RESEARCH CENTRE

Inria Saclay Centre at Université Paris-Saclay

IN PARTNERSHIP WITH:

Université Versailles Saint-Quentin

2024 ACTIVITY REPORT

Project-Team PETRUS

PErsonal & TRUSted cloud

DOMAIN

Perception, Cognition and Interaction

THEME

Data and Knowledge Representation and Processing

Inria

Contents

Project-Team PETRUS

Creation of the Project-Team: 2017 July 01

Keywords

Computer sciences and digital sciences

A1.1.8. – Security of architectures

A1.1.9. – Fault tolerant systems

A1.3. – Distributed Systems

A3.1.2. – Data management, quering and storage

A3.1.3. – Distributed data

A3.1.5. – Control access, privacy

A3.1.6. – Query optimization

A3.1.9. – Database

A3.1.11. – Structured data

A4.7. – Access control

A4.8. – Privacy-enhancing technologies

Other research topics and application domains

B2.5.3. – Assistance for elderly

B6.4. – Internet of things

B6.6. – Embedded systems

B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

- Luc Bouganim [Team leader, INRIA, Senior Researcher]
- Nicolas Anciaux [INRIA, Senior Researcher, until May 2024]

Faculty Members

- Philippe Pucheral [UVSQ, Professor]
- Iulian Sandu Popa [UVSQ, Associate Professor, until Sep 2024]

Post-Doctoral Fellow

- Mariem Habibi [INRIA, Post-Doctoral Fellow, until Jun 2024]

PhD Students

- Xinqing Li [INRIA, until Sep 2024]
- Julien Mirval [INRIA, until Feb 2024]
- Ali Ncibi [INRIA]
- Haoying Zhang [INSA CENTRE VDL, until May 2024]

Technical Staff

- Ludovic Javet [INRIA, Engineer]

Interns and Apprentices

- Nathan Champeil [INRIA, Intern, from May 2024 until May 2024]
- Marine Conor [INRIA, Intern, from Jul 2024 until Nov 2024]
- Mayeul Docq [INRIA, Intern, from May 2024 until Aug 2024]
- Abdel-Malik Fofana [INRIA, Apprentice, from Sep 2024]
- Fabien Girard [INRIA, Intern, from May 2024 until May 2024]
- Hairiya Guidado Aissatou [INRIA, Intern, from May 2024 until May 2024]
- Ivan Krivokuca [INRIA, Apprentice, from Sep 2024]
- Lazare Ricour-Dumas [INRIA, Intern, from Jun 2024 until Aug 2024]
- Jasmine Watissee [INRIA, Intern, until Mar 2024]

Administrative Assistant

- Katia Evrat [INRIA]

Visiting Scientist

- Jose Maria De Fuentes [UNIV CARLOS III, until Aug 2024]

External Collaborators

- Xavier Bultel [INSA CENTRE VDL, from Feb 2024 until May 2024]
- Cedric Eichler [INSA CENTRE VDL, from Feb 2024 until May 2024]
- Luis Ibanez Lissen [UNIV CARLOS III, until Jun 2024]
- Benjamin Nguyen [INSA CENTRE VDL, until May 2024]

2 Overall objectives

We are witnessing an exponential accumulation of personal data on central servers: data automatically gathered by administrations and companies but also data produced by individuals themselves (e.g., photos, agendas, data produced by smart appliances and quantified-self devices) and deliberately stored in the cloud for convenience. The net effect is, on the one hand, an unprecedented threat on data privacy due to abusive usage and attacks and, on the other hand, difficulties in providing powerful user-centric services (e.g. personal big data) which require crossing data stored today in isolated silos. The Personal Cloud paradigm holds the promise of a Privacy-by-Design storage and computing platform, where each individual can gather her complete digital environment in one place and share it with applications and users, while preserving her control. However, this paradigm leaves the privacy and security issues in user's hands, which leads to a paradox if we consider the weaknesses of individuals' autonomy in terms of computer security, ability and willingness to administer sharing policies. The challenge is however paramount in a society where emerging economic models are all based - directly or indirectly - on exploiting personal data.

