IN PARTNERSHIP WITH:
CNRS

Université Versailles Saint-Quentin

Activity Report 2012

Project-Team SMIS

Secured and Mobile Information Systems

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

RESEARCH CENTER
Paris - Rocquencourt

THEME
Knowledge and Data Representation and Management

Table of contents

Project-Team SMIS

Keywords: Databases, Privacy, Ubiquitous Computing, Distributed System, Information Indexing And Retrieval

Creation of the Project-Team: September 01, 2004 .

1. Members

Research Scientists

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

Faculty Members

Philippe Pucheral [Team leader, Professor - UVSQ, HdR]
Benjamin Nguyen [Associate Professor - UVSQ]
Iulian Sandu Popa [Associate Professor - UVSQ]

Engineer

Alexei Troussov [Senior Software Engineer]

PhD Students

Quoc-Cuong To [UVSQ, CORDI, since Nov 2012]
Saliha Lallali [UVSQ, CORDI, since Nov 2012]
Lionel Le Folgoc [UVSQ, CORDI]
Niv Dayan [U. of Copenhagen, co-supervision]
Matias Bjørling [U. of Copenhagen, co-supervision]

Administrative Assistant

Laurence Bourcier [Administrative Assistant]

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. Scientific Foundations

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 [39] and Hippocratic databases [30]. In the last years, many other policy languages have been proposed for different application scenarios, including EPAL [44], XACML [41] and WSPL [34]. 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 behaviour 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) [36]. 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 (R. Agrawal : IBM Almaden, C.C. Aggarwal: IBM Watson), 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 [40]. 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 (R.N.Wright : Stevens Institute of Tech. / Rutgers Univ, NJ, V. Shmatikov : Univ Austin Texas), on Secure Multi-party Computing for Privacy Preserving Data Mining (J. Vaidya : Rutgers Univ, C. Clifton : Purdue Univ) and on distributed PPDP algorithms (D. DeWitt : Univ Wisconsin, K. Lefevre : Univ Michigan, J. Vaidya : Rutgers Univ, C. Clifton : Purdue Univ) while none of them share the same architectural hypothesis as us.

4. Application Domains

4.1. Application Domains

data privacy, personal information management, healthcare, ambient intelligence

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, one application is today more specifically targeted by the SMIS project. This application deals with privacy preservation in EHR (Electronic Health Record) systems. Several countries (including France) launched recently ambitious EHR programs where medical folders will be centralized and potentially hosted by private Database Service Providers. Centralization and hosting increase the risk of privacy violation. In 2007, we launched two projects (PlugDB and DMSP) tackling precisely this issue, with the final objective to experiment our technologies in the field. In 2011, we launched a new project (KISS) capitalizing on the previous ones and extending their scope towards the protection of any personal data delivered to individuals in an electronic form.

5. Software

5.1. Introduction

In our research domain, developing software prototypes is mandatory to validate research solutions and is an important vector for publications, demonstrations at conferences and exhibitions as well as for cooperations with industry. This prototyping task is however difficult because it requires specialized hardware platforms (e.g., new generations of smart tokens), themselves sometimes at an early stage of development.

For a decade, we have developed successive prototypes addressing different application domains, introducing different technical challenges and relying on different hardware platforms. PicoDBMS was our first attempt to design a full-fledged DBMS embedded in a smart card [42] [32]. Chip-Secured Data Access (C-SDA) embedded a reduced SQL query engine and access right controller in a secure chip and acted as an incorruptible mediator between a client and an untrusted server hosting encrypted data [37]. Chip-Secured XML Access (C-SXA) was an XML-based access rights controller embedded in a smart card [38]. Prototypes of C-SXA have been the recipient of the e-gate open 2004 Silver Award and SIMagine 2005 Gold award, two renowned international software contests. The next subsections details the two prototypes we are focusing on today.

5.2. PlugDB engine

Participants: Nicolas Anciaux [correspondent], Luc Bouganim, Philippe Pucheral, Shaoyi Yin, Yanli Guo, Lionel Le Folgoc, Alexei Troussov.

