IN PARTNERSHIP WITH:
CNRS

Université de Lorraine

Activity Report 2017

Project-Team COAST

Web Scale Trustworthy Collaborative Service Systems

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

RESEARCH CENTER
Nancy - Grand Est

THEME
Distributed Systems and middleware

Table of contents

Project-Team COAST

Creation of the Team: 2014 July 01, updated into Project-Team: 2015 July 01

Keywords:

Computer Science and Digital Science:

A1.1.6. - Cloud
A1.1.7. - Peer to peer
A1.3. - Distributed Systems
A2.5.1. - Software Architecture & Design
A2.6.2. - Middleware
A3.1.3. - Distributed data
A3.1.5. - Control access, privacy
A5.1.1. - Engineering of interactive systems
A5.1.2. - Evaluation of interactive systems
A7.1. - Algorithms

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
B9.1.1. - E-learning, MOOC
B9.5.1. - Psychology
B9.6. - Reproducibility
B9.8. - Privacy

1. Personnel

Research Scientist

Claudia-Lavinia Ignat [Inria, Researcher]

Faculty Members

François Charoy [Team leader, Univ de Lorraine, Professor, HDR]
Khalid Benali [Univ de Lorraine, Associate Professor, HDR]
Gérôme Canals [Univ de Lorraine, Associate Professor]
Claude Godart [Univ de Lorraine, Professor]
Gérald Oster [Univ de Lorraine, Associate Professor]
Olivier Perrin [Univ de Lorraine, Professor, HDR]
Pascal Urso [Univ de Lorraine, Associate Professor]
Samir Youcef [Univ de Lorraine, Associate Professor]

Post-Doctoral Fellows

Mohammed Riyadh Abdmeziem [Univ de Lorraine]
Kahina Bessai Ép. Youcef [Univ de Lorraine, until Oct 2017]

PhD Students

Amina Ahmed Nacer [Univ de Lorraine, until Aug 2017]

Anis Ahmed Nacer [Inria, from Mar 2017]
Quang Vinh Dang [Univ de Lorraine, until Nov 2017]
Victorien Elvinger [Univ de Lorraine]
Abir Ismaili-Alaoui [Univ de Lorraine]
Quentin Laporte Chabasse [Univ de Lorraine]
Béatrice Linot [Univ de Lorraine]
Jordi Martori Adrian [Inria, until May 2017]
Hoai Le Nguyen [Univ de Lorraine]
Hoang Long Nguyen [Univ de Lorraine]
Matthieu Nicolas [Univ de Lorraine, from Oct 2017]
Guillaume Rosinosky [Univ de Lorraine, From Sep 2017]

Technical staff

Jean-Philippe Eisenbarth [Univ de Lorraine]
Phillippe Kalitine [Univ de Lorraine]

Interns

Laurent Bidet [Univ de Lorraine, from Jun 2017 until Aug 2017]
Camille Menou [Univ de Lorraine, from Jun 2017 until Aug 2017]
Quentin Tardivon [Univ de Lorraine, from Jun 2017 until Aug 2017]

Administrative Assistants

Sophie Drouot [Inria]
Delphine Hubert [Univ de Lorraine]
Martine Kuhlmann [CNRS]
Sylvie Musilli [Univ de Lorraine]

Visiting Scientist

Valerie Shalin [Invited Professor, from May 2017 until Jul 2017]

2. Overall Objectives

2.1. Overall Objectives

The advent of the Cloud, of smart mobile devices and of service-based architecture has opened a field of possibilities 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 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. Thus, clients consume applications through service API (web services), available on delivery platforms, called stores or markets. This approach of the distribution of software outstrips the traditional software distribution channels, in both 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 likely insufficient. We conduct a dedicated research effort to tackle the problems arising from the evolution of contemporary technologies and of those we can anticipate.

[1] See http://blog.programmableweb.com/2011/09/16/open-api-growth-a-visualization/

For this purpose, 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. Although this research has a long history referring back to [19], new challenges arise regarding needs that are occurring with the acculturation of users to collaboration like the number of participants to a collaboration (a crowd), sharing among different organisations and the nature of documents that are shared and produced. The problem here 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 API oblige us to reconsider the problem [18]. Our goal here is, taking into account this evolution, like the advent of the Cloud, the availability at a large scale of public API based on the REST [2] architectural style, 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 in the two first topics, our main research direction aims at providing support to build **trustworthy collaborative applications** based on the knowledge that we can gather from the underlying algorithms, from the composition of services and from the quality of services that can be deduced and monitored. The complexity of the context in which applications are executed does not allow 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. It is very surprising to see to what extent people rely today on services with very little knowledge about the amount of confidence they put in these services. As soon as these services are based on composition of other unknown services, it becomes very difficult to understand the consequences of the failure of a component of the composition for instance.

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 organisations. We see there that we aim at combining our results and expertise to achieve a new leap forward toward the understanding and the mastering of methods and techniques that allow the engineering and the use of large scale collaborative systems.

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 research area. 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 towards a decentralized collaboration where users have full control over their data that they can store 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.