While many research works tackle the organization of the user's workspace, the semantic unification of personal information, the personal data analytics problems, the objective of the PETRUS project-team is to tackle the privacy and security challenges from an architectural point of view. More precisely, our objective is to help providing a technical solution to the personal cloud paradox. More precisely, our goals are (i) to propose new architectures (encompassing both software and hardware aspects) and administration models (decentralized access and usage control models, data sharing, data collection and retention models) for secure personal cloud data management, (ii) to propose new secure distributed database indexing models, privacy preserving query processing strategies and data anonymization techniques for the personal cloud, and (iii) study economic, legal and societal issues linked to secure personal cloud adoption.

3 Research program

To tackle the challenge introduced above, we identify three main lines of research:

- (Axis 1) Personal cloud server architectures and administration models. Based on the intuition that user control, security and privacy are key properties in the definition of trusted personal cloud solutions, our objective is to propose new architectures (encompassing both software and hardware aspects) for secure personal cloud data management. We also focus in this axis on administration models and their enforcement in relation to the architecture of the system, so that the exclusive control of a non expert individual can be ensured.
- (Axis 2) Global query evaluation. The goal of this line of research is to provide capabilities for crossing data belonging to multiple individuals (e.g., performing statistical queries over personal data, computing queries on social graphs or organizing participatory data collection) in a fully decentralized setting while providing strong and personalized privacy guarantees. This means proposing new secure distributed database indexing models and query processing strategies. In addition, we concentrate on locally ensuring to each participant the good behaviour of the processing, such that no collective results can be produced if privacy conditions are not respected by other participants.

- (Axis 3) Technical, legal and economical issues linked to PDMS adoption. This research axis is more transverse and entails multidisciplinary research, addressing the links between economic, legal, societal and technological aspects. We are particularly interested in some specific issues related to the design, implementation and deployment of real PDMS solutions.

Our contributions also rely on tools (algorithms, protocols, proofs, etc.) from other communities, namely security (cryptography, secure multiparty computations, differential privacy, etc.) and distributed systems (distributed hash tables, gossip protocols, etc.). Beyond the research actions, we structure our software activity around advanced platforms integrating our main research contributions. These platforms are cornerstones to help validating our research results through accurate performance measurements, a common practice in the DB community, and target the best conferences. It is also a strong vector to federate the team, simplify the bootstrapping of new PhD or master students, conduct multidisciplinary research and open the way to industrial collaborations and technological transfers. Our main platform is called PlugDB and has reached a high level of maturity. It runs on a microcontroller and is integrated in a real PDMS home box solution deployed in the field for social-medical care. In addition, we are developing a second platform (which is only in the research prototype stage), to provide the user with a PDMS solution that can be hosted on a cloud platform using trusted execution environments (such as Intel SGX) to ensure data privacy and security.

4 Application domains

4.1 Personal cloud, home care, IoT, sensing, surveys

As stated in the software section, the Petrus research strategy aims at materializing its scientific contributions in an advanced hardware/software platform with the expectation to produce a real societal impact. Hence, our software activity is structured around a common Secure Personal Cloud platform rather than several isolated demonstrators. This platform will serve as the foundation to develop a few emblematic applications.

Several privacy-preserving applications can actually be targeted by a Personal Cloud platform, like: (i) smart disclosure applications allowing the individual to recover her personal data from external sources (e.g., bank, online shopping activity, insurance, etc.), integrate them and cross them to perform personal big data tasks (e.g., to improve her budget management) ; (ii) management of personal medical records for care coordination and well-being improvement; (iii) privacy-aware data management for the IoT (e.g., in sensors, quantified-self devices, smart meters); (iv) community-based sensing and community data sharing; (v) privacy-preserving studies (e.g., cohorts, public surveys, privacy-preserving data publishing). Such applications overlap with all the research axes described above but each of them also presents its own specificities. For instance, the smart disclosure applications will focus primarily on sharing models and enforcement, the IoT applications require to look with priority at the embedded data management and sustainability issues, while community-based sensing and privacy-preserving studies demand to study secure and efficient global query processing.

