RESEARCH CENTRE

Inria Centre at Université de Lorraine

IN PARTNERSHIP WITH:

Université de Lorraine, CNRS

2023 ACTIVITY REPORT

Project-Team COAST

Web Scale Trustworthy Collaborative Service Systems

IN COLLABORATION WITH: Laboratoire lorrain de recherche en informatique et ses applications (LORIA)

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Distributed Systems and middleware

Inria

Contents

Project-Team COAST

Creation of the Project-Team: 2015 July 01

Keywords

Computer sciences and digital sciences

A1.3. – Distributed Systems

A1.3.1. – Web

A1.3.3. – Blockchain

A1.3.4. – Peer to peer

A1.3.5. – Cloud

A1.3.6. – Fog, Edge

A2.5. – Software engineering

A2.6.2. – Middleware

A3.1.3. – Distributed data

A3.1.5. – Control access, privacy

A3.1.8. – Big data (production, storage, transfer)

A5.1.1. – Engineering of interactive systems

A5.1.2. – Evaluation of interactive systems

Other research topics and application domains

B6.1.1. – Software engineering

B6.3.1. – Web

B6.5. – Information systems

B8.4. – Security and personal assistance

B8.4.1. – Crisis management

B9.6.1. – Psychology

B9.8. – Reproducibility

B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientist

- Claudia-Lavinia Ignat [Team leader, INRIA, Researcher, HDR]

Faculty Members

- Khalid Benali [UL, Associate Professor, HDR]
- Gérôme Canals [UL, Associate Professor]
- François Charoy [UL, Professor, HDR]
- Thomas Lambert [UL, Associate Professor]
- Gérald Oster [UL, Associate Professor]
- Olivier Perrin [UL, Professor, HDR]

PhD Students

- Quentin Acher [INRIA, from Sep 2023]
- Clélie Amiot [INRIA]
- Alexandre Bourbeillon [INRIA, until Oct 2023]
- Lisa Formentini [INRIA, from Oct 2023]
- Linda Ouchaou [UL, ATER, until Aug 2023]
- Ludovic Paillat [hive, CIFRE, from Oct 2023]
- Pierre-Antoine Rault [INRIA]
- Mohammad Rizk [INRIA, from Nov 2023]

Technical Staff

- Victorien Elvinger [INRIA, Engineer, until Feb 2023]
- Baptiste Hubert [INRIA, Engineer]

Interns and Apprentices

- Quentin Acher [INRIA, Intern, from Feb 2023 until Jul 2023]
- Habibatou Ba [INRIA, Intern, from Apr 2023 until Aug 2023]
- Lyra Bastian [UL, Intern, from Mar 2023 until Sep 2023]
- Hua Junrui [hive, Intern, from Nov 2023]
- Ludovic Paillat [hive, Intern, from Mar 2023 until Aug 2023]

Administrative Assistants

- Sophie Drouot [INRIA]
- Nathalie Fritz [UL, until Aug 2023]
- Delphine Hubert [UL, from Sep 2023]

2 Overall objectives

The advent of the Cloud, smart mobile devices and service-based architecture has opened a field of possibilities as wide as the invention of the Web 30 years ago. Software companies now deliver applications and services using the Web as a platform. From text to video editing, from data analytics to process management, they distribute business applications to users within their web browser or on some mobile appliance [1]. These services are deployed on sophisticated infrastructures that can cope with very demanding loads. The Software as a Service approach (SaaS) highlights their cooperative nature, by enabling the storage of data in cloud infrastructures that can be easily shared among users.

Clients consume applications through service APIs (web services), available on delivery platforms, called stores or markets. This approach of software distribution outstrips the traditional software distribution channels, in scale and opportunity. Scale has different dimensions: the number of users (communities rather than groups), the size of data produced and managed (billions of documents), the number of services and of organizations (tens of thousands). Opportunity refers to the infinite number of combinations between these services and the many ways to consume and use them.

This fast-paced evolution challenges research because the creation of applications from the composition of services must incorporate new content and context based constraints. From a socio-technical perspective, the behaviour of users is evolving constantly as they get acculturated to new services and ways to cooperate. Mere enhancement of current existing solutions to cope with these challenges is insufficient.

We conduct a dedicated research effort to tackle the problems arising from the evolution of contemporary technologies and of those we can anticipate. We explore three directions: large scale collaborative data management, data centred service composition and above all, a foundation for the construction of trustworthy collaborative systems.

Large scale collaborative data management concerns mostly the problem of allowing people to collaborate on shared data, synchronously or not, on a central server or on a peer to peer network. This research has a long history referring back to a paper by Ellis [18]. User acculturation to online collaboration triggers new challenges. These refer to the number of participants in a collaboration (a crowd), to the number of different organizations and to the nature of the documents that are shared and produced. The problem is to design new algorithms and to evaluate them under different usage conditions and constraints and for different kinds of data.

Data centred service composition deals with the challenge of creating applications by composing services from different providers. Service composition has been studied for some time now but the technical evolution and the growing availability of public APIs require us to reconsider the problem [16]. Our goal here is, taking into account this evolution, like the advent of the Cloud, the availability at a large scale of public APIs based on the REST [2] architecture, to design models, methods and tools to help developers to compose these services in a safe and effective way.

Based on the work that we do on the two first topics, our main research direction aims at providing support to build **trustworthy collaborative applications.** We base it on the knowledge that we can gather from the underlying algorithms, the composition of services and the quality of services that we can deduce and monitor. The complexity of the context in which applications are executed does not allow us to provide proven guarantees. Our goal is to base our work on a contractual and monitored approach to provide users with confidence in the service they use. Surprisingly, people rely today on services with very little knowledge about the amount of confidence they can put in these services. They are based on composition of other unknown services. Thus, it becomes very difficult to understand the consequences of the failure of a component of the composition. We follow a path that portrays a ruptured continuum, to underscore both the endurance of the common questions along with the challenge of accommodating a new scale. We regard collaborative systems as a combination of supportive services, encompassing safe data management and data sharing. Trustworthy data centred services are an essential support for collaboration at the scale of communities and organizations. We will combine our results and expertise to achieve a new leap forward toward the design of methods and techniques to enable the construction of usable large scale collaborative systems.

