contextual attribute-based access control model. In: OTM
30. Danezis G, Gürses S (2010) A critical review of 10 years of privacy technology. In: Surveillance cultures: a global surveillance society?, UK
31. Danezis G, Mittal P (2009) SybilInfer: detecting sybil nodes using social networks. In: NDSS
32. de Montjoye Y-A, Hidalgo CA, Verleysen M, Blondel V (2013) Unique in the crowd: the privacy bounds of human mobility. *Nat Sci Rep* 3:1376
33. De Capitani di Vimercati S, Samarati P (2011) PrimeLife project: next generation policies. <http://primelife.ercim.eu/results/documents/150-523d> Accessed 15 Feb 2013
34. Diaz C (2005) Anonymity privacy in electronic services. PhD thesis, Cath. Univ. Leuven
35. Dingledine R, Mathewson N, Syverson PF (2004) Tor: the second-generation onion router. In: 13th USENIX security symposium, San Diego
36. Dwork C (2006) Differential privacy. In: International colloquium on automata, languages and programming (ICALP) Springer, Venice
37. Filho JB (2010) A family of context-based access control models for pervasive environments. PhD thesis, MSTII Doctoral School, Joseph Fourier University, Grenoble
38. Filho JB, Agoulmine N (2011) A quality-aware approach for resolving context conflicts in context-aware systems. *IEEE/IFIP conference on embedded and ubiquitous computing*
39. Flinn J (2012) Cyber foraging: bridging mobile and cloud computing. *Synth Lect Mob Pervasive Comput* 7(2):1–103
40. Freytag J-C (2009) Context quality and privacy—friends or rivals? In: First international workshop on quality of context. Lecture notes in computer science, vol 5786. Springer, Berlin
41. Google (2013) The Google geocoding API. <https://developers.google.com/maps/documentation/geocoding/> Accessed 15 Feb 2013
42. Graf C, Busch M, Schulz T, Hochleitner C, Skeide Fuglerud K (2012) D2.7 updated design guidelines on the security feedback provided by the “Things”. Technical report, uTRUSTIt project
43. Graf C, Hochleitner C et al (2011) Towards usable privacy enhancing technologies: lessons learned from the PrimeLife project. <http://primelife.ercim.eu/results/documents/149-416d> Accessed 15 Feb 2013
44. Han Q, Hakkarinen D, Boonma P, Suzuki J (2010) Quality-aware sensor data collection. *Int J Sens Netw* 7(3):127–140
45. Hedbom H (2009) A survey on transparency tools for enhancing privacy. In: The future of identity in the information society, vol 298. Springer, Berlin
46. Henriksen K, Indulska J (2004) Modelling and using imperfect context information. In: IEEE PERCOM 1st workshop CoMoRea
47. Huebscher MC, McCann JA (2005) A learning model for trustworthiness of context-awareness services. In: Third IEEE international conference on PerCom workshops, Hawaii
48. Inglesant P, Sasse MA, Chadwick D, Shi LL (2008) Expressions of expertness: the virtuous circle of natural language for access control policy specification. In: SOUPS
49. Accuracy ISO (2011) (Trueness and precision) of measurement methods and results—part 1: introduction and basic principles. ISO/WD 15725-1 document
50. ITU Internet Reports (2005) The Internet of Things, 7th edn. ITU, Geneva
51. Kelly D, Raines R, Grimaila M, Baldwin R, Mullins B (2008) A survey of state-of-the-art in anonymity metrics. In: 1st ACM workshop on network data anonymization. ACM
52. Kim Y, Lee K (2006) A quality measurement method of context information in ubiquitous environments, vol 2. In: ICHIT '06 proceedings of the 2006 international conference on hybrid information technology
53. Korpipää P, Mäntyjärvi J, Kela J, Keränen H, Malm EJ (2003) Managing context information in mobile devices. *IEEE Pervasive Comput* 2(3):42–51
54. Krause M, Hochstatter I (2005) Challenges in modelling and using quality of context (QoC). In: Mobility aware technologies and applications, vol 3744. Springer, Berlin
55. Kulkarni D, Tripathi A (2008) Context-aware role-based access control in pervasive computing systems. In: Proceedings of the 13th ACM symposium on access control models and technologies, SACMAT '08, ACM, New York, pp 113–122
56. Kumar A, Karnik NM, Chafle G (2002) Context sensitivity in role-based access control. *Oper Syst Rev* 36(3):53–66
57. Langheinrich M (2009) Privacy in ubiquitous computing. In: Krumm J (ed) Ubiquitous computing. CRC, Boca Raton, pp 95–160