[2] representational state transfer

Service oriented Computing [21] 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 and more generally to the emerging discipline of Service Science has been and is still to focus on the issue of the efficient and flexible construction of reliable and secure high level services 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 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) [25] 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 allows improving 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 [23] 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:

- the operational transformation (OT) algorithms [19]
- the algorithms based on commutative replicated data types (CRDT) [22].

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 parametrized 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 is a new class of algorithms initiated by WOOT [20] a 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 studies.

Much efforts have been devoted in the past years to establish standard business process models founded on well grounded theories (e.g. Petri Nets) that meet the needs of both business analysts but also of software engineers and software integrators. This has lead to heated debate in the BPM community as the two points of view are very difficult to reconcile. On one side, the 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 semantic in order to be able transform them into executable artefacts. Part of our work has been an attempt to reconcile these point of views. It resulted in the development of the Bonita Business process management system and more recently on our work in crisis management where the same people are designing, executing and monitoring the process as it executes. But more generally, and at a larger scale, we have been considering the problem of processes spanning the barriers of organisations and thus more general problem of service composition as a way to coordinate inter organisational construction of applications providing value based on the composition of lower level services [17].

3.5. Service Composition

We are considering processes as pieces of software whose execution traverse the boundaries of organisations. This is especially true with service oriented computing where processes compose services produced by many organisations. We tackle this problem from very different perspectives, trying to find the best compromise between the need for privacy of internal processes from organisations and the necessity to publicize large part of them, proposing to distribute the execution and the orchestration of processes among the organisations themselves, and attempting to ensure non-functional properties in this distributed setting [16].

Non-functional aspects of service composition relate to all the properties and service agreements that one wants to ensure and that are orthogonal to the actual business but that are important when a service is selected and integrated in a composition. This includes transactional context, security, privacy, and quality of service in general. Defining and orchestrating services on a large scale while providing the stakeholders with some strong guarantees on their execution is a first class problem for us. For a long time, we have proposed models and solutions to ensure that some properties (e.g. transactional properties) were guaranteed on process execution, either through design or through the definition of some protocols. Our work has also been extended to the problems of security, privacy and service level agreement among partners.

We extended some of our previous work around authorization policies for Enterprise Social Networks, and we propose two directions: a first one is dedicated to formal verification using PlusCal-2[11], while the other one is a formal approach for the Verification of AWS IAM Access Control Policies..

Recently, we started a work on service composition for software architects where services are coming from different providers with different plans (capacity, degree of resilience,...). The objective is to help the architects to select the most accurate services (wrt. to thier requirements, both functional and non functional) and plans for building their software. We also compute the properties that can be guarantee for the composition of these services.

4. New Software and Platforms

4.1. BeGoood

FUNCTIONAL DESCRIPTION: BeGoood is a generic system for managing non-regression tests on knowledge bases. BeGoood allows to define test plans in order to monitor the evolution of knowledge-bases. Any system answering queries by providing results in the form of set of strings can be tested with BeGoood. BeGoood has been developed following a REST architecture and is independent of any application domain. BeGoood is a part of the Kolflow infrastructure.

- Participant: Gérôme Canals
- Contact: Gérôme Canals
- URL: https://github.com/kolflow/begoood

4.2. MUTE

Multi-User Text Editor
FUNCTIONAL DESCRIPTION: MUTE (Multi-User Text Editor) is a web-based text editing tool that allows to edit documents collaboratively in real-time. It implements our recent work on collaborative editing algorithms and more specifically the LogootSplit+ approach. Compared to existing web-based collaborative text editing tool this editor does not require a powerful central server since the server is not performing any computation and acts as a simple broadcast server. Our editor offers support for working offline while still being able to reconnect at a later time.

- Participants: Claudia-Lavinia Ignat, François Charoy, Gérald Oster and Luc André
- Contact: Gérald Oster
- URL: https://github.com/coast-team/mute-demo/

4.3. Replication Benchmarker

FUNCTIONAL DESCRIPTION: The Replication Benchmarker is a performance evaluation framework for optimistic replication mechanisms used in collaborative applications. It contains a library of implementation of several CRDT (Commutative Replicated Data Type) and OT (Operational Transformation) algorithms for different data types: text, set, trees. The framework is able to evaluate the performance of comparable algorithms on different corpus of events traces. These events traces can be produced randomly according to different parameters, can be extracted from real real-time editing session that have been recorded, or can be automatically extracted from distributed version control repositories such as the one produced with Git. Performances of the algorithms are measured in term of execution time, memory footprint and merge result quality (compared to manual merge history stored in git repositories).

- Participants: Gérald Oster, Mehdi Ahmed-Nacer and Pascal Urso
- Contact: Pascal Urso
- URL: https://github.com/score-team/replication-benchmarker/

4.4. Rivage

Real-tIme Vector grAphic Group Editor
FUNCTIONAL DESCRIPTION: Rivage is a real-time collaborative graphical editor. Several users can edit at the same time and in real-time a graphical document, user changes being immediately seen by the other users. The editor relies on a peer-to-peer architecture where users can join and leave the group at any time. Each user has a copy of the shared document and user changes on the document copies are merged in real-time by using a CRDT (Commutative Replicated Data Type) algorithm.