[1] See "Open API Growth: a Visualization" (ProgrammableWeb 2011)
[2] representational state transfer

3 Research program

3.1 Introduction

Our scientific foundations are grounded on distributed collaborative systems supported by sophisticated data sharing mechanisms and on service oriented computing with an emphasis on orchestration and on non-functional properties. Distributed collaborative systems enable distributed group work supported by computer technologies. Designing such systems requires an expertise in Distributed Systems and in Computer-supported collaborative Work. Besides theoretical and technical aspects of distributed systems, the design of distributed collaborative systems must take into account the human factor to offer solutions suitable for users and groups.

The Coast team vision is to move away from a centralized authority based collaboration toward a decentralized collaboration. Users will have full control over their data. They can store them locally and decide with whom to share them. The Coast team investigates the issues related to the management of distributed shared data and coordination between users and groups.

Service oriented Computing [25] is an established domain on which the ECOO, Score and now the Coast teams have been contributing for a long time. It refers to the general discipline that studies the development of computer applications on the web. A service is an independent software program with a specific functional context and capabilities published as a service contract (or more traditionally an API). A service composition aggregates a set of services and coordinates their interactions. The scale, the autonomy of services, the heterogeneity and some design principles underlying Service Oriented Computing open new research questions that are at the basis of our research. They span the disciplines of **distributed computing**, **software engineering** and **computer supported collaborative work** (CSCW). Our approach to contribute to the general vision of Service Oriented Computing is to focus on the issue of the efficient and flexible construction of reliable and secure high-level services. We aim to achieve it through the coordination/orchestration/composition of other services provided by distributed organizations or people.

3.2 Consistency Models for Distributed Collaborative Systems

Collaborative systems are distributed systems that allow users to share data. One important issue is to manage the consistency of shared data according to concurrent access. Traditional consistency criteria such as serializability, linearizability are not adequate for collaborative systems. Causality, Convergence and Intention preservation (CCI) [28] are more suitable for developing middleware for collaborative applications. We develop algorithms for ensuring CCI properties on collaborative distributed systems. Constraints on the algorithms are different according to the kind of distributed system and to the data structure. The distributed system can be centralized, decentralized or peer-to-peer. The type of data can include strings, growable arrays, ordered trees, semantic graphs and multimedia data.

3.3 Optimistic Replication

Replication of data among different nodes of a network promotes reliability, fault tolerance, and availability. When data are mutable, consistency among the different replicas must be ensured. Pessimistic replication is based on the principle of single-copy consistency while optimistic replication allows the replicas to diverge during a short time period. The consistency model for optimistic replication [27] is called eventual consistency, meaning that replicas are guaranteed to converge to the same value when the system is idle. Our research focuses on the two most promising families of optimistic replication algorithms for ensuring CCI:

- operational transformation (OT) algorithms [18]
- algorithms based on commutative replicated data types (CRDT) [26].

Operational transformation algorithms are based on the application of a transformation function when a remote modification is integrated into the local document. Integration algorithms are generic, being parametrised by operational transformation functions which depend on replicated document types. The advantage of these algorithms is their genericity. These algorithms can be applied to any data type

and they can merge heterogeneous data in a uniform manner. Commutative replicated data types are a new class of algorithms initiated by WooT [24], the first algorithm designed WithOut Operational Transformations. They ensure consistency of highly dynamic content on peer-to-peer networks. Unlike traditional optimistic replication algorithms, they can ensure consistency without concurrency control. CRDT algorithms rely on natively commutative operations defined on abstract data types such as lists or ordered trees. Thus, they do not require a merge algorithm or an integration procedure.

3.4 Process Orchestration and Management

Process Orchestration and Management is considered as a core discipline behind Service Management and Computing. It includes the analysis, the modelling, the execution, the monitoring and the continuous improvement of enterprise processes and is for us a central domain of study. Many efforts have been devoted to establishing standard business process models founded on well-grounded theories (e.g. Petri Nets) that meet the needs of business analysts, software engineers and software integrators. This has led to heated debate in the Business Process Management (BPM) community as the two points of view are very difficult to reconcile. On one side, business people in general require models that are easy to use and understand and that can be quickly adapted to exceptional situations. On the other side, IT people need models with an operational semantics in order to be able transform them into executable artifacts. Part of our work has been an attempt to reconcile these points of view. This resulted in the development of the Bonita BPM system. It also resulted more recently in our work on crisis management where the same people are designing, executing and monitoring the process as it executes. More generally, and at a larger scale, we have been considering the problem of processes spanning the barriers of organizations. This leads to the more general problem of service composition as a way to coordinate inter organizational construction of applications. These applications provide value, based on the composition of lower level services [15].

Recently, we also proposed new approaches for augmenting business processes. By relying mainly on data analysis, machine learning algorithms, and complex event processing, we exploited the data generated by business process execution (event data, event logs) and improved these processes from different perspectives such as instance scheduling and event management in an IoT environment [20].

4 Application domains

4.1 Crisis Management

Crisis management research investigates all the dimensions regarding the management of unexpected catastrophic events like floods, earthquakes, terrorist attacks or pandemics. All the phases of a crisis, from preparedness to recovery require collaboration between people from many organizations. This provides opportunities to study inter-organizational collaboration at a large scale and to propose and evaluate mechanisms that ensure secure and safe collaboration. The PhD thesis of Béatrice Linot supervised by François Charoy and Jérôme Dinet and defended in 2021 provided us with a deep understanding of the factors that encourage collaboration and help to maintain trustworthy collaboration between stakeholders. This work is continued by the PhD thesis of Clélie Amiot who studies the effects of human chat-bot collaboration in this kind of setting and with the PhD thesis of Lisa Formentini who will study the impact of the COVID lockdown on firefighters work practices and ecology of artefacts.

