Acronym of the proposal	STREAMS		
Title of the proposal (in French)	Solutions pair-à-pair pour le web social temps-réel		
Title of the proposal (in English)	**S**olutions for peer-**T**o-peer **REA**l-ti**M**e **S**ocial web		
Theme	☐ 1 ☒ 2 ☐ 3 ☐ 4 ☒ 5 2. Infrastructure for Internet, supercomputing or services 5. Safety and Security		
Type of research	☒ Basic Research ☐ Industrial Research ☐ Experimental Development		
Type of scientific project	☐ Platform		
Total requested funding	545,875 €	**Project Duration**	36 months

1 Executive Summary 3

2 Context and Relevance to the Call 5
2.1 Context, Economic and Societal Issues . . . 5
2.2 Relevance of the Proposal . . . 7
2.2.1 Relevance regarding the ARPEGE Call . . . 7
2.2.2 Positioning regarding national projects . . . 8
2.2.3 Positioning regarding european projects . . . 9
2.2.4 Positioning regarding industrial and general public platforms . . . 9

3 Scientific and Technical Description 9
3.1 State of the Art . . . 9
3.1.1 Peer-to-peer infrastructure . . . 9
3.1.2 Replication, consistency maintenance and collaborative editing . . . 10
3.1.3 Security and privacy . . . 11
3.2 Scientific and Technical Objectives, Progress Beyond the State of the Art . . . 12

4 Scientific and Technical Objectives, Project Description 13
4.1 Scientific programme, Project Structure . . . 13
4.2 Project management . . . 15
Task 1: Project co-ordination . . . 15

4.3 Detailed description of the work organised by task . . . 16
Task 2: Requirements and Infrastructure . . . 16
Task 3: Efficient Replication and Consistency Maintenance . . . 18
Task 4: Security and Privacy . . . 20
Task 5: Experimentations and Performance Analysis . . . 22
4.4 Tasks Schedule, deliverables and milestones . . . 24

5 Dissemination and Exploitation of Results, Management of Intellectual Property 26
5.1 Dissemination of results . . . 26
5.2 Consortium agreement . . . 26

6 Consortium description 27
6.1 Partners description & relevance, complementarity . . . 27
6.2 Relevant experience of the project coordinator . . . 29

7 Scientific justification for the mobilisation of the resources 30
7.1 Partner 1: SCORE - LORIA / Université Henri Poincaré Nancy I . . . 31
7.2 Partner 2: REGAL - LIP6 / INRIA Paris - Rocquencourt . . . 33
7.3 Partner 3: ASAP - IRISA / INRIA Rennes - Bretagne Atlantique . . . 34
7.4 Partner 4: CASSIS - LORIA / INRIA Nancy - Grand Est . . . 35
7.5 Partner 5: XWiki SAS . . . 36

8 Appendices 38
8.1 References . . . 38
8.2 CVs, Resume . . . 40
8.2.1 Gérald Oster, Associate Professor, Université Nancy 1 / Partner 1. SCORE . . . 40
8.2.2 Claudia Ignat, Chargée de Recherche 1, INRIA / Partner 1. SCORE . . . 41
8.2.3 Pascal Urso, Associate Professor, Université Nancy 1 / Partner 1. SCORE . . . 41
8.2.4 Pascal Molli, Associate Professor, HdR, Université Nancy 1 / Partner 1. SCORE 42
8.2.5 Marc Shapiro, Research Director, HdR, INRIA / Partner 2. REGAL . . . 43
8.2.6 Achour Mostefaoui, Associate Professor, HdR, IRISA / Partner 3. ASAP . . . 43
8.2.7 Marin Bertier, Associate Professor, HdR, IRISA / Partner 3. ASAP . . . 43
8.2.8 Michel Raynal, Professor, HdR, IRISA / Partner 3. ASAP . . . 43
8.2.9 Abdessamad Imine, Associate Professor, Université Nancy 2 / Partner 4. CASSIS 44
8.2.10 Michaël Rusinowitch, Research Director, HdR, INRIA / Partner 4. CASSIS . . . 44
8.2.11 Ludovic Dubost / Partner 5. XWIKI . . . 45
8.2.12 Vincent Massol / Partner 5. XWIKI . . . 45
8.2.13 Fabio Mancinelli / Partner 5. XWIKI . . . 45
8.2.14 Sergiu Dumitriu / Partner 5. XWIKI . . . 46
8.2.15 Thomas Eveilleau / Partner 5. XWIKI . . . 46
8.2.16 Jean-Vincent Drean / Partner 5. XWIKI . . . 46
8.2.17 Thomas Mortagne / Partner 5. XWIKI . . . 46
8.2.18 Jérôme Velociter / Partner 5. XWIKI . . . 46
8.3 PhD. and Post-doctoral Topics . . . 47
8.3.1 PhD. about Efficient Consistency and Replication . . . 47
8.3.2 PhD. about Security and Privacy . . . 48
8.3.3 Post-doctoral Position about Reconciliation of Structured Documents . . . 49
8.3.4 Post-doctoral Position about Peer-to-Peer Infrastructure . . . 50

