INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team adept

Algorithms for Dynamic Dependable Systems

Rennes - Bretagne-Atlantique

Theme : Distributed Systems and Services

Table of contents

1. Team **1**
2. Overall Objectives **1**
3. Scientific Foundations **2**
4. Application Domains **3**
4.1. Space Domain Applications 3
4.2. Telecommunication Applications 3
5. Software **3**
6. New Results **4**
6.1. Evaluation of Consensus Protocols 4
6.2. Mobile Agent and Transactions 5
6.3. Framework for Proving Self-Organization 5
6.4. Persistent Feedback 6
6.5. Induced Churn to Face Malicious Behaviors 6
6.6. Privacy-preserving Identification System 8
6.7. Privacy in Social Networking Sites 8
6.8. Geo-privacy 8
7. Contracts and Grants with Industry **8**
8. Other Grants and Activities **9**
8.1. International Initiatives 9
8.2. Exterior research visitors 9
9. Dissemination **9**
9.1. Animation of the scientific community 9
9.2. Teaching 11
10. Bibliography **11**

ADEPT *is a common project with CNRS and the University of Rennes I.*

1. Team

Research Scientists

Michel Hurfin [Team leader, Junior Researcher (CR) INRIA, HdR]
Emmanuelle Anceaume [Junior Researcher (CR) CNRS]

Faculty Members

Sébastien Gambs [Associate Professor (MdC) Université de Rennes 1, INRIA research chair in Security of Information Systems]
Frédéric Majorczyk [ATER Université de Rennes 1, since September 2010]

Technical Staff

Romaric Ludinard [Technical Staff, INRIA, until March 2010]

PhD Students

Ahmed Gmati [PhD Student, Université de Rennes 1, since December 2010]
Izabela Moise [PhD Student, Université de Rennes 1]
Heverson Borba Ribeiro [PhD Student, Université de Rennes 1, until September 2010]

Visiting Scientist

Linda Zeghache [PhD Student, USTBH-CEDRIC (université des Sciences et de la Technologie Houari Boumédiène, Algeria), Décember 2010]

Administrative Assistant

Lydie Mabil [Administrative Assistant, INRIA]

Other

Jean-Pierre Le Narzul [External Collaborator, Associate Professor, Institut Télécom / Telecom Bretagne, until July 2010]

2. Overall Objectives

2.1. Overall Objectives

Information technologies are evolving and maturing at a very high pace. Networks and connected entities have progressed so much that their improvements have induced radical changes in the very nature of distributed applications. Many information systems are now based on massively networked devices that support a large population of interacting and cooperating entities. While computer-based systems become increasingly open and complex, accidental and intentional failures tend to get considerably more frequent and severe. In the context of large-scale distributed and dynamic systems, interacting with unknown entities becomes an unavoidable habit despite the induced risk.

In the field of distributed systems and algorithms, the ADEPT team is focusing on dependability and security issues (namely reliability, availability, integrity, confidentiality, and privacy). Our main objective is to study and design services based on detection and protection mechanisms for open environments.

The design of dependable mechanisms mainly depends on the types of faults that might occur during the computation. Benign faults (crash, omission, ...) are distinguished from the arbitrary faults (Byzantine faults). In the former case, processes behave according to their specification but after some time they may omit some (or all) computation steps. In the latter case, processes involved in the computation may arbitrarily deviate from their specification. Such faults can be the consequence of malicious intents of individuals. While an active adversary may trigger either benign or malign faults, a passive adversary which just observes the protocol behavior has also to be considered in order to protect the privacy of the interacting entities.

Our contributions focus on the three following themes:

- **Dependability and group of nodes.** Many groups of nodes include only a limited number of cooperating processes (e.g. sets of replicas) or are the results of a decomposition of the whole system into several sub-systems (e.g. hierarchies, clusters, neighborhoods, or communities of interests). In this context we aim to consider both accidental and intentional faults and to design algorithms and methods to detect or to mask such faults which are sometimes transient (another dynamic aspect). An important part of our activity is dedicated to the study of synchronization and agreement problems and to their use in group communication services.
- **Reputation in large scale distributed systems.** We consider different types of large scale systems and study the main dependability issues that are associated with. To reduce the risk to rely on dishonest entities, a *reputation mechanism* is an essential prevention tool that aims at measuring the capacity of a remote node to provide a correct service. Such a mechanism should allow to overcome ill-founded suspicions and to be aware of established misbehaviors. It can be used to punish nodes displaying a malicious behavior.
- **Privacy enhancing technologies.** The protection of privacy is now recognized as a fundamental right of individuals. Yet, very few systems tackle the issue of guaranteeing its respect. We investigate the preservation of privacy in various contexts such as social networks, geo-privacy, privacy-preserving data mining and identity management systems.

3. Scientific Foundations

3.1. Dependable Distributed Computing

Economic activities and human lives are now heavily dependent on distributed systems and applications. When computing resources and stored data can be affected by the occurrence of failures, dependability becomes a crucial issue.

When a low level of dynamicity (also called churn) is assumed or when the system size is rather small, a process involved in a distributed computation may know and observe all the other participants. Distributed applications often rely on the identification of such sets of interacting entities. These small sets are either called groups, clusters, collections, neighborhoods, spheres, or communities according to the criteria that define the membership. The adopted criteria may for instance reflect the fact that its members are administrated by a unique person, that they share a unique security policy, that they are located in closed physical places, that they need to be strongly synchronized, that they cooperate together, or that they share mutual interests. When all the participants can share a common knowledge of the group of interacting processes, various fundamental problems (related to observation and synchronization) can be defined. Adaptive algorithms can be proposed to detect a modification of the whole execution context and react globally to this modification (reconfiguration, execution of another code, ...). In particular, to cope with the dynamic evolution of a distributed system, the Group paradigm (and the associated concept of membership service) allows to efficiently address dependability issues. Solutions to agreement problems (such as the consensus problem) can be used as basic building blocks for designing solutions to higher level protocols that are in charge of maintaining global properties at the group level despite the occurrence of faults within the group. Due to the increasing adversity of the system (asynchrony and failures), the design of efficient solutions that are simple to deploy and easy to adapt remains a difficult issue.