More than a stand-alone prototype, PlugDB is part of a complete architecture dedicated to a secure and ubiquitous management of personal data. PlugDB aims at providing an alternative to a systematic centralization of personal data. To meet this objective, the PlugDB architecture lies on a new kind of hardware device called Secure Portable Token (SPT). Roughly speaking, a SPT combines a secure microcontroller (similar to a smart card chip) with a large external Flash memory (Gigabyte sized). The SPT can host data on Flash (e.g., a personal folder) and safely run code embedded in the secure microcontroller. PlugDB engine is the cornerstone of this embedded code. PlugDB engine manages the database on Flash (tackling the peculiarities of NAND Flash storage), enforces the access control policy defined on this database, protects the data at rest against piracy and tampering, executes queries (tackling low RAM constraint) and ensures transaction atomicity. Part of the on-board data can be replicated on a server (then synchronized) and shared among a restricted circle of trusted parties through crypto-protected interactions. PlugDB engine has been registered at APP (Agence de Protection des Programmes) in 2009 [33] and a new version is registered each year. The underlying Flash-based indexing system has also been patented by Inria and Gemalto [43]. It has been demonstrated in a dozen of national and international events including JavaOne and SIGMOD. It is being experimented in the field to implement a secure and portable medical-social folder helping the coordination of medical care and social services provided at home to dependent people. In 2012, we have ported PlugDB-engine on a new hardware platform to 1) become completely independent from Gemalto, 2) have a plug-and-play implementation on Android, 3) serve as a basement to port it on other custom hardware implementation. We have already discussed with hardware companies located in “île de France” to produce new hardware tokens to host future versions of PlugDB-engine. Link: http://www-smis.inria.fr/_DMSP/home.php.

5.3. uFLIP Benchmark

Participants: Luc Bouganim [correspondent], Philippe Bonnet, Bjorn Jónsson, Lionel Le Folgoc.

It is amazingly easy to produce meaningless results when measuring flash devices, partly because of the peculiarity of flash memory, but primarily because their behavior is determined by layers of complex, proprietary, and undocumented software and hardware. uFLIP is a component benchmark for measuring the response time distribution of flash IO patterns, defined as the distribution of IOs in space and time. uFLIP includes a benchmarking methodology which takes into account the particular characteristics of flash devices. The source code of uFLIP, available on the web (700 downloads, 4000 distinct visitors), was registered at APP in 2009 [35]. It has been demonstrated at SIGMOD.

Link: http://www.uflip.org.

6. New Results

6.1. Embedded Data Management

Participants: Nicolas Anciaux, Luc Bouganim, Lionel Le Folgoc, Yanli Guo, Saliha Lallali, Philippe Pucheral, Iulian Sandu Popa, Shaoyi Yin.

Inspired by low cost economic models, this work draws the idea of a one-dollar database machine, with the objective to disseminate databases everywhere, up to the lightest smart objects. In contrast to traditional database machines relying on massively parallel architectures, the one-dollar database machine considers the cheapest form of computer available today: a microcontroller equipped with GBs size (external) Flash storage. Designing such a database machine is very challenging due to a combination of conflicting RAM and NAND Flash constraints. To tackle this challenge, this work proposes a new paradigm based on database serialization (managing all database structures in a pure sequential way) and stratification (restructuring them into strata when a scalability limit is reached).We show that a complete DBMS engine can be designed according to this paradigm and demonstrate the effectiveness of the approach through a performance evaluation. This work capitalizes on previous results related to the indexing of Flash resident data [16] and has also obvious connections with the more general study we are conducting on Flash-based data management (see Section 6.2). Partial elements of this solution have been demonstrated at [13]. In 2012, we have extended our previous results on indexation of flash resident data [IS] and we have proposed the design of a complete DBMS engine [DAPD] complying by nature with the conflicting RAM and NAND Flash constraints we are facing. Currently, we work at the extension of the embedded DBMS engine to support document data (e.g., text documents or any type of documents that are tagged)) and spatio-temporal data (e.g., vehicle trajectory data or any type of time-stamped and/or geo-located data).

6.2. Flash-based Data Management

Participants: Matias Bjørling, Philippe Bonnet, Luc Bouganim, Niv Dayan.

Solid State Drives (SSDs) are replacing magnetic disks as secondary storage for database management, as they offer orders of magnitude improvement in terms of bandwidth and latency. In terms of system design, the advent of SSDs raises considerable challenges. First, the storage chips, which are the basic component of a SSD, have widely different characteristics – e.g., copy-on-write, erase-before-write and page-addressability for flash chips vs. in-place update and byte-addressability for PCM chips. Second, SSDs are no longer a bottleneck in terms of I/O latency forcing streamlined execution throughout the I/O stack. Finally, SSDs provide a high degree of parallelism that must be leveraged to reach nominal bandwidth. This evolution puts database system researchers at a crossroad. The first option is to hang on to the current architecture where secondary storage is encapsulated behind a block device interface. This is the mainstream option both in industry and academia. This leaves the storage and OS communities with the responsibility to deal with the complexity introduced by SSDs in the hope that they will provide us with a robust, yet simple, performance model. We showed that this option amounts to building on quicksand. We illustrated our point by debunking some popular myths about flash devices and by pointing out mistakes in the papers we have published throughout the years. The second option is to abandon the simple abstraction of the block device interface and reconsider how database storage managers, operating system drivers and SSD controllers interact. We gave our vision of how modern database systems should interact with secondary storage. This approach requires a deep re-design of the database system architecture, which is the only viable option for database system researchers to avoid becoming irrelevant. This work started at the end of 2011 and was published at CIDR'13 [20], in cooperation with the IT University of Copenhagen.

