

RESEARCH CENTRE

Inria Nancy - Grand Est Center

IN PARTNERSHIP WITH:

Université de Lorraine, CNRS

2022

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

The Inria logo is a stylized, red, cursive script of the word "Inria".

Contents

Project-Team COAST	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
3 Research program	4
3.1 Introduction	4
3.2 Consistency Models for Distributed Collaborative Systems	4
3.3 Optimistic Replication	4
3.4 Process Orchestration and Management	5
3.5 Service Composition	5
4 Application domains	5
4.1 Crisis Management	5
4.2 Collaborative Editing	5
5 Highlights of the year	6
6 New software and platforms	6
6.1 New software	6
6.1.1 MUTE	6
6.1.2 Synql	6
7 New results	7
7.1 Overcoming Identifier Overhead in Sequence CRDTs	7
7.2 Distributed Access Control using CRDTs	8
7.3 Analysis of Social Networks as Collaboration Support	8
7.4 Ethereum's Peer-to-Peer Network Monitoring	8
7.5 Composite Service Selection based on API Call Limit	9
7.6 Straggler Detection in Big Data Analytic Systems	9
7.7 Privacy-Preserving Graph Processing in Geo-Distributed Data Centers	9
7.8 Proactive IoT Business Process Management with Data Analytics and Event Processing	10
8 Bilateral contracts and grants with industry	10
8.1 Bilateral contracts with industry	10
9 Partnerships and cooperations	11
9.1 National initiatives	11
9.1.1 Inria Challenge	11
10 Dissemination	11
10.1 Promoting scientific activities	12
10.1.1 Scientific events: organisation	12
10.1.2 Scientific events: selection	12
10.1.3 Journal	13
10.1.4 Research administration	13
10.2 Teaching - Supervision - Juries	13
10.2.1 Teaching	13
10.2.2 Supervision	14
10.2.3 Juries	14
10.3 Popularization	14
10.3.1 Education	14

11 Scientific production	14
11.1 Major publications	14
11.2 Publications of the year	15
11.3 Cited publications	16

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]
- G er ome Canals [UL, Associate Professor]
- Fran ois Charoy [UL, Professor, (Team Leader until August 2022), HDR]
- Claude Godart [UL, Professor, until Oct 2022, HDR]
- Thomas Lambert [UL, Associate Professor]
- Gerald Oster [UL, Associate Professor]
- Olivier Perrin [UL, Professor, HDR]
- Samir Youcef [UL, Associate Professor, until Aug 2022]

PhD Students

- Cl elie Amiot [CNRS]
- Alexandre Bourbeillon [INRIA]
- Abir Ismaili-Alaoui [UL]
- Matthieu Nicolas [UL, ATER, until Aug 2022]
- Linda Ouchaou [UL, ATER, from Sep 2022]
- Pierre-Antoine Rault [INRIA]

Technical Staff

- Victorien Elvinger [INRIA, Engineer]
- Baptiste Hubert [INRIA, Engineer]

Interns and Apprentices

- Sami Djouadi [Inria, from Mar 2022 until Jul 2022]

Administrative Assistants

- Sophie Drouot [INRIA]
- Nathalie Fritz [UL]

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 25 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 appliance¹. 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 [21]. 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 [19]. 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 REST² 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.

¹See "Open API Growth: a Visualization" (ProgrammableWeb 2011)

²representational 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 [26] 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) [30] 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 [28] 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 [21]
- algorithms based on commutative replicated data types (CRDT) [27].

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 [25], 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 [18].

3.5 Service Composition

Recently, we started a study on service composition for software architects where services are coming from different providers with different plans (capacity, degree of resilience...). The objective is to support the architects to select the most accurate services (w.r.t. to their requirements, both functional and non-functional) and plans for building their software. We also compute the properties that we enforce for the composition of these services.

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.

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, it requires us to consider the problem of access control of participants[15].

5 Highlights of the year

The Inria Challenge Alvearium (<https://project.inria.fr/alvearium/>) between Inria and HIVE coordinated by Coast started on December 1, 2022.

6 New software and platforms

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 2022 we refactored the base code in order to use the InterPlanetary File System (IPFS) protocol as our new underlying peer-to-peer communication layer.

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.

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 Overcoming Identifier Overhead in Sequence CRDTs

Participants: Matthieu Nicolas, Gérald Oster, Olivier Perrin.