When the system has a very high level of churn, implementing a global observation mechanism that allows to reconfigure the whole system in a single step is no more realistic. Only local observations and progressive adaptations to changes can be performed on cohesive subsets of nodes. Such a radical gap on the scale and dynamicity of systems militates in favor of a paradigm shift for designing solutions to the problems raised by these new systems. Several partial and inconsistent views of the system may coexist (each participant may have its own view). All classical distributed computing problems (for example, dependability issues,

communication problems, resource allocation, and data management) require new solutions that address these challenges in the new settings. In the context of large-scale distributed and dynamic systems, interaction with unknown entities becomes an unavoidable habit despite the induced risk. In this general context, we consider mainly reputation mechanisms in P2P systems and privacy protection issues.

Our scientific contributions aim to reach a deeper understanding of some fundamental problems that arise in dynamic distributed systems prone to accidental/intentional failures. We consider mainly problems corresponding to middleware services that need to be correctly and continuously provided to the upper-layer entities despite the occurrence of faults.

During the study of a particular problem, we aim to design, for a particular execution environment (characterized by a set of assumptions on the computation model, the failure model, the dynamicity, the scalability, ...), efficient algorithmic solutions that are optimal and generic if possible. If no solution exists, we aim at exhibiting impossibility results. To validate and to promote the use of these algorithmic solutions, we conduct in parallel experimental evaluations by developing flexible and adaptive middleware services that integrate our know-how and experience in distributed computing. This prototyping activity leads us to consider technical and operational problems as well as methodological issues. The feed-back that we receive helps us to define new directions in our research activity.

4. Application Domains

4.1. Space Domain Applications

To cope with more and more complex requirements, this sector of activity shows a growing interest in distributed computing. More precisely, the adequacy between the properties ensured by their applications (that are getting increasingly stronger) and the assumptions about their systems (that are getting inexorably weaker) becomes questionable. In particular, regarding fault tolerance, a large number of entities (software and hardware entities) of the embedded computer-based system interact with each other. To make interaction robust, a broad range of failures (from benign failures up to malicious failures) have to be tolerated. Regarding flexibility and adaptability, the new generation of distributed services has to be adaptive. To achieve this goal, algorithmic solutions have to benefit from the recent advances in software engineering (componentware approach) and a provable methodology to specify, design and prove the distributed algorithms is needed.

4.2. Telecommunication Applications

The telecommunication domain is currently very interested in peer-to-peer computing. Nowadays, people are not just satisfied with the ability that they can hear a person from another side of the earth. "Instead, the demands of clearer voice in real-time are increasing globally. Just like the TV network, there are already cables in place, and it's not very likely for companies to change all the cables. Many of them turn to use the Internet, more specifically P2P networks. For instance, Skype, one of the most widely used Internet phone applications is using P2P (peer-to-peer) technology" [excerpt from *Wikipedia*]. By relying on a P2P paradigm, the telecommunication industry is enlarging its panel of innovating applications ranging from video on demand to massively-shared and user-generated unbounded digital universe. A prerequisite for these applications to meet quality of service requirements of their users is the effective and honest participation of these very same users. In absence of any large centralized enforcement institution in charge of controlling users behavior, the only viable alternative for encouraging trustworthy behavior is to rely on informal social mechanisms collecting, and aggregating information about user behaviors, a.k.a., reputation mechanisms.

5. Software

5.1. Prometeus

Participants: Michel Hurfin, Izabela Moise, Jean-Pierre Le Narzul.

The PROMETEUS project, part of the Inria Gforge, is a software environment for reliable programming developed by the Adept team. The basic elements of PROMETEUS are Eva, a component-based framework and Adam, a set of group communication services.

EVA is an implementation of a component model that aims at supporting the development of distributed abstractions and high-level communication protocols. EVA implements a publish/subscribe communication environment to structure components composing high level protocols. In the EVA model, protocols are regarded as a number of cooperating components that communicate via an event channel. Communication is achieved via the production of events (output data) by supplier components, and the consumption of these events (input data) by consumer components. A supplier component uses the service of an event channel to route the events it produces to any consumer component that has registered with the event channel it is interested in consuming that particular type of event. The event channel decouples suppliers from consumers yielding an interesting flexibility. Synchronous interactions between components is also supported in EVA. Special attention has been devoted to optimize the implementation. For example, potential sources of overheads (in the management or transmission of events) have been limited or eliminated in the design and implementation of EVA.

ADAM is a library of agreement components, based on the component model implemented by Eva. The central element of the ADAM library is GAC (Generic Agreement Component). It implements a generic and adaptive fault-tolerant consensus algorithm that can be customized to cope with the characteristics of the environment. Moreover, thanks to a set of versatile methods, its behavior can be tuned to fit the exact needs of a specific agreement problem. A range of fundamental ADAM components are implemented as specializations of this GAC component. The ADAM library currently includes the most important components for reliable distributed programming (Group Membership, Atomic Broadcast). Based on their (local and inconsistent) observations of the system, all members are obliged to continuously update and share a unique view of the system. This common perception of the state of the group has to be consistent with i) the decided sequence of view changes that have to be installed (membership service), ii) the decided sequence of messages that have to be consumed (total order broadcast) , and iii) the decided interleaving of the view change notifications with the flow of ordered messages (view synchrony property).

6. New Results

6.1. Evaluation of Consensus Protocols

Participants: Michel Hurfin, Izabela Moise, Jean-Pierre Le Narzul.