6.3. Minimal Exposure

Participants: Nicolas Anciaux, Walid Bezza, Danae Boutara, Benjamin Nguyen, Michalis Vazirgiannis.

When users request a service, the service provider usually asks for personal documents to tailor its service to the specific situation of the applicant. For example, the rate and duration of consumer's loans are usually adapted depending on the risk based on the income, assets or past lines of credits of the borrower. In practice, an excessive amount of personal data is collected and stored. Indeed, a paradox is at the root of this problem: service providers require users to expose data in order to determine whether that data is needed or not to achieve the purpose of the service. We explore a reverse approach, where service providers would publicly describe the data they require to complete their task, and where software (placed, depending on the context, on the client, on the server, or in a trusted hardware component) would use those descriptions to determine a minimum subset

of information to expose. In 2012, we have presented our general framework called Minimum Exposure [14], we have modelled the underlying problem (for simple tasks) and proposed resolution algorithms [19], [24], and we have addressed the case of multi-label classifiers [18]. In the short term, we plan to adapt the minimum exposure architecture to support hidden decision rules using smart cards. Then, we will investigate new privacy metrics to capture the degree of exposure of sets of personal data items better.

6.4. Secure Global Computing on Asymmetric Architecture

Participants: Tristan Allard, Benjamin Nguyen, Philippe Pucheral, Quoc-Cuong To.

This research direction is based on the asymmetric architecture, composed of a powerful, available and untrusted computing infrastructure (server or cloud), and a large set of low powered, highly disconnected trusted devices. Trust is assumed ad hoc and can be justified by the use of secure tokens, open source software, friend relationships etc. In our work, we use tamper resistant secure tokens running trusted software, which provide a high degree of trust, due to the overwhelming cost of hardware tampering. The main difficulty on such an architecture is global processing i.e. constructing aggregate data from the individual records, because the entity in charge of executing the global computation is untrusted. Given our large scale data centric applications (e.g. nationwide surveys), we also discard solutions based on secure multi-party computation, which do not scale. We have studied 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. This work is an extension of [31] We are now studying more generally the execution of SQL "Group by" queries on this architecture, which is the topic of Quoc-Cuong To's Ph.D. thesis started in sept. 2012. We have published preliminary results on this novel problem in [23], which adapts the techniques proposed in [31].

6.5. Trusted Cell Data Management

Participants: Nicolas Anciaux, Philippe Bonnet, Luc Bouganim, Benjamin Nguyen, Philippe Pucheral, Iulian Sandu Popa.

With the convergence of mobile communications, sensors and online social networks technologies, we are witnessing an exponential increase in the creation and consumption of personal data. Such data is volunteered by users, automatically captured by sensors or inferred from existing data. Today, there is a wide consensus that individuals should have increased control on how their personal data is collected, managed and shared. Yet there is no appropriate technical solution to implement such personal data services: centralized solutions sacrifice security for innovative applications, while decentralized solutions sacrifice innovative applications for security. In this work, we argue that the advent of secure hardware in all personal IT devices, at the edges of the Internet, could trigger a sea change. We propose the vision of trusted cells: personal data servers running on secure smart phones, set-top boxes, secure portable tokens or smart cards to form a global, decentralized data platform that provides security yet enables innovative applications. We motivate our approach, describe the trusted cells architecture and define a range of challenges for future research in a paper published at CIDR'13 (Int. Conf on Innovative Data Systems Research) [17].

6.6. Experiment in the medical field

Participants: Nicolas Anciaux, Luc Bouganim, Philippe Pucheral, Alexei Troussov.