Among these applications domains, one is already receiving a particular attention from our team. Indeed, we gained a strong expertise in the management and protection of healthcare data through our past DMSP (Dossier Medico-Social Partagé) experiment in the field. This expertise is being exploited to develop a dedicated healthcare and well-being personal cloud platform. We are currently deploying 10000 boxes equipped with PlugDB in the context of the DomYcile project. In this context, we are currently setting up an Inria Innovation Lab with the Hippocad company to industrialize this platform and deploy it at large scale (see Section the bilateral contract OwnCare II-Lab).

5 New software, platforms, open data

5.1 New software

5.1.1 PlugDB

Keywords: Databases, Personal information, Privacy, Hardware and Software Platform

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 prototype version of 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 tamper-resistant device. 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). Then, PlugDB was extended to run both on secure devices provided by Gemalto and on specific secure devices designed by PETRUS and assembled by electronic SMEs. Mastering the hardware platform opens up new research and experiment opportunities (e.g., support for wireless communication, secure authentication, sensing capabilities, battery powered ...).

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. PlugDB was used in an educational platform that we set up : SIPD (Système d'Information Privacy- by-Design). SIPD was used at ENSIIE, INSA CVL and UVSQ through the Versailles Sciences Lab fablab, to raise students awareness of privacy protection problems and embedded programming.

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.

PlugDB is currently industrialized in the context of the OwnCare Inria Innovation Lab (II-Lab). In OwnCare, PlugDB acts as a secure personal cloud to manage medical/social data for people receiving care at home. It is currently being deployed over 10.000 patient in the Yvelines district. The industrialization process covers the development of a complete testing environment, the writing of a detailed documentation and the development of additional features (e.g., embedded ODBC driver, TPM support, flexible access control model and embedded code upgrade notably). It has also required the design of a new hardware platform equipped with a battery power supply, introducing new energy consumption issues for the embedded software.

URL: https://project.inria.fr/plugdb/

Contact: Luc Bouganim

5.2 New platforms

Participants: Nicolas Anciaux, Luc Bouganim, Iulian Sandu Popa *(correspondent)*.

Personal Data Management Systems (PDMS) arrive at a rapid pace boosted by smart disclosure initiatives and new regulations such as GDPR. However, our survey [1] indicates that the existing PDMS solutions cover partially the PDMS data life-cycle and, more importantly, focus on specific privacy threats depending on the employed architecture. To address this issue, we proposed in [1] a logical reference architecture for an extensive (i.e., covering all the major functionalities) and secure (i.e., circumventing all the threats specific to the PDMS context) PDMS. We also discussed several possible physical instances fo the architecture and showed that TEEs (Trusted Execution Environments) are a prime option for building a trustworthy PDMS platform [2].

Hence, based on our previous studies, we have developed a first prototype of an extensive and secure PDMS (ES-PDMS) platform using the state-of-the-art TEE technology available today, i.e., Intel Software Guard eXtension (SGX). The originality of our approach is to achieve extensibility through a set of isolated data-oriented tasks potentially untrusted by the PDMS owner, running alongside a trusted module which controls the complete workflow and limits data leakage. Our ES-PDMS software stack can be deployed on any SGX-enabled machine (i.e., any relatively recent computer having an Intel CPU). This prototype was presented in a demonstration paper [6] focusing on security properties of the platform with the help of several concrete scenarios and interactive games.

Xinqing Li's PhD thesis, which started in Oct 2023, takes up some of the plateforms's elements and adapts them to the context of a cloud database service (going beyond that of PDMS) with code extensions potentially vulnerable to certain attacks.

6 New results

6.1 A new PET for Data Collection via Forms with Data Minimization, Full Accuracy and Informed Consent (Axis 1)

Participants: Nicolas Anciaux *(correspondent)*, Benjamin Nguyen.