- Participant: Claudia-Lavinia Ignat
- Contact: Claudia-Lavinia Ignat
- URL: https://github.com/stephanemartin/rivage/

5. New Results

5.1. Design and Analysis of Collaborative Editing Approaches

Participants: Matthieu Nicolas, Victorien Elvinger, Hoai Le Nguyen, Quentin Laporte Chabasse, Claudia-Lavinia Ignat [contact], Gérald Oster, François Charoy, Olivier Perrin.

Since the Web 2.0 era, the Internet is a huge content editing place on which users collaborate. Such shared content can be edited by thousands of people. However, current consistency maintenance algorithms seem not to be adapted to massive collaborative updating involving large amount of contributors and a high velocity of changes. This year we designed new optimistic replication algorithms for maintaining consistency for complex data such as wikis. We also designed a peer-to-peer web-based real-time collaborative editor relying on our proposed algorithms as well as a mechanism that balances awareness and disturbance in this kind of systems. We also started to study collaborative editing user behavior.

Wikis are one of the most important tools of Web 2.0 allowing users to easily edit shared data. However, wikis offer limited support for merging concurrent contributions on the same pages. Users have to manually merge concurrent changes and there is no support for an automatic merging. Real-time collaborative editing reduces the number of conflicts as the time frame for concurrent work is very short. We proposed extending wiki systems with real-time collaboration and designed an automatic merging solution adapted for rich content wikis [2]. Our merging solution is based on an operational transformation approach for which we defined operations with high-level semantic capturing user intentions when editing wiki content such as move, merge and split. Our solution is the first one that deals with high level operations, existing approaches being limited to operations of insert, delete and update on textual documents.

Existing real-time collaborative editors rely on a central authority that stores user data which is a perceived privacy threat.We designed MUTE [8], a peer-to-peer web-based real-time collaborative editor that eliminates the disadvantages of central authority based systems. Users share their data with the collaborators they trust without having to store their data on a central place. MUTE features high scalability and supports offline and ad-hoc collaboration. MUTE relies on LogootSplit, a CRDT-based consistency maintenance algorithm for strings [15]. MUTE collaborative editor will be integrated in the virtual desktop of OpenPaaS::NG project [8].

When people work collaboratively on a shared document, they have two contradictory requirements on their editors that may affect the efficiency of their work. On the one hand, users would like to be aware of other users work on a particular part of the document. On the other hand, users would like to focus their attention on their own current work, with as little disturbance from the concurrent activities as possible. We designed a mechanism that lets users handle a balance between disturbance and awareness of concurrent updates [10]. Users can define focus regions and concentrate on the work in these regions without being disturbed by work of other users. Occasionally, users can preview concurrent updates and select a number of these updates to be integrated into the local copy.

We are interested in analysing user behavior during collaborative editing. This year we studied concurrency and conflicts in asynchronous collaboration [7]. We chose to study collaboration traces of distributed version control systems such as Git. We analysed Git repositories of four projects: Rails, IkiWiki, Samba and Linux Kernel. We analyzed the collaboration process of these projects at specific periods revealing how changes integration evolves during project development. We also analyzed how often users decide to rollback to previous document version when the integration process results in conflict. Finally, we studied the mechanism adopted by Git to consider changes made on two continuous lines as conflicting.

5.2. Trust-based Collaboration

Participants: Quang Vinh Dang, Claudia-Lavinia Ignat, Francois Charoy, Olivier Perrin, Mohammed Riyadh Abdmeziem, Hoang Long Nguyen.

Trust between users is an important factor for the success of a collaboration. Users might want to collaborate only with those users they trust. We are interested in assessing users trust according to their behaviour during collaboration in a large scale environment. In order to compute the trust score of users according to their contributions during a collaborative editing task, we need to evaluate the quality of the content of a document that has been written collaboratively. We investigated how to automatically assess the quality of Wikipedia articles in order to provide guidance for both authors and readers of Wikipedia. Most existing approaches for quality classification of Wikipedia articles rely on traditional machine learning with manual feature engineering, which requires a lot of expertise and effort and is language dependent. We proposed an

approach that addresses the trade-off between accuracy, time complexity and language independence for the prediction models [5]. Our approach relying on Recurrent Neural Networks (RNN) eliminates disadvantages of feature engineering, i.e. it learns directly from raw data without human intervention and is language-neutral. Experimental results on English, French and Russian Wikipedia datasets show that our approach outperforms state-of-the-art solutions.

Rating prediction is a key task of e-commerce recommendation mechanisms. Recent studies in social recommendation enhance the performance of rating predictors by taking advantage of user relationships. However, these prediction approaches mostly rely on user personal information which is a privacy threat. We proposed dTrust [6], a simple social recommendation approach that avoids using user personal information. It relies uniquely on the topology of an anonymized trust-user-item network that combines user trust relations with user rating scores for items. This topology is fed into a deep feed-forward neural network. Experiments on real-world data sets showed that dTrust outperforms state-of-the-art in terms of Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) scores for both warm-start and cold-start problems.