The PlugDB engine is being experimented in the field since September 2011 to implement a secure and portable medical-social folder. The objective is to improve the coordination of medical care and social services provided at home for dependent people. Details related to this experiment are given in Section 7.2. While this action did not generate new academic results (though it helped us validating some previous results), it imposed us a strong investment in terms of test and optimization for our prototype and in terms of communication to promote this experiment at the regional level.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

The SMIS project has a long lasting cooperation with Axalto, recently merged with Gemplus to form 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. Meanwhile, we are developing partnerships with SMEs capable of building ad-hoc hardware prototypes conforming to our own design. We cooperate also with Santeos, an Atos Origin company developing software platforms of on-line medical services. Santeos is member of the consortium selected by the French Ministry of Health to host the French DMP (the national Personal Medical Folder initiative). This cooperation helps us tackling one of our targeted applications, namely the protection of medical folders.

7.2. DMSP Yvelines District grant (Nov 2010 - Apr. 2012)

Partners: Inria-SMIS (coordinator), Gemalto, UVSQ, Santeos
SMIS funding : 75k€
http://www-smis.inria.fr/_DMSP/accueil.php
Electronic Health Record (EHR) projects have been launched in most developed countries to increase the quality of care while decreasing its cost. Despite their unquestionable benefits, patients are reluctant to abandon their control of highly sensitive data to a distant server. The objective of the DMSP project is to complement a traditional EHR server with a secure and mobile personal medical folder (1) to protect and share highly sensitive data among trusted parties and (2) to provide a seamless access to the data even in disconnected mode. The DMSP architecture builds upon the technology designed in the PlugDB project (see above). It is currently experimented in the context of a medical-social network providing care and services at home for elderly people. The experiment in the field started in September 2011 with a targeted population of 120 volunteer patients and practitioners in the Yvelines district.

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. ANR DEMOTIS (Feb. 2009 - Feb. 2012)

Partners: SopinSpace (coordinator), Inria (SMIS, SECRET), CECOGI
SMIS funding: 85k€
http://www.demotis.org/
The design and implementation of large-scale infrastructure for sensitive and critical data (e.g., electronic health records) have to face a tangle of legal provisions, technical standards, and societal concerns and expectations. DEMOTIS project aims to understand how the intrication between legal and technical domains constrains the design of such data infrastructures. DEMOTIS consists of two interdependent facets: legal (health law, privacy law, intellectual property law) and computer science (database security, cryptographic techniques). Combining expertise of researchers in Law and computer scientists should help to better assess whether law statements can be actually put in practice, to characterize the related technological challenges when mismatches are detected and, when possible, to suggest preliminary solutions.

8.1.2. 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.

8.1.3. ARC CAPPRIS (Dec. 2011 - Dec. 2015)

Inria Large Scale Initiative
Inria Partners: PRIVATICS (coordinator), SMIS, PLANETE, CIDRe, COMETE
External partners: Univ. of Namur, Eurecom, LAAS
Funding: not associated to individual project-teams
An ARC 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 will be 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 signicance of the research results, they will be confronted to three classes of case studies CAPPRIS partners are involved in, namely Online Social Networks, Location Based Services and Electronic Health Record Systems.

8.2. European Initiatives

8.2.1. Collaborations in European Programs, except FP7

Program: Danish Council for Independent Research (FTP call)

Project acronym: CLyDE

Project title: Cross-LaYer optimized Database Engine

Duration: 10/2011 - 10/2014

Coordinator: Philippe Bonnet (ITU of Copenhagen)

Other partners: IT University of Copenhagen - Denmark, SMIS

Abstract: The goal is to explore how flash devices, operating system and database system can be designed together to improve overall performance. Such a co-design is particularly important for the next generation database appliances, or cloud-based relational database systems for which well-suited flash components must be specified. More generally, our goal is to influence the evolution of flash devices and commodity database systems for the benefit of data intensive applications. The project should result in two complementary open-source software systems: (i) a bimodal flash device software component based on the idea from [6], and (ii) a database system optimized for bimodal flash devices. The project funding will be managed by the IT University of Copenhagen and will cover the expenses for two co-supervised PhD students (including regular visits to and from Denmark).

8.2.2. Collaborations with Major European Organizations

The SMIS members have developed tight european cooperations with the following persons/teams:

- P. Bonnet (Associate Professor at the University of Copenhagen, Denmark): collaboration on Flash-based data management for high-end servers. The study of flash devices started during a short sabbatical of Luc Bouganim (from April to August 2008) in Copenhagen. The uFLIP study has been conducted in close cooperation with Philippe Bonnet from IT University of Copenhagen and Björn Þór Jónsson from Reykjavík University. The cooperation with Copenhagen is very active and led to new studies on flash devices and on the Trusted Cell architecture Two PhD students are currently co-supervised by Luc Bouganim and Philippe Bonnet. Philippe Bonnet got a Marie-Curie IEF grant and will visit SMIS for one year in 2013-2014.
- Michalis Vazirgiannis (Athens University of Economics and Business): collaboration on Minimal Exposure in the context of Michalis' Digiteo Chair at LIX (Ecole Polytechnique).

