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

Research Scientist

Michel Hurfin [Team leader, Research Associate (CR) INRIA, HdR]
Emmanuelle Anceaume [Research Associate (CR) CNRS]

Faculty Member

Sébastien Gambs [Associate Professor (MdC) Université de Rennes 1, INRIA research chair in Security of Information Systems - since September 2009]
Frédéric Majorczyk [ATER Université de Rennes 1, until August 2009]

Technical Staff

Romaric Ludinard [Technical Staff, INRIA]

PhD Student

Izabela Moise [PhD Student, Université de Rennes 1]
Heverson Ribeiro [PhD Student, Université de Rennes 1]

Administrative Assistant

Elodie Besnard [Administrative Assistant, INRIA, until February 2009]
Elise Guilloux [Administrative Assistant, INRIA, from February to June 2009]
Lydie Mabil [Administrative Assistant, INRIA, since June 2009]

Other

Jean-Pierre Le Narzul [External Collaborator, Faculty Member, Institut Télécom / Telecom Bretagne]

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 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.

Our contributions focus on the three following themes:

- **Dependability and group communication.** 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 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. 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 user right. Yet, very few systems tackle the issue of guaranteeing its respect. We investigate the preservation of privacy in various contexts: privacy-preserving identification systems, data mining & privacy, and geo-privacy.

3. Scientific Foundations

3.1. Introduction

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 solved easily. 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.

3.2. Dependability and Group Communication

3.2.1. Agreement Problems and Group Communication Services

To cope with dependability and security issues (namely reliability, availability, integrity, confidentiality, and privacy) when both accidental and intentional faults may occur, we promote the use of group communication services. We target mainly small distributed systems that contain less than a few tens of nodes. Obviously, the size of the system and the evolution speed of the group composition are two main factors that inherently may increase the cost of the proposed solutions. Yet, despite these limitations, the group concept is a very attractive approach even in the case of large scale dynamic networks. Indeed, many groups 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 a system which is composed of numerous heterogeneous, transient and unfamiliar entities, the group concept is sometimes a palliative approach. It compensates for these negative factors by identifying long lasting sub-systems of small size that unify their transient members while forcing them to know each other and to synchronize their activities.

Providing group communication services within a system is essential. Thanks to these general purpose services [38], entities located at different nodes of a distributed network can remain tightly synchronized despite failures and the asynchrony of the underlying distributed system. The membership service tracks changes in the group composition that result from explicit join and leave operations, as well as implicit (and unpredictable) leave operations due to failures [46]. The membership service ensures that all the processes share a *consistent* view of the group composition and allows to synchronize the activities of the processes with regard to the successive evolutions of the group composition (view installation, view synchrony, ...). Fundamental communication abstractions, called broadcast, are also provided. When an entity (belonging to the group or not) broadcasts a message using a group reference, the message is forwarded to all the entities belonging to the current view. When using a reliable broadcast, the message is received by all the non faulty members of the group or by none of them. The broadcast of a message can satisfy various ordering constraints: FIFO order, Causal order or Total order [30]. In the case of a reliable total order broadcast (also called atomic broadcast), all the messages addressed to a group are delivered in the same order by the group members even if these messages have been received in a different order. All these services facilitate the task of an application designer since they guarantee strong properties regarding the delivery of the messages to the recipients and the order in which these messages are delivered. For instance, to increase the overall reliability of a system, both critical data and functionalities may be replicated on a group of nodes. Ensuring consistency within such a set of replicas becomes trivial if group communication services are available.

Many group communication services can be classified as agreement problems. In our work, we design and develop an homogeneous set of services that rely on a solution to the basic consensus problem [29], [26]. More precisely, we propose to build all the group communication services on top of a generic and adaptive solution to the consensus problem, that can be customized to cope with the characteristics of the environment, as well as the properties of the reliable distributed abstractions that have to be ensured (see the description of the *Prometeus* software). From an algorithmic point of view, several design choices (definition of tunable consensus protocol parameters, consensus algorithms with multiple round participations, continual execution

of consensus instances, use of clock synchronization algorithms to fix the round duration, ...) lead to obtain an original software whose performance differs from that obtained by other group communication projects (*Ensemble* - Cornell University and the Hebrew University, *Appia* - University of Lisbon [43], *Samoa* - EPFL [44], ...). As these services are used very often, efficiency is a key issue when designing solutions to such agreement problems. Our main goal is to have an even better understanding of these problems while considering various levels of adversity (various failure models but also various computational models ranging from the purely synchronous one to the purely asynchronous one). Agreement protocols usually used in group communication only take into account an active adversary which may trigger either crashes (benign faults) or arbitrary behaviors (malign faults). Passive adversary which just observe the protocol's behavior have also to be considered in order to protect the privacy of each group's member.

From a software engineering point of view, the use of a componentware approach helps to implement the group concept in a modular way. Code tangling is a major concern when designing group communication services. Even if a consensus-based solution allows for a clean separation between agreement related code and protocol specific code, many concerns that crosscuts the various protocol codes remain. Part of the difficulty lies in the identification of all the hidden and tangled synchronizations which exist between the numerous protocols. We aim at promoting the "separation of concerns" principle to overtake existing toolkits in terms of adaptability. Indeed current existing solutions are poorly flexible and possible tunings usually require a deep expertise. To reach this objective, we need to revisit several protocols and their interaction schemes. Conducting a performance evaluation of our proposal is also a part of our activities.