In order to achieve high availability, large-scale distributed systems often replicate data and minimize coordination among nodes. One approach that has gained attention in both literature and industry is the use of Conflict-free Replicated Data Types (CRDTs) to design these systems. CRDTs are new specifications of existing data types, such as sets or sequences, that maintain the same behavior as previous specifications in sequential executions, but excel in distributed settings by natively supporting concurrent updates. To accomplish this, CRDTs embed conflict resolution mechanisms into their specifications. These mechanisms typically rely on identifiers attached to elements of the data structure to resolve conflicts in a deterministic and coordination-free manner. However, these identifiers must comply with certain constraints, such as uniqueness or belonging to a dense total order, which can increase the size of the identifiers over time and cause performance issues. To address this problem, we propose a novel Sequence CRDT [10] that incorporates a renaming mechanism, allowing nodes to reassign shorter identifiers to elements in an uncoordinated manner.

We conducted an experimental evaluation to assess the performance of our proposed renaming mechanisms in various areas. Specifically, we evaluated the impact of the renaming mechanism on: (i) the size of the data structure; (ii) the integration time of the rename operation; and (iii) the integration time of insert and remove operations. For comparison, we used LogootSplit [17] (see section 6.1.1) as the baseline data structure in cases (i) and (iii). The results we obtained were very promising, as we observed a significant reduction in integration time when using the renaming mechanism, even when taking into account the time required to execute the rename operation.

7.2 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. Lack of a central authority raises issues of group management such as joining and leaving the group as well as rights revocation. 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.

We proposed a distributed access control mechanism for collaborative applications where access rights as well as data are replicated [15]. We illustrated by means of examples the challenges faced by replication algorithms for obtaining consistency over both data and access control policies in distributed applications such as Google Docs and POSIX file systems. We provided a replication algorithm for access control policies based on Conflict-free Replicated Data Types (CRDTs). We also gave an overview of our proposed solution for the composition of the CRDT on access control policies with a CRDT on data.

7.3 Analysis of Social Networks as Collaboration Support

Participants: Quentin Laporte Chabasse, Gérald Oster, François Charoy.

Creating secure peer-to-peer collaborative services requires a trusted peer-to-peer network. In order to achieve this goal, we examined how to utilize the underlying social networks of inter-organizational collaboration to support such collaboration. To accomplish this, we analyzed collaborative graphs, which provide valuable insights into the behavior of groups of individuals. Exponential Random Graph Models (ERGMs) are commonly used to analyze social processes and dependencies among group members. Our approach uses a modified version of ERGMs, modeling the problem as an edge-labeling one. The main challenge is inferring the model, as the normalizing constant involved in traditional Markov Chain Monte Carlo approaches is not available in closed-form.

We proposed [9] to use the ABC Shadow algorithm [29], which can sample from posterior distributions while avoiding this limitation. The method was demonstrated using real data sets provided by the HAL platform and offers new insights into self-organized collaborations among researchers. We applied this method in a longitudinal study to identify patterns in the evaluation of collaboration over multiple years for the same teams. We also applied it to a dataset from a social study in a French primary school.

7.4 Ethereum's Peer-to-Peer Network Monitoring

Participants: Olivier Perrin, Thibault Cholez (RESIST team) , Jean-Philippe Eisenbarth (RESIST team) .

The two main blockchains, Bitcoin and Ethereum, rely on peer-to-peer networks to ensure their functionalities. In the work carried out with Jean-Philippe Eisenbarth and Thibault Cholez (from the RESIST team), we first proposed an in-depth study of these networks using a supervision mechanism. In particular, we have studied a number of criteria related to reliability, such as the number of peers, their geographical distribution, their distribution over the IP network, the churn, the proportion of clients with known vulnerabilities, the existence of daily connection patterns or the ability to infer topology. It appears that both networks show good properties on all these points.

Second, based on the observation that, on the one hand, the distributed hash table (DHT) of the Ethereum P2P network is largely untapped, and that, on the other hand, the storage of blockchain

data is only growing (which will pose problems in the long run), we have developed a new distributed storage architecture for Ethereum that takes advantage of the DHT. The proposed solution is backward compatible with current clients and can reduce the disk space used for long-term storage by 95% (58% of total storage) without impacting the guarantees or performance of the Ethereum blockchain.