Data minimization is a privacy principle of the GDPR, which states that the collection of personal data must be minimized according to the purpose of the intended processing. We have proposed a "privacy-enhancing technology" (PET) for the collection of personal data, specifically targeting public sectors such as social services, where forms are filled in by applicants to apply for certain benefits. Our approach, based on classical logic and game theory, aims to minimize data collection while ensuring that users make informed choices and obtain all due benefits. It offers a practical solution for preserving privacy without compromising service accuracy. We also demonstrated the implementation of our data minimization model in two real-life scenarios concerning French social benefits. This work resulted in an article in EDBT'24 [3]. This work was done in collaboration with Sabine Frittella (INSA-VDL) and Guillaume Scerri (ENS Paris Saclay).

6.2 Decentralized and Reliable Secure Aggregation Protocols (Axis 2)

Participants: Luc Bouganim, Julien Mirval, Iulian Sandu Popa *(correspondent)*
.

The development and adoption of personal data management systems (PDMS) has been fueled by legal and technical means such as smart disclosure, data portability and data altruism. By using a PDMS, individuals can effortlessly gather and share data, generated directly by their devices or as a result of their interactions with companies or institutions. In this context, federated learning appears to be a very promising technology, but it requires secure, reliable, and scalable aggregation protocols to preserve user privacy and account for potential PDMS dropouts. Despite recent significant progress in secure aggregation for federated learning, we still lack a solution suitable for the fully decentralized PDMS context. In this work, we proposed a family of fully decentralized protocols that are scalable and reliable with respect to dropouts. We focused in particular on the reliability property which is key in a peer-to-peer system wherein aggregators are system nodes and are subject to dropouts in the same way as contributor nodes. We showed that in a decentralized setting, reliability raises a tension between the potential completeness of the result and the aggregation cost. We proposed a set of strategies that deal with dropouts and offer different trade-offs between completeness and cost. We extensively evaluated the proposed protocols and showed that they cover the design space allowing to favor completeness or cost in all settings. This work was published in TLDKS [12] and is part of Julien Mirval's Phd [12].

6.3 Edgelet Computing: Opportunistic Queries on Secure Edges (Axis 3)

Participants: Nicolas Anciaux, Luc Bouganim, Ludovic Javet *(correspondent)*, Philippe Pucheral.

The convergence of Opportunistic Networks and Trusted Execution Environments at the network edge presents a compelling opportunity for fully decentralized privacy-preserving data processing. Based on this convergence, we defined the concept of Edgelet computing, a new paradigm for executing powerful and privacy-preserving distributed queries on personal devices. Our objective was to establish a robust, secure, and scalable execution framework with strong individual privacy guarantees. In this work, we first propose a liability model tailored to decentralized executions on crowd members' devices, along with a query evaluation model that differs from the traditional database closed-world assumption. Second, we defined essential properties for ensuring the security, resiliency and validity of executions, and subsequently presents several methods and strategies for their enforcement. Through a comprehensive qualitative analysis and extensive evaluations, we showcased the relevance and effectiveness of the approach, demonstrating that Edgelet Computing holds potential for the emergence of novel and important classes of applications. This work has been published in the Personal and Ubiquitous Computing journal [11].

6.4 PETSCRAFT creation

Participants: Nicolas Anciaux, Mariem Brahem, Nathan Champeil, José Maria de Fuentes, Fabien Girard, Hairiya Guidado Aissatou, Benjamin Nguyen *(correspondent)*, Jasmine Watissee, Haoying Zhang.

Since July 2023, a PETRUS spin-off has been incubating around the development of Privacy-Enhancing Technologies (PETs). This project-team, called PETSCRAFT, is a collaboration between Inria Saclay and INSA-CVL, with the primary goal of modeling and implementing PETs to protect personal data in various contexts. During the begining of year 2024, building the scientific methodology used in our work on a new PET for Data Collection [3] (see Section 6.1), we initiated new collaborative research directions between Inria and Insa-CVL, focusing on (1) privacy control in new data structures such as Matrix Profile, (2) the use of large language models (LLMs) to preserve text anonymity (LLM-Privacy), and (3) copyright and privacy issues arising from the proliferation of LLMs (Comply-LLM) in partnership with legal experts. For detailed results on these research directions after June 1, 2024, including related publications [14, 15], please refer to the PETSCRAFT evaluation report.

