7. SMIS - Secured and Mobile Information Systems

Leader : Philippe Pucheral

Research center : Paris - Rocquencourt

Field : Perception, Cognition and Interaction

http://www.inria.fr/en/domains/perception-cognition-and-interaction

Theme : Data and Knowledge Representation and Processing

http://www.inria.fr/en/domains/perception-cognition-and-interaction/(theme)/97608

CNRS,Université Versailles Saint-Quentin, Parallelisme, réseaux, systèmes, modélisation (PRISM) (UMR8144)

Evaluating experts: Anastasia Ailamaki (Chair), Yannick Cras, Fosca Giannotti

Note: The reviewers unanimously congratulate the team for their preparation, collaborative spirit, and candidness during this review, which made possible a thorough evaluation of their achievements.

7.1 Primary Topics and Objectives

SMIS focuses on several aspects of IoT data management challenges, including flash-related storage devices (the team has a seminal paper in the area), and security (malicious attacks). The overarching vision of the Personal Cloud includes and connects two parallel lines of work on embedded data management and on secure distributed computation.

More specifically, the top goals of the project are:
1. **personal data management in very constrained hardware through techniques of pipelined indexing on document stores,**
2. storage of massively-indexed, log-only databases, using flash as the most convenient technology for storage, indexing, and handling of malicious attacks on the flash device
3. design and implementation of privacy-preserving query execution using both hardware and software methods
4. cross-layer optimisations for optimising data management performance on flash devices.

These research topics are particularly timely: the growing market of personal devices, while from one side is demanding for increased processing capability and higher intelligence levels at the device level, and on the other side it poses two doubly interleaved challenges: to provide trusted processing framework, and to create the context for which the persons can take advantage of such enormous collection of their personal traces on the various dimensions of their lives.

From an initial focus on the development of data management capabilities under constrained hardware capabilities, the team has indeed tackled and provided contributions to a broader range of important areas – from low-level technology issues regarding the optimization of data structures adapted to the inherent properties of Flash storage to socially impactful issues of privacy and ethics, also considering the architectural impact of achieving scalability in conjunction with secure computing. This range of domains is nevertheless kept coherent under the umbrella vision for personal clouds being developed by the team.

7.2 International Standing and Reputation

The key SMIS members (Bouganim, Pucheral, Anciaux) and the external collaborator Bonnet are top-level researchers and known internationally as authorities in the field. Their paper on flash devices at CIDR 2009 is considered seminal in the area, and they have regular publications at top-level venues (CIDR, VLDB).

The overall team has displayed an excellent record of participation in major conferences and journals (including the most competitive database conferences SIGMOD, VLDB, and CIDR). This demonstrates that the vision, and the consequent research activities, elaborated in this last five years are truly disruptive, and the importance is higher because their results are based on system implementation. The team have also proven academic collaboration through an interesting example and a live presentation by the partner.

7.3 Major Achievements and Impact (Theory, Research Software, etc.)

This is a fascinating team project with a lot of very influential work. Except for the award-winning work on data management on top of flash storage, they have a principled approach on the Personal Cloud with clear requirements and applicable results. Their work on flash has already found its way to products. Another impressive factor is the assembly of skills demonstrated by the team in the service of their vision, and the variety of the achieved results which broaden their impact. It is our suggestion to the team to design and propose a seminar to teach the other teams in INRIA on the challenges and opportunities on combining cutting-edge research and active, pipelined technology transfer, which is exemplary in this team.

The team has invented novel data structures and methods to achieve efficient querying under constrained RAM and avoid the problem of random writes in Flash. The reviewers particularly appreciated the care that they took to illustrate why the problems they tackle – low ratio of RAM to storage, performance of random writes – are not tactical and won't be solved only by mechanical evolution of the underlying technologies. In addition, such techniques might also find applications in other areas, for instance in the case of in-memory columnar storage where it is desired to limit RAM usage outside of the storage itself.