To be able to coordinate efficiently the activities of replicas, a significant body of work on replication techniques, group communication services and agreement problems has been done. The Consensus service has been recognized as a fundamental building block for fault-tolerant distributed systems. Many different protocols to implement such a service have been proposed, however, little effort has been placed in evaluating their performance. In [21], we present a protocol designed to solve several consecutive consensus instances in an asynchronous distributed system prone to crash failures and message omissions. The protocol follows the Paxos approach and integrates two different optimizations to reduce the latency of learning a decision value. During an execution of Paxos [32], [30], an external proposer is linked to an external learner by a communication path of length six in the worst case (external proposer ⟵ leader ⟶ acceptors ⟶ leader ⟶ acceptors ⟶ leader ⟶ external learner). A first well-known strategy allows to reduce the latency to four (external pro- poser ⟶ leader ⟶ acceptors ⟶ leader ⟶ external learner) when the elected leader remains stable (i.e. during long lasting failure-free synchronous periods). This optimization consists in removing of a phase called the Prepare phase. A second strategy, presented by Lamport in a protocol called Fast Paxos (with a blank) [31], tries to take advantage from a low throughput of the flow of initial values provided by the proposers. It aims at reducing the number of communication steps to three (external proposer ⟶ acceptors ⟶ leader ⟶ external learner), in favorable circumstances. If all proposers provide the same initial value, a gain can be obtained during a new consensus instance by anticipating some part of the computation. Rather than

being idle, a leader prepares the next consensus instance by sending a special value, called an Any value, to the acceptors. Once an acceptor receives it, it is allowed to adopt a value directly provided by a proposer. As such an initial value does not pass in transit through the leader, the decision latency is reduced. As the values adopted by different acceptors during an attempt are not necessarily equal, larger quorums have to be used. Moreover, when different values are proposed simultaneously, this second optimization (denoted O2) can be counterproductive as it may require the execution of a time consuming recovery procedure. In the proposed solution [21], the decision to use the second optimization is adopted by the leader during the computation. More precisely, as the use of the second optimization is risky, we check at runtime if the context seems to be favorable or not: the current leader chooses to use the second optimization only if it has been waiting for an initial value for some time. The proposed protocol is adaptive as it tries to obtain the best performance gain depending on the current context. Moreover, it guarantees the persistence of all decision values. Our experimentation results focus on the impact and the prediction of collisions.

6.2. Mobile Agent and Transactions

Participants: Michel Hurfin, Izabela Moise, Linda Zeghache.

Mobile devices are now equipped with multiple sensors and networking capabilities. They can gather information about their surrounding environment and interact both with nearby nodes, using a dynamic and self-configurable ad-hoc network, and with distant nodes via the Internet. While the concept of mobile agent is appropriate to explore the ad-hoc network and autonomously discover service providers, it is not suitable for the implementation of strong distributed synchronization mechanisms. Moreover, the termination of a task assigned to an agent may be compromised if the persistence of the agent itself is not ensured. In the case of a transactional mobile agent, we identify two services, *Availability of the Sources* and *Atomic Commit*, that can be supplied by more powerful entities located in a cloud. In [22], we propose a solution where these two services are provided in a reliable and homogeneous way. To guarantee reliability, the proposed solution relies on a single agreement protocol that orders continuously all the new actions whatever the related transaction and service. In [23] we show that these services can be integrated within the cloud.

6.3. Framework for Proving Self-Organization

Participant: Emmanuelle Anceaume.

Self-organization is an evolutionary process that appears in many disciplines. Physics, biology, chemistry, mathematics, economics, and more recently distributed systems (peer-to-peer systems, ad-hoc networks, sensors networks, cooperative robotics) just to cite a few of them, show many examples of self-organizing systems. In all these systems, self-organization is described as a process from which properties emerge at a global level of the system. These properties are solely due to local interactions among components of the system, that is with no explicit control from outside the system. Influence of the environment is present but not intrusive, in the sense that it does not disturb the internal organization process.

The main focus of our work presented in [27] is to propose a formal specification of the self-organization notion which, for the best of our knowledge, has never been formalized in the area of scalable and dynamic systems, in spite of an overwhelming use of the term. Specifically, we introduce the notion of *local self-organization*. Intuitively, a locally self-organizing system should reduce locally the entropy of the system. For example, a locally self-organized P2P system forces components to be adjacent to components that improve, or at least maintain, some property or evaluation criterion. The second contribution of the paper is the proposition of different classes of self-organization through safety and liveness properties that both capture information regarding the entropy of the system. Basically, we propose three classes of self-organization. The first one characterizes dynamic systems that converge toward sought global properties only during stable periods of time (these properties can be lost due to instability). The second one depicts dynamic systems capable of infinitely often increasing the convergence towards global properties (despite some form of instability). Finally, the last one describes dynamic systems capable of continuously increasing that convergence. We show that complex emergent properties can be described as a combination of local and independent properties.

6.4. Persistent Feedback

Participants: Emmanuelle Anceaume, Heverson Borba Ribeiro.

In our quest of building robust reputation systems, we are currently addressing the issue of feedback persistency. Reputation systems are clearly concerned with this issue in dynamic systems for the simple reason that nodes can freely join and leave the system at any time. This continuous churn must not prevent the reputation system to have continuous access to feedback the system has succeeded to gather, evaluate and store at multiple nodes in the system. Beyond churn, the unavoidable presence of malicious nodes may also be an obstacle to guarantee durable access to feedback. Reasons are numerous, and among them one can cite whitewashing, transaction repudiations, bad mouthing or even ballot stuffing. In [17], [18] we have proposed *DataCube*, a persistent storage functionality for implementing long-lived data objects despite massive and targeted attacks. The solution we have proposed to face adversarial environment was by designing a hybrid redundancy schema (a compound of light replication and rateless erasure coding) on top of a clustered DHT-based overlay. We have shown that this schema guarantees durable access and integrity of data despite adversarial attacks. We have evaluated the efficiency of DataCube, a P2P persistent data storage platform guaranteeing durable access and integrity of data despite collusion of malicious nodes. The evaluation we have performed has shown that parameters selection impacts codes performance. In particular we have shown the benefit of judiciously collecting check blocks. Regarding our evaluation of DataCube in a severe adversarial environment we have shown the benefit of hybrid replication over full replication in terms of data availability, storage overhead and bandwidth usage.