STREAMS:

Solutions for peer-**T**o-peer **REA**l-ti**M**e **S**ocial web

1 Executive Summary

The web is continuously evolving faster and faster. Starting with Web2.0 era, the web became easily writable and changeable, and nowadays, it is getting more social and more real-time. We are faced with an explosion of real-time and social software such as wikis, blogs, micro-blogs, and social networks services (Facebook, LinkedIn, MySpace, etc.). This tendency is acknowledged by a recent statistical study[1] from data collected by Alexa.com which shows that Facebook website has now more page views than Google website. Real-time information delivery is emerging as one of the most important elements of our online experience. Rather than requiring that users or their software check a source periodically for updates, real-time web is a paradigm based on the principle of pushing information to users as soon as it is available. The shift toward real-time social web is certainly due to the fact that nowadays network connectivity is provided almost everywhere at every time. High speed connection has become more affordable and you get it not only at the office but also at home or at the hotel, not only when you are steady but also when you are moving by train or airplane, not only by means of your desktop computer but also via your mobile devices (your laptop, your mobile phone, etc). Real-time information delivery is likely to become ubiquitous, a requirement for almost any website or service.

The most emblematic examples of this evolution from social web and groupware towards real-time collaboration are the micro-blogging services such as Twitter and Google products such as Google Docs. For instance, the trendy new communication service Google Wave combines real-time editing of a document with conversation within that document, all these functionalities being connected to your social network.

Unfortunately, all existing real-time social web services rely on a central authority. Indeed, in order to benefit from this real-time social services, users are obliged to provide and to store their data to vendors of these services and to trust that they will preserve privacy of their data. Moreover, centralisation of the platforms hosting these services makes their scalability and reliability very costly and limit their deployment by customers for their own needs. Although this centralisation complies with the business models used by real-time social service vendors, users do not have any other alternative. Decentralised alternatives do not exist in the state of the art of collaborative algorithms and peer-to-peer architectures and protocols.

This project proposes to design peer-to-peer solutions that offer underlying services required by real-time social web applications and that eliminate the disadvantages of centralised architectures. Theses solutions are meant to replace a central authority-based collaboration with a distributed collaboration that offers support for decentralisation of services.

In real-time social software, high-responsiveness is a key factor. It means users should be aware of changes (other user status changes, contribution to a document, etc.) as soon as possible, and without requiring any interaction from the user. Therefore the underlying peer-to-peer architecture must offer fast propagation mechanisms of events and changes on shared data between the socially connected

[1] *Web becomes more social, Facebook gets more hits than Google*. `http://tinyurl.com/moresocialweb`.

peers. These mechanisms should take care of the real-time[1] aspect of communication. This aspect will be considered and addressed in task 2 of this project.