7 Bilateral contracts and grants with industry

7.1 Bilateral contracts with industry

OwnCare-2 IILab (Jan 2022 - Dec 2025) - Partners: PETRUS, Hippocad

Participants: Nicolas Anciaux, Luc Bouganim, Ludovic Javet, Philippe Pucheral *(correspondent)*, Laurent Schneider.

The OwnCare IILab – Inria Innovation Lab - (Jan 2018-Dec 2021) aimed at conceiving a secured personal medical folder facilitating the organization of medical and social care provided at home to elderly people and at deploying it in the field. This IILab has been built in partnership with the Hippocad company which won, in association with Inria and UVSQ, a public call for tender launched by the Yvelines district to deploy this medical folder on the whole distinct (10.000 patients). This solution, named DomYcile in the Yvelines district, is based on a home box combining the PlugDB hardware/software

technology developed by the Petrus team (to manage and secure the medical folder) and additional technology developed by Hippocad. The primary result of the OwnCare IILab has been to build a concrete industrial solution based on PlugDB and deploy it so far among 3000 patients in the Yvelines district, despite the Covid pandemia. In 2022, Hippocad has become a subsidiary of the La Poste group opening new opportunities in terms of deployment. Hence, Inria, UVSQ and Hippocad have launched a follow up of the OwnCare IILab for the period Jan 2022-Dec 2025. The goal of the OwnCare2 IILab is (1) to integrate our solution in the MaSanté 2022 national roadmap by making it interoperable with external services (without hurting the security provided by the box), (2) to handle, in a privacy-preserving way, new usages like actimetrics, teleassistance and global statistics based on IoT techniques, machine learning and decentralized computations and (3) try to deploy it at the national/international level. In 2023, a new district (Hauts de Seine) has decided to deploy the DomYcile solution on its own territory, leading to an extended partnership.

8 Partnerships and cooperations

8.1 International research visitors

8.1.1 Visits of international scientists

Other international visits to the team

José M. de Fuentes:

Status Associate Professor

Institution of origin: Carlos III University of Madrid (UC3M)

Country: Spain

Dates: from 1 Sept. 2023 to 31 Aug. 2024

Context of the visit: The aim of Prof. José M. de Fuentes' stay in the Petrus project-team is to combine our respective expertise to propose new protection mechanisms that are adapted to today's real-world personal data clouds. Two lines of collaborative research are being investigated: (1) the combination of database techniques (such as matrix profile) and artificial intelligence for authenticating data sources, and (2) privacy protection techniques for sharing and exploiting personal data in specific contexts such as home surveillance and remote working. The one-year research stay is fully covered by a UC3M grant.

Mobility program/type of mobility: Research stay

8.2 National initiatives

8.2.1 iPoP, interdisciplinary Project on Privacy, PEPR Cybersécurité (July 2022 - June 2028)

Partners: Inria, CNRS, EDHEC, INSA CVL, INSA Lyon, UGA, Université de Lille, Université Rennes 1, UVSQ, CNIL

Digital technologies provide services that can greatly increase quality of life (e.g. connected e-health devices, location based services or personal assistants). However, these services can also raise major privacy risks, as they involve personal data, or even sensitive data. Indeed, this notion of personal data is the cornerstone of French and European regulations, since processing such data triggers a series of obligations that the data controller must abide by. This raises many multidisciplinary issues, as the challenges are not only technological, but also societal, judiciary, economic, political and ethical. The objectives of this project are thus to study the threats on privacy that have been introduced by these new services, and to conceive theoretical and technical privacy-preserving solutions that are compatible with French and European regulations, that preserve the quality of experience of the users. These solutions will

be deployed and assessed, both on the technological and legal sides, and on their societal acceptability. In order to achieve these objectives, we adopt an interdisciplinary approach, bringing together many diverse fields: computer science, technology, engineering, social sciences, economy and law.