3.2.2. Group Communication Services to Secure a Web Access

To address the confidentiality, integrity, and availability of an information system, a security policy has to be defined and enforced. In the case of large public applications deployed on Internet, prevention techniques are necessary but not sufficient. At runtime, additional mechanisms have to be used to detect any violation of the above properties and to limit their consequences. Indeed, some weaknesses and vulnerabilities often remain in the executed applications. The fact that these faults have never been identified and corrected before is partly due to both the ever increasing complexity of the information systems and the ever decreasing time-to-market of the new applications and services. Until they are discovered and eliminated, these design faults may prevent an application from behaving according to its specification. Hence, whether it may be accidental or intentional, when a user activates such a bug, security rules can be violated. Intentional faults are produced by malicious attackers who try to take advantage of residual vulnerabilities of the information system. Assuming that an intrusion can succeed, we want to be able to detect it, to confine damage and to clean and recover corrupted entities from errors.

Research on Intrusion Detection Systems (IDS) has been carried out following two distinct approaches: *misuse detection* and *anomaly detection*. Misuse detection, also called pattern-based detection, consists in recognizing attacks using a signature database which contains descriptions of already known attacks. The control focuses on the contents of the incoming requests. Unfortunatelly unknown attacks or variants of known ones may succeed. Once a vulnerability is discovered, new security advisories are published. Until they are taken into account, an information system remains unsecured. Even if the database is continuously refreshed and updated (which is a tremendous task in itself), this countermeasure fails to effectively react against an attack that may spread over Internet in just a few minutes, as some recent worms did. Anomaly detection, also called profile-based detection, consists in analyzing deviations from an expected normal behavior. The control focuses on the computing activities induced by the incoming requests. The accuracy of the detection relies on the two following assumptions. First, any intrusion should have a noticeable and unexpected impact on the activity of the system. This is almost always the case as an attack implies an abnormal use of the system. Second, a model that characterizes all the normal behavior patterns should be available. When both assumptions are satisfied, the occurrence of an intrusion tallies with the observation of a significant deviation from the expected behavior and vice versa. In this approach, new or unknown attacks can be detected. Of course, the definition of a model is far from being trivial. If the model is too general and permissive or too precise and restrictive, the IDS will probably make mistakes: it may for instance ignore real attacks (*false negative*) or raise an alert although the suspicious request is actually not an attack (*false positive*). The model representing normal behaviors is

usually explicitly defined. In that case, after a static construction, it can be dynamically improved using normal training data sets during a preliminary learning phase.

In our work, we consider a radically different approach called *implicit intrusion detection*. In the case of a Web server that delivers dynamic contents, we show that the use of diversified COTS (Components-Off-The-Self) servers allows to detect intrusions. To secure a web access to a set of data, we assume that data is replicated and accessible through different systems that may have residual vulnerabilities but hopefully not necessarily the same ones. Consequently, an attack can succeed on a particular copy but not on all the redundant servers. By checking the values returned by the different copies to the malicious attacker we can identify differences and detect anomalies. Of course, the difficult part is to provide replication and detection mechanisms that are safe and will not become an even more simple target for the attacker. Our aim is to study how group services can be used and adapted to achieve this objective. This approach can detect even previously unknown attacks. Similar studies (leading however to different algorithmic solutions) have been conducted by the LAAS (Delta-4 [37], [28] and DIT architectures [41]) and by the university of Texas at Austin [47].

3.3. Reputation in Dynamic Large Scale Systems

Digital reputation has recently emerged as a promising approach to cope with the specificities of large scale and dynamic systems. Briefly, reputation stimulates the development of relationship among trustworthy entities, while discouraging them in presence of untrustworthy entities. EBay, Amazon, Slashdot, ePinions, Yahoo! auctions, just to cite a few of them rely on a reputation mechanism to foster a trust relationship among entities that do not know each other *a priori*, and may possibly interact only once. Specifically, similarly to real world, a reputation mechanism expresses a collective opinion about some target entity by gathering and aggregating feedbacks about the past behavior of that target entity. The derived reputation score is used to help entities to decide whether an interaction with that entity is conceivable or not. By encouraging trust or distrust, reputation helps in finding new resources by using trusted entities as sources of knowledge. It is also a powerful tool to incite entities to behave correctly. Indeed, a well behaving entity maintains a good reputation score so that entities are interested in interacting with it. On the other hand, reputation can be also used as a punishment mechanism. By lowering reputation score of misbehaving entities, establishment of relationships with other entities is made harder.

