IN PARTNERSHIP WITH:
CNRS

Université Versailles Saint-Quentin

Activity Report 2015

Project-Team SMIS

Secured and Mobile Information Systems

IN COLLABORATION WITH: Parallelisme, réseaux, systèmes, modélisation (PRISM)

RESEARCH CENTER
Paris - Rocquencourt

THEME
Data and Knowledge Representation and Processing

Table of contents

Project-Team SMIS

Creation of the Project-Team: 2004 September 01

Keywords:

Computer Science and Digital Science:

1.1.8. - Security of architectures
1.4. - Ubiquitous Systems
3.1.2. - Data management, quering and storage
3.1.3. - Distributed data
3.1.5. - Control access, privacy
3.1.6. - Query optimization
3.1.8. - Big data (production, storage, transfer)
3.1.9. - Database
4.7. - Access control
4.8. - Privacy-enhancing technologies

Other Research Topics and Application Domains:

2.1. - Well being
2.3. - Epidemiology
2.5.3. - Assistance for elderly
6.4. - Internet of things
6.5. - Information systems
8.1. - Smart building/home
8.5. - Smart society
9.10. - Ethics
9.8. - Privacy

1. Members

Research Scientists

Nicolas Anciaux [Inria, Researcher, HdR]
Luc Bouganim [Inria, Senior Researcher, HdR]

Faculty Members

Philippe Pucheral [Team leader, Univ. Versailles, Professor, HdR]
Iulian Sandu Popa [Univ. Versailles, Associate Professor]

Engineers

Aydogan Ersoz [Inria]
Quentin Lefebvre [Inria, until Aug 2015, granted by ANR KISS project]

PhD Students

Chao Chen [Inria, from Dec 2015]
Athanasia Katsouraki [Inria]
Saliha Lallali [Inria, granted by ANR KISS project]
Quoc Cuong To [Inria, until Oct 2015]
Matias Bjørling [Univ. of Copenhagen, co-supervision, until Aug 2015]
Niv Dayan [Univ. of Copenhagen, co-supervision, until Aug 2015]

Paul Tran Van [CozyCloud, CIFRE]

Visiting Scientist

Benjamin Nguyen [INSA CVL, Professor, HdR]

Administrative Assistant

Laurence Bourcier [Inria]

2. Overall Objectives

2.1. Overall Objectives

The research work within the project-team is devoted to the design and analysis of core database techniques dedicated to the definition of secured and mobile information systems.

Ubiquitous computing and ambient intelligence entail embedding data in increasingly light and specialized devices (chips, sensors and electronic appliances for smart buildings, telephony, transportation, health, etc.). These devices exhibit severe hardware constraints to match size, security, power consumption and also production costs requirements. At the same time, they could highly benefit from embedded database functionalities to store data, analyze it, query it and protect it. This raises a first question "Q_1: *How to make powerful data management techniques compatible with highly constrained hardware platforms?*". To tackle this question, SMIS contributes to the design and validation of new storage and indexing models, query execution and optimization techniques, and transaction protocols. The relevance of this research goes beyond embedded databases and may have potential applications for database servers running on advanced hardware.

By making information more accessible and by multiplying –often transparently– the means of acquiring it, ubiquitous computing involves new threats for data privacy. The second question addressed by the project-team is then "Q_2: *How to make smart objects less intrusive?*". New access and usage control models have to be devised to help individuals keep a better control on the acquisition and sharing conditions of their data. This means integrating privacy principles like user's consent, limited collection and limited retention in the access and usage control policy definition. This also means designing appropriate mechanisms to enforce this control and provide accountability with strong security guarantees.

In parallel, thanks to a high degree of decentralization and to the emergence of low cost tamper-resistant hardware, ubiquitous computing contains the seeds for new ways of managing personal/sensitive data. The third question driving the research of the project-team is therefore "Q_3: *How to build privacy-by-design architectures based on trusted smart objects?*". The objective is to capitalize on embedded data management techniques, privacy-preserving mechanisms, trusted devices and cryptographic protocols to define an integrated framework dedicated to the secure management of personal/sensitive data. The expectation is showing that credible alternatives to a systematic centralization of personal/sensitive data on servers can be devised and validating the approach through real case experiments.

3. Research Program

3.1. Embedded Data Management

The challenge tackled is this research action is twofold: (1) to design embedded database techniques matching the hardware constraints of (current and future) smart objects and (2) to set up co-design rules helping hardware manufacturers to calibrate their future platforms to match the requirements of data driven applications. While a large body of work has been conducted on data management techniques for high-end servers (storage, indexation and query optimization models minimizing the I/O bottleneck, parallel DBMS, main memory DBMS, etc.), less research efforts have been placed on embedded database techniques. Light versions of popular DBMS have been designed for powerful handheld devices; yet DBMS vendors have never addressed the complex problem of embedding database components into chips. Proposals dedicated to databases embedded

on chip usually consider small databases, stored in the non-volatile memory of the microcontroller –hundreds of kilobytes– and rely on NOR Flash or EEPROM technologies. Conversely, SMIS is pioneering the combination of microcontrollers and NAND Flash constraints to manage Gigabyte(s) size embedded databases. We present below the positioning of SMIS with respect to international teams conducting research on topics which may be connected to the addressed problem, namely work on electronic stable storage, RAM consumption and specific hardware platforms.

Major database teams are investigating data management issues related to hardware advances (EPFL: A. Ailamaki, CWI: M. Kersten, U. Of Wisconsin: J. M. Patel, Columbia: K. Ross, UCSB: A. El Abbadi, IBM Almaden: C. Mohan, etc.). While there are obvious links with our research on embedded databases, these teams target high-end computers and do not consider highly constrained architectures with non traditional hardware resources balance. At the other extreme, sensors (ultra-light computing devices) are considered by several research teams (e.g., UC Berkeley: D. Culler, ITU: P. Bonnet, Johns Hopkins University: A. Terzis, MIT: S. Madden, etc.). The focus is on the processing of continuous streams of collected data. Although the devices we consider share some hardware constraints with sensors, the objectives of both environments strongly diverge in terms of data cardinality and complexity, query complexity and data confidentiality requirements. Several teams are looking at efficient indexes on flash (HP LABS: G. Graefe, U. Minnesota: B. Debnath, U. Massachusetts: Y. Diao, Microsoft: S. Nath, etc.). Some studies try to minimize the RAM consumption, but the considered RAM/stable storage ratio is quite large compared to the constraints of the embedded context. Finally, a large number of teams have focused on the impact of flash memory on database system design (we presented an exhaustive state of the art in a VLDB tutorial [7]). The work conducted in the SMIS team on bi-modal flash devices takes the opposite direction, proposing to influence the design of flash devices by the expression of database requirements instead of running after the constantly evolving flash device technology.