4.2 Collaborative Editing

Collaborative editing is a common application of optimistic replication in distributed settings. The goal of collaborative editors, irrespective of the kind of document, is to allow a group of users to update a document concurrently while ensuring that they eventually get all the same copy at the end. Our algorithm allows us to implement a collaborative editor in a peer to peer way. It avoids the need for a central server ensuring a higher level of privacy among collaborators. In this context, in order to decrease the overhead of the replicated data structure, we proposed a CRDT for Sequence which embeds a renaming mechanism [23]. The domain of collaborative editing requires us to consider the problem of access control of participants[11] and group key management[10].

4.3 Peer-to-peer storage

Peer-to-peer storage systems use the combined capacity of the peers to provide storage functionality to end users. Peer-to-peer storage systems are designed to provide persistence and availability of the stored content despite unreliability of the individual autonomous peers in a decentralised environment. We started to apply our work on data replication algorithms [7] and on group key management [10] for IPFS (InterPlanetary File System) peer-to-peer storage and we will transfer it to hive.

5 Highlights of the year

Claudia-Lavinia Ignat and François Charoy are the co-coordinators together with Myriam Lewkowicz (Université de Technologie de Troyes) of the PILOT project (2023-2030) of PEPR eNSEMBLE. The 13 partners of this project are the top French institutions with an interdisciplinary expertise on collaborative systems.

6 New software, platforms, open data

6.1 New software

6.1.1 MUTE

Name: Multi-User Text Editor

Keywords: Collaborative systems, Peer-to-peer, Replication and consistency, Privacy, Distributed systems, CRDT

Scientific Description: MUTE is a peer-to-peer collaborative editing platform that is used to evaluate the performance of replication algorithms in editing situations and to understand how it affects user experience.

Functional Description: MUTE (Multi-User Text Editor) is a web-based real-time collaborative editor that overcomes the limitation of existing collaborative systems which generally rely on a service provider that stores and has control over user data which is a threat for privacy. MUTE uses a peer-to-peer architecture and therefore it is highly scalable and resilient to faults and attacks. Several users may edit in real-time a shared document and their modifications are immediately sent to the other users without transiting through a central server. Our editor offers support for working offline while still being able to reconnect at a later time, which gives it a unique feature. Data synchronisation is achieved by using the LogootSplit algorithm developed by Coast.

News of the Year: In 2023 we replaced the underlying architecture which was mainly based on WebRTC (Web Real-time Communication) with libP2P from IPFS (InterPlanetary File System) project.

URL: https://github.com/coast-team/mute

Publications: hal-00903813, hal-01655438, hal-03772633

Contact: Gerald Oster

Participants: Claudia-Lavinia Ignat, François Charoy, Gerald Oster, Luc André, Matthieu Nicolas, Victorien Elvinger, Baptiste Hubert

6.1.2 Synql

Name: Conflict-free replicated relational database for SQLite

Keywords: Relational database, CRDT, Replication and consistency, Integrity constraints

Scientific Description: Synql allows to replicate an existing relational database without modifying the database engine or the application. To do this, Synql relies on a Git-like model. First the administrator has to initialize an existing database in order to obtain a replicated database. The initialization creates new relations and new triggers that store and maintain replicated metadata. Metadata allows us to synchronize several database replicas and to resolve potential conflicts. An administrator can add replicas by cloning an existing replica. The replicas can be concurrently updated without any coordination. The application reads and updates its database in the usual way by submitting SQL requests. The database triggers automatically update the replicated metadata. The replicas are synchronized in background.

Our replication mechanism is defined by the composition of CRDT (Conflict-free Replicated Data Types) primitives. We identify every inserted tuple with a globally unique identifier consisting of a monotonically increasing timestamp and a replica identifier. Each replica maintains a causal context that maps every replica identifier to the latest timestamp generated by the replica. The causal context allows fine-grained synchronization between any pair of replicas. The state of the database is computed from the replicated state by deterministically resolving all integrity violations.

Functional Description: Many applications use an embedded relational database, such as SQLite, to manage their local data. The replication of the database eases the addition of collaborative features to its applications. Most of the approaches for replicating a relational database require coordination.

Synql is a proof of concept of a coordination-less replication for relational databases that allows offline work and that respects commonly used integrity constraints such as uniqueness and referential integrity. The current implementation relies on SQLite. Synql is written in SQL. It can be used in existing database instances without changing the SQLite engine.

News of the Year: In 2023 we improved the code base and the continuous integration workflow.

URL: https://github.com/coast-team/synql

Publication: hal-02983557

Contact: Claudia-Lavinia Ignat

Participants: Victorien Elvinger, Claudia-Lavinia Ignat

7 New results

7.1 Distributed Access Control using CRDTs

Participants: Claudia-Lavinia Ignat, Olivier Perrin, Pierre-Antoine Rault.

Existing access control mechanisms mainly based on a central authority feature several difficulties in the context of collaborative systems. In the case of a federation of organizations, agreeing on an authority that manages the access rights is almost impossible. The lack of a central authority raises issues of group management such as joining and leaving the group as well as rights revocation. Indeed, it should be possible for a partner to revoke granted rights without contacting an external authority. Moreover, current access control mechanisms feature performance issues that are critical for real-time collaboration when the number of updates is high. Indeed, delays are too high for sending at each user action an access request and waiting for its answer from a trusted central authority which maintains the security policies. In order to avoid the use of a central server that stores all data, we propose that access rights as well as data are replicated.