8.3. International Initiatives

8.3.1. Inria International Partners

The SMIS members have developed tight international cooperations with the following persons/teams:

- Dennis Shasha (Professor at the University of New-York, USA): collaboration on tamper-resistant data management issues. Dennis Shasha has done a one year sabbatical stay in SMIS (July 2006 to June 2007).
- I. Ray and I.Ray (Professors at Colorado State University, USA): collaboration on data privacy and usage control (Indrajit and Indrakshi Ray have visited SMIS from September 2009 up to February 2010).
- Cristian Borcea (Associate Professor at New Jersey Institute of Technology, USA): collaboration on spatio-temporal data management issues.

8.4. International Research Visitors

8.4.1. Visits of International Scientists

8.4.1.1. Internships

- Quoc-Cong To (Vietnam): Internship on distributed query processing in the PDS architecture. September 2012 - February 2013.
- Athanasia Katsouraki (Greece): Internship on usage control models. January-September 2012.
- Danae Boutara (Greece): Internship on Minimal exposure. April-August 2012.

9. Dissemination

9.1. Scientific Animation

- Philippe Pucheral
  - Area Editor of the Information Systems international journal (2007-now).
  - PC member of MOBIWIS'12, CODAPSY'12, EDBT'12 (demo).
  - Member of the recruiting committees of UVSQ and Paris VI.
  - Co-founder of the bi-annual French Summer School "Masses de Données Distribuées" and co-organiser of this school in 2010 and 2012.
- Luc Bouganim
  - President of the BDA Board, the steering committee of the national French conference on Databases.
  - PC member of MOBIWIS'12, BDA'12.

  - Member of the recruiting committees of INSA Lyon.
  - Reviewer for the Information Systems international journal and ACM TKDE journal.
  - President of the Inria Post-Doc and Delegation Commission.
  - Member of the Inria “Bureau du Comité des Projets” (BCP).
  - Member of the Inria "Cordi-S (Inria PhD grant)" commission.
  - Member of the Commission PES (Prime d’Excellence Scientifique) for computer science at UVSQ (since 2010).
- Nicolas Anciaux
  - Demonstration chair at BDA 2012.
  - Co-organizer of the 4th edition of Atelier Protection de la Vie Privée (APVP 2013)
  - Member of Commission de Developpement Technologique (CDT) at Inria Rocquencourt
  - Member of the Editorial Board of TSI Journal (Technique et Science Informatiques) (2007 – until 2012).
- Benjamin Nguyen
  - Member of the Advisory Committee of the W3C for the UVSQ, of the W3C XQuery Working Group (Test Suite Editor) and of the W3C Social Web Interest Group.
  - PC member of VLDB 2012, VLDB 2012 PhD Workshop, ICCSAMA 2013, BDA 2012.
  - Member of the editorial committee of TSI (Techniques et Sciences Informatiques), Eds. Lavoisier.
  - Organizer of Atelier Protection de la Vie Privée (APVP) 2013.
  - Member of the Selection Committee of Paris-VI and INSA Lyon.
  - Elected member of the PRiSM Laboratory Council.
  - Elected member of the UVSQ Science faculty Scientific Committee.
  - In charge of the UVSQ Computer Science Masters diploma. In charge of the 2014/2015 proposal.
  - Responsible for the UVSQ continuous training of High School teachers in computer science.
- Iulian Sandu Popa
  - Tutorial chair of MOBILWARE 2012 (5th International Conference on MOBILe Wireless MiddleWARE, Operating Systems, and Applications).
  - Reviewer for the Earth Science Informatics Journal and the Information Systems Journal.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

SMIS is a joint project-team with University of Versailles St-Quentin (UVSQ) and CNRS. Hence SMIS members are naturally deeply involved in teaching.

- P. Pucheral: (130h/y)
  - Full professor at UVSQ.
  - Director of the research Master COSY (UVSQ).
  - Member of the HDR committee of the STV doctoral school.
  - Courses on databases, DBMS architecture and security in Master1, Master2 and engineer school ISTY.
- B. Nguyen: (107h/y)

  - Associate professor at UVSQ.
  - Courses on object programing, databases, XML in undergraduate and Master2.
  - Courses on Teaching CS courses for High School teachers.
- I. Sandu Popa: (192h/y)
  - Associate professor at UVSQ.
  - Courses on C programing and databases in undergraduate and Master2.
- L. Bouganim: (90h/y)
  - Courses on DBMS architecture, data security, database technology in Master1 and Master2 (AFTI, Orsay) and in engineering school (ENST Paris).
- N. Anciaux: (90h/y)
  - Courses on DBMS internal mechanisms, database technology in Master1 and Master2 (UVSQ), and in engineering school (ENSTA Paris).
- L. Le Folgoc: (64h/y)
  - Courses on Relational Database Concepts and SQL, system programming in Masters 1 (UVSQ).