The coding schema initially used in DataCube was based on Online codes [17], [18]. This coding reveal to be efficient however we quickly realized that performance were considerably dependent on the tuning of the different parameters of the coding. This led us to compare the different rateless erasure coding that have been so far designed. For the best of our knowledge, no experimental study comparing the two classes of fountain codes had ever been performed. Two classes of rateless erasure codes exist. Representative of the first class is the LT coding proposed by Luby [33], and representatives of the second one are Online codes proposed by Maymounkov [34] and Raptor codes by Shokrollahi [35]. The latter class differs from the former one by the presence of a pre-coding phase. Thus the objective of [16] has been to provide such a comparison. Specifically, we have proposed to compare the experimental performance of both LT and Online codes. Note that Raptor codes [35] could have been analyzed as a representative of the second class of fountain codes; however because part of their coding process relies on LT codes, we have opted for Online codes. This evaluation sought not only to compare the performance of both codes in different adversarial environments, and in different application contexts (which is modeled through different size of data), but also to understand the impact of each coding parameter regarding the space and time complexity of the coding process. As such, this work should be considered, for the best of our knowledge, as the first comprehensive guideline that should help application designers to configure these codes with optimal parameters values. Experiments have confirmed that it is more efficient to collect random check blocks as theoretically predicted than favoring only small degree check blocks. Clearly, we expect that this study should allow a good insight into the properties of these codes. In particular we have confirmed the good behavior of both coders/decoders when the number of input blocks increases. This is of particular interest for multimedia and storage applications.

6.5. Induced Churn to Face Malicious Behaviors

Participants: Emmanuelle Anceaume, Sébastien Gambs, Romaric Ludinard.

Ensuring durable access to feedback is a first barrier against simple attacks as discussed above. However it is still possible for the adversary to to use several distinct identities so as to bias the reputation mechanism. Recall that trustworthiness of the reputation mechanism we are considering in our work is solely based on statistical measurements. Consequently, an attacker that would be powerful enough to create a significant number of distinct identities could violate such an assumption. Our contribution is centered around the study of robust mechanisms that can resist such attacks.

Toward this goal, we have first investigating the problem of uniform sampling in large scale open systems in presence of adversarial nodes. Uniform sampling ensures that any individual in a population has the same probability to be selected as sample. Uniform sampling finds its root in many problems such as data collection, dissemination, load balancing, and data-caching. Achieving uniform sampling in large scale open systems has been shown to be difficult. Reasons are two-fold. Firstly, the population which is typically very large (*e.g.*, thousand or millions of nodes) shows a very high churn (recent studies on the eDonkey file-sharing network have shown that in average $500,000$ nodes connect and disconnect per day). Secondly, openness makes unavoidable the presence of adversarial parties that control arbitrarily many nodes. These nodes strategize to isolate honest nodes within the system by violating randomness assumptions, these assumptions being are at the very foundation of these systems (*i.e.*, scalability of structured peer-to-peer systems rely on the assumption that nodes are uniformly distributed over the structured communication graph, while connectivity of unstructured peer-to-peer systems result from the assumption that nodes choose their neighbors arbitrarily). This is achieved by poisoning local views of honest nodes with malicious node identifiers. In unstructured graphs, a number of push operations logarithmic in the size of local views is sufficient to fully eclipse honest nodes from the local view of a node, while in structured graphs, a linear number of join operations is needed.

By relying on the topological properties of structured peer-to-peer systems, it has been shown that it is possible to guarantee that with high probability any node is equally likely to appear in the local view of each other honest node in a number of rounds polynomial in the size of the system. This is achieved by imposing nodes to frequently depart from their position and move to another random position in the system. Indeed, in a recent paper that will appear in 2011 [11], we have shown that an adversary can very quickly subvert overlays based on distributed hash tables by simply never triggering leave operations. We have also demonstrated that when all nodes (honest and malicious ones) are imposed on a limited lifetime, the system eventually reaches a stationary regime where the ratio of polluted clusters is bounded, independently from the initial amount of corruption in the system. This work has been conducted in collaboration with the INRIA project team Dionysos and SSIR team of Supelec.

In unstructured peer-to-peer systems, nodes cannot rely on the topological nature of structured graphs to detect undesirable behaviors. To circumvent this issue, Bortnikov *et al.* [29] rely on the properties of min-wise independent permutations, which are fed by the streams of gossiped node ids, to eventually converge towards uniform sampling on the node ids. However this random sample is definitive, in the sense that no other node ids received in the input stream can ever appear in the random sample, which makes their sampling uniform but not ergodic. Informally, this property guarantees that each received node id infinitely often has a non-null probability to locally appear as a sample [13]. This lack of adaptivity seems to be the only defense against adversarial behavior when considering bounded resources (memory and bandwidth). A preliminary step in determining conditions under which uniform and ergodic sampling is achievable in unstructured peer-to-peer systems potentially populated with a large proportion of Byzantine nodes has been presented in [13]. Briefly, it is shown that imposing strict restrictions on the number of messages gossiped by malicious nodes during a given period of time and providing each honest node with a very large memory (in the size of the system) are necessary and sufficient conditions to solve this problem.