Finally, we analyzed Ethereum peers for patterns that could reflect Sybil attacks and showed the existence of thousands of suspicious nodes grouping a large number of identifiers for the same IP address (up to 10,000/IP). Following this analysis, we designed and implemented a protection architecture against Sybil attacks. It is based on a crawler detecting suspicious nodes in real time, a smart contract structuring the information and distributing it to all peers, and finally a completely distributed revocation mechanism, each peer noticing itself the attack and cutting its connections to Sybils nodes. The implementation on an Ethereum test network has shown the efficiency of the proposed architecture [7].

7.5 Composite Service Selection based on API Call Limit

Participants: François Charoy.

Today, services, also known as APIs, have different types of call limits. These limits can include the number of requests that can be made to the API within a certain time period, the number of concurrent connections allowed, or the amount of data that can be transferred. These call limits are put in place by the API providers to ensure the stability and performance of their services, and to manage the cost of providing them. Different APIs may have different call limits, depending on their intended usage and the resources required to provide the service. This can make it challenging for customers, particularly developers, to select the most appropriate services for their needs, and to effectively compose different services together. In [12] we proposed an approach for selecting the most relevant compositions of APIs based on this notion of call limit. Specifically, we showed how the call limits of the individual services can be aggregated to obtain the call limits of a given composition. We introduced the notion of minimal budget skyline, which comprises the most interesting compositions that fit within the customer's budget. In addition, we developed two algorithms, based on effective pruning strategies, to efficiently compute the minimal budget skyline. We validated the approach experimentally.

7.6 Straggler Detection in Big Data Analytic Systems

Participants: Thomas Lambert.

Speculative execution can significantly improve the performance of Big Data applications by launching other copies of stragglers (slow tasks). Straggler detection plays an important role in the effectiveness of speculative execution. Not detecting real stragglers (false negatives) may slow down the whole computation. On the contrary, launching speculative copies of non-straggler tasks (false positives) can lead to a waste of resources. Most state-of-the-art methods employed to detect stragglers use the information extracted from the last received heartbeats which may be outdated when triggering detection. This, in turn, can mislead Big Data analytics systems to make inaccurate detection.

In [14] we shed light on this issue by carrying out extensive simulations to identify how heartbeat arrival, task starting times, and detection methods impact the accuracy of straggler detection in Big Data analytic systems. In particular we investigated two families of state-of-the-art detection methods: one based on progression score and the other one based on progression speed. We showed that they can both lead to rather large inaccuracy, by either overlooking real stragglers or by marking normal tasks as stragglers. We also highlighted the fact that this inaccuracy is not only due to outdated information, but also to the asynchrony of starting times and heartbeat arrivals.

7.7 Privacy-Preserving Graph Processing in Geo-Distributed Data Centers

Participants: Thomas Lambert.

Graph processing is a popular computing model for big data analytics. Very roughly, in such model, data are stored on different nodes of a graph and are updated iteratively with data transferred from the neighboring nodes. The famous Page Rank algorithm, which is very popular for ranking web pages or users, is an example of distributed analytics that can use this model. Meanwhile, emerging big data applications are often maintained in multiple geographically distributed (geo-distributed) data centers (DCs) to provide low-latency services to global users. Graph processing in geo-distributed DCs suffers from costly inter-DC data communications. Furthermore, due to increasing privacy concerns, geo-distribution imposes diverse, strict, and often asymmetric privacy regulations that constrain geo-distributed graph processing. For example, the European General Data Protection Regulation (GDPR) does not allow the movement of private data from a DC in Europe to a DC in the USA. Existing graph processing systems fail to address these two challenges.

In [16], we designed and implemented PGPreGel, an end-to-end system that provides privacy-preserving graph processing in geo-distributed DCs with low latency and high utility. The design of PGPreGel is based on three techniques. The first technique is differential privacy, a popular way to preserve privacy by adding noise to transferred data. The second technique is sampling, i.e. only part of the data is sent from one DC to another. The objective is to reduce both inter-DC communication and the amount of added noise in order to keep the same level of privacy (which results in a possible improvement of accuracy). The last technique is combiners, i.e. aggregated messages. Similarly to sampling, combiners reduce inter-DC communication and mitigate the impact of noise on accuracy.