3.2. Access and Usage Control Models

Access control management has been deeply studied for decades. Different models have been proposed to declare and administer access control policies, like DAC, MAC, RBAC, TMAC, and OrBAC. While access control management is well established, new models are being defined to cope with privacy requirements. Privacy management distinguishes itself from traditional access control is the sense that the data to be protected is personal. Hence, the user's consent must be reflected in the access control policies, as well as the usage of the data, its collection rules and its retention period, which are principles safeguarded by law and must be controlled carefully.

The research community working on privacy models is broad, and involves many teams worldwide including in France ENST-B, LIRIS, Inria LICIT, and LRI, and at the international level IBM Almaden, Purdue Univ., Politecnico di Milano and Univ. of Milano, George Mason Univ., Univ. of Massachusetts, Univ. of Texas and Colorado State Univ. to cite a few. Pioneer attempts towards privacy wary systems include the P3P Platform for Privacy Preservation [40] and Hippocratic databases [37]. In the last years, many other policy languages have been proposed for different application scenarios, including EPAL [43], XACML [42] and WSPL [38]. Hippocratic databases are inspired by the axiom that databases should be responsible for the privacy preservation of the data they manage. The architecture of a Hippocratic database is based on ten guiding principles derived from privacy laws.

The trend worldwide has been to propose enhanced access control policies to capture finer behavior and bridge the gap with privacy policies. To cite a few, Ardagna *et al.* (Univ. Milano) enables actions to be performed after data collection (like notification or removal), purpose binding features have been studied by Lefevre *et al.* (IBM Almaden), and Ni *et al.* (Purdue Univ.) have proposed obligations and have extended the widely used RBAC model to support privacy policies.

The positioning of the SMIS team within this broad area is rather (1) to focus on intuitive or automatic tools helping the individual to control some facets of her privacy (e.g., data retention, minimal collection) instead of increasing the expressiveness but also the complexity of privacy models and (2) to push concrete models enriched by real-case (e.g., medical) scenarios and by a joint work with researchers in Law.

3.3. Tamper-resistant Data Management

Tamper-resistance refers to the capacity of a system to defeat confidentiality and integrity attacks. This problem is complementary to access control management while being (mostly) orthogonal to the way access control policies are defined. Security surveys regularly point out the vulnerability of database servers against external (i.e., by intruders) and internal (i.e., by employees) attacks. Several attempts have been made in commercial DBMSs to strengthen server-based security, e.g., by separating the duty between DBA and DSA (Data Security Administrator), by encrypting the database footprint and by securing the cryptographic material using Hardware Security Modules (HSM) [39]. To face internal attacks, client-based security approaches have been investigated where the data is stored encrypted on the server and is decrypted only on the client side. Several contributions have been made in this direction, notably by U. of California Irvine (S. Mehrotra, Database Service Provider model), IBM Almaden (R. Agrawal, computation on encrypted data), U. of Milano (E. Damiani, encryption schemes), Purdue U. (E. Bertino, XML secure publication), U. of Washington (D. Suciu, provisional access) to cite a few seminal works. An alternative, recently promoted by Stony Brook Univ. (R. Sion), is to augment the security of the server by associating it with a tamper-resistant hardware module in charge of the security aspects. Contrary to traditional HSM, this module takes part in the query computation and performs all data decryption operations. SMIS investigates another direction based on the use of a tamper-resistant hardware module on the client side. Most of our contributions in this area are based on exploiting the tamper-resistance of secure tokens to build new data protection schemes.

While our work on Privacy-Preserving data Publishing (PPDP) is still related to tamper-resistance, a complementary positioning is required for this specific topic. The primary goal of PPDP is to anonymize/sanitize microdata sets before publishing them to serve statistical analysis purposes. PPDP (and privacy in databases in general) is a hot topic since 2000, when it was introduced by IBM Research (IBM Almaden: R. Agrawal, IBM Watson: C.C. Aggarwal), and many teams, mostly north American universities or research centres, study this topic (e.g., PORTIA DB-Privacy project regrouping universities such as Stanford with H. Garcia-Molina). Much effort has been devoted by the scientific community to the definition of privacy models exhibiting better privacy guarantees or better utility or a balance of both (such as differential privacy studied by C. Dwork: Microsoft Research or D. Kifer: Penn-State Univ and J. Gehrke: Cornell Univ) and thorough surveys exist that provide a large overview of existing PPDP models and mechanisms [41]. These works are however orthogonal to our approach in that they make the hypothesis of a trustworthy central server that can execute the anonymization process. In our work, this is not the case. We consider an architecture composed of a large population of tamper-resistant devices weakly connected to an untrusted infrastructure and study how to compute PPDP problems in this context. Hence, our work has some connections with the works done on Privacy Preserving Data Collection (Stevens Institute of Tech. / Rutgers Univ,NJ: R.N.Wright, Univ Austin Texas: V. Shmatikov), on Secure Multi-party Computing for Privacy Preserving Data Mining (Rutgers Univ: J. Vaidya, Purdue Univ: C. Clifton) and on distributed PPDP algorithms (Univ Wisconsin: D. DeWitt, Univ Michigan: K. Lefevre, Rutgers Univ: J. Vaidya, Purdue Univ: C. Clifton) while none of them share the same architectural hypothesis as us.

4. Application Domains

4.1. Application Domains

Our work addresses varied application domains. Typically, data management techniques on chip are required each time data-driven applications have to be embedded in ultra-light computing devices. This situation occurs for example in healthcare applications where medical folders are embedded into smart tokens (e.g., smart cards, secured USB keys), in telephony applications where personal data (address book, agenda, etc.) is embedded into cellular phones, in sensor networks where sensors log row measurements and perform local computation on them, in smart-home applications where a collection of smart appliances gather information about the occupants to provide them a personalized service, and more generally in most applications related to ambient intelligence.