In addition to consider malign attacks, we have investigated in which extend large scale systems, and in particular cluster-based peer-to-peer systems, are robust to churn, churn being classically defined as the rate of turnover of peers in the system. In [14], we have studied the ability of the system to continue to operate correctly despite high churn. This has been achieved by accurately predicting the minimal number of join and leave events that need to be globally triggered in the system to give rise to the first topological operation in the system. We have shown that $\Theta(N)$ join/leave events are required before any of these topological operations occur, where N is the number of peers currently in the system. This is actually an important result as it shows that due to the rare occurrence of those relatively costly topological operations, it is simple for the network to correctly update routing tables in due time and thus maintain the graph structure. This solves the Achilles heel of DHTs, i.e., the cost induced to maintain nodes routing tables consistency in presence of high churn. From a practical point of view this is interesting as it shows the appropriateness of these overlay networks as substrate for large scale applications demanding in terms of routing latency and topology stability such as multimedia streaming platforms, and persistent data storage.

6.6. Privacy-preserving Identification System

Participant: Sébastien Gambs.

Principles and Techniques used to preserve privacy are discussed in [24].

We aim at studying privacy-preserving identification systems. In a joint work with Yves Deswarte (LAAS) [12], we propose to replace the national identity card, currently used in many countries, by a personal device that allows its user to prove some binary statements about himself while minimizing personal information leakage. The privacy of the user is protected through the use of anonymous credentials which allows him to prove binary statements about himself to another entity without having to disclose his identity or any unnecessary information. The proposed scheme also prevents the possibility of tracing the user, even if he proves several times the same statement (unlinkability property). A tamper-proof smartcard is used to store the personal information of the user thus protecting his privacy and preventing the risks of forgery at the same time. The user identifies himself to the card via biometrics thus forbidding an unauthorized use in the situation where the card is stolen or lost. Two practical implementations of the privacy-preserving identity card are described and discussed.

6.7. Privacy in Social Networking Sites

Participant: Sébastien Gambs.

Social Networking Sites (SNS), such as Facebook and LinkedIn, have become the established place for keeping contact with old friends and meeting new acquaintances. As a result, a user leaves a big trail of personal information about him and his friends on the SNS, sometimes even without being aware of it. This information can lead to privacy drifts such as damaging his reputation and credibility, security risks (for instance identity theft) and profiling risks. In an ongoing collaboration [15] with Ai Thanh Ho and Esma Aïmeur (Université de Montréal), we first highlight some privacy issues raised by the growing development of SNS and identify clearly three privacy risks. While it may seem a priori that privacy and SNS are two antagonist concepts, we also identified some privacy criteria that SNS could fulfill in order to be more respectful of the privacy of their users. Finally, we introduce the concept of a Privacy-enhanced Social Networking Site (PSNS) and we describe Privacy Watch, our first implementation of a PSNS.

6.8. Geo-privacy

Participant: Sébastien Gambs.

A geolocalised system generally belongs to an individual and as such knowing its location reveals the location of its owner, which is a direct threat against his privacy. To protect the privacy of users, a sanitization process, which adds uncertainty to the data and removes some sensible information, can be performed but at the cost of a decrease of utility due to the quality degradation of the data. In a joint work with Marc-Olivier Killijian (LAAS) [19], we introduce GEPETO (for GEoPrivacy-Enhancing TOolkit), a flexible open source software which can be used to visualize, sanitize, perform inference attacks and measure the utility of a particular geolocalised dataset. The main objective of GEPETO is to enable a user to design, tune, experiment and evaluate various sanitization algorithms and inference attacks as well as visualizing the following results and evaluating the resulting trade-off between privacy and utility.

In [20], we report on our preliminary experiments with GEPETO for comparing different clustering algorithms and heuristics that can be used as inference attacks, and evaluate their efficiency for the identification of point of interests, as well as their resilience to sanitization mechanisms such as sampling and perturbation.

7. Contracts and Grants with Industry

7.1. Grants with Industry

The P2Pim@ge project (2007-2010) is supported by the Direction Générale des Entreprises. This project aims at studying, prototyping and testing legal advanced streaming technology on peer-to-peer systems. Different applications are addressed such as video on demand, immediate or differed download, access to scare content, etc. Partners of the project are Thomson R&D, Thomson Broadcast & Multimedia, Mitsubishi Electric ITE/TCL, Devoteam, France Telecom, ENST Bretagne, Marsouin, IRISA, IPdiva, and TMG.

In such large-scale dynamic systems, users may have a strategic behavior that is neither obedient nor malicious, but just rational. Tracking such behavior is complex since it requires taking into account a large set of features: large population, asymmetry of interest, collusion, "zero-cost identity", high turnover, and rationality. Techniques from the security domain (e.g. intrusion detection), and new fault tolerant distributed algorithms inspired from social theories will be investigated to deal with these undesirable behaviors.

8. Other Grants and Activities

8.1. International Initiatives

Participants: Sébastien Gambs, Emmanuelle Anceaume.

CANADA: Sébastien Gambs is co-supervising Ai Thanh Ho, a PhD student from the Université de Montréal with whom he has been actively collaborating for 2 years on the subject of privacy issues in social networking sites. The main supervisor of Ai Thanh Ho is Esma Aïmeur (full professor, Université de Montréal). In 2010, this cooperation has led to a joint publication [15].

BRAZIL: Francisco Brasileiro, Professor at the Federal University of Paraiba (Campina Grande) was involved with us in a four years Capes/Cofecube project (2005-2009). We still cooperate on the dependability evaluation of cluster-based systems [11].

8.2. Exterior research visitors

Participants: Michel Hurfin, Izabela Moise.

ALGERIA: Linda Zeghache, Phd student at USTBH-CEDRIC (université des Sciences et de la Technologie Houari Boumédiène, Algeria) visited us during one month in december 2010/january 2011. This cooperation has led to two joint publications [22], [23].