We implemented our design in Apache Giraph(<https://giraph.apache.org/>) and evaluated it in real cloud DCs. We targeted four graph processing applications. Results show that PGPreGel can preserve the privacy of graph data with low overhead and good accuracy, beating strategies based on differential privacy, no-inter-DC communication, or application-specific privacy-preserving implementation alone.

7.8 Proactive IoT Business Process Management with Data Analytics and Event Processing

Participants: Khalid Benali, Abir Ismaili-Alaoui.

IoT is becoming a hot-spot area of technological innovations and economic development promises for many industries and services. This new paradigm shift affects all enterprise architecture layers, from infrastructure to business. Business Process Management (BPM) is affected by this new technology.

To tackle the data and event explosion resulting, among others, from IoT, data analytic processes are combined with event processing techniques. These techniques examine large data sets to uncover hidden patterns, and unknown correlations between collected events, either at a very technical level (incident/anomaly detection, predictive maintenance) or at the business level (customer preferences, market trends, revenue opportunities). They provide improved operational efficiency, better customer service, and competitive advantages over rival organizations.

In [8] and [13] we 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.

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, François Charoy.

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 [20] and its validation methodology [22] will be adapted for the Fair&Smart personal data management platform for computing trust between the different users of this platform. Our decentralised mechanism for identity certification relying on a blockchain [23, 24] will be transferred 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 (<https://project.inria.fr/alvearium/>) between Inria and HIVE (<https://www.hivenet.com/>)

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

Dates: 2022-2026

Inria coordinator: Claudia-Lavinia Ignat

Inria teams: Coast, Coati, Myriads, 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.

10 Dissemination

Participants: Khalid Benali, G r me Canals, Fran ois Charoy, Claude Godart, Claudia-Lavinia Ignat, Thomas Lambert, G rald Oster, Olivier Perrin.

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

Member of the organizing committees

- Claudia-Lavinia Ignat was a member of the Scientific Organisation Committee for the Inria Prospective Seminar on Distributed Systems and Middleware in 2022.

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 2022.

Member of the conference program committees

- Khalid Benali was a PC member of WorldCIST'22 (10th World Conference on Information Systems and Technologies), I3E 2022 (IFIP Conference on e-Business, e-Services and e-Society), INFORSID 2022 (INformatique des ORganisations et Systèmes d'Information et de Décision), ICCCI 2022 (14th International Conference on Computational Collective Intelligence), and MEDES 2022 (14th International Conference on Management of Digital EcoSystems).
- François Charoy was a PC Member of ICSSOC 2022 (International Conference on Service Oriented Computing) and of several workshops.
- Claudia-Lavinia Ignat was an associate chair at the ACM CHI Conference on Human Factors in Computing Systems (CHI) 2022. She was a PC member of the European Conference on Computer-Supported Cooperative Work (ECSCW) 2022, the International Conference on Cooperative Design, Visualization and Engineering (CDVE) 2022, the International Conference on Collaboration Technologies and Social Computing (CollabTech) 2022 and the International Conference on Intelligent Computer Communication and Processing (ICCP) 2022.
- Thomas Lambert was a PC member of ICPP 2022 (International Conference on Parallel Processing).
- Gérald Oster was a PC member of the International Conference on Collaboration Technologies and Social Computing (CollabTech) 2022, and the International Conference on Intelligent Computer Communication and Processing (ICCP) 2022.
- Olivier Perrin was a PC Member of ICSSOC 2022 (International Conference on Service Oriented Computing) and of several workshops.

Reviewer

- Claudia-Lavinia Ignat reviewed articles for the International Workshop on Distributed Infrastructure for Common Good (DICG 2022) co-located with ACM/IFIP Middleware 2022.
- Thomas Lambert reviewed articles for CCGrid22 (International Symposium on Cluster, Cloud and Internet Computing).
- Gérald Oster reviewed articles for CSCW 2022 (ACM Conference On Computer-Supported Cooperative Work And Social Computing) and CODASPY 2023 (ACM Conference on Data and Application Security and Privacy).

10.1.3 Journal

Member of the editorial boards

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

Reviewer - reviewing activities

- Claudia-Lavinia Ignat reviewed articles for Computer Supported Cooperative Work (CSCW) journal.
- Thomas Lambert reviewed an article for the IEEE Transactions on Parallel and Distributed Systems (TPDS).
- Gérald Oster reviewed an article for Computer Supported Cooperative Work (CSCW).
- Olivier Perrin reviewed papers for the International Journal of Information Management.