Safeguarding data confidentiality has become a primary concern for citizens, administrations and companies, broadening the application domains of our work on access control policies definition and enforcement. The threat on data confidentiality is manifold: external and internal attacks on the data at rest, on the data on transit, on the data hosted in untrusted environments (e.g., Database Service Providers, Web-hosting companies) and subject to illegal usage, insidious gathering of personal data in an ambient intelligence surrounding. Hence, new access control models and security mechanisms are required to accurately declare and safely control who is granted access to which data and for which purpose.

While the application domain mentioned above is rather large, two applications are today more specifically targeted by the SMIS team. The first one deals with privacy preservation in EHR (Electronic Health Record) systems and PCEHR (Personally Controlled EHR). We are developing technologies tackling this issue and experiment them in the field. The second application area deals with privacy preservation in the context of personal Cloud, that is personal data hosted in dedicated servers staying under the holder's control (e.g., in a personal internet box or in a home automation box).

5. New Software and Platforms

5.1. PLUG-DB ENGINE

FUNCTIONAL DESCRIPTION

PlugDB is a complete platform dedicated to a secure and ubiquitous management of personal data. It aims at providing an alternative to a systematic centralization of personal data. The PlugDB engine is a personal database server capable of storing data (tuples and documents) in tables and BLOBs, indexing them, querying them in SQL, sharing them through assertional access control policies and enforcing transactional properties (atomicity, integrity, durability). The PlugDB engine is embedded in a tamper-resistant hardware device combining the security of smartcard with the storage capacity of NAND Flash. The personal database is hosted encrypted in NAND Flash and the PlugDB engine code runs in the microcontroller. Complementary modules allow to pre-compile SQL queries for the applications, communicate with the DBMS from a remote Java program, synchronize local data with remote servers (typically used for recovering the database in the case of a broken or lost devices) and participate in distributed computation (e.g., global queries). PlugDB runs both on secure devices provided by Gemalto and on specific secure devices designed by SMIS and assembled by electronic SMEs. Mastering the hardware platform opens up new research and experiment opportunities (e.g., we have recently integrated a Bluetooth module to communicate wirelessly with PlugDB and a fingerprint module to strongly authenticate users) and allows us to engage ourselves in an open-source/open hardware initiative. Open-SW/open-HW contributes to the trust the community of users can put in any privacy preserving solution and is key to enable a diversity of solutions, hence decreasing the risk of class attacks. PlugDB engine has been registered first at APP (Agence de Protection des Programmes) in 2009 - a new version being registered every two years and the hardware datasheets in 2015. PlugDB has been experimented in the field - notably in the healthcare domain - and we recently set up an educational platform to raise students awareness of privacy protection problems and embedded programming. As a conclusion, PlugDB combines several research contributions from the team, at the crossroads of flash data management, embedded data processing and secure distributed computations. It then strongly federates all members of our team (permanent members, PhD students and engineers). It is also a vector of visibility, technological transfer and dissemination and gives us the opportunity to collaborate with researchers from other disciplines around a concrete privacy enhancing platform.

- Participants: Nicolas Anciaux, Luc Bouganim, Philippe Pucheral, Shaoyi Yin, Yanli Guo, Kevin Jacquemin, Aydogan Ersoz and Quentin Lefebvre
- Contact: Nicolas Anciaux
- URL: https://project.inria.fr/plugdb/

6. New Results

6.1. Embedded Data Management

Participants: Nicolas Anciaux, Saliha Lallali, Philippe Pucheral, Iulian Sandu Popa [correspondent].

Embedded keyword indexing: In this work, we revisit the traditional problem of information retrieval queries over large collections of files in an embedded context. A file can be any form of document, picture or data stream, associated with a set of terms. A query can be any form of keyword search using a ranking function (e.g., TF-IDF) identifying the top-k most relevant files. The proposed search engine can be used in sensors to search for relevant objects in their surroundings, in cameras to search pictures by using tags, in personal smart dongles to secure the querying of documents and files hosted in an untrusted Cloud, or in a personal cloud securely managed using a tamper resistant smart object. A search engine is usually based on a (large) inverted index and queries are traditionally evaluated by allocating one container in RAM per document to aggregate its score, making the RAM consumption linear with the size of the document corpus. To tackle this issue, we designed a new form of inverted index which can be accessed in a pure pipeline manner to evaluate search queries without materializing any intermediate result. Successive index partitions are written once in Flash and maintained in the background by timely triggering merge operations while files are inserted or deleted from the index. By combining this new index and the corresponding evaluation techniques, our embedded search engine is capable of reconciling high insert/delete/update rate and query scalability. We have demonstrated the search engine on a secure USB token in the context of a personal cloud, and have conducted in depth performance evaluations on a development board representative for different smart objects characteristics. The experimental results demonstrate the scalability of the approach and its superiority compared to state of the art methods. This work was published at VLDB'15 [21] and demonstrated at SIGMOD'15 [24]. It constitutes the main contribution of the PhD thesis of Saliha Lallali

Spatio-temporal indexing in Flash storage: The convergence of mobile computing, wireless communications and sensors has raised the development of many applications exploiting massive flows of spatio-temporal data such as in location-based services, participatory sensing, or traffic management [15]. Spatio-temporal data indexing is among the most active research topics in this area. Nevertheless, since a few years a new fundamental parameter has made its entry on the database scene: the NAND flash storage. The peculiar characteristics of flash memory require redesigning the existing data storage and indexing techniques that were devised for magnetic hard-disks. TRIFL, proposed in [16] is an efficient and generic TRajectory Index for FLash, designed around the key requirements of both trajectory indexing and flash storage. TRIFL is generic in the sense that it is efficient for both simple flash storage devices such as the SD cards and more powerful devices such as the solid state drives. In addition, TRIFL includes an online self tuning algorithm that allows adapting the index structure to the workload and the technical specifications of the flash storage device to maximize the index performance. Moreover, TRIFL achieves good performance with relatively low memory requirements, making it appropriate for many application scenarios. The experimental evaluation shows that TRIFL outperforms the representative indexing methods on flash disks but also on magnetic disks. This work [15] [16] is part of Dai Hai Ton That's Ph.D. thesis, co-supervised by Iulian Sandu Popa.

6.2. Secure Global Computing on Asymmetric Architecture

Participants: Benjamin Nguyen [correspondent], Philippe Pucheral, Quoc Cuong To.