We proposed CRDT algorithms for the synchronisation of access rights that were composed with CRDT algorithms for data synchronisation. In the face of concurrent edits on the access rights and the document, conflicts are likely to occur. For instance, users might execute operations on the document while their rights of executing these operations are concurrently revoked. An important feature of

collaborative applications is to allow multiple dynamic administrators that can modify users access rights (e.g. read or write) to the shared documents. In the face of dynamic groups of administrators conflicts might get complex. For instance, an administrator might assign an access right to a user, while concurrently this administrator right is removed. We proposed a replicated access control mechanism that manages a collaborative document with multiple, dynamic administrators. Besides maintaining consistency over the replicated document state and access rights, our CRDT solution preserves document integrity and prevents unauthorized modifications. An a posteriori enforcement is provided in order to correct the document state by compensating the effect of unauthorized modifications. Based on causality mechanisms, operations on document are mapped to intervals of authorisations where a right is enabled or disabled. This allows us to generate the appropriate compensations to the document state even after unauthorized operations have been integrated [11].

7.2 CRDTs for Mutable Data in Peer-to-peer Storage Systems

Participants: Quentin Acher, Claudia-Lavinia Ignat, Shadi Ibrahim(Inria, Myriads team) .

The continuous growth in data volume increases the interest in using peer-to-peer systems not only to store static data (i.e., immutable data) but also to store and share mutable data – data that are updated and modified by multiple users. Unfortunately, current peer-to-peer systems are mainly optimized to manage immutable data. Thus, each modification creates a new copy of the file, which leads to a high useless network usage.

Conflict-free Replicated Data Types (CRDTs) are specific data types built in a way that mutable data can be managed without the need for consensus-based concurrency control. In [7] we demonstrated the potential benefits of integrating CRDTs in the InterPlanetary File System (IPFS), an open-source widely used peer-to-peer content sharing system, meant to become the storage layer of the decentralized web. In an attempt to fill the gap between theory and practice, we provided the first quantitative measurement of the performance of CRDTs in IPFS. We introduced IM-CRDT (InterPlanetary Merkle-CRDT), an implementation of Merkle-CRDTs in IPFS that focuses on set data types. Experiments on Grid'5000 showed that IM-CRDT reduces the data transfer of an update by up to 99.96% and the convergence time by 54.6%-62.6%. More importantly, we find that IM-CRDT can sustain low convergence time under concurrent updates.

7.3 CRDTs for Replicated Relational Databases with Integrity Constraints

Participants: Victorien Elvinger, Ba Habibatou, Claudia-Lavinia Ignat.

Many offline-first applications use an embedded relational database, such as SQLite, to manage their data. The replication of the database eases the addition of collaborative features to its applications. Most of the approaches for replicating a relational database require coordination at some extent. A few approaches propose a coordination-less replication to allow offline work. These approaches are limited in two ways: (i) They do not respect Strong Eventual Consistency that states that two replicas converge as soon as they integrate the same set of modifications; (ii) They fail to preserve the combined effect of operations' intent in complex scenarios.

In [14] we proposed a new CRDT for replicating relations and maintaining integrity constraints in face of concurrent modifications. In contrast to previous approaches, our proposal enforces Strong Eventual Consistency and respects combined effect of operations' intent in complex scenarios. Its replicated state consists of the composition of CRDT primitives. The state of the database is computed from the replicated state by deterministically resolving all integrity violations. Local modifications are compensated in a way that ensures the preservation of combined effect of operations' intent.

7.4 Distributed Delivery Service for Group Key Agreement Protocols

Participants: Davide Frey(Inria, WIDE team) , Claudia-Lavinia Ignat, Amine Ismail(hive) , Ludovic Paillat (hive) , Mathieu Turuani(Inria, PESTO team) .

End-to-end encrypted messaging applications such as Signal and Whatsapp became widely popular thanks to their capability to ensure the confidentiality and integrity of online communication. While the highest security guarantees were long reserved to two-party communication, solutions for n-party communication remained either inefficient or less secure until the standardization of the MLS Protocol (Messaging Layer Security). The MLS protocol relies on a Group Key Agreement Protocol that allows members of a group to derive a common secret called group key which serves as a basis to secure group communications. It is scalable in terms of the number of operations modifying the group such as adding/removing members and it supports periodic group-key renewals preventing compromised communication. The MLS Protocol offers an efficient solution to guarantee the confidentiality and integrity of communication. However, the availability of the protocol depends on the centralized Delivery-Service component. The centralization of this component makes it an ideal target for attackers who wish to disrupt communication. Notably, with the help of a compromised Delivery Service, an attacker can prevent group members from refreshing their keys and resolving the compromise.

In order to overcome these limitations we proposed a fully distributed Delivery Service [10]. It combines two distributed communication mechanisms adapted to the need of the messages exchanged by the protocol. We used a Probabilistic Reliable Broadcast mechanism to reliably deliver messages allowing users to propose changes to the group (i.e. Proposal messages) and a Cascade Consensus Protocol to deliver the messages that actually modify the group (i.e. Commit messages) and thus require an agreement between members.

7.5 Ethereum's storage enhancement

Participants: Jean-Philippe Eisenbarth(University of Luxembourg) , Olivier Perrin, Thibaut Cholez(Inria, RESIST team) .

Blockchains face many challenges in time, among which the ever-growing storage needs for blockchain data. In particular, Ethereum is quickly approaching the 1 TB storage limit for a node. This may significantly reduce the candidates able to run an Ethereum fullnode.

In 2023 we proposed an elegant way to reduce the storage needs of an Ethereum node by leveraging Ethereum's DHT (Distributed Hash Table) to distribute the storage of old blocks and receipts (what we called cold data). Indeed, Ethereum's peer-to-peer (P2P) network is based on a well-known architecture named Kademlia that can efficiently retrieve data at the P2P network level with a logarithmic complexity. However, while fully functional, Ethereum's DHT is under-exploited so far and only used to list contacts (known nodes) but not to store any blockchain data. To understand why, we investigated the Ethereum's past and current storage and synchronization strategies and we explained why the distribution of the storage makes much sense today, blocks being no more executed since genesis block by a new node. We then proposed and evaluated a new synchronization strategy that takes advantage of Ethereum's DHT wasted potential, is fully backward compatible with current clients and does not affect any core mechanism of the Ethereum blockchain. We provided an implementation of our new synchronization mechanism exploiting Ethereum's DHT in Geth. Our evaluation on a private Ethereum instance proved the validity of the solution. When applied to an Ethereum peer, our solution saves roughly 60% of the storage of a node (360 GB) which represents a total of 12 PB of data at the network scale [8, 9].