We argue that reputation is a clear added-value to tackle some security issues, and envision to use it as a building block for the deployment of security policies. Those policies will be dynamically set according to the level of hostility perceived by each machine. However, to be considered as a valuable tool for trust assessment, a reputation mechanism has to be itself robust against adversity. In other words, reputation must have the ability to self-heal or at least to self-protect against undesirable behavior that may jeopardize users security. Moreover, attacks in open systems are numerous and can be magnified through collusion. Just to name a few, reputation mechanism must be able to face:

- whitewashing (badly scored entities leave and rejoin the system to renew their reputation score);
- masquerading (badly scored entities pretend to be another entity to acquire its good reputation score);
- bad mouthing (collusion to discredit the reputation of a service provider to lately benefit from it);
- ballot stuffing (collusion to advertise the quality of service of a service provider more than its real value to increase its reputation to push users to be involved in fraudulent transactions);
- sybil attack (generation of numerous fake entities to manipulate the reputation score);
- transaction repudiation (an entity can deny the existence of a transaction).

Increasing the robustness of reputation mechanisms encompasses robustness both at the reputation mechanism itself as previously described, but also at the underlying network level. Specifically appropriate mechanisms should prevent message corruption, rerouting, and denial of service during the feedback collect phase.

In this context, we envision to contribute at the different phases of the reputation mechanism construction. Regarding feedback aggregation, we propose to extend existing works (e.g., [22], [27]) by enlarging the behavioral assumptions of interacting entities (e.g., variation of the effort exerted by providing entities according to the entities with which they interact, according to their welfare, or their level of hostility), by minimizing the number of relevant feedback needed to build a fair enough score estimation so that reputation could quickly react to highly dynamic environments. An interesting approach would be to combine credibility-based reputation function with endogenous techniques, well adapted for massive churn [24]. Regarding feedback availability, a classical solution amounts in replicating feedback at different entities hence guaranting that despite disconnections and malicious behavior, feedback information remain available within the system. However this type of solution relies on entities propensity to fully and honestly cooperate. Such assumptions are ideal ones, and cannot be enforced without relying on incentive mechanisms. Finally, it has been shown that peer-to-peer overlay networks can only survive severe (Byzantine) attacks if malicious peers are not able to predict what is going to be the topology of the network for a given sequence of join and leave operations. Induced churn, by which peers are required to rejoin (leave and, immediately after, join again) the system seems to be an appealing solution for the construction of Byzantine-resilient overlays.

3.4. Preservation of Privacy

In forthcoming years, the *protection of privacy* is one of the greatest challenge that lies ahead and also an important condition for the development of the "Society of Information". In the ubiquitous world where we live, each individual constantly leaves "numerical traces" related to his activities and interests which can be linked to his identity. For example, when a person surfs the Internet and accesses a website, his IP address can be linked at the same time to his localization, his centers of interests and his identity. Sometimes it may even happen that someone disseminates information without being aware of it, such as when a RFID chip is inserted in its pull-over and diffuses information in a passive manner without even its holder being aware of it. If all these numerical traces are collected by an unauthorized entity, this can lead to a privacy breach and may be used against the individual itself. A company might for instance use this information to send targeted spam or a malicious person could perpetrate an identity theft for fraudulent purposes. Moreover, due to legality and confidentiality issues, problematics linked to privacy emerge naturally for applications working on sensitive data, such as medical records of patients or proprietary datasets of enterprises.

Privacy Enhancing Technologies [11] (PETs) are generally designed to respect both the principles of *data minimization* and *data sovereignty*. The *data minimization principle* states that only the information necessary to complete a particular application should be disclosed (and no more). This principle is a direct application of the legitimacy criteria defined by the European data protection directive (Article 7, [45]). The *data sovereignty principle* states that the data related to an individual belongs to him and that he should stay in control of how this data is used and for which purpose. This principle can be seen as an extension of many national legislations on medical data that consider that a patient record belongs to the patient, and not to the doctors that create or update it, nor to the hospital that stores it.

In our works, we investigate PETs that are generally based on a mix of different foundations such as cryptographic techniques, security policies and access control mechanisms just to name a few. Examples of domains that we are investigating and where privacy and utility aspects collide include: identity and privacy, geo-privacy, distributed computing and privacy, privacy-preserving data mining and privacy issues in social networks.

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: a Group Communication Service

Participants: Michel Hurfin, 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. Dependability and Group Communication

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

6.1.1. Agreement Problems and Group Communication Services

We consider an asynchronous distributed system which is prone to message losses and crash failures. Within a group, several important services, such as total order broadcast or group membership, can be solved by relying on repeated calls to a Consensus service. The classical specification of the Consensus problem requires that each participant proposes an initial value during an invocation of the *Propose* primitive and, despite failures, all the correct processes have to decide on a single value selected out of these proposed values. In a pure asynchronous system, this problem is impossible to solve [29]. Yet under some well-identified additional synchrony properties which can be indirectly exploited by a failure detector or a leader election service, several consensus protocols have been proposed. Among them, the Paxos Protocol presented by Lamport [34], [32] is probably the most famous one. Lamport has identified four basic roles: proposer, learner, coordinator, and acceptor. Each participant may take on multiple roles or just a single one. Proposers are entities that may provide initial values. The learners are in charge of detecting that the protocol has successfully converged toward a decision value. Proposers and learners are not involved in the convergence procedure which is only driven by the interactions between coordinators and acceptors. Coordinators and acceptors play a central role in ensuring that eventually a single value is selected to become the decision value. A leader election service is used to grant eventually a privilege to a single coordinator. If a correct coordinator becomes the unique leader forever (or at least, till the current consensus instance ends), it is able to impose a selected value to a majority of acceptors and to detect the successful termination of its attempt. Acceptors are used to implement quorums as majority sets. Therefore, by assumption, a majority of acceptors should never crash during the computation.