Asymmetric Architecture Computing: This research direction studies the secure execution of various algorithms on data stored in an unstructured network of Trusted Cells (i.e., personal trusted device) so that each user can keep control over her data. The data could be stored locally in a trusted cell or encrypted on some external cloud. Execution takes place on a specific infrastructure called the Asymmetric Architecture: the network of trusted cells, supported by an untrusted cloud supporting IaaS or PaaS. Our objective is to show that many different algorithms and computing paradigms can be executed on the Asymmetric Architecture, thus achieving secure and private computation. Our first contribution in this area was to study the execution of Privacy Preserving Data Publishing (PPDP) algorithms on such an architecture, and provided generic

protocols to deal with all kinds of PPDP algorithms, which are robust against honest-but-curious and malicious adversaries [2][3]. Our second contribution was to study general SQL queries in this same execution context. We concentrated on the subset of SQL queries without joins, but including Group By and aggregates, and show how to secure their execution in the presence of honest-but-curious attackers [9]. This work was part of Quoc-Cuong To's Ph.D defended in 2015 [13]. We are extending this general framework through a collaboration with INSA Centre Val de Loire, LIFO Lab and University of Paris Nord, LIPN lab, to study the secure execution of Map/Reduce on the Asymmetric Architecture. Computing MapReduce processes on the Asymmetric Architecture means maintaining the flexibility and efficiency of MapReduce, while adding security into the mix. We have shown in [25] that it is possible to achieve seamless integration of distributed MapReduce processing using trusted cells, while maintaining reasonable performance.

Secure spatio-temporal distributed processing: Mobile participatory sensing could be used in many applications such as vehicular traffic monitoring, pollution tracking, or even health surveying (e.g., to allow measuring in real-time the individual exposure to environmental risk factors or the propagation of an epidemic). However, its success depends on finding a solution for querying a large number of users which protects user location privacy and works in real-time. We addressed these issues and proposed PAMPAS, a privacy-aware mobile distributed system for efficient data aggregation in mobile participatory sensing. In PAMPAS, mobile devices enhanced with secure hardware, called secure probes, perform distributed query processing, while preventing users from accessing other users' data. Secure probes exchange data in encrypted form with help from an untrusted supporting server infrastructure. PAMPAS uses two efficient, parallel, and privacy-aware protocols for location-based aggregation and adaptive spatial partitioning of secure probes. Our experimental results and security analysis demonstrate that these protocols are able to collect, aggregate and share statistics or derived data in real-time, without any privacy leakage. This work is part of Dai Hai Ton That's Ph.D. thesis, co-supervised by Iulian Sandu Popa. The system implementation was demonstrated in [26], and a paper describes the technical details of the system [31].

6.3. Personal Cloud

Participants: Nicolas Anciaux [correspondent], Luc Bouganim, Athanasia Katsouraki, Benjamin Nguyen, Philippe Pucheral, Iulian Sandu Popa, Paul Tran Van.

We are witnessing an exponential increase in the acquisition of personal data about the individuals or produced by them. Today, this information is managed using Web applications, centralizing this data in cloud data servers, under the control of few Web majors [4]. However, it has now become clear that (1) centralizing millions of personal records exposes the data to very sophisticated attacks, linked to a very high potential benefit in case of success (millions of records being revealed), and (2) delegating the management of personal records without any tangible guarantee for the individuals leads to privacy violations, the data being potentially made accessible to other organizations (e.g., governments, commercial partners) and being subject to lucrative secondary usages (not advertised to the individuals). To face this situation, many recent initiatives push towards the emergence of the Personal Cloud paradigm. A personal cloud can be viewed as a personal server, owned by a given individual, which gives to its owner the ability to store her complete digital environment, synchronize it among various devices and share it with other individuals and applications under control. In the SMIS team, we claim the need of a Secure Personal Cloud, and promote the introduction of a secure (tamper resistant) data engine in the architecture [1]. On this basis, we investigate new data sharing and dissemination models, where usage and access control rules endorsed by the individuals could be enforced and have presented this vision at EDBT'14 and at ADBIS'15 [18]. We have started a cooperation with the startup CozyCloud at the end of 2014. A contract was signed at the end of 2014 to integrate PlugDB in a CozyCloud instance and the PhD of Paul Tran Van (CIFRE SMIS-CozyCloud) has started to explore new data sharing techniques which could be enforced in the secure personal cloud model. A second PhD CIFRE SMIS-CozyCloud is being submitted to explore privacy-preserving distributed computations over personal clouds. Athanasia Katsouraki is working on privacy issues and on adoption of the secure data engine [29] in cooperation with the economists (CERDI) in the context of the Digital Society Institute (DSI). A paper written by jurists, economists and computer scientists from DSI has been invited for publication in Legicom'2016 to present our common vision of Privacy-by-Design principles in the context of Open Data and Internet of Things.

6.4. Applications

Participants: Nicolas Anciaux [correspondent], Luc Bouganim, Philippe Pucheral.

In 2014, we proposed a new paradigm, that we call Folk-enabled Information System (Folk-IS), based on a fully decentralized and participatory approach, where each individual implements a small subset of a complete information system without the need for a shared networked infrastructure [5]. Folk-IS builds upon the emergence of highly secure, portable and low-cost storage and computing devices, called hereafter Smart Tokens. Here however, the focus is on low-cost of ownership, deployment and maintenance, and on the absence of a networked infrastructure. With Folk-IS and thanks to their smart tokens, people will transparently and opportunistically perform data management and networking tasks as they physically move, so that IT services are truly delivered by the crowd. Following this work, we collaborate with researchers and doctors from Cameroon to study the specific case of diabete follow-up. Indeed, there are currently more than half a million diabetes cases in Cameroon and the deaths caused by diabetes complications will double before 2030. Diabetes complications mostly occur due to a bad follow-up of patients. Based on an analysis of the current situation, we proposed a new IT architecture for diabetes follow-up and introduce the bases of a new distributed computation protocol for this architecture. Our approach does not require any preexisting support communication infrastructure, can be deployed at low cost, and provides strong privacy and security guarantees. This work, published in AFRICOM [20] envisions an experiment in the field we plan to conduct under the authority of the Cameroonian National Center for Diabetes and Hypertension, with a potential for generalization to other diseases.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

The SMIS project has a long lasting cooperation with Gemalto, the world's leading providers of microprocessor cards. Gemalto provides SMIS with advanced hardware and software smart card platforms which are used to validate numbers of our research results. In return, SMIS provides Gemalto with requirements and technical feedbacks that help them adapting their future platforms towards data intensive applications. While no bilateral contract exists between Gemalto and SMIS, we are partners in several projects. Meanwhile, we are developing partnerships with SMEs capable of building ad-hoc hardware prototypes conforming to our own design.