One dimension of our work is dedicated to ensure consistency of the key server. We design Trusternity, which is a secure, scalable auditing mechanism using a blockchain to replace the gossiping mechanism of transparent log system. We have implemented Trusternity as a proof-of-concept, and we have led some evaluation about the detection of malicious behavior on the blockchain network.

Securing P2P collaborative system remains a critical issue for its widespread adoption. One of our goals is to ensure that communication between collaborating partner is secure from end to end. We need to encrypt exchange of operations among partners. For that we propose to rely on group keys management. One of the issue is that the composition of the partnership can change and this require to change the group key. Since we don't want a central server to manage keys, that would break the p2p nature of our approach we need to propose group key management protocols that are resilient to change in groups, even in group of large size. [3]

5.3. Cloud Provisioning for Elastic BPM

Participants: François Charoy, Samir Youcef, Guillaume Rosinosky.

Cloud computing provider do not help consumer to use optimally the available resources. For this, several approaches have been proposed [24] that take benefit from the elasticity of the Cloud, starting and stoppping virtual machines on demand. They suffer from several shortcomings. Often they consider only one objective, the reduction of the cost, or a level of quality of service. We proposed to optimize two conflicting objectives, the number of migrations of tenants that is helpful to reach the optimal cost and the cost incurred considering a set of resources. Our approach allows to take into account the multi-tenancy property and the Cloud computing elasticity, and is efficient as shown by an extensive experimentation based on real data from Bonita BPM customers [9].

5.4. Risk Management for the Deployment of a Business Process in a Multi-Cloud Context

Participants: Amina Ahmed-Nacer, Claude Godart, Samir Youcef.

The lack of trust in cloud organizations is often seen as braking forces to SaaS developments. This work proposes an approach which supports a trust model and a business process model in order to allow the orchestration of trusted business process components in the cloud.

The contribution is threefold and consists in a method, a model and a framework. The method categorizes techniques to transform an existing business process into a risk-aware process model that takes into account security risks related to cloud environments. These techniques are partially described in the form of constraints to automatically support process transformation. The model formalizes the relations and the responsibilities between the different actors of the cloud. This allows to identify the different information required to assess and quantify security risks in cloud environments.

The framework is a comprehensive approach that decomposes a business process into fragments that can automatically be deployed on multiple clouds. The framework also integrates a selection algorithm that combines the security information of cloud offers and of the process with other quality of service criteria to generate an optimized configuration. It is implemented in a tool to assess cloud providers and decompose processes.

Rooted in past years work, we are contributing this year at the methodological and framework levels in two directions:

- At the methodological level, while our risk computing model rested previously only on data provided by cloud providers (provider-side risk model), we are developing a risk model integrating client-side knowledge (client-side risk model) [4].
- Additionally are developing a simulation tool for supporting designer decision with the ability to balance risk with cost when selecting the best cloud configuration.

5.5. Scheduling and Resource Allocation in Business Processes

Participants: Khalid Benali, Abir Ismaili-Alaoui.

Business Process Management (BPM) is concerned with continuously enhancing business processes by adapting a systematic approach that enables companies to increase the performance of their existing business processes and achieve their business goals. Business processes are generaly considered as blind and stateless, which mean that in each business process execution results from past process instances are not taken into consideration.

The main objective of our current research is to exploit the data generated from previous instances in order to enhance business processes in regards with several aspects, such as improvement of process business logical correctness, optimization of business process modeling issues, or improvemment of resource allocation and scheduling procedure in order to particularly optimize costs and time (among other factors).

We focus currently on this last aspect, i.e. scheduling and resource allocation in business processes. Business Processes may contain automatic tasks and non automatic tasks, so managing resources depends on the type of those resources (human or machine) In this context, our work use machine learning techniques to analyze data generated from previous business process execution to improve business process scheduling. This step ensure the assignment of the most critical business process instance task to a qualified (and may be costly) human resource while minimizing global execution costs through assignement of “dummy” tasks to machine agents.

6. Bilateral Contracts and Grants with Industry

6.1. Bilateral Contracts with Industry

6.1.1. Industrial funding Groupe Open (2016–2019)

Groupe Open is a leading french company specialised in digital services and operations. The goal of the project is to propose an industrial composition model for APIs that takes into account the new constraints imposed by this new way to distribute and operate software. It will be based on a formal API contract along with trust and reputation attributes in order to allow consumers to anticipate risks regarding the quality and the safety of services. A PhD student is under recruitment for this project. Coast funding : 237,000 €

6.2. Bilateral Grants with Industry

6.2.1. CIFRE Grant with Bonitasoft

Participants: François Charoy, Samir Youcef, Guillaume Rosinosky.

Bonitasoft is a leading software company in the domain of open source Business Process Management Systems. The objective of this grant is to help Bonitasoft to support effective elastic BPM operation in the Cloud by leveraging the business knowledge, the process models and the execution history of process instances and correlate them with cloud resource consumption. Guillaume Rosinosky has been recruited as a PhD Student to work on this project. We will define models that will be validated based on a detailed analysis of existing use cases that we have started to collect from Bonitasoft and its clients.