In [20], we revisit the interaction scheme between proposers, learners, coordinators, and acceptors. We formally define the Multiple Integrated Consensus problem and consider a protocol in charge of the whole sequence of consensus instances. Consensus instances are still executed sequentially but not in a complete isolation from each other. We extend the remit of the sub-group of coordinators and acceptors so that they also have to ensure the availability of the past decisions and they have to control when a new consensus instance can start. In the context of a long lasting computation performed by a (potentially large) collection of (possibly ephemeral) processes, the core of dedicated processes formed by the coordinators and acceptors is able, on one hand, to provide all the decision values already computed (or only the most recent ones) to any current member of the collection and, on the other hand, to ensure the progress of the successive consensus instances while regulating the activity of the proposers that may dynamically join and leave the collection. By definition, the k^{th} decision value corresponds to the outcome of the k^{th} consensus instance which selects an initial value v, proposed by at least one member and generates a decision $< v, k >$. A member of the collection may ignore this outcome but, instead of this decision, it cannot consider another couple $< v', k >$ with $v' \neq v$. To regulate the rate of consensus instances, a classical constraint is used: a participant is not allowed to act as a proposer during consensus instance k if it is not able to access the $k-1$ previous decisions.

The repeated and intensive use of a consensus building block militates in favor of an optimization of the performance of this basic agreement protocol. In a recent past, two different protocols, namely FastPaxos (without space) by Boichat *et al.* [25] and Fast Paxos (with a space) by Lamport [33], have been designed to reduce the latency of learning a decision value to respectively, three and two communication steps, in favorable circumstances. The first strategy which is also adopted in [35] (where the notion of view is proposed) and in [36] (where the concept of regency is introduced) tries to benefit from the stability of an elected leader during long lasting failure-free synchronous periods. The second strategy tries to take advantage from a low throughput of the flow of initial values provided by the proposers. To solve efficiently the Multiple Integrated Consensus problem, we present in [20] a protocol called Paxos-MIC that integrates, for the first time to our knowledge, within a single simple framework the two best known methods for reducing decision latency in Paxos-like protocols. Our protocol unifies these two different strategies, in order to obtain the best performance gain.

6.1.2. Group Communication Services to Secure a Web Access

In addition to the classical prevention security tools, Intrusion Detection Systems (IDS) are nowadays widely used by security administrators to detect attack occurrences against their systems. Anomaly detection is often viewed as the only approach to detect new forms of attack. The main principle of this approach consists in building a reference model of the behavior for a given entity (user, machine, service, or application) in order to compare it with the current observed behavior. If the observed behavior diverges from the model, an alert is raised to report the anomaly.

Intrusion detection is traditionally based on the definition of an explicit reference model. In the context of a joint work with Supelec, we consider an implicit model. We propose a solution to protect Web applications which is based on the concepts of diversity and redundancy. A set of COTS (Components-Off-The-Self) servers executed on different nodes and different operating systems constitutes the core of the generic architecture: they provide simultaneously the service to the client. Design diversity is used to build at runtime the reference model. As an attack takes advantage of a vulnerability which is specific to either an operating system or a running software (*i.e.*, a web server), an attack will succeed on at most one node. If at least three nodes are used, the normal behavior is the one adopted by a majority of servers. To ensure integrity and confidentiality, any request is forwarded to the different servers which implement the same functionality but through diverse designs. Any difference between results that are returned can be interpreted as a possible attack and a possible corruption of one node. This approach can detect even previously unknown attacks.

Furthermore to ensure also availability, replication techniques implemented on top of agreement services are used to avoid any single point of failure. Secured and robust group communication mechanisms (see the PROMETEUS software description) are used to maintain consistency at various stages of the architecture. Our system has been deployed on a intrusion detection platform that is based on a set of diversified Web servers running on top of three different operating systems (Windows, Linux, Mac OS X). Performance evaluations are currently conducted. We aim at evaluating the relevance of our solution along two axes. On one hand, we have to show that diversification of COTS servers can improve the detection of attacks with respect to false positives. On the other hand, we have to show that the cost of the atomic broadcast service is reasonable enough to be used in real applications where dependability is a key requirement.

In addition to this main activity, we have proposed a solution to protect web applications running on top of a diversified architecture against code injection. Our solution consists in creating diversity in the web applications scripts by randomizing the language understood by all the redundant servers. The automatization of this process called *Instruction-Set Randomization* is presented in [17].

6.2. Reputation in Dynamic Large Scale Systems

Participants: Emmanuelle Anceaume, Romaric Ludinard, Heverson Ribeiro.