7.1.1. Cozy Cloud bilateral contract (Dec 2014 - Nov. 2015)

Partners: Cozy Cloud, Inria-SMIS
SMIS funding: 50k€.
Many personal data end up today on servers where they can be scrutinized by companies and governmental agencies. To face this situation, the most emblematic initiative is the Personal Cloud paradigm. Roughly speaking, the Personal Cloud is an architecture which gives users the ability to store their complete digital environment, synchronize it among various devices and share it with other users and applications under their control. It reflects the expectation of the individuals for the emergence of privacy-by-design next-generation storage and computing services. Cozy Cloud is a French startup providing such a personal Cloud platform. The Cozy product is a software stack that anyone can deploy to run his personal server in order to host his personal data and web services. Cozy defines itself as the "Android of personal servers". While centralizing all personal data in the holder's hand is a natural way to reestablish his control on his privacy, this represents an unprecedented threat in case of attacks by an intruder, especially for individuals who are not security experts. The objective of this bilateral contract is typically to address this issue by integrating the PlugDB solution into the Cozy stack. Roughly speaking, the Cozy data system will be modified in such a way to store only encrypted files and each file access will be intercepted and routed to PlugDB. PlugDB will act as a doorkeeper for the whole individual dataspace by managing the files' metadata, the access control rules defined on these metadata, the decryption keys and the user/application authentication.

7.1.2. Cozy Cloud CIFRE contract (Oct 2014 - Sept 2017)

Partners: Cozy Cloud, Inria-SMIS
SMIS funding: 30k€.
In relation with the bilateral contract mentioned above, a CIFRE PhD thesis has been started by Paul Tran Van. The objective is to capitalize on the Cozy-PlugDB platform to devise new access and usage control models to exchange data among devices of the same user (devices may have different levels of trustworthiness) and among different users. A particular focus will be put on the enforcement of the access and usage control rules in this thesis.

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. ANR KISS (Dec. 2011 - Dec. 2015)

Partners: Inria-SMIS (coordinator), Inria-SECRET, LIRIS, Univ. of Versailles, CryptoExperts, Gemalto, Yvelines district.
SMIS funding: 230k€.

The idea promoted in KISS is to embed, in trusted devices, software components capable of acquiring, storing and managing securely various forms of personal data (e.g., salary forms, invoices, banking statements, geolocation data, depending on the applications). These software components form a Personal Data Server which can remain under the holder's control. The scientific challenges include: embedded data management issues tackling regular, streaming and spatio-temporal data (e.g., geolocation data), data provenance-based privacy models, crypto-protected distributed protocols to implement private communications and secure global computations.

8.1.2. CAPPRIS Project-Lab (Dec. 2011 - Dec. 2015)

Inria Partners: PRIVATICS (coordinator), SMIS, PLANETE, CIDRE, COMETE.
External partners: Univ. of Namur, Eurecom, LAAS.
Funding: not associated to individual project-teams.

An Inria Project Lab (IPL) is a long-term multi-disciplinary project launched by Inria to sustain large scale risky research actions in line with its own strategic plan. CAPPRIS stands for "Collaborative Action on the Protection of Privacy Rights in the Information Society". The key issues that are addressed are: (1) the identification of existing and future threats to privacy, (2) the definition of formally grounded measures to assess and quantify privacy, (3) the definition of the fundamental principles underlying privacy by design and methods to apply them in concrete situations and (4) The integration of the social and legal dimensions. To assess the relevance and significance of the research results, they are confronted to three classes of case studies CAPPRIS partners are involved in: namely Online Social Networks, Location Based Services and Electronic Health Record Systems.

8.1.3. CityLab@Inria, Inria Project Lab (May 2014 -).

Inria Partners: ARLES-MIMOVE, CLIME, DICE, FUN, MYRIADS, OAK, SMIS, URBANET, WILLOW.
External partners: UC Berkeley.
Funding: not associated to individual project teams.

CityLab@Inria studies ICT solutions toward smart cities that promote both social and environmental sustainability. A strong emphasis of the Lab is on the undertaking of a multi-disciplinary research program through the integration of relevant scientific and technology studies, from sensing up to analytics and advanced applications, so as to actually enact the foreseen smart city Systems of Systems. SMIS contributes to Privacy-by-Design architectures for trusted smart objects so as to ensure privacy to citizens, which is critical for ensuring that urbanscale sensing contributes to social sustainability and does not become a threat. https://citylab.inria.fr/

8.1.4. VALDO (Valorisation et monétisation des données personnelles à l'ère du Big Data), Digital Society Institute (DSI) (May 2015 - Sept. 2016).

Partners: DANTE and SMIS (co-organizers), CERDI, RITM.
SMIS funding: 50K€.

The objective of this project is to study with a multidisciplinary approach (i.e., computer science, law and economics) the impact of putting a certain (e.g., monetary) value on personal data, over the behavior of individuals (that are the rightful owners of the data) and market companies (that make usage of the personal data) in terms of data protection practices and data usage.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific events organization

9.1.1.1. General chair, scientific chair

- Philippe Pucheral: Co-founder of the bi-annual French Summer School 'Masses de Données Distribuées' and co-organizer of this school in 2016
- Benjamin Nguyen: Organizer and chair of 'Atelier sur la Protection de la Vie Privée' (APVP) in 2015

9.1.2. Scientific events selection

9.1.2.1. Member of the conference program committees

- Philippe Pucheral: EDBT'15, MOBIWIS'15, DATA'15, BDA'15
- Luc Bouganim: EDBT'15, EDBT'16
- Nicolas Anciaux: BDA'15, Int. Conference on Sustainable Energy and Environmental Engineering (SEEE'15)
- Benjamin Nguyen: BDA 2015, ACOMP'15, ECML-PKDD'15, EDA 2015
- Iulian Sandu Popa: MOBILWARE 2015, IEEE Mobile Cloud 2016, APVP 2015

9.1.2.2. Reviewer

- Iulian Sandu Popa: EDBT'15, ACM SIGSPATIAL'15

9.1.3. Journal

9.1.3.1. Member of the editorial boards

- Nicolas Anciaux: Area Editor of the VLDB Journal (since 2015)
- Benjamin Nguyen: Member of the editorial committee of TSI (Techniques et Sciences Informatiques), French Journal, Eds. Lavoisier since 2012