9.2.2. Supervision

Lionel Le Folgoc. Serveurs personnels de données : réflexions sur leur conception et leurs performances. PhD Thesis University of Versailles Saint-Quentin-en-Yvelines (UVSQ), December 2012, co-supervised by Luc Bouganim and Nicolas Anciaux.

Bogdan Butnaru, Optimizations of XQuery in peer-to-peer distributed XML databases. PhD Thesis University of Versailles Saint-Quentin-en-Yvelines (UVSQ), April 2012, co-supervised by Benjamin Nguyen and Georges Gardarin.

PhD in progress : Quoc-Cuong To, starting in November 2012, Co-supervised by Philippe Pucheral and Benjamin Nguyen.

PhD in progress : Saliha Lallali, starting in November 2012, Co-supervised by Philippe Pucheral, Nicolas Anciaux and Iulian Sandu Popa.

PhD in progress : Matias Bjørling, starting in December 2011, co-supervised by Philippe Bonnet and Luc Bouganim.

PhD in progress : Niv Dayan, starting in December 2011, co-supervised by Philippe Bonnet and Luc Bouganim.

9.2.3. Juries

- L. Bouganim
  - Referee for the PhD thesis of Toufik Sarni: Vers une Mémoire Transactionnelle Temps Réel, University of Nantes, October 2012.
  - Referee for the PhD thesis of Stéphane Jacob: Protection Cryptographique des Bases de Données : Conception et Cryptanalyse, University of Paris 6, March 2012.
  - Jury member of the PhD defense of Lionel Le Folgoc: Personal Data Server Engine: Design and Performance Considerations, University of Versailles, December 2012.
  - Jury member of the PhD defense of Hien Thi Thu Truong: A Contract-based and Trust-aware Collaboration Model, University of Lorraine, December 2012.
- N. Anciaux
  - Jury member of the PhD defense of Lionel Le Folgoc: Personal Data Server Engine: Design and Performance Considerations, University of Versailles, December 2012.
- B. Nguyen

- Jury member of the PhD defense of Bogdan Butnaru : Optimizations of XQuery in peer-to-peer distributed XML databases, University of Versailles, April 2012
- Jury member of the PhD defense of Noor Malla : Partitioning XML data, towards distributed and parallel management, University of Paris-XI, September 2012.

9.3. Popularization

SMIS members have an important dissemination activity, motivated both by the popularity of the addressed research domain (security/privacy) and of the targeted applications (e.g., personal medical folder), leading to interactions with different institutions (e.g., the French Deputy Chamber, the Parliamentary Office for Evaluation of Scientific and Technological Options (OPECST), the French Network and Information Security Agency), interviews resulting in articles in large audience magazine (e.g., "La Recherche", the CNRS Journal or BBC news) and talks and demonstrations targeting industrials in wide audience conferences (e.g., JavaOne, e-Smart, "Les rendez-vous Carnot")

- Organization of the French Summer School "Masses de Données Distribuées", Philippe Pucheral, June 2012
- Presentation at the "1/2 heure de science", Database cryptography, Luc Bouganim, March 2012, http://www.inria.fr/centre/paris-rocquencourt/recherche/la-demi-heure-de-science/2012/database-cryptography.
- Presentation at the "Colloque du labex DIGICOSME", Orsay, Architectures pour des grandes bases de données sécurisées, Luc Bouganim, September 2012.
- Participation to the realization of a pedagogic video on Security and Safety, Luc Bouganim, October 2012.
- Participation to the Round Table at the 5th Open World Forum (OWF), Paris,Logiciels libres, culture et enseignement de l'informatique, Benjamin Nguyen, 2012.
- Invited talk at the "Les mardis du CRIDS", Le projet CAPPRIS, Benjamin Nguyen, July 2012.

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, p. 1-8, Best Paper Award, http://hal.inria.fr/hal-00624043/en.

[3] N. ANCIAUX, M. BENZINE, L. BOUGANIM, P. PUCHERAL, D. SHASHA. *GhostDB: querying visible and hidden data without leaks*, in "26th International Conference on Management of Data (SIGMOD)", June 2007.