9. Dissemination

9.1. Animation of the scientific community

Emmanuelle Anceaume served as a program committee member of:

- Algotel 2010, 12èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications, Belle Dune, France, May 31- June 3, 2010
- HPCC 2010, the 12th edition of the highly successful International Conference on High Performance and Communications (HPCC), Melbourne, Australia, September 1-3, 2010.
- NSS2010, the 4rd International Conference on Network and System Security (NSS 2010) Melbourne, Australia, September 1-3, 2010.
- SMPE2010, the 4th International Symposium on Security and Multimodality in Pervasive Environment 2010 (SMPE 2010)Perth, Australia, April 20-23, 2010.
- SSDU2010, the 2010 International Symposium on Service, Security and its Data management technologies in Ubi-com (SSDU'2010, http://www.ftrai.org/ssdu2010), Xian, China, October 26-29, 2010.
- Trust-Com2010, the six IEEE/IFIP International Symposium on Trusted Computing and Communications (TrustCom-10, http://trust.csu.edu.cn/conference/trustcom2010/), to be held in Hong Kong during December 11-13, 2010.

Emmanuelle Anceaume acted as a reviewer for the ANR (programme blanc international), and for the ACM Transactions on Autonomous and Adaptive Systems journal (ACM TAAS) and for the Journal of Computers (JCP).

Sébastien Gambs was co-organizing:

- the 6th Workshop on Computer Privacy in Electronic Commerce (http://www.iro.umontreal.ca/~prive10/) with Esma Aïmeur and Gilles Brassard (Université de Montréal). The workshop was held in May 2010 in Montréal (one day) and had about 100 participants.
- the 1st workshop on Games, Logic and Security called GIPSy (http://www.irisa.fr/prive/Sophie.Pinchinat/GIPSy/gipsy10.html) along with Nathalie Bertrand (INRIA) and Sophie Pinchinat (Université de Rennes 1). The workshop was held in November 2010 in Rennes and had about 40 participants.

Sébastien Gambs served as a program committee member of:

- CMS'2010, the 11th Joint IFIP TC6 and TC11 Conference on Communications and Multimedia Security.

Sébastien Gambs is member of the editorial board of International Journal of Data Mining, Modelling and Management (http://www.inderscience.com/browse/index.php?journalID=342#board).

Sébastien Gambs was an external reviewer for the User Modeling and User-Adapted Interaction (UMUAI), the 36th International Conference on Very Large Data Bases (VLDB'2010) and the 40th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'2010).

Since January 2010, Sébastien Gambs has given 9 talks on subjects such as:

- "What do your queries, social network and movements reveal about yourself and can we protect them?", (TAMALE seminar, Ottawa University + LITQ's seminar, Université de Montréal + Annual security summit, Technicolor, Rennes + ASAP seminar, INRIA Rennes Bretagne Atlantique + Seminar, ENS Cachan Bretagne),
- "Private similarity computation in distributed systems: from cryptography to differential privacy", (Gossple meeting, INRIA Rennes Bretagne Atlantique),
- "Towards a privacy-preserving national identity card" (IRMAR cryptography seminar, Université de Rennes 1),
- "Retour sur le workshop PQSM'10", (LITQ's seminar, Université de Montréal),
- "Show me how you move and I will tell you who you are", (SPRINGL'10, San José).

Michel Hurfin served as a program committee member of:

- CARI'2010, the 10th African Conference on Research in Computer Science and Applied Mathematics, October 18-21, 2010, Yamoussoukro, Côte d'Ivoire. http://www.cari-info.org
- SRDS 2010, the 29th IEEE Int. Symposium on Reliable Distributed Systems, November 1-3, 2010, Delhi, Inde. http://www.scs.ryerson.ca/iwoungan/SRDS2010
- SSS 2010, the 12th Int. Symposium on Stabilization, Safety, and Security of Distributed Systems, September 20-22, 2010, New York, USA. http://www.cs.bgu.ac.il/~dolev/SSS10
- CloudCom 2010, the 2nd Int. Conf. on Cloud Computing Technology and science, December 1-3, 2010, Indianapolis, USA. http://2010.cloudcom.org
- IWSN 2011, the 2nd Int. workshop on Interconnections of Wireless Sensor Networks, June 27-29, 2011, Barcelona, Spain. http://iwsn2011.gforge.uni.lu/index.html

Michel Hurfin is is member of the editorial board of the Springer Journal of Internet Services and Applications. http://www.springer.com/computer/communications/journal/13174

Michel Hurfin has reporting on Mr Hassan Omar PhD work entitled "Privacy preserving reputation systems for decentralized environments". The thesis defense has been held at INSA Lyon on september 2010.

Michel Hurfin acted as a external reviewer for the ANR (programme blanc international), for the ANRT (bourse Cifre) for the Elsevier Journal "Parallel Computing", for EDCC 2010 (Eighth European Dependable Computing Conference) and for PODC 2010 (Twenty-Ninth Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing).

Michel Hurfin is Member of the COST-GTAI and acts as a reviewer for both ARCs and ADT submissions.

9.2. Teaching

Emmanuelle Anceaume participates in the Master research of the University of Rennes 1. She is responsible of the BIB module.

This academic year (2010-2011), Sébastien Gambs will be teaching two graduate courses (Master 2), one on Security and Authentication and the other on the Protection of Privacy, which is a new course he has created.

Michel Hurfin gave lectures on fault tolerance and distributed computing to students of two engineering schools: Telecom Bretagne (Rennes, 5 hours) and Supelec (Rennes, 8 hours).

Since November 2009, Sébastien Gambs is co-supervising Ai Thanh Ho, a PhD student from the Université de Montréal with whom he has been actively collaborating for 2 years on the subject of privacy issues in social networking sites. The main supervisor of Ai Thanh Ho is Esma Aïmeur (full professor, Université de Montréal).