9.1.3.2. Reviewer - Reviewing activities

- Iulian Sandu Popa: ACM Transactions on Storage 2015

9.1.4. Invited talks

- Invited tutorial: Towards an Era of Trust in Personal Data Management, Tutorial at ADBIS'15 conference, N. Anciaux, B. Nguyen, I.Sandu Popa, 2015
- Invited talk: Gestion de données personnelles respectueuse de la vie privée. "1/2 heure de science", N. Anciaux, June 2015
- Invited talk: Gestion de données embarquées dans des calculateurs sécurisés, une solution pour la protection des données personnelles ?, Journée ASF, COMPAS conference, L. Bouganim, June 2015

- Invited talks and demonstrations: PlugDB: towards the Secure Personal Cloud. Smart City & Mobility Innovations, Cities, IoT and Analytics mobility, RII, San Francisco, N. Anciaux, 2015

9.1.5. Research administration

- Philippe Pucheral: Member of the HDR committee of the STV doctoral school (UVSQ) since 2014
- Philippe Pucheral: Philippe Pucheral: Member of the steering committee of the ED STIC doctoral school of University Paris-Saclay, 'Data, Knowledge and Interactions' committee (about 250 PhD students) since 2014
- Luc Bouganim: Co-president of the Inria 'Emplois Scientifiques' commission (includes PhD grant, Post-Doc and Délégation attributions), 2015
- Nicolas Anciaux: Co-director of the 'Privacy and digital identity' WG at Digital Society Institute (DSI), since January 2015
- Nicolas Anciaux: President of the 'Comité de Suivi Doctoral' (CSD) at Inria Paris-Rocquencourt, since 2015
- Nicolas Anciaux: Member of Commission de Developpement Technologique (CDT) at Inria Rocquencourt since 2012
- Benjamin Nguyen: Director of Digital Affairs (INSA CVL)
- Benjamin Nguyen: Member of the Executive Committee of INSA CVL

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Master : Philippe Pucheral, responsible of the DataScale master, courses in M1 and M2, UVSQ, France

Master: Nicolas Anciaux, Courses on database internal mechanisms and database security, 80, in Master1 and Master2 (AFTI, Orsay) and in engineering school (ENSTA ParisTech, Telecom Paristech)

Master : Luc Bouganim, Bases de données, 39, niveau M2, CFA AFTI/UVSQ, France

Master : Luc Bouganim, Systèmes d'information Privacy by Design, 30, niveau M2, ENSIIE, France

Master : Luc Bouganim, Sécurité des bases de données, 15, niveau M2, Télécom ParisTech, France

Licence : Iulian Sandu Popa, Fondaments de l'informatique, 37, niveau L1, UVSQ, France

Licence : Iulian Sandu Popa, Initiation aux bases de données, 66, niveau L2, UVSQ, France

Master : Iulian Sandu Popa, Bases de données relationnelles, 60, niveau M1, UVSQ, France

Master : Iulian Sandu Popa, Mécanismes internes des bases de données, M2, niveau M2, UVSQ, France

Master : Iulian Sandu Popa, Sécurité des bases de données, 9, niveau M2, UVSQ, France

Master : Iulian Sandu Popa, Projet bases de données, 27, niveau M2, UVSQ, France

Master : Benjamin Nguyen, Databases, IA, Security, 192, INSA CVL, France

9.2.2. Supervision

PhD : Cuoc-Quong To, Privacy-Preserving Query Execution using Tamper Resistant Hardware : Design and Performance Considerations, UVSQ, Octobre 2015, Benjamin Nguyen and Philippe Pucheral

PhD : Mathias Bjorling, Operating System Support for High-Performance Solid State Drives, IT University of Copenhagen, August 2015, Philippe Bonnet and Luc Bouganim

PhD : Niv Dayan, Modelling and Managing SSD Write Amplification, IT University of Copenhagen, August 2015, Philippe Bonnet and Luc Bouganim

PhD in progress : Saliha Lallali, A Scalable Search Engine for the Personal Cloud, October 2012, Nicolas Anciaux, Philippe Pucheral, and Iulian Sandu Popa

PhD in progress : Dai Hai Ton That, Efficient Management and Secure Sharing of Mobility Traces, November 2012, Iulian Sandu Popa and Karine Zeitouni (UVSQ)

PhD in progress : Paul Tran Van, Partage de documents sécurisé dans le Cloud Personnel , October 2014, Nicolas Anciaux and Philippe Pucheral

PhD in progress : Chao Chen, A Pivacy-by-Design Middleware for Urban-scale Mobile Crowdsensing, December 2015, Nicolas Anciaux and Valérie Issarny (MiMOVE)

PhD in progress : Axel Michel, Secure Distributed Computations, October 2015, Benjamin Nguyen and Philippe Pucheral

9.2.3. Juries

Philippe Pucheral: member of the PhD jury of Benjamin Billet (UVSQ, 19/03/2015)

Benjamin Nguyen: Reviewer of the PhD of Regina PAIVA MELO MARIN (Supélec Rennes, 7/9/2015)

Benjamin Nguyen: Reviewer of the PhD of Maria Laura NECULA MAAG (Université Pierre et Marie Curie (Paris-VI), 8/4/2015)

Benjamin Nguyen: President of the PhD jury of Mouhamadou Lamine BA (Telecom ParisTech, 30/3/2015)

9.3. Popularization

- Round table: Privacy & Personal Data. Who owns the data? International Conference on Digital Assets, Data Philanthropy, and Public Benefit, CITRIS, Inria, EIT ICT Labs, San Francisco. N. Anciaux, 2015
- General public (large audience magazines, television, videos): Gestion sécurisée de données personnelles. 1024 - Bulletin de la Société Informatique de France, numéro 5, pp. 17-41, N. Anciaux and B. Nguyen, 2015
- Institutions, decision makers and industrials: Présentation du Pôle Vie Privée de l'ISN (DSI), Cabinet de Mme Axelle Lemaire, Secrétaire d'Etat chargée du numérique, Ministère de l'économie, Paris, N. Anciaux and L. Bouganim, May 2015
- Round table: Peut-on être seul maître de son espace numérique ? Comprendre et construire la société numérique. Conférence de l'Institut de la Société Numérique (ISN/DSI), Hôtel Potocki, Paris, N. Anciaux, March 2015
- Round table: P. Pucheral, "Données personnelles, approche interdisciplinaire". Busines Convention on Big Data, Campus HEC, 24-25 Nov. 2015