[4] N. ANCIAUX, M. BENZINE, L. BOUGANIM, P. PUCHERAL, D. SHASHA. *Revelation on Demand*, in "Distributed and Parallel Database Journal (DAPD)", April 2009, vol. 25, n^{o} 1-2.

[5] N. ANCIAUX, L. BOUGANIM, P. PUCHERAL, P. VALDURIEZ. *DiSC: Benchmarking Secure Chip DBMS*, in "IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE)", October 2008, vol. 20, n^{o} 10.

[6] P. BONNET, L. BOUGANIM. *Flash Device Support for Database Management*, in "5th Biennial Conference on Innovative Data Systems Research (CIDR)", Asilomar, California, USA, January 2011, p. 1-8.

[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, F. DANG-NGOC, P. PUCHERAL. *Dynamic Access-Control Policies on XML Encrypted Data*, in "ACM Transactions on Information and System Security (ACM TISSEC)", January 2008, vol. 10, n^{o} 4.

[9] 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.

[10] S. YIN, P. PUCHERAL, X. MENG. *A Sequential Indexing Scheme for Flash-Based Embedded Systems*, in "Proc. of the International Conference on Extending Database Technology (EDBT)", Saint-Petersburg, Russia, March 2009.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] B. BUTNARU. *Optimizations of XQuery in peer-to-peer distributed XML databases*, University of Versailles St-Quentin-en-Yvelines, April 2012.

[12] L. LE FOLGOC. *Personal Data Server Engine: Design and Performance Considerations*, University of Versailles St-Quentin-en-Yvelines, December 2012.

Articles in International Peer-Reviewed Journals

[13] 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", 2013, to appear, http://hal.archives-ouvertes.fr/hal-00768355.

[14] N. ANCIAUX, B. NGUYEN, M. VAZIRGIANNIS. *The Minimum Exposure Project : Limiting Data Collection in Online Forms*, in "ERCIM News", July 2012, n^{o} 90, p. 41-42, Digiteo LeTeVoNe, http://hal.archives-ouvertes.fr/hal-00768369.

[15] N. ANCIAUX, J.-M. PETIT, P. PUCHERAL, K. ZEITOUNI. *Personal Data Server: Keeping Sensitive Data under the Individual's Control*, in "ERCIM News", July 2012, p. 39-40, http://hal.archives-ouvertes.fr/hal-00768354.

[16] S. YIN, P. PUCHERAL. *PBFilter: a Flash-Based Indexing Scheme for Embedded Systems*, in "Information Systems", 2012, vol. 37, n^{o} 7, p. 634-653, http://hal.archives-ouvertes.fr/hal-00768380.

International Conferences with Proceedings

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

[18] N. ANCIAUX, D. BOUTARA, B. NGUYEN, M. VAZIRGIANNIS. *Limiting Data Exposure in Multi-Label Classification Processes*, in "20th International Symposium on Methodologies for Intelligent Systems", Macao, 2012, 12, Digiteo LeTeVoNe, http://hal.archives-ouvertes.fr/hal-00768373.

[19] N. ANCIAUX, B. NGUYEN, M. VAZIRGIANNIS. *Limiting Data Collection in Application Forms : A real-case application of a Founding Privacy Principle*, in "PST '12 Proceedings of the 2012 Tenth Annual International Conference on Privacy, Security and Trust (PST)", France, 2012, p. 59-66, Digiteo LeTeVoNe [*DOI :* 10.1109/PST.2012.6297920], http://hal.archives-ouvertes.fr/hal-00768363.

[20] M. BJORLING, P. BONNET, L. BOUGANIM, N. DAYAN. *The Necessary Death of the Block Device Interface*, in "6th Biennial Conference on Innovative Data Systems Research (CIDR)", Asilomar, États-Unis, 2013, 4, http://hal.archives-ouvertes.fr/hal-00768386.

[21] S. J. PAN, M. KHAN, I. SANDU-POPA, K. ZEITOUNI, C. BORCEA. *Proactive vehicle re-routing strategies for congestion avoidance*, in "8th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS 2012)", Chine, May 2012, p. 265-272, EPI SMIS EPI SMIS, http://hal.archives-ouvertes.fr/hal-00768334.

[22] I. SANDU-POPA, K. ZEITOUNI. *Modeling and Querying Mobile Location Sensor Data*, in "The Fourth International Conference on Advanced Geographic Information Systems, Applications, and Services (GEOProcessing 2012)", Espagne, February 2012, 18, EPIS SMIS EPIS SMIS, http://hal.archives-ouvertes.fr/hal-00768335.