To offer high-availability of data and reliability of the system, peer-to-peer networks traditionally rely on replication of data. Most of peer-to-peer overlays only support immutable data which means it is not possible to update the content of stored data. The few peer-to-peer overlays that support mutable data consider these data as coarse grained-data, meaning it is possible to update the whole value of the data (for instance, the whole content of a document) but not perform fine-grained changes on these data (for instance, only inserting a paragraph in a document). Therefore, currently it is not possible to support real-time collaboration on the same document in a peer-to-peer network. This weakness of existing peer-to-peer approaches is due to the current scientific state of the art of replication. It should be noted that existing solutions do not consider at all the real-time aspect. In this project, we propose to address these issues in task 3. The consortium brings together experts in the domains of data replication and real-time collaborative editing, making us confident that outcomes of the project will advance the state of the art on this topic.

Real-time social web applications provide computer support for users to freely post their data (text, files, images, etc.) and share it with their friends, friends to friends, or anyone else owning access to the web application. This project is aimed at supporting flexible real-time web applications, in the sense where data availability is ensured by replicating the shared objects at every user site. So, balancing users' desire of sharing their data with users' desire to keep data private against some undesirable use is a challenging problem. The user might decide to give access to parts of his personal data to another user or let them publicly available. At the meanwhile, he can keep the other parts strictly private. Real-time social web applications typically grant access to view this personal data to three entities: the user, the user's friends or to all users. These entities need to coordinate a data sharing process that raises challenging questions: How can the personal data space be securely replicated between users in order to avoid information leakage? How can we precisely define the access control rules? Furthermore, due to the real-time feature of our social web applications, locally updating the shared objects requires high responsiveness. However adding standard access control mechanisms to real-time social web applications may degrade responsiveness since updates must be granted by some authorisation coming from a distant user (such as a central server). Furthermore, these applications have to allow for dynamic change of access permissions since users may join and leave the group in an ad-hoc manner and other shared objects may be created and destroyed. Besides access control solutions suitable for real-time collaboration, we are interested in analyzing violation of data privacy due to data disclosure to malicious peers which misuse data. In order to prevent data misuse a user has to trust the users with whom he collaborates. Therefore reputation techniques have to be used to help users in their decisions about trusting or not a person. Levels of trust have to be adjusted if a person is spotted as a cheater. In task 4 we address the above mentioned issues concerning access control and data privacy. The consortium gathers experts in the security field as well as experts in the real-time collaboration and therefore we expect that a suitable solution will be proposed in the context of STREAMS.

The solutions proposed in this project concerning efficient replication algorithms, propagation mechanisms and optimistic access control coupled with privacy-aware models will be evaluated for their suitability for real-time peer-to-peer environments in task 5. Both a theoretical and practical evaluation of these solutions will be performed. The theoretical evaluation will be performed in terms of time and space complexities of the proposed algorithms, while the practical evaluation will be made

[1]In this proposal, *real-time* term does not refer to the definition commonly agreed in computer science, i.e. there is no operational deadlines from event to system response. We refer to a soft real-time constraint which is related to users' feeling of a quasi-real-time response of the system.

on traces of human produced data modifications provided by our industrial partner. Moreover, as a proof of concept, we plan to build a demonstrator that integrates all the proposed solutions on top of which we can connect the XWiki collaborative platform of our industrial partner. In this way we will obtain the first peer-to-peer real-time wiki system.

As scientific results, the STREAMS project will present significant advances in the knowledge of distributed systems, computer-supported cooperative work and security focussing on peer-to-peer real-time web topic. These results will be assessed by a number of publications in major conferences and journals related to those areas.

The technical results of the project will be research prototypes that demonstrate the suitability and the viability of our proposals. In particular, data collected and brought into the project by the involved industrial partner will enable us to simulate real use cases on our proposed solutions. Moreover, we expect to successfully deploy the industrial partner collaborative platform (XWiki) on our prototyped implementation. This should demonstrate the full potential of our solutions.

2 Context and Relevance to the Call

Real-time social web applications are now central to our economies and every day life.