10.1.4 Research administration

- François Charoy is an elected member of the CNU (Conseil National des Universités) 27. He is 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 is a member of the Inria Evaluation Commission. She is member of the Inria Nancy-Grand Est "Bureau du Comité de Projets" (BCP) and of the Inria Nancy-Grand Est COMIPERS committee. In 2022, she was a member of the CRCN recruitment jury at Inria Grenoble and Inria Sophia Antipolis and of Inria secondment jury.

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.
- Claude Godart was responsible for the Computer Science Department of the Polytech 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 defended: Abir Ismaïli-Alaoui, Methodology for an Augmented Business Process Management in IoT Environment", defended in December 2022, Khalid Benali and Karim Baïna (Université Mohammed V, Rabat, Morocco)
- PhD defended: Matthieu Nicolas, Coordination-free re-identification in Conflict-free Replicated Data Types, defended in December 2022, Olivier Perrin and Gérald Oster
- PhD defended: Jean Philippe Eisenbarth (RESIST team), Analysis and protection of public blockchains, defended in December 2022, Olivier Perrin and Thibault Cholez (RESIST team)
- PhD in progress: Clélie Amiot, Trust and Human/Chatbot collaboration, started in October 2019, Jérôme Dinet and François Charoy
- PhD in progress: Alexandre Bourbeillon, Trust among users in collaborative systems, started in November 2020, Claudia-Lavinia Ignat
- PhD in progress: Pierre-Antoine Rault, Security mechanisms for decentralised collaborative systems, started in October 2020, Claudia-Lavinia Ignat and Olivier Perrin

10.2.3 Juries

- Julien Coche, PhD, Ecole des Mines d'Albi, May 2022 (François Charoy, Jury President)
- Clément Cormi, PhD, Université Technologique de Troyes, December 2022 (François Charoy, Reporter)

10.3 Popularization

10.3.1 Education

In March 2022 Claudia-Lavinia Ignat presented her research works to the first year students at École des Mines de Nancy while they were visiting Loria.

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](https://doi.org/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](https://doi.org/10.1145/3329250). URL: <https://hal.inria.fr/hal-02307981>.
- [3] Q. Laporte-Chabasse, R. S. Stoica, M. Clausel, F. Charoy and G. Oster. 'Morpho-statistical description of networks through graph modelling and Bayesian inference'. In: *IEEE Transactions on Network Science and Engineering* 9.4 (2022), pp. 2123–2138. DOI: [10.1109/TNSE.2022.3155359](https://doi.org/10.1109/TNSE.2022.3155359). URL: <https://hal.science/hal-03744409>.
- [4] 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](https://doi.org/10.1007/s10606-018-9323-3). URL: <https://hal.science/hal-01917249>.
- [5] 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](https://doi.org/10.1109/TPDS.2022.3172570). URL: <https://hal.inria.fr/hal-03772633>.

- [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 journals

- [7] J.-P. Eisenbarth, T. Cholez and O. Perrin. 'Ethereum's Peer-to-Peer Network Monitoring and Sybil Attack Prevention'. In: *Journal of Network and Systems Management*. Special Issue on Blockchains and Distributed Ledgers in Network and Service Management 30.4 (July 2022), p. 65. DOI: [10.1007/s10922-022-09676-2](https://doi.org/10.1007/s10922-022-09676-2). URL: <https://hal.inria.fr/hal-03777454>.
- [8] A. Ismaili-Alaoui, K. Baïna and K. Benali. 'IoDEP: Towards an IoT-Data Analysis and Event Processing Architecture for Business Process Incident Management'. In: *International journal of advanced computer science and applications (IJACSA)* 13.4 (2022). DOI: [10.14569/IJACSA.2022.01304104](https://doi.org/10.14569/IJACSA.2022.01304104). URL: <https://hal.inria.fr/hal-03660862>.
- [9] Q. Laporte-Chabasse, R. S. Stoica, M. Clausel, F. Charoy and G. Oster. 'Morpho-statistical description of networks through graph modelling and Bayesian inference'. In: *IEEE Transactions on Network Science and Engineering* 9.4 (2022), pp. 2123–2138. DOI: [10.1109/TNSE.2022.3155359](https://doi.org/10.1109/TNSE.2022.3155359). URL: <https://hal.science/hal-03744409>.
- [10] 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](https://doi.org/10.1109/TPDS.2022.3172570). URL: <https://hal.inria.fr/hal-03772633>.