[23] C. Q. TO, B. NGUYEN, P. PUCHERAL. *Secure global protocol for computing aggregate functions*, in "First AVSE Doctoral Workshop", Cachan, France, September 2012, 29, http://hal.archives-ouvertes.fr/hal-00768378.

Conferences without Proceedings

[24] N. ANCIAUX, B. NGUYEN, M. VAZIRGIANNIS. *Cas d'usage d'un principe fondamental de protection de la vie privée*, in "Bases de Données Avancées (BDA)", Clermont-Ferrand, France, October 2012, Digiteo LeTeVoNe, http://hal.archives-ouvertes.fr/hal-00768367.

[25] N. ANCIAUX, B. NGUYEN, M. VAZIRGIANNIS. *Limiting Data Collection in Application Forms*, in "Atelier Protection de la Vie Privée (APVP)", Groix, France, June 2012, Digiteo LeTeVoNe, http://hal.archives-ouvertes.fr/hal-00768365.

[26] J.-P. ARCHAMBAULT, L. BLOCH, L. CHÉNO, B. NGUYEN. *Logiciels libres, culture et enseignement de l'informatique*, in "5th Open World Forum (OWF)", France, October 2012, http://hal.archives-ouvertes.fr/hal-00768374.

[27] P. BONNET, L. BOUGANIM. *Flash Device Support for Database Management*, in "28èmes journées Bases de Données Avancées (BDA)", Clermont-Ferrand, France, October 2012, http://hal.archives-ouvertes.fr/hal-00768387.

Scientific Books (or Scientific Book chapters)

[28] I. BEDINI, B. NGUYEN, C. MATHEUS, F. PATEL-SCHNEIDER, A. BORAN. *Mining XML Schemas to Extract Conceptual Knowledge*, in "Semi-Automatic Ontology Development: Processes and Resources", M.

T. PAZIENZA, A. STELLATO (editors), IGI Global, February 2012, p. 79-105, http://hal.archives-ouvertes.fr/hal-00768346.

[29] L. BOUGANIM. *Bases de Données Avancées - Modélisation, Interrogation et Optimisation*, Lavoisier, 2012, 138, Ingénierie des Systèmes d'Information (ISI), vol. 17, n°5, http://hal.archives-ouvertes.fr/hal-00768383.

References in notes

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

[31] T. ALLARD. *Sanitizing Microdata Without Leak: A Decentralized Approach*, University of Versailles, 2011.

[32] N. ANCIAUX, C. BOBINEAU, L. BOUGANIM, P. PUCHERAL, P. VALDURIEZ. *PicoDBMS : Validation and Experience*, in "Proc. of the Int. Conf. on Very Large Data Bases (VLDB)", 2001.

[33] N. ANCIAUX, L. BOUGANIM, P. PUCHERAL, S. YIN, M. BENZINE, K. JACQUEMIN, D. SHASHA, C. SALPERWYCK, M. E. KHOLY. *Logiciel PlugDB-engine version 2, enregistré à l'Agence pour la Protection des Programmes (APP) sous le numéro IDDN.FR.001.280004.000.S.C.2008.0000.10000 en date du 27 avril 2009*, April 2009.

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

[35] L. BOUGANIM. *Logiciel uFLIP version 2.1, enregistré à l'Agence pour la Protection des Programmes (APP) sous le numéro IDDN.FR.001.110020.000.S.P.2009.0000.10000 en date du 10 mars 2009*, March 2009.

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

[37] L. BOUGANIM, P. PUCHERAL. *Chip-Secured Data Access : Confidential Data on Untrusted Servers*, in "Proc. of the 28th Int. Conf. on Very Large Data Bases (VLDB)", 2002.

[38] L. BOUGANIM, F. DANG-NGOC, P. PUCHERAL. *Client-Based Access Control Management for XML Documents*, in "Proc. of the 30th Int. Conf. on Very Large Databases (VLDB)", 2004.

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

[40] 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.

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

[42] P. PUCHERAL, L. BOUGANIM, P. VALDURIEZ, C. BOBINEAU. *PicoDBMS : Scaling down Database Techniques for the Smartcard*, in "Very Large Data Bases Journal (VLDBJ), Best Paper Award VLDB'2000", 2001, vol. 10, n^{o} 2-3.

[43] P. PUCHERAL, S. YIN. *System and Method of Managing Indexation of Flash Memory*, May 2007, Dépôt par Gemalto et Inria du brevet européen nr 07290567.2.

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