58. Lazouski A, Martinelli F, Mori P (2010) Usage control in computer security: a survey. *Elsevier Comput Sci Rev* 4(2):81–99

59. Lederer S, Hong J, Dey A, Landay J (2004) Personal privacy through understanding and action: five pitfalls for designers. *Personal Ubiquit Comput* 8(6):440–454
60. Machanavajjhala A, Kifer D, Gehrke J, Venkatasubramanian M (2007) L-diversity: privacy beyond k -anonymity. *ACM Trans Knowl Discov Data* 1(3):3:1–3:52
61. Machara Marquez S, Chabridon S, Taconet C (2013) Models@Run.time for privacy and quality of context level agreements in the Internet of Things. Technical report, UMR SAMOVAR, Télécom SudParis
62. Manzoor A (2010) Quality of context in pervasive systems; models, techniques, and applications. PhD thesis, School of Computer Science, Wien TU
63. Manzoor A, Truong H-L, Dustdar S (2012) Quality of context: models and applications for context-aware systems in pervasive environments. *Knowl Eng Rev*. doi:10.1017/S0000000000000000. Special issue on web and mobile information services
64. Manzoor A, Truong HL, Dustdar S (2008) On the evaluation of quality of context. In: *Smartsensing and context*, Springer, Berlin
65. Marie P, Desprats T, Chabridon S, Sibilla M (2013) A meta-model for the management of the quality of context information. Technical report, University, Toulouse, IRIT
66. Marx G (2001) Murky conceptual waters: the public and the private. *Ethics Inf Technol* 3(3):157–169
67. Matos A (2012) Privacy in next generation networks. PhD thesis. <http://ria.ua.pt/handle/10773/8697> Accessed 15 Feb 2013
68. McKeever S, Ye J, Coyle L, Dobson S (2009) A context quality model to support transparent reasoning with uncertain context. In: *First international workshop on quality of context*. Lecture notes in computer science, vol 5786. Springer, Berlin
69. Mehta B (2007) Learning from what others know: privacy preserving cross system personalization. In: *11th conferences on user modeling*
70. Miorandi S, an Sicari D, De Pellegrini F, Chlamtac I (2012) Survey Internet of things: vision, applications and research challenges. *Ad Hoc Netw* 10(7):1497–1516
71. Narayanan A, Shmatikov V (2008) Robust De-anonymization of large sparse datasets. In: *IEEE symposium security and privacy*
72. Neisse R (2012) Trust and privacy management support for context-aware service platforms. PhD thesis, CTIT School, University of Twente, NL
73. Neisse R, Pretschner A, Di Giacomo V (2011) A trustworthy usage control enforcement framework. In: *6th international conference on ARES*
74. Neisse R, Wegdam M, van Sinderen M (2008) Trustworthiness and quality of context information. In: *9th conference for young computer scientists*, Hunan
75. Nguyen DH, Mynatt ED (2002) Privacy mirrors: understanding and shaping socio-technical ubiquitous computing systems. Technical report GIT-GVU-02-16, Georgia Techniques, Atlanta
76. OASIS (2012) Extensible access control markup language (XACML). <http://www.oasis-open.org/committees/xacml/> Accessed 15 Feb 2013
77. OECD (1980) Guidelines on the protection of privacy and trans-border flows of personal data
78. Official Journal of the European Communities (2002) EU Directive 2002/58/ec on the processing of personal data and the protection of privacy in the electronic communications sector
79. Official Journal of the European Communities (1995) EU Directive 95/46/ec on the protection of individuals with regard to the processing of personal data and the free movement of such data
80. Oglaza A, Laborde R, Zarate P (2013) Authorization policies: using decision support system for context-aware protection of user's private data. In: *IEEE international symposium on UbiSafe computing*
81. Okagawa T, Nishida K, Miura A (2003) A proposed routing procedure in IP2. In: *IEEE 58th VTC*, vol 3
82. OpenID (2013) Foundation website. <http://openid.net/> Accessed 15 Feb 2013
83. Paquin C (2011) U-prove technology overview. <http://research.microsoft.com/apps/pubs/default.aspx?id=166980>
84. Pearson S (2012) Privacy management in global organisations. In: *Communications and multimedia security*, Springer
85. Pearson S, Casassa Mont M (2011) Sticky policies: an approach for managing privacy across multiple parties. *IEEE Comput* 44(9):60–68
86. Perera C, Zaslavsky A, Christen P, Georgakopoulos D (2013) Context aware computing for the internet of things: a survey. *Commun Surv Tutor*, IEEE PP(99):1–41. doi:10.1109/SURV.2013.042313.00197. ISSN 1553-877X