At the other end of the spectrum, the practical experiments the team have conducted with PlugDB have led to interesting insights and questions regarding the ethics and social acceptability aspects of the solution – first of all the question of who should be assigned the responsibility to administer access

permissions of the personal device in the case of a dependent person. It would be of major interest to further formalize these issues and, as the teams intends, to work with a pluridisciplinary group. Since a project lab (CAPRISS) exists around privacy, it might be interesting to make sure this lab as a whole establishes relevant links with the community of jurists, specialists in ethic, and decision makers.

The secure calculation protocols exemplified by the team are also of great interest in the perspective of personal cloud architectures.

The demonstration of a working piece of hardware and professionally designed application on top provided for a vivid and compelling illustration of the vision.

7.4 Industry Transfer and Partnership

The team showed an impressive demo on the medical application enabled by the technology. It is commendable that the team are working hard on their research (and publishing at top venues) and at the same time they liaise with companies, which takes at least as much time as the research itself. The practical experiments they did with PlugDB, which clearly have required intense interaction with local policy makers and healthcare and social assistance professional, as well as the ATOS subsidiary specialized in healthcare, is also very noteworthy. In the case of the medical app demo, SanteOS develops the application front-end and they also have to interact with doctors and other hospital personnel for feedback on the use of the login/fingerprint reading device. They have partnered closely with the CozyCloud startup, and done so in a way that preserves their autonomy as a research team while providing measurable value to the partner. There is also a tight collaboration / partnership with Gemalto, the world's leading smart card provider. Technology transfer is therefore clearly an area of focus for the SMIS project and happens almost automatically. It seems as if people are conducting their research with the constant feedback of the users and target applications, which is an excellent idea.

7.5 Training of Personnel

The PlugDB software is used as a basis for training new graduate students. This means that the incoming PhD students do not need to spend time developing infrastructure, and they can "hit the ground running" with a platform that has been tested in research as well as in real application settings. This helps alleviate the inefficiencies resulting from the universally applied (and very limiting) rule that PhDs have a maximum duration of 3 years.

7.6 Principal Strengths and Weaknesses of the Project

The project should involve more people in deployment of the ideas and reach out / create industrial exploitation opportunities.

Strenghts:

- The project has a focused vision which translates into a very elaborate and well-managed research program.
- The large breadth of approaches and deployed solutions is exciting and of high educational and practical value.

- The work is at the top of the fields of research it covers and more importantly takes all ideas to production. The combination of these properties is rare in research projects which explore deep modelling algorithms.
- The use cases are practical and concrete, and are complete with working implementations.
- SMIS has a collaboration with with the SECRET project team which works on cryptography (in relation to the goal #2 above).

Weaknesses (none of which are the project team's fault):

- There is limited availability of engineering resources including qualified PhD students
- There are inherent difficulties to spread such technology in the healthcare space.

7.7 Plan for the next period (4 years)

As the project concludes this term, there is no plan for the next period per se. Nevertheless, the team presented us with a coherent and compelling perspective of where a new project could find fruitful research topics.

7.8 Opportunities and risks/difficulties faced by the project

Risks:

- the legal, ethical and social acceptability issues may well dominate the scientific and technology issues when it comes to transfer. Special attention should be paid to this.
- The industry will move fast in this area which is hot – the team will need to be reactive to what happens on that front and identify the right messages/differentiators to use when they speak to an industrial audience.

Difficulties:

- As SMIS is about very applied work, the team need engineers. However, they can only give contracts to engineers for six months, one year, or a maximum two years. This is detrimental because (a) it breaks continuity of the work (b) highly qualified engineer candidates look for contracts with a bit longer renewal prospect so that they can develop their skills along with the team.

7.9 Recommended actions and suggested measures of success

Please see suggestions in previous sections, and also

- Engage into coordinated discussions with experts and policy makers about the ethical and privacy concepts behind the personal cloud. E.G, "conseil du Numérique". Maybe CAPRISS or some other cross-team INRIA entity should engage on this aspect, which is not only sensitive in the case of healthcare.
- An opportunity may lay in a european initiative: EIT-ICT lab in which INRIA is involved that has launched an High Impact action aimed at defining the requirements of a trusted European Cloud: https://www.eitdigital.eu/innovation-entrepreneurship/future-cloud/. The group might have interesting contributions for a related scenario.