The project's scientific program focuses on new forms of personal information collection, on the learning of Artificial Intelligence (AI) models that preserve the confidentiality of personal information used, on data anonymization techniques, on securing personal data management systems, on differential privacy, on personal data legal protection and compliance, and all the associated societal and ethical considerations. This unifying interdisciplinary research program brings together internationally recognized research teams (from universities, engineering schools and institutions) working on privacy, and the French Data Protection Authority (CNIL).

8.2.2 YPPOG, Youth Privacy Protection in Online Gaming, DATAIA project (Sept. 2021 - Sept. 2024)

Partners: CERDI (Université Paris Saclay), LITEM (IMT-BS), PETRUS (Inria-UVSQ).

Youth Privacy Protection in Online Gaming is a crucial topic in the field of Privacy. Approximately 94% of minor children play video games, and 60% of 10-17 year olds play online, but only 12% of parents actively monitor their children's activities. The YPPOG project, involving researchers in law, economics, and computer science, explores multidisciplinary questions such as minor consent and GDPR enforcement online. In 2022, PETRUS developed a technical prototype for managing the data of League of Legends players, and in 2023, this prototype was extended to other games and platforms. The CNIL is particularly interested in this project and is closely monitoring its progress.

9 Dissemination

9.1 Promoting scientific activities

9.1.1 Scientific events: organisation

General chair, scientific chair

- Luc Bouganim: Co-organizer "École thématique BDA Masses de Données Distribuées", Ceillac (2024)
- Iulian Sandu Popa: Proceedings Chair IEEE ICDCS'24

9.1.2 Scientific events: selection

Member of the conference program committees

- Nicolas Anciaux: EDBT'24, WISE'24, PaveTrust @ FM'24 (int. symposium on Formal Methods)
- Iulian Sandu Popa: BDA'24, DATA'24, SSDBM'24, ACM Sigmod'25.

9.1.3 Journal

Reviewer - reviewing activities

- Iulian Sandu Popa: DAPD'24.

9.1.4 Scientific expertise

- Nicolas Anciaux: Member of the jury of the 9th edition of CNIL-Inria Privacy Award 2024.
- Nicolas Anciaux: Member of the Recruitment Jury for DR2 positions at Inria in 2024
- Nicolas Anciaux: Vice-president of the Recruitment Jury for CRCN-ISFP positions at Inria Saclay in 2024

- Iulian Sandu Popa: Member of the selection committee (COS) for the position of associate professor - UVSQ.
- Philippe Pucheral: Member of the jury of the SPRINGCS call from the ISN Graduate School

9.1.5 Research administration

- Nicolas Anciaux : Deputy Scientific Delegate (DSA) at Inria Saclay, and de facto Member of the Inria Evaluation Committee (CE) and Member of the Inria Scientific Committee (CoSi)
- Nicolas Anciaux : Member of the Inria Saclay Technology Development Committee (CDT)
- Nicolas Anciaux : Member of the UPSaclay Research Committee (CR du CAC, CR HDR)
- Nicolas Anciaux : Member of the UPSaclay Research and Valorization Direction Committee (Co-direV)
- Iulian Sandu Popa: Member of the ATER committee at UVSQ.
- Iulian Sandu Popa: Member of the Bureau of David lab at UVSQ.
- Luc Bouganim: PhD thesis referent for the Doctoral School of University Paris-Saclay
- Philippe Pucheral: Member of the Scientific Commission (CS) of the ISN Graduate School of Paris-Saclay University.