> *Real-time information delivery is fast emerging as one of the most important elements of our online experience. No more waiting for the Pony Express to deliver a parcel cross-country, no more waiting for web services to communicate from one polling instance to another. This is information being available to you at nearly the moment it's produced, whether you're watching for it or not.*
>
> **Source:** Introduction to the Real-Time Web, ReadWriteWeb[28]

Real-time social web communities hosted on Twitter or Facebook are becoming larger and larger. It is impressive to observe how much the announcements of GoogleWave or of the recent GoogleBuzz are attracting attention from media and public. This tendency is also impacting enterprise information systems as illustrated by the recent announcements of Salesforce Chatter[9], a real-time social network for enterprise, and SAP's 12sprints[10], the new SAP collaborative decision making solution.
Unfortunately, available solutions on the market are making use of *cloud computing* which is drawing more and more attention due to the privacy issues it raises as we will explain latter in this section.
To address this issue, we propose to deploy real-time web applications not in *clouds* but on peer-to-peer networks that might be composed of computers belonging to participants. Therefore, we will offer a decentralised alternative which is meant to replace a central authority-based solution with a distributed solution that offers support for social real-time services.
The STREAMS project aims to leverage the scientific locks that prevent the deployment of peer-to-peer real-time social web applications. It involves scientific literature survey and studies of challenges, design and experimentations of new suitable solutions.

2.1 Context, Economic and Societal Issues

With the rising popularity of social networking applications, violation of privacy and security requirements has become more and more frequent. Privacy and security have different concerns. Privacy represents *what* must be protected. It encompasses the collection, use, and disclosure of personal

[9] Salesforce Chatter: *Social Application Platform*. `http://www.salesforce.com/chatter/platform/`.
[10] SAP 12sprints: *Discuss, Decide, Deliver. Collaborative Decision-Making*. `http://12sprints.com/`.

information. Security represents *how* information must be protected. It covers methods for accessing and protecting information. It is possible to secure information without keeping it private, but it is not possible to protect privacy without having security. But both security and privacy are often intertwined on social networks as anyone who violates a site's security network offers an easy access to private information belonging to any user.

Placing personal information in the hands of large corporations is a perceived privacy threat. These corporations allow a profile to be produced based on the individual behaviour and therefore detrimental decisions to an individual may be taken. Moreover, due to large amounts of information social network sites process every day, it makes easier to exploit a single flaw in the system. For instance on Facebook features such as messages, invitations and photos help users gain access to private information. Moreover, flaws in the Facebook's third-party API have been found which allow for easy theft of private information. In [8] it is highlighted that social networking sites such as Facebook and MySpace might give users the impression that they can act in the privacy of a social circle as they can make explicit their social network by specifying their set of friends. Authors point out that this impression is false as user privacy might be compromised due to lack of control over activity streams, unwelcome linkage and deanonymisation through merging of social graphs. The lack of control over activity streams means that either the user is not aware of all events that are fed into the activity stream or that he exposes activity streams to an unanticipated audience. Unwelcome linkage occurs when links reveal information that they had not intended to reveal. The deanonymisation through merging of social graphs comes from the fact that social networking sites extract a lot of personal information about people such as birth date or favourite hobbies. It is possible to de-anonymise users by comparing this information across social networking sites, even if the information is partially obfuscated in each networking site.

Most of the privacy issues mentioned above are due to a single authority hosting the whole amount of data shared by all users of the system. Our peer-to-peer solution where peers are social participants is promising as it allows for better control over personal data by their owners. The infrastructure that we propose including solutions for peer-to-peer collaborative editing and content sharing as well as for privacy and security will allow us to build peer-to-peer social networks where changes are transmitted in real-time as in GoogleWave.