7.6 Impact of Chatbots on Virtual Teamwork Dynamics and Performance

Participants: Clélie Amiot, François Charoy.

We studied the role of chatbots as a pivotal element in enhancing virtual teamwork. We delved into the effects of chatbots on group dynamics and performance within an online collaborative setting. To this end, a unique collaborative online activity was developed, completed with an integrated platform and a custom-designed chatbot assistant. The study involved 72 participants, systematically arranged into teams of four. These teams were further allocated into four distinct experimental conditions based on the nature of chatbot assistance provided: no assistance, private chat assistance, group chat assistance, or a combination of both.

The core findings of this investigation revealed a pronounced enhancement in team performance metrics attributable to the chatbot intervention. Teams with chatbot assistance exhibited not only improved performance but also experienced a notable reduction in response times for information requests during the collaborative activity. This improvement underscores the efficacy of chatbots in streamlining communication and information dissemination within team settings.

A particularly compelling aspect of our findings was the significant correlation observed between the chatbot's communication capabilities and the cognitive workload of team members. Teams interacting with chatbots demonstrating higher communication proficiency reported reduced cognitive strain, suggesting that the quality of chatbot interaction plays a crucial role in the overall team experience[12, 13].

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

Fair & Smart

Company: Fair & Smart

Dates: 2020-2024

Participants: Claudia-Lavinia Ignat *(contact)*, Gérald Oster, Olivier Perrin
.

The goal of this project is the development of a platform for the management of personal data according to the General Data Protection Regulation (GDPR). The other partners of this project are CryptoExperts and team READ from LORIA. The computational personal trust model that we proposed for repeated trust game [17] and its validation methodology [19] was adapted for the Fair&Smart personal data management platform for computing trust between the different stakeholders of this platform. Our decentralised mechanism for identity certification relying on a blockchain [21, 22] was transfered to Fair& Smart for user identification for their personal data management platform.

9 Partnerships and cooperations

9.1 National initiatives

9.1.1 Inria Challenge

Alvearium between Inria and hive

Title: Large Scale Secure and Reliable Peer-to-Peer Cloud Storage

Dates: 2022-2026

Inria coordinator: Claudia-Lavinia Ignat

Inria teams: Coast, Coati, Myriads, Pesto, Wide

Participants: Claudia-Lavinia Ignat *(contact)*, Thomas Lambert, Gérald Oster.

The project aims to propose an alternative peer-to-peer cloud which provides both computing and data storage via a peer-to-peer network rather than from a centralised set of data centers. HIVE proposes to exploit the unused capacity of computers and to incentivize users to contribute their computer resources to the network in exchange for similar capacity from the network and/or monetary compensation. By exchanging similar computer resources and network capacity users can benefit from all cloud services. Peers store encrypted fragments of the data of other peers. This proposed peer-to-peer cloud solution addresses users concerns about the privacy of their data and the dependency on centralised cloud providers. In this collaboration with HIVE we will apply our work on replication mechanisms for sharded encrypted data, data placement, Byzantine fault tolerance and security mechanisms in peer-to-peer environments.

9.1.2 PEPR

PILOT project of PEPR eNSEMBLE (Future of digital collaboration)

Title: Practices and infrastructure lor Long-term collaboration

Dates: 2023-2030

Coordinators: François Charoy (Université de Lorraine), Claudia-Lavinia Ignat (Inria), Myriam Lewkowicz (Université de Technologie de Troyes)

Partners: Inria (coordinator), CNRS, Université Grenoble Alpes, Université Paris-Saclay, Sorbonne Université, IMT, Université de Technologie de Troyes, INSA Lyon, Université Claude Bernard, Nantes Université, ENSAM, Université de Lille, Université de Toulouse III

Participants: François Charoy, Claudia-Lavinia Ignat *(contact)*, Gérald Oster, Olivier Perrin.

The project aims to design and engineer collaborative platforms that build upon regulatory challenges, organizational theories, and field descriptions. The project seeks to anticipate technological and societal evolutions and enable a French (or European) exception on digital platforms that guarantee individual actors' autonomy and foster care, trust, and digital well-being. The project's key challenges stem from revisiting the socio-technical stack, which includes novel conceptual models and design frameworks for long-term collaborative practices and enabling fluid collective experiences that support interoperability and evolution.

TRUSTINCloudS project of PEPR Cloud

Title: Cybersecurity of cloud infrastructures

Dates: 2023-2030

Coordinator: CEA (Aymen Boudguiga)

Partners: AMU, IMT, UL, EURECOM, UT3, CEA, INRIA

Participants: Claudia-Lavinia Ignat *(contact)*.

TRUSTINCloudS project develops solutions for the major cybersecurity challenges specific to Cloud environments, in order to ensure the confidentiality, integrity and availability of data, applications and services. The work carried out in this project aims at adapting traditional security mechanisms to the characteristics of the Cloud in order to address the specific threats of the different types of Clouds (IaaS, PaaS,...). The main objective of TRUSTINCloudS is to study and develop new methodologies to strengthen Cloud security and implement them in platforms in order to build a sovereign and trusted Cloud. In the context of this project, Coast team will work on the security of peer-to-peer clouds for storage.

10 Dissemination

Participants: Khalid Benali, Gérôme Canals, François Charoy, Claudia-Lavinia Ignat, Thomas Lambert, Gérald Oster, Olivier Perrin, Pierre-Antoine Rault.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

- Claudia-Lavinia Ignat and François Charoy organized the kick-off meeting of PILOT project of PEPR eNSEMBLE in November 2023.