6.2.1. Persistent Feedbacks

We aim at providing mechanisms that will guarantee the persistence of the feedback about entities within a structured peer to peer overlay [18], [19]. Persistent feedbacks clearly leverage the level of difficulty for an attacker to mount withwashing attacks, or transaction repudiations. The use of erasing and rateless codes [40] is a promising way to provide strong persistence mechanims [31], [39] at a reasonnable cost. We have proposed solutions that partially fulfill these requirements in the context of a structured peer to peer overlay [21]. As currently designed, these mechanisms are so powerful that it can be almost impossible to erase wrongly attributed feedbacks. Because we are also interested with privacy issue, we plan to extend these solutions so as to offer the possibility for a right to oblivion (*i.e.*, a right to erasure of data).

6.2.2. Induced Churn to Face Malicious Users

Persistent feedbacks are a first barrier against the simpler attacks. However, it is still quite easy for a malicious user to use several distinct identities so as to bias the reputation mechanism. Recall that the trustworthiness of the reputation mechanism we are considering here, is solely based on statistical measurements. Consequently, an attacker that could create a statistically significant number of different identities could make collapsing this hypothesis. Our contribution is centered around the study of robust mechanisms that can resist such attacks. It has been shown [23] that peer-to-peer overlay networks can only survive severe (Byzantine) attacks if malicious peers are not able to predict what is going to be the topology of the network for a given sequence of join and leave operations. Designing a P2P overlay that makes such previsions extremely difficult is a key feature to obtain a robust routing [42] mechanism within the overlay. In turn, a robust routing mechanism will guarantee that malicious users will not be able to corrupt information transmission and by the way persistence of collected and aggregated feedbacks.

A promising way to reach this goal is to have limited lifetime identities within the system. This limitation forces malicious and honest users to regularly leave and reenter the system. This helps spreading uniformly malicious users within the overlay making the task of mounting an attack extremely difficult. However, this solution induces a permanent churn for the overlay. As a consequence, it is of the uttermost importance to evaluate the cost of such solutions in terms of extra communication messages. In collaboration with the Inria project team Dionysos and Supelec, we have started investigating this approach with interesting results. We consider adversarial strategies by following specific games. Our analysis demonstrates first that an adversary can very quickly subvert DHT-based overlays by simply never triggering leave operations. We then show 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 results, obtained by using Markov models, are shown in [12] and [13].

6.3. Preservation of Privacy

Participant: Sébastien Gambs.

We describe herein only the research activities corresponding to papers co-authored by Sébastien Gambs and published since september 2009 (date of his arrival in the project team).

6.3.1. Privacy-preserving Identification System

We aim at studying privacy-preserving identification systems. In a joint work with Yves Deswarte (LAAS) [15], 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. This research was mainly conducted when Sébastien Gambs was a CNRS postdoctoral researcher at LAAS-CNRS, Toulouse (from October 2008 to August 2009).

6.3.2. Privacy in Social Networking Sites

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 [14] 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.

6.3.3. Geo-privacy

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) [16], 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.

7. Other Grants and Activities

7.1. National Project

7.1.1. DGE Project: P2Pim@ges (2007-2009)

Participants: Emmanuelle Anceaume, Jean-Pierre Le Narzul, Romaric Ludinard.

The P2Pim@ge project 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.

7.2. International Cooperations

7.2.1. Brazil (Federal University of Bahia and Federal University of Campina Grande)

Participants: Emmanuelle Anceaume, Michel Hurfin, Jean-Pierre Le Narzul.

A cooperation project with the Federal University of Bahia, the Federal University of Paraiba, and several French laboratories (EPI ADEPT, EPI GRAND LARGE, EPI REGAL and ENST Bretagne) was supported by Capes/Cofecub (projet 497/05) during a period of four years that ends at the beginning of 2009. Michel Hurfin is the French coordinator of this project which focuses on distributed computing and Grid computing. In 2009, this cooperation has led to two joint publications presented at SSS 2009 [12] and WRAS 2009 [13].

7.3. Visits of at least one month

Linda Zeghache, Phd student at USTBH-CEDRIC (université des Sciences et de la Technologie Houari Boumédiène, Algeria) visited us during one month in october/november 2009.

8. Dissemination

8.1. Teaching Activities

- Emmanuelle Anceaume participates in the Master research (modules BIB and META).
- Emmanuelle Anceaume is supervising the PhD of Heverson Ribeiro.
- Emmanuelle Anceaume was co-supervising with members of Supelec the research internship of a 2nd year Master student, Diego Casado Mansilla, on the subject of data availability and persistence in virtual worlds over P2P systems (From february to june 2009).
- Sébastien Gambs teaches a graduate course (Master 2) on Security and Authentification and gives practical classes in an undergraduate course on Algorithmics.
- 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).
- Sébastien Gambs is currently co-supervising with Anne-Marie Kermarrec (Inria project team Asap) the research internship of a 2nd year Master student, Mohammad Nabil Al-Aggan, on the subject of enhancing the privacy aspect in gossip-based networks such as Gossple.
- 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).
- Michel Hurfin and Jean-Pierre Le Narzul are supervising the PhD of Izabela Moise.
- Jean-Pierre Le Narzul has the responsibility for organizing several teaching units at Telecom Bretagne (RSM Department). He gives lectures on both distributed computing and object-oriented language. He is also involved in the setting of programs for continuous training.