7. Partnerships and Cooperations

7.1. Regional Initiatives

7.1.1. Region Lorraine TV Paint (2016–2017)

Participants: Claudia-Lavinia Ignat [contact], Gérald Oster, Quang Vinh Dang, Matthieu Nicolas.

Partners: TVPaint Development, Inria COAST project-team

Website: https://www.tvpaint.com/

This is a project in collaboration with TVPaint Development financed by Region Lorraine. The goal is to contribute to the creation of a collaborative system dedicated to animation movies, that allows to manipulate high quantities of digital artifacts in a collaborative way.

7.1.2. Region Grand Est TV Paint (2017–2019)

Participants: Claudia-Lavinia Ignat [contact], Gérald Oster.

Partners: TVPaint Development, Inria COAST project-team

Website: https://www.tvpaint.com/

This is a follow-up project in collaboration with TVPaint Development financed by Region Grand Est.

The goal is to contribute to the creation of a collaborative system dedicated to manage the production of animated movies. This system has to manipulate a large amount of data in a safe and secure manner. Based on the previously proposed architecture and prototype, this project intends to design and implements a commercial product. In the framework of this project, we bring our expertise in data management, business process management, distributed systems and collaborative systems.

Coast funding : 81,600 €

7.2. National Initiatives

7.2.1. OpenPaas NG (2015–2019)

Participants: Claudia-Lavinia Ignat, François Charoy [contact], Gérald Oster, Olivier Perrin, Jean-Philippe Eisenbarth, Phillippe Kalitine, Matthieu Nicolas, Mohammed Riyadh Abdmeziem, Kahina Bessai, Victorien Elvinger, Quentin Laporte Chabasse, Hoai Le Nguyen, Hoang Long Nguyen.

Partners: Linagora, XWiki SAS, Nexedi, COAST project-team (Université de Lorraine, LORIA), DaScim team (LIX).

Website: http://www.open-paas.org/

This project is financed by BpiFrance and involves French industrial leaders in open-source software development (Linagora, Nexedi, XWiki) and academic partners in collaborative work (COAST team) and recommender systems (DaScim team, LIX). The goal of the project is to develop next generation cloud enabled virtual desktop based on an Enterprise Social Network to provide advanced collaborative and recommendation services. COAST team is responsible of the work package dedicated to the design of the peer-to-peer collaborative middleware. In this context, we bring our expertise on data replication for collaborative data in peer-to-peer environments and on trust and access control and identity management in distributed collaborative information systems.

7.3. International Initiatives

7.3.1. Inria Associate Teams Not Involved in an Inria International Labs

7.3.1.1. USCOAST2

Title: User Studies on Trustworthy Collaborative Systems

International Partner (Institution - Laboratory - Researcher):

Wright State University (United States) - Department of Psychology, Knoesis - Valerie Shalin

Start year: 2016

See also: http://uscoast.loria.fr

The proposed project addresses the perception of trust by users, the appropriateness of a trust-based security approach and the role of trust metrics in the management of distributed work. The main challenge of this project is how to measure trust based on user behaviour and to verify by means of experimental studies with users that the trust-based mechanism is acceptable by users. We plan to apply this trust-based mechanism for two types of applications. The first one is collaborative editing where user trust will be computed based on the quality of user contributions for a document or project. The second type of application is in the management of work over a large group of people in order to conduct efficient, high-yield, high-density real time crowdsourcing activities.

Partners of USCOAST2 project have complementary expertise. COAST provides expertise in collaborative methods, systems and related technologies. Coast will propose algorithms that track and manipulate trust metrics. Kno.e.sis provides expertise on the analysis of human work-related behavior, including methods of data collection and data analysis, as well as a theoretical foundation for the evaluation of human performance. Knoesis will analyse trust from a psychological phenomenon point of view.

7.4. International Research Visitors

7.4.1. Visits of International Scientists

Valerie Shalin from Wright State University spent one month and a half (May-July 2017) in our team as part of the USCoast2 Inria associated team.

Weihai Yu from Arctic University of Norway spent two weeks in March 2017 in the team as invited professor.

7.4.2. Visits to International Teams

7.4.2.1. Research Stays Abroad

- Béatrice Linot spent 3 months at Wright State University as part of our collaboration with Dr Valerie Shalin and Prof. Amit Sheth, funded by her LUE PhD grant.