10.1.2 Scientific events: selection

Member of the conference steering committees

- Claudia-Lavinia Ignat was a member of the Steering Committee for the International Conference on Intelligent Computer Communication and Processing (ICCP) in 2023.
- Khalid Benali was a member of the Steering Committee for Inforsid in 2023.

Member of the conference program committees

- Claudia-Lavinia Ignat was an associate chair at the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW) in 2023. She was a PC member of the European Conference on Computer-Supported Cooperative Work (ECSCW) 2023, the International Conference on Cooperative Design, Visualization and Engineering (CDVE) 2023, the International Conference on Collaboration Technologies and Social Computing (CollabTech) 2023 and the International Conference on Intelligent Computer Communication and Processing (ICCP) 2023.
- Olivier Perrin was PC member of the ICSOC 2023 conference.
- Khalid Benali was PC member of WorldCist'23 (11st World Conference on Information Systems and Technologies), CCCI 2023 (15th International Conference on Computational Collective Intelligence), ACM MEDES 2023 (15th International Conference on Management of Digital EcoSystems), and OCTA'2023 (Organization of Knowledge and Advanced Technologies).
- François Charoy was a PC Member of ICSOC 2023 (International Conference on Service Oriented Computing) and ICWS 2023
- Gérald Oster was a PC member of the International Conference on Collaboration Technologies and Social Computing (CollabTech) 2023 and the International Conference on Intelligent Computer Communication and Processing (ICCP) 2023.

Reviewer

- Thomas Lambert was reviewer for FOCS'23 and CHEOPS'23.
- Gérald Oster reviewed papers for the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW) in 2023

10.1.3 Journal

Member of the editorial boards

- Claudia-Lavinia Ignat is an associate editor of Computer Supported Cooperative Work (CSCW): The Journal of Collaborative Computing and Work Practices.
- François Charoy is a member of the editorial board of Service Oriented Computing and Applications (Springer)

Reviewer - reviewing activities

- Olivier Perrin reviewed papers for the Computers & Security journal.
- François Charoy was reviewer for IEEE IOT and TSC
- Thomas Lambert was reviewer for IEEE Network Magazine

10.1.4 Invited talks

- Claudia-Lavinia Ignat participated to a panel discussion on the industrial and societal challenges of the PEPR eNSEMBLE in October 2023 (Video on YouTube).
- Claudia-Lavinia Ignat presented the Alvearium project to Inria's Scientific Board in December 2023.

10.1.5 Research administration

- François Charoy was an elected member of the CNU (Conseil National des Universités) 27 till August 2023. He was a member of the board as assessor. He is also co-head of the Computer Science mention of the IAEM Doctoral School (Université de Lorraine).
- Claudia-Lavinia Ignat was an elected member of the Inria's Evaluation Commitee till August 2023. She is member of the Inria Nancy-Grand Est "Bureau du Comité de Projets" (BCP). In 2023, she was a member of the hiring committees for Junior Research Scientists (CRCN/ISFP) at Inria Grenoble and Inria Lille.
- Thomas Lambert is the scientific deputy for the Nancy site of Grid5000/SLICES-FR. He is the coordinator for the Grid5000 datacenter of CPER GENI. He is the representative of the Inria Nancy center in the Inria shared computing ressources users' commitee.

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

Permanent members of the Coast project-team are leading teachers in their respective institutions. They are responsible of lectures in disciplines like software engineering, database systems, object oriented programming and design, distributed systems, service computing and more advanced topics at all levels and in different departments in the University. Most PhD Students have also teaching duties in the same institutions. Claudia-Lavinia Ignat teaches a course on data replication and consistency at Master level (M2 SIRAV) at Université de Lorraine. As a whole, the Coast team accounts for more than 2,500 hours of teaching. Members of the Coast team are also deeply involved in the pedagogical and administrative life of their departments.

- Khalid Benali is responsible for the professional Master degree speciality "Distributed Information Systems" of MIAGE (Université de Lorraine) and of its international branch in Morocco.
- Gérôme Canals is the deputy director of IUT Nancy-Charlemagne of Université de Lorraine.
- François Charoy is responsible for the Software Engineering specialisation at the TELECOM Nancy Engineering School of Université de Lorraine.
- Gérald Oster is the deputy director of TELECOM Nancy Engineering School of Université de Lorraine. He is responsible for the 3rd (last) year of study and President of the jury of the Diploma at TELECOM Nancy.

10.2.2 Supervision

- PhD in progress: Clélie Amiot, Trust and Human/Chatbot collaboration, started in October 2019, supervised by Jérome Dinet and François Charoy
- PhD in progress: Pierre-Antoine Rault, Security mechanisms for decentralised collaborative systems, started in October 2020, supervised by Claudia-Lavinia Ignat and Olivier Perrin
- PhD in progress: Quentin Acher, Management of mutable data over P2P storage, started in September 2023, supervised by Claudia-Lavinia Ignat and Shadi Ibrahim (Inria, Myriads team)
- PhD in progress: Ludovic Paillat (HIVE), Security for peer-to-peer cloud storage without central authority, started in October 2023, supervised by Claudia-Lavinia Ignat, Davide Frey (Inria, WIDE team), Mathieu Turuani (Inria, PESTO team) and Amine Ismail (HIVE)
- PhD in progress: Lisa Formentini, Artifact ecology evolution in civil security organisations, started in October 2023, supervised by François Charoy and Matthieu Tixier (Université de Technologie de Troyes)
- PhD in progress: Mohammad Rizk (Inria, Myriads team), started in November 2023, supervised by Shadi Ibrahim (Inria, Myriads team) and Thomas Lambert
- PhD discontinued: Alexandre Bourbeillon, Trust among users in collaborative systems, started in November 2020, interrupted since September 2022, discontinued in October 2023, supervised by Claudia-Lavinia Ignat