87. Pfitzmann A, Waidner M (1985) Networks without user observability: design options. In: *Advances in cryptology—EUROCRYPT*, workshop on the theory and application of cryptographic techniques, Linz, Austria. Lecture notes in computer science, vol 219. Springer, Berlin, pp 245–253
88. Preuveneers D, Berbers Y (2006) Quality extensions and uncertainty handling for context ontologies. In: *Proceedings of context and ontologies: theory practice and applications*, Italy
89. Preuveneers D, Berbers Y (2007) Architectural backpropagation support for managing ambiguous context in smart environments. In: *Fourth conference on universal access in HCI*. Lecture notes in computer science, vol 4555. Springer, Berlin
90. Privacy by Design Resolution Data protection and 32nd conference of privacy commissioners (2010) <http://www.privacybydesign.ca/content/uploads/2010/11/pbd-resolution.pdf>. Jerusalem, Israel Accessed 15 Feb 2013
91. Ranganathan A, Al-Muhtadi J, Campbell RH (2004) Reasoning about uncertain contexts in pervasive computing environments. *IEEE Pervasive Comput* 3(2):10–18
92. Roman R, Najera P, Lopez J (2011) Securing the Internet of Things. *IEEE Comput* 44(9):51–58
93. Roman R, Zhou J, Lopez J (2013) On the features and challenges of security and privacy in distributed Internet of Things. *Comput Netw* 57(10):2266–2279
94. Satyanarayanan M, Bahl P, Caceres R, Davies N (2009) The case for VM-based cloudlets in mobile computing. *IEEE Pervasive Comput* 8:14–23
95. Schmid S, Eggert L, Brunner M, Quittek J (2005) TurfNet: an architecture for dynamically composable networks. In: *Autonomic communication*. Lecture notes in computer science, vol 3457. Springer, Berlin
96. Schmidt A (2006) Ontology-based user context management: the challenges of imperfection and time-dependence. In: *Conference on ontologies, databases and applications*. Lecture notes in computer science, vol 4275. Springer, Berlin
97. Schrammel J, Hochleitner C, Tscheligi M (2011) Privacy, trust and interaction in the Internet of Things. In: *Ambient intelligence*. Lecture notes in computer science, vol 7040. Springer, Berlin, pp 378–379
98. Sheikh K, Wegdam M, Sinderen MV (2008) Quality-of-context and its use for protecting privacy in context-aware systems. *J Softw* 3(3):83–93
99. Shen Y, Pearson S (2011) Privacy enhancing technologies: a review. Technical report HPL-2011-113, HP Labs
100. Shibboleth (2013) Consortium website. <http://shibboleth.net/> Accessed 15 Feb 2013

101. Solove D. (2006) A taxonomy of privacy. *Univ Pennsylvania Law Rev* 153(3):477
102. Stepien B, Matwin S, Felty A (2011) Advantages of a non-technical XACML notation in role-based models. In: 9th annual international conference on privacy, security and trust (PST), pp 193–200
103. Sweeney L (2002) k-anonymity: a model for protecting privacy. *J Uncertain Fuzziness Knowl Based Syst* 10(5):557–570
104. Toch E, Wang Y, Cranor L (2012) Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Model User-Adap Inter* 22(1–2):203–220
105. Twidle K, Dulay N, Lupu E, Sloman M (2009) Ponder2: a policy system for autonomous pervasive environments. In: IEEE workshop on policies for distributed systems and networks
106. Van Blarckom GW, Borking JJ, Olk JGE (2003) Handbook of privacy and privacy-enhancing technologies: the case of intelligent softwares. College Bescherming Persoonsgegevens, The Hague
107. W3C (2011) The platform for Privacy Preferences (P3P) Project. <http://www.w3.org/P3P/> Accessed 15 Feb 2013
108. Wang Y, Kobza A (2008) Handbook of research on social and organizational liabilities in information security, chapter privacy enhancing technology. IGI Publishing, Hershey
109. Warren SD, Brandeis LD (1890) The right to privacy. *Harvard Law Rev* 4(5):193–220
110. Wishart R, Henriksen K, Indulska J (2005) Context obfuscation for privacy via ontological descriptions. In: 1st international workshop on location and context-awareness (LoCA). Lecture notes in computer science, vol 3479. Springer, Berlin
111. Yasar A-U-H, Paridel K, Preuveneers D, Berbers Y 2011
112. Ylitalo J, Nikander P (2006) BLIND: a complete identity protection framework for end-points. In: Security protocols. Lecture notes in computer science, vol 3957. Springer, Berlin