8. Dissemination

8.1. Promoting Scientific Activities

8.1.1. Scientific Events Organisation

8.1.1.1. Member of the Organizing Committees

- Claudia-Lavinia Ignat was member of the organisation committee of The Fifteenth International Workshop on Collaborative Editing Systems in conjunction with CSCW 2017

8.1.2. Scientific Events Selection

8.1.2.1. Member of the conference program committees

- Claude Godart was PC member of the conference program committee of BPMDS (Business Process Modeling, Development and Support), EDOC (The enterprise computing conference), ICSOC (International Conference on Services Oriented Computing), IEEE CLOUD Computing, ICWS (IEEE International Conference on Web Services), SCC (IEEE International Conference on Services Computing), S2 ICIOT (S2 International Conference on Internet of Things), IEEE/WIC/ACM WI (Web Intelligence conference), WISE (Web Information Systems Engineering) conferences.
- Claudia-Lavinia Ignat was PC member of CDVE (International Conference on Cooperative Design, Visualization and Engineering) 2017 and The Fifteenth International Workshop on Collaborative Editing Systems in conjunction with CSCW 2017
- Olivier Perrin was PC Member of ICSOC 2017, CoopIS 2017, ATC 2017, MoLS 2017 and some workshops.
- François Charoy was PC Member of ICEBE (International Conference on Business Engineering) 2017, CTS 2017 (International Symposium on Collaborative Technologies and Systems), DG.O (International Conference on Digital Government Research) 2017, IEEE WETICE 2017, ICSOC 2017, IEEE International Conference on Business Information Systems, ISCRAM 2017 and of several workshops.
- Gérald Oster was a PC member of CoopIS (International Conference on Cooperative Information Systems) 2017, CSCW (21st ACM Conference on Computer-Supported Cooperative Work and Social Computing) 2018 (Online-first).
- Khalid Benali was PC Member of WorldCIST'17 (World Conference on Information Systems and Technologies), I3E 2017 (IFIP Conference on e-Business, e-Services and e-Society), ICICS 2017 (8th International Conference on Information and Communication Systems), INFORSID'2017, CloudTech'17 (Third International Conference of Cloud Computing Technologies and Applications), MEDES 2017 (9th International Conference on Management of Digital EcoSystems),

8.1.3. Journal

8.1.3.1. Member of the editorial boards

- Claude Godart is member of the editorial board of IEEE Transactions on Service Computing, International Journal of Services Computing, and member of the review board of the International Journal of Next Generation Computing.
- Claudia-Lavinia Ignat is member of the editorial board of Journal of CSCW (Computer Supported Cooperative Work).
- François Charoy is member of the editorial board of Service Oriented Computing and Applications Journal (Springer).

8.1.3.2. Reviewer - Reviewing activities

- Olivier Perrin reviewed papers for IEEE Transactions on Services Computing journal, IEEE Transactions on Parallel and Distributed Systems and Journal of Systems and Software.
- Claudia-Lavinia Ignat reviewed papers for Transactions on Parallel and Distributed Systems, Transactions on Internet Technology, Transactions on Interactive Intelligent Systems and CSCW 2018.
- Gérald Oster reviewed papers for Computer Supported Cooperative Work Journal and Concurrency and Computation: Practice and Experience journal.

8.1.4. Scientific Expertise

François Charoy was member of the HCERES committee for the CRI lab of Paris 1 Sorbonne

8.2. Teaching - Supervision - Juries

8.2.1. 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 of departments in the University. Most of the PhD Students have also teaching duties in the same institutions. As a whole, the COAST team accounts for more than 2500 hours of teaching. Members of the COAST team are also deeply involved in the pedagogical and administrative life of their departments.

- Claude Godart is responsible for the Computer Science department of the engineering school ESSTIN.
- Khalid Benali is responsible for the professional Master degree speciality "Distributed Information Systems" of MIAGE and of its international branch in Morocco.
- François Charoy is responsible of the Software Engineering specialisation at the TELECOM Nancy Engineering School of University of Lorraine.
- Pascal Urso was responsible for the "Security, Services, Systems and Network" track of the master degree in computer science at University of Lorraine from September 2013 to June 2017.

8.2.2. Supervision

- PhD (in progress): Quang Vinh Dang, Trust-based large scale collaboration, started in 10/2014, Claudia-lavinia Ignat and François Charoy
- PhD (in progress): Hoai Le Nguyen, Study of group performance and behavior in collaborative editing, started in 9/2015, Claudia-Lavinia Ignat and François Charoy
- PhD (in progress): Hoang Long Nguyen, A Trust Based Authorization Model and Framework for the Cloud, started in 11/2015, Claudia-Lavinia Ignat and Olivier Perrin
- PhD (in progress): Victorien Elvinger, Secured Replication for Peer-to-Peer Collaborative Infrastructures, started in 10/2015, François Charoy and Gérald Oster
- PhD (defended): Jordi Martori i Adrian, Data constraints for large-scale collaboration, started in 10/2013, defense in 5/2017, François Charoy and Pascal Urso
- PhD (in progress): Guillaume Rosinoski, Elastic BPM and the Cloud, started in 10/2014, François Charoy and Samir Youssef
- PhD (in progress): Quentin Laporte-Chabasse, Federation of Organisations over Peer to Peer Collaborative Network, started in 10/2016, François Charoy and Gérald Oster
- PhD (in progress): Béatrice Linot, Trust in cooperative systems, Jérome Dinet et François Charoy, started 11/2016
- PhD (in progress): Anis Ahmed Nacer, Safe Service Composition, Olivier Perrin and François Charoy, started 3/2017
- PhD (in progress): Matthieu Nicolas, Optimisation of Replication Algorithms, Olivier Perrin and Gérald Oster, started 10/2017