10. Bibliography

Major publications by the team in recent years

[1] T. ALLARD, N. ANCIAUX, L. BOUGANIM, Y. GUO, L. LE FOLGOC, B. NGUYEN, P. PUCHERAL, I. RAY, I. RAY, S. YIN. *Secure Personal Data Servers: a Vision Paper*, in "Proc. of the 36th Int. Conf. on Very Large Databases (VLDB)", 2010

[2] T. ALLARD, B. NGUYEN, P. PUCHERAL. *Safe Realization of the Generalization Privacy Mechanism*, in "Privacy, Security and Trust", Montreal, Canada, 2011, pp. 1-8, Best Paper Award, http://hal.inria.fr/hal-00624043/en

[3] T. ALLARD, B. NGUYEN, P. PUCHERAL. *MetaP: Revisiting Privacy-Preserving Data Publishing using Secure Devices*, in "Distributed and Parallel Databases", June 2014, vol. 32, n° 1, pp. 191-244 [*DOI :* 10.1007/s10619-013-7122-x], https://hal.archives-ouvertes.fr/hal-00934586

[4] N. ANCIAUX, P. BONNET, L. BOUGANIM, B. NGUYEN, P. PUCHERAL, I. SANDU POPA. *Trusted Cells : A Sea Change for Personnal Data Services*, in "CIDR 2013 - 6th Biennal Conference on Innovative Database Research", Asilomar, United States, 2013, 4 p. , http://hal.inria.fr/hal-00768379

[5] N. ANCIAUX, L. BOUGANIM, T. DELOT, S. ILARRI, L. KLOUL, N. MITTON, P. PUCHERAL. *FolkIS: Opportunistic Data Services in Least Developed Countries*, in "40th International Conference on Very Large Data Bases (VLDB)", Hangzhou, China, Zhejiang University, September 2014, https://hal.inria.fr/hal-00906204

[6] N. ANCIAUX, L. BOUGANIM, P. PUCHERAL, Y. GUO, L. LE FOLGOC, S. YIN. *MILo-DB: a personal, secure and portable database machine*, in "Distributed and Parallel Databases", March 2014, vol. 32, n° 1, pp. 37-63 [*DOI :* 10.1007/s10619-012-7119-x], https://hal.archives-ouvertes.fr/hal-00768355

[7] P. BONNET, L. BOUGANIM, I. KOLTSIDAS, S. VIGLAS. *System Co-Design and Data Management for Flash Devices*, in "Very Large Data Bases (Tutorial)", 2011

[8] L. BOUGANIM, B. JÓNSSON, P. BONNET. *uFLIP: Understanding Flash IO Patterns*, in "4th Biennial Conference on Innovative Data Systems Research (CIDR)", Asilomar, California, USA, January 2009, best paper award

[9] Q. C. TO, B. NGUYEN, P. PUCHERAL. *Privacy-Preserving Query Execution using a Decentralized Architecture and Tamper Resistant Hardware*, in "17th International Conference on Extending Database Technology (EDBT)", Athens, Greece, March 2014 [*DOI :* 10.5441/002/EDBT.2014.44], https://hal.inria.fr/hal-01096639

[10] S. YIN, P. PUCHERAL. *PBFilter: a Flash-Based Indexing Scheme for Embedded Systems*, in "Information Systems", 2012, vol. 37, n° 7, pp. 634-653 [*DOI :* 10.1016/j.is.2012.02.002], http://hal.archives-ouvertes.fr/hal-00768380

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] M. BJØRLING. *Operating System Support for High-Performance Solid State Drives*, IT University of Copenhagen, August 2015

[12] N. DAYAN. *Modelling and Managing SSD Write Amplification*, IT University of Copenhagen, August 2015

[13] C. Q. TO. *Privacy-Preserving Query Execution using Tamper Resistant Hardware. Design and Performance Considerations*, Universite´ de Versailles Saint-Quentin-en-Yvelines, September 2015, https://hal.archives-ouvertes.fr/tel-01253759

Articles in International Peer-Reviewed Journals

[14] N. ANCIAUX, D. BOUTARA, B. NGUYEN, M. VAZIRGIANNIS. *Limiting Data Exposure in Multi-Label Classification Processes*, in "Fundamenta Informaticae", 2015, vol. 137, n° 2, pp. 219-236, https://hal.inria.fr/hal-01176445

[15] I. SANDU POPA, K. ZEITOUNI, V. ORIA, A. KHARRAT. *Spatio-temporal compression of trajectories in road networks*, in "GeoInformatica", 2015, vol. 19, n° 1, pp. 117-145 [*DOI :* 10.1007/s10707-014-0208-4], https://hal.inria.fr/hal-01096623

[16] D. H. TON THAT, I. SANDU POPA, K. ZEITOUNI. *TRIFL: A Generic Trajectory Index for Flash Storage*, in "ACM Transactions on Spatial Algorithms and Systems", July 2015, vol. 1, n° 2, 44 p. [*DOI :* 10.1145/2786758], https://hal.inria.fr/hal-01176563

Articles in National Peer-Reviewed Journals

[17] N. ANCIAUX, B. NGUYEN. *Gestion sécurisée de données personnelles*, in "1024, le bulletin", March 2015, n° 5, pp. 17-41, https://hal.archives-ouvertes.fr/hal-01179047

Invited Conferences

[18] N. ANCIAUX, B. NGUYEN, I. SANDU POPA. *Towards an Era of Trust in Personal Data Management*, in "Proceedings of the 19th East-European Conference on Advances in Databases and Information Systems (ADBIS '15). Tutorial", Poitiers, France, 2015, https://hal.inria.fr/hal-01176512

[19] L. BOUGANIM. *Gestion de données embarquées dans des calculateurs sécurisés, une solution pour la protection des données personnelles ?*, in "Atelier ASF "Aux frontières du Système", Conférence en Parallélisme, Architecture et Système", Lille, France, 2015, https://hal.inria.fr/hal-01178257

International Conferences with Proceedings

[20] N. ANCIAUX, S. GUILLOTON, L. BOUGANIM, I. SERGIO, A. KAMGANG, A. NGAMI, C. NOUEDI, P. PUCHERAL, M. TCHUENTÉ. *Managing Personal Health Records in an Infrastructure-weak Environment*, in "Proceedings of the 7th EAI International Conference on e-Infrastructure and e-Services for Developing Countries", Cotonou, Benin, December 2015, https://hal.archives-ouvertes.fr/hal-01254961