10.2.3 Juries

- Corentin Stalder, PhD defense jury, "Jumeau numérique support aux pratiques coopératives : application au système ferroviaire", Université de Technologie de Troyes, May 2023 (Claudia-Lavinia Ignat, Reviewer)
- William Aboucaya, PhD defense jury, "Collaborative systems for large-scale online citizen participation", Inria Paris, November 2023 (Claudia-Lavinia Ignat, Reviewer)
- Omar Hasan, HDR, "Privacy preservation in trust-deficient decentralized systems", INSA Lyon, June 2023 (François Charoy, Reviewer)
- Debashisha Mishra, PhD defense jury, Exploiting the synergies of unmanned aerial vehicles (UAVs) and 5G network, Université de Lorraine, June 2023 (François Charoy, President)
- Linda Ouchaou, PhD defense jury, "Système de découverte des services dans un environnement Cloud Computing", Université de Sciences et de la Technologie Houari Boumediene, September 2023 (François Charoy, guest)
- Amina Hamila Brahem, PhD defense jury, "Ensuring trustworthy and adaptable business process choreographies in open and dynamic environment", Université de Tour, December 2023 (François Charoy, Reviewer and President)
- Mohamed Oulaafaart, PhD Defense jury, "Automating Security Enhancement for Cloud Services", Université de Lorraine (François Charoy, President)

10.3 Popularization

10.3.1 Articles and contents

- an interview of Claudia-Lavinia Ignat on the Alvearium Inria challenge (interview link) that was also published in Le Monde (binaire) (article link)

10.3.2 Interventions

- Claudia-Lavinia Ignat presented her work on "Large scale trustworthy distributed collaboration" and Pierre-Antoine Rault presented his work on "Distributed access control for collaborative applications" at FADEX (French American Doctoral EXchanges) 2023 event on cybersecurity organised in October 2023 in Nancy and gathering doctoral students from various universities in USA, Inria (Nancy, Rennes and Paris) and Campus Cyber.
- Claudia-Lavinia Ignat participated in [#Visages d'Inria] "L'interview" in 2023 (video link).

11 Scientific production

11.1 Major publications

[1] C.-L. Ignat, L. André and G. Oster. 'Enhancing rich content wikis with real-time collaboration'. In: *Concurrency and Computation: Practice and Experience* 33.8 (25th Apr. 2021). DOI: 10.1002/cpe.4110. URL: https://hal.inria.fr/hal-01404024.

[2] C.-L. Ignat, Q.-V. Dang and V. Shalin. 'The Influence of Trust Score on Cooperative Behavior'. In: *ACM Transactions on Internet Technology* 19.4 (19th Sept. 2019), pp. 1–22. DOI: 10.1145/3329250. URL: https://hal.inria.fr/hal-02307981.

[3] H. Le Nguyen and C.-L. Ignat. 'An Analysis of Merge Conflicts and Resolutions in Git-based Open Source Projects'. In: *Computer Supported Cooperative Work* 27.3-6 (June 2018), pp. 741–765. DOI: 10.1007/s10606-018-9323-3. URL: https://hal.science/hal-01917249.

[4] M. Nicolas, G. Oster and O. Perrin. 'Efficient Renaming in Sequence CRDTs'. In: *IEEE Transactions on Parallel and Distributed Systems* 33.12 (1st Dec. 2022), pp. 3870–3885. DOI: 10.1109/TPDS.2022.3172570. URL: https://hal.inria.fr/hal-03772633.

[5] L. Paillat, C.-L. Ignat, D. Frey, M. Turuani and A. Ismail. 'Design of an Efficient Distributed Delivery Service for Group Key Agreement Protocols'. In: *Lecture Notes in Computer Science (LNCS)*. FPS 2023 - 16th International Symposium on Foundations & Practice of Security. Bordeaux, France, 11th Dec. 2023, pp. 1–16. URL: https://inria.hal.science/hal-04337821.

[6] G. Rosinosky, S. Youcef and F. Charoy. 'A Genetic Algorithm for Cost-Aware Business Processes Execution in the Cloud'. In: *Lecture Notes in Computer Science*. ICSOC 2018 - The 16th International Conference on Service-Oriented Computing. ICSOC 2018: Service-Oriented Computing 11236. Hangzhou, China: Springer, 7th Nov. 2018, p. 14. URL: https://hal.science/hal-01870828.

11.2 Publications of the year

International peer-reviewed conferences

[7] Q. Acher, C.-L. Ignat and S. Ibrahim. 'Quantifying the Performance of Conflict-free Replicated Data Types in InterPlanetary File System'. In: *Middleware 2023 Companion Proceedings*. DICG 2023 - 4th International Workshop on Distributed Infrastructure for Common Good. Bologna, Italy, 2023, pp. 1–6. DOI: 10.1145/3631310.3633488. URL: https://inria.hal.science/hal-04337761.

[8] J.-P. Eisenbarth, T. Cholez and O. Perrin. 'Avoiding the 1 TB Storage Wall: Leveraging Ethereum's DHT to Reduce Peer Storage Needs'. In: The 5th ACM International Symposium on Blockchain and Secure Critical Infrastructure (BSCI 2023). Melbourne, Australia, 10th July 2023, p. 10. URL: https://inria.hal.science/hal-04163897.

[9] J.-P. Eisenbarth, T. Cholez and O. Perrin. 'Valorisation de la DHT d'Ethereum pour réduire les besoins de stockage des pairs'. In: CoRes 2023 - 8èmes Rencontres Francophones sur la Conception de protocoles, l'évaluation de performances et l'expérimentation de Réseaux de communication. Cargèse (Corse), France, 22nd May 2023. URL: https://hal.science/hal-04080219.

[10] L. Paillat, C.-L. Ignat, D. Frey, M. Turuani and A. Ismail. 'Design of an Efficient Distributed Delivery Service for Group Key Agreement Protocols'. In: *Lecture Notes in Computer Science (LNCS)*. FPS 2023 - 16th International Symposium on Foundations & Practice of Security. Bordeaux, France, 11th Dec. 2023, pp. 1–16. URL: https://inria.hal.science/hal-04337821.