8.2.3. Juries

- Claudia-Lavinia Ignat was member of CR recruitement jury at Inria Nancy-Grand Est

COAST members were members of the following PhD and HdR defense committees:

- Hala Skaf Molli, HdR, Université de Nantes, October 2017 (François Charoy)
- Mourad Bouneffa, HdR, Université du Littoral Côte d'Opale (Claude Godart)
- Elian Aubry, PhD, Université de Lorraine, December 2017 (François Charoy)
- Mohsen Sayed, PhD, Université de Lorraine, July 2017 (François Charoy)
- Fatma Slaimi, PhD, Université d'Aix Marseille (Claude Godart)
- Cheick Salmi, PhD, Ecole Nationale Supérieure de Mécanique et d'Aéronautique (Claude Godart)
- Emma Hachicha Belghith, PhD, Université Paris-Saclay (Claude Godart)
- Jihane Lakhrouit, PhD, ENSIAS Rabat (Claude Godart)
- Hafida Naim, PhD, Université d'Aix Marseille (Claude Godart)

8.3. Popularization

- In June 2017 Claudia-Lavinia Ignat organised the meeting Research@Inria at Inria Nancy-Grand Est for presenting to internship students, PhD students and postdocs various Inria programs and the main research topics at Inria with a focus on activities of research teams at Inria Nancy-Grand Est. She also presented the main activities of a researcher and briefly described her research work.

8.4. Institutional commitment

- Claudia-Lavinia Ignat is in charge of European affairs for Inria Nancy Grand-Est. She is the Delegate of International Relations for Inria Nancy-Grand Est and member of COST-GTRI commission. She is member of the Inria Nancy-Grand Est COMIPERS committee. She is member of Inria CAP Chercheurs commission. She is responsible with the activity kindergarten at AGOS Inria Nancy-Grand Est.

9. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journals

[1] M. R. ABDMEZIEM, D. TANDJAOUI, I. ROMDHANI. *Lightweighted and energy-aware MIKEY-Ticket for e-health applications in the context of internet of things*, in "International Journal of Sensor Networks", 2017, forthcoming, https://hal.archives-ouvertes.fr/hal-01589967

[2] C.-L. IGNAT, L. ANDRÉ, G. OSTER. *Enhancing rich content wikis with real-time collaboration*, in "Concurrency and Computation: Practice and Experience", March 2017 [*DOI :* 10.1002/CPE], https://hal.inria.fr/hal-01404024

International Conferences with Proceedings

[3] M. R. ABDMEZIEM, F. CHAROY. *Fault-tolerant and Scalable Key Management Protocol for IoT-based Collaborative Groups*, in "SecureComm 2017 : 13th EAI International Conference on Security and Privacy in Communication Networks", Niagara falls, Canada, October 2017, pp. 1-20, https://hal.inria.fr/hal-01588490

[4] A. AHMED NACER, C. GODART, S. YOUCEF, A. TARI. *A Metric for Evaluating the Privacy Level of a Business Process Logic in a Multi-Cloud Deployment*, in "IEEE 21st International Enterprise Distributed Object Computing Conference (EDOC)", Quebec, Canada, October 2017, https://hal.archives-ouvertes.fr/hal-01620023

[5] Q.-V. DANG, C.-L. IGNAT. *An end-to-end learning solution for assessing the quality of Wikipedia articles*, in "OpenSym 2017 - International Symposium on Open Collaboration", Galway, Ireland, August 2017 [*DOI :* 10.1145/3125433.3125448], https://hal.inria.fr/hal-01559693

[6] Q.-V. DANG, C.-L. IGNAT. *dTrust: a simple deep learning approach for social recommendation*, in "The 3rd IEEE International Conference on Collaboration and Internet Computing (CIC-17)", San Jose, United States, October 2017, https://hal.inria.fr/hal-01578316

[7] H. LE NGUYEN, C.-L. IGNAT. *Parallelism and conflicting changes in Git version control systems*, in "IWCES'17 - The Fifteenth International Workshop on Collaborative Editing Systems", Portland, Oregon, United States, February 2017, https://hal.inria.fr/hal-01588482

[8] M. NICOLAS, V. ELVINGER, G. OSTER, C.-L. IGNAT, F. CHAROY. *MUTE: A Peer-to-Peer Web-based Real-time Collaborative Editor*, in "ECSCW 2017 - 15th European Conference on Computer-Supported Cooperative Work", Sheffield, United Kingdom, Proceedings of 15th European Conference on Computer-Supported Cooperative Work - Panels, Posters and Demos, EUSSET, August 2017, vol. 1, n^{o} 3, pp. 1-4 [*DOI :* 10.18420/ECSCW2017_P5], https://hal.inria.fr/hal-01655438

[9] G. ROSINOSKY, S. YOUCEF, F. CHAROY. *Efficient migration-aware algorithms for elastic BPMaaS*, in "15th International Conference on Business Process Management", Barcelone, Spain, J. CARMONA, G. ENGELS, A. KUMAR (editors), Business Process Management, Springer, September 2017, vol. LNCS, Volume 10445, pp. 147-163, https://hal.archives-ouvertes.fr/hal-01501705