9.2 Teaching - Supervision - Juries

9.2.1 Teaching

- Philippe Pucheral: head of the M1 and M2 DataScale master program at University Paris-Saclay.
- Master: Iulian Sandu Popa, Bases de données relationnelles (niveau M1), Gestion des données spatiotemporelles (niveau M2), Sécurité des bases de données (niveau M2), 96, UVSQ, France. Philippe Pucheral, courses in M1 and M2 in databases and in security, introductory courses for jurists,UVSQ, France.
- Licence: Iulian Sandu Popa, Bases de données (niveau L2), 96, UVSQ, France.
- Engineers school: Nicolas Anciaux, Databases (ENSTA, module IN206, M1), 32, and Advanced databases (ENSTA, module ASI13, level M2), 32. Luc Bouganim, Bases de données relationnelles (ENSTA, module IN207, M1), 32.
- MOOC: Défis technologiques des villes intelligentes participatives. Co-Auteurs: Nicolas Anciaux, Stéphane Grumbach, Valérie Issarny, Nathalie Mitton, Christine Morin, Animesh Pathak et Hervé Rivano. Sessions sur la plateforme FUN en 2022. The Mooc will go into "archived open for registration" mode from April 2024. To date, the Mooc has 21320 registered users, with a total of 2041 badges and certificates of achievement issued since April 2019.

9.2.2 Supervision

- Defended PhD (March 11, 2024): Julien Mirval, DISSEC-ML: towards distributed and secured machine learning in the personal cloud, Luc Bouganim and Iulian Sandu Popa.
- PhD in progress: Ali Ncibi, Secure machine Learning on IOT traces for daily activity discovery, Inria, since March 2023, Luc Bouganim and Philippe Pucheral.
- PhD in progress: Xinqing Li, Securing Algorithms for Classification and Machine Learning on Personal Data using Trusted Execution Environments, Inria, since October 2023, Nicolas Anciaux and Iulian Sandu Popa.

- PhD in progress: Haoying Zhang, Privacy-Enhancing Technologies for telework data sharing: an approach based on informed user consent, INSA-CVL, since Sept 2023, Nicolas Anciaux and Benjamin Nguyen.

9.2.3 Juries

- Nicolas Anciaux: Member of the HDR Jury of Sami Zhioua (IPParis), June 2024
- Nicolas Anciaux: Member (invited) of the PhD jury of S. Tanigassalame (IPParis), May 2024
- Luc Bouganim: Reviewer of the PhD of Niclas Hedam (ITU Copenhagen), october 2024.

10 Scientific production

10.1 Major publications

[1] N. Anciaux, P. Bonnet, L. Bouganim, B. Nguyen, P. Pucheral, I. Sandu-Popa and G. Scerri. 'Personal Data Management Systems: The security and functionality standpoint'. In: *Information Systems* 80 (2019), pp. 13–35. DOI: 10.1016/j.is.2018.09.002. URL: https://hal.archives-ouvertes.fr/hal-01898705 (cit. on p. 5).

[2] N. Anciaux, L. Bouganim, P. Pucheral, I. S. Popa and G. Scerri. 'Personal Database Security and Trusted Execution Environments: A Tutorial at the Crossroads'. In: *Proceedings of the VLDB Endowment (PVLDB)* (Aug. 2019). DOI: 10.14778/3352063.3352118. URL: https://hal.inria.fr/hal-02269292 (cit. on p. 5).

[3] N. Anciaux, S. Frittella, B. Joffroy, B. Nguyen and G. Scerri. 'A new PET for Data Collection via Forms with Data Minimization, Full Accuracy and Informed Consent'. In: *EDBT*. 27th International Conference on Extending Database Technology, EDBT 2024. Paestum, Italy, 25th Mar. 2024. URL: https://inria.hal.science/hal-04149000 (cit. on pp. 6, 7).

[4] L. Bouganim, J. Loudet and I. Sandu Popa. 'Highly distributed and privacy-preserving queries on personal data management systems'. In: *The VLDB Journal* 32.2 (1st Mar. 2023), pp. 415–445. DOI: 10.1007/s00778-022-00753-1. URL: https://inria.hal.science/hal-03814840.

[5] M. Brahem, N. Anciaux, V. Issarny and G. Scerri. 'Consent-driven Data Reuse in Multi-tasking Crowdsensing Systems: A Privacy-by-Design Solution'. In: *Pervasive and Mobile Computing* 83 (2022). DOI: 10.1016/j.pmcj.2022.101614. URL: https://hal.science/hal-03775759.