8.2. Presentations of Research Works

- Since January 2000, Emmanuelle Anceaume co-organizes with Bruno Tuffin the seminars entitled "Networks and Systems" that are held in our institute.
- In 2009, Sébastien Gambs has given 9 talks on subjects such as:
  - "Towards a privacy-preserving national identity card",
    (LITQ's seminar, Université de Montréal)
  - "Privacy-preserving data mining",
    (DIRO's colloquium, Université de Montréal)
  - "Cryptography in a quantum world",
    (journée pôle SINC du LAAS, Toulouse)

- "UPP : User Privacy Policy for social networking sites",
  (ICIW'09, Venice)
- "Extensions to anonymous quantum communication",
  (LITQ's seminar, Université de Montréal)
- "Geolocalisation and privacy",
  (DIWALL's seminar, Brittany)
  (Cryptis' seminar, Université de Limoges)
  (LITQ's seminar, Université de Montréal)
  (C&ESAR 2009, Rennes)

8.3. Integration within the Scientific Community

- Emmanuelle Anceaume :
  - She was the co-chair of the 1st Workshop on Dependability and Security in P2P (WeeP 2009) organized in conjunction with LADC 2009, September 1-4, 2009, João Pessoa, Brazil. http://www.sbc.org.br/ladc
  - She served as a program committee member of :
    * The 3rd International Symposium on Service, Security and its Data management technologies in Ubi-comp (SSDU-09), May 4-8, 2009, Geneva, Switzerland. http://www.sersc.org/SSDU2009
    * The 1st International Workshop on Service-Oriented P2P Networks (ServP2P 09) organized in conjunction with CCGRID 09, May 18-21, 2009, Shangai, China. http://www.computing.surrey.ac.uk/personal/pg/S.Stafrace/ServP2P09
    * The 3rd International Symposium on Security and Multimodality in Pervasive Environments (SMPE-09) organized in conjunction with AINA-09, May 26-29, 2009, Bradford, UK. http://www.sersc.org/SMPE2009
    * The 11th IEEE International Conference on High Performance Computing and Communications (HPCC-09), June 25-27, 2009, Seoul, Korea. http://www.sersc.org/HPCC2009
    * The International Workshop on Ubiquitous Computing Security (UC-Sec 09), July 13-16, 2009, Las Vegas, USA. http://www.sersc.org/UCSEC2009
    * The 3rd International Conference on Network and System Security (NSS 09), October 19-21, 2009, Gold Coast, Australia. http://nss2007.cqu.edu.au
    * The 2009 IEEE/IFIP International Symposium on Trusted Computing and Communications (TrustCom 2009), August 29-31, 2009, Vancouver, Canada. http://trust.csu.edu.cn/conference/trustcom2009
    * La $4^{\text{ième}}$ Conférence sur la Sécurité des Architectures Réseaux et des Systèmes d'Information (SARSSI 2009), June 22-26, 2009, Luchon, France. http://sarssi.enseeiht.fr/FR/home.php
  - She was member of the PhD jury of Francois Lesueur (november 27th). His thesis is entitled "Autorité de certification distribuée pour des réseaux pair-à-pair structurés: mise en oeuvre et exemples d'applications".
  - She acted as a reviewer for the Swiss National Science Foundation.
- Sébastien Gambs :
  - He is member of the editorial board of the International Journal of Data Mining, Modelling and Management (Inderscience Publishers). http://www.inderscience.com/browse/index.php?journalCODE=ijdmmm

  - He was member of the organizing committee of the 5th workshop on Computer Privacy in Electronic Commerce, May 2nd, 2009, Université de Montréal, Montréal, Canada. He is member of this organizing committee for the next edition of the workshop in May 2010. http://www.iro.umontreal.ca/~prive09
  - He was in the local organization committee and the treasurer of the 4th International Conference on Risks and Security of Internet and Systems (CRiSIS 2009), October 19-22, 2009, Toulouse, France. http://www.crisis2009.org/
  - He serves as a member of the program committee of the 11th Joint IFIP TC6 and TC11 Conference on Communications and Multimedia Security (CMS'2010), May 31st-June 2nd, 2010, Linz, Austria. http://www.cms2010.net/
  - In 2009, he acted as an external reviewer for UMUAI (User-Modeling and User Adapted Interaction - the Journal of Personalization Research), WiMob 2009 (5th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications), DSN-DCCS 2009 (39th Annual IEEE/IFIP International Conference on Dependable Systems and Networks - The Dependable Computing and Communication Symposium), and SEC 2009 (24th IFIP International Information Security Conference).