[10] W. YU, G. OSTER, C.-L. IGNAT. *Handling Disturbance and Awareness of Concurrent Updates in a Collaborative Editor*, in "CDVE 2017 - 14th International Conference on Cooperative Design, Visualization, and Engineering", Mallorca, Spain, Y. LUO (editor), LNCS - Lecture Notes in Computer Science, Springer, September 2017, vol. 10451, https://hal.inria.fr/hal-01652656

[11] E. ZAHOOR, S. AKTAR, O. PERRIN. *Formal Verification of Authorization Policies for Enterprise Social Networks using PlusCal-2*, in "CollaborateCom 2017 - 13th EAI International Conference on Collaborative Computing: Networking, Applications and Worksharing", Edimburg, United Kingdom, December 2017, https://hal.inria.fr/hal-01657116

Books or Proceedings Editing

[12] F. CHAROY, K. DRIRA, H. WANG, Q. YU, Y. WANG, Y. YAN, J. MENDLING, M. MOHAMED, Z. WANG, S. BHIRI (editors). *Service-Oriented Computing – ICSOC 2016 Workshops*, Springer, 2017, 222 p. [*DOI :* 10.1007/978-3-319-68136-8], https://hal.inria.fr/hal-01651943

[13] M. S. MACFADDEN, A. NG, C.-L. IGNAT, N. GU, C. SUN (editors). *The Fifteenth International Workshop on Collaborative Editing Systems*, ACM, Portland, United States, 2017 [*DOI :* 10.1145/3022198.3022653], https://hal.inria.fr/hal-01652759

Other Publications

[14] G. ROSINOSKY, S. YOUCEF, F. CHAROY. *Optimisation of business process tenant distribution in the Cloud with a genetic algorithm*, June 2017, working paper or preprint, https://hal.archives-ouvertes.fr/hal-01558202

References in notes

[15] L. ANDRÉ, S. MARTIN, G. OSTER, 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, October 2013, http://hal.inria.fr/hal-00903813

[16] S. BHIRI, W. GAALOUL, C. GODART. *Mining and Improving Composite Web Services Recovery Mechanisms*, in "International Journal of Web Services Research", 03 2008, vol. 3, http://hal.inria.fr/inria-00438424/en/

[17] S. BHIRI, O. PERRIN, W. GAALOUL, C. GODART. *An Object-Oriented Metamodel For Inter-Enterprises Cooperative Processes Based on Web Services*, in "Journal of Integrated Design and Process Science", 2004, vol. 8, pp. 37–55, http://hal.inria.fr/inria-00099953/en/

[18] F. CASATI. *Promises and Failures of Research in Dynamic Service Composition*, in "Seminal Contributions to Information Systems Engineering", J. BUBENKO, J. KROGSTIE, O. PASTOR, B. PERNICI, C. ROLLAND, A. SØLVBERG (editors), Springer Berlin Heidelberg, 2013, pp. 235-239, http://dx.doi.org/10.1007/978-3-642-36926-1_18

[19] C. A. ELLIS, 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, http://doi.acm.org/10.1145/67544.66963

[20] G. OSTER, P. URSO, P. MOLLI, A. IMINE. *Data Consistency for P2P Collaborative Editing*, in "ACM Conference on Computer-Supported Cooperative Work - CSCW 2006", Banff, Alberta, Canada, ACM Press, 11 2006, pp. 259 - 268, http://hal.inria.fr/inria-00108523/en/

[21] M. P. PAPAZOGLOU, P. TRAVERSO, S. DUSTDAR, F. LEYMANN. *Service-Oriented Computing: State of the Art and Research Challenges*, in "Computer", 2007, vol. 40, pp. 38-45

[22] N. PREGUIÇA, J. M. MARQUÈS, M. SHAPIRO, M. LETIA. *A commutative replicated data type for cooperative editing*, 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/

[23] Y. SAITO, M. SHAPIRO. *Optimistic Replication*, in "Computing Surveys", March 2005, vol. 37, n^{o} 1, pp. 42–81, http://doi.acm.org/10.1145/1057977.1057980

[24] S. SCHULTE, C. JANIESCH, S. VENUGOPAL, I. WEBER, P. HOENISCH. *Elastic Business Process Management: State of the art and open challenges for BPM in the cloud*, in "Future Generation Computer Systems", 2015, vol. 46, pp. 36 - 50 [*DOI :* 10.1016/J.FUTURE.2014.09.005], http://www.sciencedirect.com/science/article/pii/S0167739X1400168X

[25] C. SUN, X. JIA, Y. ZHANG, Y. YANG, D. CHEN. *Achieving Convergence, Causality Preservation, and Intention Preservation in Real-Time Cooperative Editing Systems*, in "ACM Transactions on Computer-Human Interaction", March 1998, vol. 5, n^{o} 1, pp. 63–108, http://doi.acm.org/10.1145/274444.274447