By its objectives the STREAMS project is targeting the *social software* market where main actors are currently proprietary software American companies. Rather than keeping their information in their own hands, clients of this software have to hand it to a third party. With the rise of Web 2.0 applications more and more people upload personal information such as emails, photographs and their works on internet-accessible servers owned by companies such as Google. As mentioned previously, this raises privacy concerns. Furthermore, most of existing social software is based on a central server. In order to deal with massive collaboration of users using social software, proprietary companies including Google, Microsoft and Amazon adopted the *cloud computing* solution. Cloud computing is a paradigm where IT power is delivered over Internet as users need it rather than drawn from a desktop computer. In this way, users and developers are enabled to use services without knowledge of, expertise with, nor control over the technology infrastructure that supports them. But there is a growing concern that mainstream adoption of cloud computing could present a mixture of privacy and ownership issues, with users potentially being locked out of their own files. Richard Stallman, founder of the Free Software Foundation and creator of the computer operating system GNU, argues that cloud computing was simply a trap aimed at forcing more people to buy into locked, proprietary systems that would cost them more and more over time. He advised users to stay local and stick with their own computers.

> *One reason you should not use web applications to do your computing is that you lose control. It's just as bad as using a proprietary program. Do your own computing on your own computer with*

> *your copy of a freedom-respecting program. If you use a proprietary program or somebody else's web server, you're defenseless. You're putty in the hands of whoever developed that software.*
>
> **Source:** Richard Stallman, *Free Software Foundation*

We share the free software initiative and rather than adopting a cloud computing solution we propose a peer-to-peer architecture that offers scalability, cost sharing and does not require a central authority. One of the most popular functionalities offered by the moment by social software is the real-time delivery of information. One of the representative social software offering the real-time feature is Google Wave, the new communication service developed by Google, that defines the wave as being equal parts conversation and document, where people can communicate and work together with richly formatted text, photos, videos and maps. While GoogleWave seems more personal productivity and entertainment oriented than group and enterprise productivity oriented, the solution that we offer is general and suitable for both categories of needs. Moreover, the algorithms that we propose for maintaining real-time consistency of documents are more efficient for the real-time functionality as the one used in GoogleWave as they rely on a peer-to-peer architecture and not on a client-server one. In this way changes are directly sent from senders to receivers and do not need to transit via a central server.

2.2 Relevance of the Proposal

2.2.1 Relevance regarding the ARPEGE Call

Considering the ARPEGE call, the work to be done is mainly focused on topic 2 *Infrastructure for Internet, Supercomputing or Services* and topic 5 *Safety and Security*.
The topic related to *Infrastructure for Internet* is considered in the project since it addresses the requested research subjects presented in section 2.2 of the call:

- **2010' Priorities:** *cloud computing, future of Internet*
- *architectures for dealing with massive data and distributed information systems*
- *management of data, access and communication*
- *peer-to-peer environments*
- *high availability of data and services*
- *collaborative platforms*

The STREAMS project contributes also to the topic *Security and Safety*, since it pays attention to the highlighted research challenges:

- **2010' Priority:** *Management of identities*
- *Privacy preservation, access control*
- *Detection of intrusion, security policies*
- *Identity management and access rights*

Considering the type of work to be done, the consortium considers the project to be part of the *basic research* class, which is defined as follows:

> **Basic research:** *experimental or theoretical work undertaken primarily to acquire new knowledge on the foundations of phenomena or observable facts, without requiring any practical application or use to be provided.*
>
> [Free english translation of an extract from the French ANR call ARPEGE]

The goal of the proposal is clearly to gain new knowledge on algorithms required for building a peer-to-peer middleware that offer support for secured collaboration over real-time social web. This addresses the first part of the definition. The software created throughout the project will include research prototypes for demonstrating the suitability and performance of the proposed solutions. This goes beyond the theoretical aspect required and the project addresses also some aspects of practical work.

2.2.2 Positioning regarding national projects

The ARC Recall (2006-2007) [7] funded by INRIA was the first project to target consistency on peer-to-peer networks for collaborative editing. One of the results of this project was the Woot[22] approach, a consistency maintenance algorithm for replicated data. Woot is suitable for peer-to-peer networks but does not consider the real-time constraint required real-time social web applications.