- Michel Hurfin :
  - He is member of the editorial board of the Springer Journal of Internet Services and Applications. http://www.springer.com/computer/communications/journal/13174
  - He served as a program committee member of :
    * The 3rd International Workshop on Latin American Grid (LAGrid 2009), October 28-31, 2009, Sao Paulo, Brazil. http://lagrid09.lncc.br
    * The 1st International Conference on Cloud Computing (CloudCom 2009), December 1-4, 2009, Beijing, China. http://2009.cloudcom.org
    * The 2nd IEEE International Symposium on Ubisafe Computing (UbiSafe-09), December 12-14, 2009, Chengdu, China. http://cs.okstate.edu/ubisafe09/
    * The Colloquium of Computation: Brazil / INRIA, Cooperations, Advances and Challenges (Colibri) organized in conjunction with CSBC, July 22-23, 2009, Bento Goncalves, Brazil. http://gppd.inf.ufrgs.br/colibri/index.en.html
    * La $4^{\text{ième}}$ Conférence sur la Sécurité des Architectures Réseaux et des Systèmes d'Information (SARSSI 2009), June 22-26, 2009, Luchon, France. http://sarssi.enseeiht.fr/FR/home.php
    * The 10th African Conference on Research in Computer Science and Applied Mathematics (CARI'2010), October 18-21, 2010, Yamoussoukro, Côte d'Ivoire. http://www.cari-info.org
  - He acted as a reviewer for the ANR (programme blanc) and for Digiteo.
  - In 2009, he acted as an external reviewer for DISC 2009 (23rd International Symposium on Distributed Computing) , DSN-DCCS 2009 (39th Annual IEEE/IFIP International Conference on Dependable Systems and Networks - The Dependable Computing and Communication Symposium), OPODIS 2009 (13th International Conference On Principle Of DIstributed Systems) and EDCC 2010 (8th European Dependable Computing Conference).
  - Since September 2009, he is a member of the INRIA COST-GTAI (Comité d'Orientation Scientifique et Technologique - Groupe de Travail Actions Incitatives).

9. Bibliography

Major publications by the team in recent years

[1] E. ANCEAUME, A. DATTA, M. GRADINARIU, 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, vol. 227, ACM, June 2007, 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", vol. 14, n^{o} 1, 2007, 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", vol. 51, n^{o} 4, April 2002, 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", vol. 51, n^{o} 8, August 2002, p. 900–915.

Year Publications

Articles in National Peer-Reviewed Journal

[11] Y. DESWARTE, S. GAMBS. *Protection de la vie privée: principes et technologies*, in "Cahiers du CRID (Centre de Recherches Informatique et Droit)", vol. 32, January 2010, To appear.

International Peer-Reviewed Conference/Proceedings

[12] E. ANCEAUME, F. BRASILEIRO, R. LUDINARD, B. SERICOLA, F. TRONEL. *Brief Announcement: Induced Churn to Face Adversarial Behavior in Peer-to-Peer Systems*, in "Proc. of the 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2009), Lyon, France", LNCS 5873, Springer Verlag, November 2009, p. 773–774, http://hal.archives-ouvertes.fr/hal-00420559/en/BR.

[13] E. ANCEAUME, R. LUDINARD, B. SERICOLA, F. TRONEL, F. BRASILEIRO. *Analytical Study of Adversarial Strategies in Cluster-based Overlays*, in "Proc of the 2nd International Workshop on Reliability, Availability, and Security (WRAS 2009), Hiroshima, Japan", December 2009, 6 pages BR .

[14] E. AÏMEUR, S. GAMBS, A. HO. *Towards a Privacy-enhanced Social Networking Site*, in "Proc. of the 5th International Conference on Availability, Reliability and Security (ARES 2010), Krakow, Poland", February 2010.

[15] Y. DESWARTE, S. GAMBS. *Towards a Privacy-Preserving National Identity Card*, in "Proc. of the 4tn International Workshop on Data Privacy Management (DPM 2009) in conjunction with ESORICS'09, Saint Malo, France", LNCS 5939, Springer Verlag, September 2009, p. 30–43.

[16] S. GAMBS, M.-O. KILLIJIAN, M. NUNEZ DEL PRADO. *GEPETO: a GEoPrivacy-Enhancing TOolkit*, in "Proc. of the International Workshop on Advances in Mobile Computing and Applications: Security, Privacy and Trust, held in conjunction with International Conference on Advanced Information Networking and Applications (AINA 2010), Perth, Australia", 2010.

[17] F. MAJORCZYK, J.-C. DEMAY. *Automated Instruction-Set Randomization for Web Applications in Diversified Redundant Systems*, in "Proc. of the 4th IEEE International Conference on Availability, Reliability and Security, (ARES - WAIS 2009), Fukuoka, Japan", March 2009, p. 978-983.

[18] H. RIBEIRO, E. ANCEAUME. *DataCube: a P2P persistent Storage Architecture based on Hybrid Redundancy Schema*, in "Proc. of the 18th Euromicro International Conference on Parallel, Distributed and Network-Based Computing (PDP 2010), Pisa, Italy", 2010, 5 pages, Short paper.

[19] H. RIBEIRO, E. ANCEAUME. *Exploiting Rateless Coding in Structured Overlays to Achieve Data Persistence*, in "Proc. of the International Conference on Advanced Information Networking and Applications (AINA 2010), Perth, Australia", 2010.

Research Reports

[20] M. HURFIN, I. MOISE. *A Multiple Integrated Consensus Protocol based on Paxos, FastPaxos, and Fast Paxos*, n^{O} 1941, IRISA, 2009, 25 pages, Rapport de recherche.

References in notes

[21] E. ANCEAUME, F. BRASILEIRO, R. LUDINARD, A. RAVOAJA. *PeerCube: an Hypercube-based P2P Overlay Robust against Collusion and Churn*, in "Proceedings of the International Conference on Self- Autonomous and Self-Organizing Systems (SASO)", 2008.