[6] R. Carpentier, F. Thiant, I. Sandu Popa, N. Anciaux and L. Bouganim. 'An Extensive and Secure Personal Data Management System Using SGX'. In: EDBT 2022 - 25th International Conference on Extending Database Technology. Edinburgh / Virtual, United Kingdom, 29th Mar. 2022. URL: https://inria.hal.science/hal-03580286 (cit. on p. 6).

[7] L. Javet, N. Anciaux, L. Bouganim and P. Pucheral. 'Edgelet Computing: Pushing Query Processing and Liability at the Extreme Edge of the Network'. In: CCGrid 2022. Taormina, Italy, 16th May 2022. URL: https://hal.inria.fr/hal-03666895.

[8] R. Ladjel, N. Anciaux, P. Pucheral and G. Scerri. 'Trustworthy Distributed Computations on Personal Data Using Trusted Execution Environments'. In: *TrustCom 2019 - The 18th IEEE International Conference on Trust, Security and Privacy in Computing and Communications / BigDataSE 2019 - 13th IEEE International Conference on Big Data Science and Engineering*. Rotorua, New Zealand, Aug. 2019. DOI: 10.1109/TrustCom/BigDataSE.2019.00058. URL: https://hal.archives-ouvertes.fr/hal-02269207.

[9] J. Loudet, I. Sandu-Popa and L. Bouganim. 'SEP2P: Secure and Efficient P2P Personal Data Processing'. In: *EDBT 2019 - 22nd International Conference on Extending Database Technology*. Lisbon, Portugal, Mar. 2019. URL: https://hal.inria.fr/hal-01949641.

[10] I. Sandu Popa, D. H. T. That, K. Zeitouni and C. Borcea. 'Mobile participatory sensing with strong privacy guarantees using secure probes'. In: *Geoinformatica* 25.3 (July 2021), pp. 533–580. DOI: 10.1007/s10707-019-00389-4. URL: https://hal.science/hal-03329908.

10.2 Publications of the year

International journals

[11] L. Javet, N. Anciaux, L. Bouganim and P. Pucheral. 'Edgelet Computing: Enabling Privacy-Preserving Decentralized Data Processing at the Network Edge'. In: *Personal and Ubiquitous Computing* (14th June 2024). DOI: 10.1007/s00779-024-01821-9. URL: https://inria.hal.science/hal-04594280 (cit. on p. 7).

[12] J. Mirval, L. Bouganim and I. Sandu Popa. 'Handling Dropouts in Federating Learning with Personal Data Management Systems'. In: *Transactions on Large-Scale Data- and Knowledge-Centered Systems* (17th May 2024). URL: https://inria.hal.science/hal-04598486. In press (cit. on p. 6).

International peer-reviewed conferences

[13] N. Anciaux, S. Frittella, B. Joffroy, B. Nguyen and G. Scerri. 'A new PET for Data Collection via Forms with Data Minimization, Full Accuracy and Informed Consent'. In: *EDBT*. 27th International Conference on Extending Database Technology, EDBT 2024. Paestum, Italy, 25th Mar. 2024. URL: https://inria.hal.science/hal-04149000.

[14] L. Ibanez-Lissen, J. Maria de Fuentes, L. Gonzales-Manzano and N. Anciaux. 'Continuous Authentication Leveraging Matrix Profile'. In: ARES 2024 - The 19th International Conference on Availability, Reliability and Security. Vienne, Austria, 30th July 2024. URL: https://inria.hal.science/hal-04663471 (cit. on p. 7).

Conferences without proceedings

[15] X. Li, I. Sandu Popa and N. Anciaux. 'Extensive and Secure Data Management System with Vulnerable Extension Code'. In: APVP 2024 - 14ème Atelier sur la Protection de la Vie Privée. Lyon, France, 24th June 2024. URL: https://inria.hal.science/hal-04598521 (cit. on p. 7).

Doctoral dissertations and habilitation theses

[16] J. Mirval. 'DISSEC-ML : towards distributed and secured machine learning in the personal cloud'. Université Paris-Saclay, 11th Mar. 2024. URL: https://theses.hal.science/tel-04889175.