[11] P.-A. Rault, C.-L. Ignat and O. Perrin. 'Access control based on CRDTs for Collaborative Distributed Applications'. In: *Proceedings of the 22nd IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom-2023)*. The International Symposium on Intelligent and Trustworthy Computing, Communications, and Networking (ITCCN-2023), in conjunction with the 22nd IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom-2023). Exeter, United Kingdom, 1st Nov. 2023. URL: https://inria.hal.science/hal-04224855.

Conferences without proceedings

[12] C. Amiot, F. Charoy and J. Dinet. 'Chatbots as decision aids: investigating reliance in Human-Chatbot collaboration vs. Human-Human collaboration'. In: Workshop on Trust and Reliance in AI-Assisted Tasks, at CHI 2023. Hamburg, Germany, 23rd Apr. 2023. URL: https://hal.univ-lorraine.fr/hal-04229467.

Reports & preprints

[13] C. Amiot, F. Charoy and J. Dinet. *Whom Do We Trust ?: A Comparative Study on Reliance between Chatbot and Human Assistance*. 5th Oct. 2023. URL: https://hal.univ-lorraine.fr/hal-04229730.

[14] V. Elvinger and C.-L. Ignat. *Synql: Replicated Relations and Integrity Maintenance*. Inria, 21st Feb. 2023. URL: https://inria.hal.science/hal-03999168.

11.3 Cited publications

[15] S. Bhiri, O. Perrin, W. Gaaloul and C. Godart. 'An Object-Oriented Metamodel For Inter-Enterprises Cooperative Processes Based on Web Services'. Anglais. In: *Journal of Integrated Design and Process Science* 8 (2004), pp. 37–55. URL: http://hal.inria.fr/inria-00099953/en/.

[16] F. Casati. 'Promises and Failures of Research in Dynamic Service Composition'. English. In: *Seminal Contributions to Information Systems Engineering*. Ed. by J. Bubenko, J. Krogstie, O. Pastor, B. Pernici, C. Rolland and A. Sølvberg. Springer Berlin Heidelberg, 2013, pp. 235–239. URL: http://dx.doi.org/10.1007/978-3-642-36926-1_18.

[17] Q. V. Dang and C. Ignat. 'Computational Trust Model for Repeated TrustGames'. In: *Proceedings of the 15th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom 2016)*. Tianjin, China, Aug. 2016. URL: https://hal.inria.fr/hal-01351250.

[18] C. A. Ellis and S. J. Gibbs. 'Concurrency Control in Groupware Systems'. In: *Proceedings of the ACM SIGMOD Conference on the Management of Data - SIGMOD 89*. Portland, Oregon, USA, May 1989, pp. 399–407. URL: http://doi.acm.org/10.1145/67544.66963.

[19] C.-L. Ignat, Q.-V. Dang and V. Shalin. 'The Influence of Trust Score on Cooperative Behavior'. In: *ACM Transactions on Internet Technology* 19.4 (Sept. 2019), pp. 1–22. DOI: 10.1145/3329250. URL: https://hal.inria.fr/hal-02307981.

[20] A. Ismaili-Alaoui. 'Methodology for an Augmented Business Process Management in IoT Environment'. Theses. Université de Lorraine ; Université Mohammed V (Rabat), Dec. 2022. URL: https://hal.univ-lorraine.fr/tel-04062128.

[21] H. Nguyen, J.-P. Eisenbarth, C.-L. Ignat and O. Perrin. 'Blockchain-Based Auditing of Transparent Log Servers'. In: *The 32nd Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec 2018)*. Proceeding of Data and Applications Security and Privacy XXXII - 32nd Annual IFIP WG 11.3 Conference. Bergamo, Italy, July 2018, pp. 21–37. URL: https://hal.archives-ouvertes.fr/hal-01917636.

[22] H. Nguyen, C.-L. Ignat and O. Perrin. 'Trusternity: Auditing Transparent Log Server with Blockchain'. In: *Companion of the The Web Conference 2018*. Lyon, France, Apr. 2018, pp. 79–80. DOI: 10.1145/3184558.3186938. URL: https://hal.inria.fr/hal-01883589.

[23] M. Nicolas. 'Ré-identification sans coordination dans les types de données répliquées sans conflits (CRDTs)'. Theses. Université de Lorraine, Dec. 2022. URL: https://theses.hal.science/tel-04312434.

[24] G. Oster, P. Urso, P. Molli and A. Imine. 'Data Consistency for P2P Collaborative Editing'. Anglais. In: *ACM Conference on Computer-Supported Cooperative Work - CSCW 2006*. Banff, Alberta, Canada: ACM Press, Nov. 2006, pp. 259–268. URL: http://hal.inria.fr/inria-00108523/en/.

[25] M. P. Papazoglou, P. Traverso, S. Dustdar and F. Leymann. 'Service-Oriented Computing: State of the Art and Research Challenges'. In: *Computer* 40 (2007), pp. 38–45.

[26] N. Preguiça, J. M. Marquès, M. Shapiro and M. Letia. 'A commutative replicated data type for cooperative editing'. Anglais. In: *29th IEEE International Conference on Distributed Computing Systems (ICDCS 2009)*. Montreal, Québec Canada: IEEE Computer Society, 2009, pp. 395–403. DOI: 10.1109/ICDCS.2009.20. URL: http://hal.inria.fr/inria-00445975/en/.

[27] Y. Saito and M. Shapiro. 'Optimistic Replication'. In: *Computing Surveys* 37.1 (Mar. 2005), pp. 42–81. URL: http://doi.acm.org/10.1145/1057977.1057980.

[28] C. Sun, X. Jia, Y. Zhang, Y. Yang and D. Chen. 'Achieving Convergence, Causality Preservation, and Intention Preservation in Real-Time Cooperative Editing Systems'. In: *ACM Transactions on Computer-Human Interaction* 5.1 (Mar. 1998), pp. 63–108. URL: http://doi.acm.org/10.1145/274444.274447.