Since October 2010, Sébastien Gambs is co-supervising Mohammad Nabil Al-Aggan, a PhD student, on the subject of enhancing the privacy aspect in gossip-based networks such as Gossple. The main supervisor is Anne-Marie Kermarrec (senior researcher, INRIA Rennes -Bretagne Atlantique). Previously, Mohammad did a research internship as a 2nd year Master student on a related topic under the supervision of Sébastien and Anne-Marie.

Since October 2010, Sébastien Gambs is co-supervising Miguel Nunez del Prado Cortez, a PhD student from LAAS-CNRS (Toulouse), on the subject of inference attacks on geolocated data. The main supervisor is Marc-Olivier Killijian (junior researcher, LAAS-CNRS). Previously, Miguel did a research internship as a 2nd year Master student on a related topic under the supervision of Sébastien and Marc-Olivier.

Emmanuelle Anceaume is supervising the PhD of Heverson Borba Ribeiro.

Michel Hurfin is co-supervising the PhD of Izabela Moise with Jean-Pierre Le Narzul, a junior researcher from Télécom Bretagne.

Starting from December 2010, Sébastien Gambs and Michel Hurfin are co-supervising Ahmed Gmati, a PhD student, on the subject of privacy-preserving data mining.

10. Bibliography

Major publications by the team in recent years

[1] E. ANCEAUME, A. DATTA, M. GRADINARIU POTOP-BUTUCARU, G. SIMON. *Publish/Subscribe Scheme for Mobile Networks*, in "Proc. of the 2nd ACM International Workshop on Principles of Mobile Computing (POMC)", Toulouse, France, October 2002, p. 74–81.

[2] E. ANCEAUME, M. HURFIN, P. RAIPIN PARVÉDY. *An Efficient Solution to the k-set Agreement Problem*, in "Proc. of the 4th European Dependable Computing Conference (EDCC)", Toulouse, France, LNCS 2485, Springer Verlag, October 2002, p. 62–78.

[3] E. ANCEAUME, C. DELPORTE-GALLET, H. FAUCONNIER, M. HURFIN, G. LE LANN. *Designing Modular Services in the Scattered Byzantine Failure Model*, in "Proc. of the 3rd International Symposium on Parallel and Distributed Computing (ISPDC)", Cork, Ireland, July 2004, p. 262–269.

[4] E. AÏMEUR, G. BRASSARD, S. GAMBS. *Quantum clustering algorithms*, in "Proc. of the 24th international conference on Machine learning (ICML)", Corvalis, Oregon, Z. GHAHRAMANI (editor), ACM International Conference Proceeding Series, ACM, June 2007, vol. 227, p. 1-8.

[5] F. BRASILEIRO, F. GREVE, M. HURFIN, J.-P. LE NARZUL, F. TRONEL. *Eva: an Event-Based Framework for Developing Specialised Communication Protocols*, in "Proc. of the 1st IEEE International Symposium on Network Computing and Applications (NCA)", Cambridge, MA, February 2002.

[6] G. BRASSARD, A. BROADBENT, J. FITZSIMONS, S. GAMBS, A. TAPP. *Anonymous Quantum Communication*, in "Proc. of the 13th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)", Kuching, Malaysia, LNCS 4833, Springer Verlag, December 2007, p. 460-473.

[7] S. GAMBS, B. KÉGL, E. AÏMEUR. *Privacy-preserving boosting*, in "Data Mining and Knowledge Discovery", 2007, vol. 14, n^{o} 1, p. 131-170.

[8] J.-M. HELARY, M. HURFIN, A. MOSTÉFAOUI, M. RAYNAL, F. TRONEL. *Computing Global Functions in Asynchronous Distributed Systems with Process Crashes*, in "Proc. of the 20th International Conference on Distributed Computing Systems (ICDCS)", April 2000, p. 584–591, Best paper award.

[9] M. HURFIN, A. MOSTÉFAOUI, M. RAYNAL. *A Versatile Family of Consensus Protocols Based on Chandra-Toueg's Unreliable Failure Detectors*, in "IEEE Transactions on Computers", April 2002, vol. 51, n^{o} 4, p. 395–408.

[10] Y. WANG, E. ANCEAUME, F. BRASILEIRO, F. GREVE, M. HURFIN. *Solving the Group Priority Inversion Problem in a Timed Asynchronous System*, in "IEEE Transactions on Computers. Special Issue on Asynchronous Real-Time Distributed Systems", August 2002, vol. 51, n^{o} 8, p. 900–915.

Publications of the year

Articles in International Peer-Reviewed Journal

[11] E. ANCEAUME, F. BRASILEIRO, R. LUDINARD, B. SERICOLA, F. TRONEL. *Dependability Evaluation of Cluster-based Distributed Systems*, in "International Journal of Foundations of Computer Science (IJFCS)", 2011, To appear).

[12] Y. DESWARTE, S. GAMBS. *A Proposal for a Privacy-preserving National Identity Card*, in "Transactions on Data Privacy", 2010, vol. 3, n^{o} 3, p. 253-276, http://hal.inria.fr/inria-00556830/en.

International Peer-Reviewed Conference/Proceedings

[13] E. ANCEAUME, Y. BUSNEL, S. GAMBS. *Uniform and Ergodic Sampling in Unstructured Peer-to-Peer Systems with Malicious Nodes*, in "Proceedings of the 14th International Conference On Principles Of Distributed Systems (OPODIS)", Tozeur, Tunisia, Lecture Notes in Computer Science, Springer, December 2010, vol. 6490, p. 64-78, http://hal.inria.fr/hal-00554219/en.

[14] E. ANCEAUME, R. LUDINARD, B. SERICOLA. *Analytic Study of the Impact of Churn in Cluster-Based Structured P2P Overlays*, in "Proceedings of the IEEE International Conference on Communications (ICC)", Cape Town, South Africa, May 2010, p. 302-307, http://hal.inria.fr/hal-00476330/en.