[22] E. ANCEAUME, A. RAVOAJA. *Incentive-based Robust Reputation Mechanism for P2P Services*, in "Proc. of the 10th International Conference On Principles Of Distributed Systems (OPODIS 2006), Bordeaux, France", December 2006.

[23] B. AWERBUCH, C. SCHEIDELER. *Towards a Scalable and Robust DHT*, in "Proc. of the ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)", 2006.

[24] B. AWERBUCH, B. PATT-SHAMIR, D. PELEG, MARK R. TUTTLE. *Collaboration of untrusting peers with changing interests*, in "ACM Conference on Electronic Commerce", 2004.

[25] R. BOICHAT, P. DUTTA, S. FROLUND, R. GUERRAOUI. *Deconstructing Paxos*, in "ACM SIGACT News", vol. 34, n^{o} 1, March 2003, p. 47–67.

[26] TUSHAR DEEPAK. CHANDRA, S. TOUEG. *Unreliable failure detectors for reliable distributed systems*, in "Journal of the ACM", vol. 43, n^{o} 2, 1996, p. 225–267.

[27] Z. DESPOTOVIC, K. ABERER. *P2P reputation management: Probabilistic estimation vs social networks*, in "Computer Networks", vol. 50, n^{o} 4, 2006, p. 485–500.

[28] Y. DESWARTE, L. BLAIN, JEAN-CHARLES. FABRE. *Intrusion Tolerance in Distributed Computing Systems*, in "Proceedings of the IEEE Symposium on Research in Security and Privacy", May 1991, p. 110-122, ftp://ftp.laas.fr/pub/Publications/1990/90373.ps.

[29] M. J. FISCHER, N. A. LYNCH, M. S. PATERSON. *Impossibility of distributed consensus with one faulty process*, in "J. ACM", vol. 32, n^{o} 2, 1985, p. 374–382.

[30] V. HADZILACOS, S. TOUEG. *Fault-tolerant broadcasts and related problems*, in "Distributed systems (2nd Ed.)", 1993, p. 97–145.

[31] M. N. KROHN, M. J. FREEDMAN, D. MAZIERES. *On-the-fly verification of rateless erasure codes for efficient content distribution*, in "Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on", 2004.

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

[33] L. LAMPORT. *Fast Paxos*, in "Distributed Computing", vol. 19, n^{o} 2, 2006, p. 79–103.

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

[35] B. LAMPSON. *The ABCDs of Paxos*, in "Proc. of the 20th Annual ACM Symposium on Principles of Distributed Computing", 2001.

[36] JEAN-PHILIPPE. MARTIN, L. ALVISI. *Fast Byzantine consensus*, in "Proc. of the Int. Conference on Dependable Systems and Networks", June 2005, p. 402–411.

[37] D. POWELL, G. BONN, D. SEATON, P. VERISSIMO, F. WAESELYNCK. *The Delta-4 Approach to Dependability in Open Distributed Computing Systems*, in "Proceedings of Twenty-Fifth International Symposium on Fault-Tolerant Computing", IEEE, 27-30 june 1995, 56.

[38] D. POWELL. *Group Communication*, in "Communications of the ACM", vol. 39, n^{o} 4, 1996, p. 50–53.

[39] A. RAVOAJA, E. ANCEAUME. *Storm: A Secure Overlay for P2P Reputation Management*, in "Proceedings of the IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2007)", 2007.

[40] R. RODRIGUES, B. LISKOV. *High Availability in DHTs: Erasure Coding vs. Replication*, in "Proceedings of the International Workshop on Peer-to- Peer Systems (IPTPS)", 2005.

[41] A. SAIDANE, Y. DESWARTE, V. NICOMETTE. *An Intrusion Tolerant Architecture for Dynamic Content Internet Servers*, in "Proceedings of the 2003 ACM Workshop on Survivable and Self-Regenerative Systems (SSRS-03), Fairfax, VA", P. LIU, P. PAL (editors), ACM Press, October 2003, p. 110-114.

[42] E. SIT, R. MORRIS. *Security Considerations for Peer-to-Peer Distributed Hash Tables*, in "Proc. for the Int'l Workshop on Peer-to-Peer Systems (IPTPS)", 2002.

[43] THE APPIA PROJECT. *APPIA Communication Framework*, http://appia.di.fc.ul.pt.

[44] THE SAMOA PROJECT. *Samoa Project*, http://lsrwww.epfl.ch/page13488.html.

[45] E. UNION. *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*, in "Official Journal", vol. L 281, 11 1995, p. 0031-0050, http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML.

[46] R. VITENBERG, I. KEIDAR, GREGORY V. CHOCKLER, D. DOLEV. *Group Communication Specifications: A Comprehensive Study.*, in "ACM Computing Surveys", vol. 33, n^{o} 4, September 2001.

[47] J. YIN, JEAN-PHILIPPE. MARTIN, A. VENKATARAMANI, L. ALVISI, M. DAHLIN. *Separating Agreement from Execution for Byzantine Fault Tolerant Services*, in "Proceedings of the 19th ACM Symp. on Operating Systems Principles (SOSP-2003)", 2003.