[21] N. ANCIAUX, S. LALLALI, I. SANDU POPA, P. PUCHERAL. *A Scalable Search Engine for Mass Storage Smart Objects*, in "Proceedings of the 41th International Conference on Very Large Databases (VLDB)", Kohala Coast, Hawaii, United States, August 2015, vol. 8, n° 9, pp. 910-921 [*DOI :* 10.14778/2777598.2777600], https://hal.inria.fr/hal-01176458

[22] M. BJØRLING, J. MADSEN, J. GONZALEZ, P. BONNET. *Linux Kernel Abstractions for Open-Channel Solid State Drives*, in "Non-Volatile Memories Workshop", San Diego, CA, United States, 2015, https://hal.inria.fr/hal-01178263

[23] M. BJØRLING, M. WEI, J. MADSEN, J. GONZALEZ, S. SWANSON, P. BONNET. *AppNVM: A software-defined, application-driven SSD*, in "Non-Volatile Memories Workshop", San Diego, CA, United States, 2015, https://hal.inria.fr/hal-01178262

[24] S. LALLALI, N. ANCIAUX, I. SANDU POPA, P. PUCHERAL. *A Secure Search Engine for the Personal Cloud*, in "Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD

'15). Demo paper", Melbourne, Australia, 2015, pp. 1445-1450 [*DOI :* 10.1145/2723372.2735376], https://hal.inria.fr/hal-01176473

[25] C. Q. To, B. Nguyen, P. Pucheral. *TrustedMR: A Trusted MapReduce System based on Tamper Resistance Hardware* , in "Proceedings of the 23rd International Conference on Cooperative Information Systems (COOPIS)", Rhodes, Greece, October 2015, https://hal.inria.fr/hal-01254951

[26] D. H. Ton That, I. Sandu Popa, K. Zeitouni. *PPTM: Privacy-aware Participatory Traffic Monitoring Using Mobile Secure Probes*, in "Proceedings of the 16th IEEE International Conference on Mobile Data Management (MDM '15). Demo paper", Pittsburgh, United States, 2015, 4 p. , https://hal.inria.fr/hal-01176486

Conferences without Proceedings

[27] N. Anciaux, S. Lallali, I. Sandu Popa, P. Pucheral. *A scalable search engine for mass storage smart objects*, in "31èmes journées Bases de Données Avancées (BDA)", Île de Porquerolles, France, 2015, https://hal.inria.fr/hal-01176462

[28] N. Anciaux, B. Nguyen. *Managing Personal Data with Strong Privacy Guarantees. Tutoriel*, in "Rendez-Vous de la Recherche et de l'Enseignement de la Sécurité des Systèmes d'Information (RESSI 2015)", Troyes, France, May 2015, https://hal.inria.fr/hal-01178684

[29] A. Katsouraki, L. Bouganim, B. Nguyen, P. Tran-Van. *Secure Portable Tokens for Sensitive Questionnaires Surveys*, in "31èmes journées Bases de Données Avancées (BDA '15). Demo paper", Île de Porquerolles, France, 2015, https://hal.inria.fr/hal-01176544

[30] Q.-C. To, B. Nguyen, P. Pucheral. *TrustedMR: A Trusted MapReduce System based on Tamper Resistance*, in "31èmes journées Bases de Données Avancées (BDA '15)", Île de Porquerolles, France, 2015, https://hal.inria.fr/hal-01176539

[31] D. H. Ton That, I. Sandu Popa, K. Zeitouni, C. Borcea. *PAMPAS: Collecte participative respectueuse de la vie privée basée sur des mobiles securisés*, in "31èmes journées Bases de Données Avancées (BDA '15)", Île de Porquerolles, France, 2015, https://hal.inria.fr/hal-01176500

[32] D. H. Ton That, I. Sandu Popa, K. Zeitouni. *PPTM: Privacy-aware Participatory Traffic Monitoring Using Mobile Secure Probes*, in "31èmes journées Bases de Données Avancées (BDA '15). Demo paper", Île de Porquerolles, France, 2015, 4 p. , https://hal.inria.fr/hal-01176493

[33] P. Tran-Van, P. Pucheral, N. Anciaux, B. André. *Partage de documents sécurisé entre Cloud personnels*, in "APVP'15 - 6e Atelier sur la Protection de la Vie Privée", Mosnes, France, June 2015, https://hal.inria.fr/hal-01226428

Patents and standards

[34] N. Anciaux, L. Bouganim, P. Pucheral. *Plans Hardware du Token PlugDB*, 2015, n° Enregistrement APP no.IDDN.FR.001.090013.000.S.P.2015.000.20600, https://hal.inria.fr/hal-01176942

[35] N. ANCIAUX, L. BOUGANIM, P. PUCHERAL, S. YIN, Q. LEFEBVRE, A. ERSO¨Z, A. TROUSSOV. *Logiciel PlugDB-engine version 4*, September 2015, n^{o} Enregistrement APP nr.IDDN.FR.001.280004.000.S.C.2008.0000.10000, https://hal.archives-ouvertes.fr/hal-01254770

Other Publications

[36] Q.-C. TO, B. NGUYEN, P. PUCHERAL. *TrustedMR: A Trusted MapReduce System based on Tamper Resistance Hardware*, September 2015, working paper or preprint, https://hal.inria.fr/hal-01185484

References in notes

[37] R. AGRAWAL, J. KIERNAN, R. SRIKANT, Y. XU. *Hippocratic Databases*, in "Proc. of the Int. Conf. on Very Large Data Bases (VLDB)", 2002

[38] A. ANDERSON. *An introduction to the web services policy language (WSPL)*, in "IEEE Computer Society", 2004

[39] L. BOUGANIM, Y. GUO. *Database Encryption*, in "Encyclopedia of Cryptography and Security", S. JAJODIA, H. VAN TILBORG (editors), Springer, 2009, pp. 307-312

[40] L. CRANOR. *Web Privacy with P3P*, O'Reilly Media, 2002

[41] B. FUNG, K. WANG, R. CHEN, P. YU. *Privacy-preserving data publishing: A survey of recent developments*, in "ACM Computing Surveys (CSUR)", 2010, vol. 42, n^{o} 4

[42] T. MOSES. *Extensible access control markup language (XACML) version 2.0*, in "Oasis Standard 200502", 2005

[43] M. SCHUNTER, C. POWERS. *Enterprise privacy authorization language (EPAL 1.1)*, in "IBM", 2003