International peer-reviewed conferences

- [11] C. Amiot, F. Charoy and J. Dinet. 'Trustworthy automation for large-scale collaboration: a proposed exploratory study'. In: *CEUR Workshop Proceedings*. AutomationXP22: Engaging with Automation, Workshop at CHI'22. Vol-3154. New Orleans (Louisiana), United States, 16th June 2022. URL: <https://hal.univ-lorraine.fr/hal-03697875>.
- [12] K. Benouaret, J. Agoun, I. Benouaret and F. Charoy. 'Call Limit-Based Composite Service Selection'. In: ICWS 2022 - IEEE International Conference on Web Services. Barcelona, Spain, 10th July 2022. URL: <https://hal.inria.fr/hal-03896182>.
- [13] A. Ismaili-Alaoui, K. Benali and K. Baïna. 'Traitement des événements complexes pour une gestion proactive des instances d'un processus métier'. In: *Actes du Congrès INFORSID 2022*. INFORSID 2022 - INFormatique des ORganisations et Systèmes d'Information et de Décision. Dijon, France, 31st May 2022, pp. 87–102. URL: <https://hal.science/hal-03892829>.
- [14] T. Lambert, S. Ibrahim, T. Jain and D. Guyon. 'Stragglers' Detection in Big Data Analytic Systems: The Impact of Heartbeat Arrival'. In: CCGrid 2022 - 22nd International Symposium on Cluster, Cloud and Internet Computing. Taormina, Italy: IEEE, 16th May 2022, pp. 747–751. DOI: [10.1109/CCGrid54584.2022.00084](https://doi.org/10.1109/CCGrid54584.2022.00084). URL: <https://hal.inria.fr/hal-03777656>.
- [15] P.-A. Rault, C.-L. Ignat and O. Perrin. 'Distributed Access Control for Collaborative Applications using CRDTs'. In: PaPoC 2022 - 9th Workshop on Principles and Practice of Consistency for Distributed Data. Rennes, France, 5th Apr. 2022. DOI: [10.1145/3517209.3524826](https://doi.org/10.1145/3517209.3524826). URL: <https://hal.inria.fr/hal-03584553>.
- [16] A. C. Zhou, R. Qiu, T. Lambert, T. Allard, S. Ibrahim and A. El Abbadi. 'PGPregel: An End-to-End System for Privacy-Preserving Graph Processing in Geo-Distributed Data Centers'. In: *Proceedings of the 13th Symposium on Cloud Computing*. SoCC '22: ACM Symposium on Cloud Computing. San Francisco California, United States: ACM, 7th Nov. 2022, pp. 386–402. DOI: [10.1145/3542929.3563474](https://doi.org/10.1145/3542929.3563474). URL: <https://hal.science/hal-03879423>.

11.3 Cited publications

- [17] L. André, S. Martin, G. Oster and C.-L. Ignat. ‘Supporting Adaptable Granularity of Changes for Massive-scale Collaborative Editing’. In: *CollaborateCom - 9th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing - 2013*. Austin, United States, Oct. 2013. URL: <http://hal.inria.fr/hal-00903813>.
- [18] 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/>.
- [19] 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.
- [20] Q. V. Dang and C.-L. 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>.
- [21] 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>.
- [22] 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](https://hal.inria.fr/hal-02307981). URL: <https://hal.inria.fr/hal-02307981>.
- [23] H.-L. 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>.
- [24] H.-L. 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](https://hal.inria.fr/hal-01883589). URL: <https://hal.inria.fr/hal-01883589>.
- [25] 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/>.
- [26] 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.
- [27] 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](http://hal.inria.fr/inria-00445975/en/). URL: <http://hal.inria.fr/inria-00445975/en/>.
- [28] 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>.
- [29] R. S. Stoica, A. Philippe, P. Gregori and J. Mateu. ‘ABC Shadow algorithm: a tool for statistical analysis of spatial patterns’. In: *Statistics and computing* 27.5 (2017), pp. 1225–1238.
- [30] 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>.