[15] E. AÏMEUR, S. GAMBS, A. HO. *Towards a Privacy-Enhanced Social Networking Site*, in "Proceedings of the 5th IEEE International Conference on Availability, Reliability and Security (ARES)", Krakow, Poland, February 2010, p. 172-179, http://hal.inria.fr/inria-00556838/en.

[16] H. BORBA RIBEIRO, E. ANCEAUME. *A Comparative Study of Rateless Codes for P2P Persistent Storage*, in "Proceedings of the 12th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)", New York, USA, Lecture Notes in Computer Science, Springer, September 2010, vol. 6366, p. 489-503, http://hal.inria.fr/hal-00554699/en.

[17] H. BORBA RIBEIRO, E. ANCEAUME. *DataCube: A P2P Persistent Storage Architecture Based on Hybrid Redundancy Schema*, in "Proceedings of the 18th IEEE Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP)", Pisa, Italy, February 2010, p. 302-306, Short paper, http://hal.inria.fr/hal-00476286/en.

[18] H. BORBA RIBEIRO, E. ANCEAUME. *Exploiting Rateless Coding in Structured Overlays to Achieve Data Persistence*, in "Proceedings of the 24th IEEE International Conference on Advanced Information Networking and Applications (AINA)", Perth, Australia, April 2010, p. 1165-1172, http://hal.inria.fr/hal-00554701/en.

[19] S. GAMBS, M.-O. KILLIJIAN, M. NUNEZ DEL PRADO CORTEZ. *GEPETO: A GEoPrivacy-Enhancing TOolkit*, in "Proceedings of the 24th IEEE International Conference on Advanced Information Networking and Applications Workshops (WAINA)", Perth, Australia, April 2010, p. 1071-1076, http://hal.inria.fr/inria-00556835/en.

[20] S. GAMBS, M.-O. KILLIJIAN, M. NUNEZ DEL PRADO CORTEZ. *Show me how you move and I will tell you who you are*, in "Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS (SPRINGL)", San Jose, USA, November 2010, p. 34-41, http://hal.inria.fr/inria-00556833/en.

[21] M. HURFIN, I. MOISE, JEAN-PIERRE. LE NARZUL. *An Adaptive Fast Paxos for Making Quick Everlasting Decisions*, in "Proceedings of the 25th IEEE International Conference on Advanced Information Networking and Applications (AINA)", Biopolis, Singapore, March 2011, To appear.

[22] I. MOISE, M. HURFIN, L. ZEGHACHE, N. BADACHE. *Remote Reliable Services to Support Transactional Mobile Agents*, in "Proceedings of the 9th IEEE International Symposium on Network Computing and Applications - Trustworthy Network Computing Workshop (NCA - TNC)", Cambridge, USA, July 2010, http://hal.inria.fr/inria-00554730/en.

[23] I. MOISE, M. HURFIN, L. ZEGHACHE, N. BADACHE. *Cloud-Based Support for Transactional Mobile Agents*, in "Proceedings of the 7th International Symposium on Frontiers of Information Systems and Network Applications (FINA) in conjunction with the 25th IEEE International Conference on Advanced Information Networking and Applications (AINA)", Biopolis, Singapore, March 2011, To appear.

Scientific Books (or Scientific Book chapters)

[24] Y. DESWARTE, S. GAMBS. *Chapitre VII - Protection de la vie privée: principes et technologies*, in "Les technologies de l'information au service des droits : opportunités, défis, limites", D. LE MÉTAYER (editor), Bruylant, 2010, vol. 32, p. 109-134, http://hal.inria.fr/inria-00556839/en.

Research Reports

[25] E. ANCEAUME, F. BRASILEIRO, R. LUDINARD, B. SERICOLA, F. TRONEL. *Dependability Evaluation of Cluster-based Systems*, IRISA, 2010, PI 1947, http://hal.inria.fr/inria-00463468/en.

[26] E. ANCEAUME, F. CASTELLA, R. LUDINARD, B. SERICOLA. *Markov Chains Competing for Transitions: Application to Large-Scale Distributed Systems*, INRIA, May 2010, n° 1953, http://hal.inria.fr/inria-00485667/en.

[27] E. ANCEAUME, X. DÉFAGO, M. GRADINARIU POTOP-BUTUCARU, M. ROY. *A framework for proving the self-organization of dynamic systems*, INRIA, 11 2010, http://hal.inria.fr/inria-00534372/en.

[28] E. ANCEAUME, B. SERICOLA, R. LUDINARD, F. TRONEL. *Performance Analysis of Large Scale Peer-to-Peer Overlays using Markov Chains*, IRISA, December 2010, n° PI-1963, http://hal.inria.fr/inria-00546039/en.

References in notes

[29] E. BORTNIKOV, M. GUREVICH, I. KEIDAR, G. KLIOT, A. SHRAER. *Brahms: Byzantine Resilient Random Membership Sampling*, in "Computer Networks", 2009, vol. 53, p. 2340-2359, A former version appeared in the 27th ACM Symposium on Principles of Distributed Computing (PODC), 2008..

[30] L. LAMPORT. *Paxos Made Simple*, in "ACM SIGACT News", December 2001, vol. 32, n° 4, p. 51–58.

[31] L. LAMPORT. *Fast Paxos*, in "Distributed Computing", 2006, vol. 19, n° 2, p. 79–103.

[32] L. LAMPORT. *The part-time parliament*, in "ACM Transaction on Computer Systems", May 1998, vol. 16, n° 2, p. 133–169.

[33] M. LUBY. *LT codes*, in "Proceedings of the IEEE International Symposium on Foundations of Computer Science (SFCS)", 2002.

[34] P. MAYMOUNKOV. *Online codes*, in "Research Report TR2002-833, New York University", 2002.

[35] A. SHOKROLLAHI. *Raptor codes*, in "IEEE/ACM Transactions on Networking", 2006, p. 2551–2567.