The RNTL XWiki Concerto project (2006-2009) [2] was sponsored by the french national research agency (ANR), within the programme *Réseau National de recherche et d'innovation en Technologies Logicielles (RNTL)*. This project allowed the industrial transfer of Woot approach into the XWiki platform to build the first peer-to-peer wiki platform available on the market. As the proposed solution is based on Woot, the real-time aspect is not considered. Moreover, one additional drawback of the approach is that it requires a total replication of the shared content on every nodes, meaning that each participant node has to store the whole data available on the network.

The DataRing project (2009-2011) [3] is sponsored by the french national research agency (ANR), within the programme *Future Networks and Services (VERSO)*. This project addresses the general problem of peer-to-peer data sharing for online communities by offering a high-level network ring across distributed data source owners. This research work targets the design of new decentralized data management techniques that scale up while addressing the autonomy, dynamic behavior and heterogeneity of both users and data sources. It also considers semantic changes and issues related to privacy, but it does not address real-time aspects of social web applications.

The ARA Respire (2006-2008) [8] was sponsored by the french national research agency (ANR). It aimed to provide a peer-to-peer environment for advanced data management applications. This project did not consider the real-time aspects nor the privacy issues that our project will investigate. The Telex middleware was funded in part by Respire.

The Wiki3.0 project (2009-2011) [11] is sponsored by the french ministry of economy, industry and employment within the call *Web Innovant*. It will address the three major evolution axes of collaborative web: real-time collaboration, social interaction integrated into the production (chat, micro-blogging, etc.) and service on-demand by means of cloud computing. The real-time aspect is considered but restricted to the context of centralised services. Moreover, it does not consider security and privacy aspects.

[7] ARC Recall (2006-2007). `http://recall.loria.fr/`.
[2] RNTL XWiki Concerto (2006-2009). `http://concerto.xwiki.com/`.
[3] ANR DataRing (2009-2011). `http://www.lina.univ-nantes.fr/projets/DataRing/`.
[8] ARA Respire (2006-2008). `http://respire.lip6.fr/`.
[11] Wiki3.0 (2009-2011). `http://wiki30.xwiki.com/` *(pending)*.

2.2.3 Positioning regarding european projects

The EU FP6 Grid4All project (2006-2009) [5] was sponsored by the european commission under the FP6 call. Grid4All aimed to develop a peer-to-peer, self-managing, federative grid system. Within Grid4All, Regal developed techniques and algorithms for optimistic and partial replication. Regal led the data storage and management work package, designing and developing a number of distributed co-operative applications, the CRDT concept and the Treedoc technology. The development of Telex was funded in part by Grid4All. This preliminary work will be continued within STREAMS.

2.2.4 Positioning regarding industrial and general public platforms

The GoogleWave platform [4]. This new communication platform announced last year by Google is the general public platform closest to our work. It combines real-time editing of a document with conversation within that document, all these functionalities being connected to your Google contact (considered as social network). As all Google applications, GoogleWave is centralised on the Google's cloud infrastructure and therefore it is subject to the same issues as any application hosted by a single vendor. The replication strategy used in GoogleWave cannot be used in a peer-to-peer context, and the consistency maintenance algorithm does not support complex structured documents as we expect to deal with in STREAMS. STREAMS project aims to provide peer-to-peer solutions that would allow to deploy real-time web applications offering the same features as GoogleWave but on a open peer-to-peer network.

3 Scientific and Technical Description

3.1 State of the Art

3.1.1 Peer-to-peer infrastructure

Peer-to-peer systems are intended to be fully decentralized, thus avoiding the bottleneck imposed by the presence of servers in traditional systems. They are highly resilient to peers arrivals and departures. Finally, individual peer behavior is based on a local knowledge of the system and yet the system converges toward global properties. However, in the context of social web the state of the art shows that it is still very centralized.
The past decade has been dominated by a major shift in scalability requirements of distributed systems and applications mainly due to the exponential growth of the Internet. A standard distributed system today is related to thousands or even millions of computing entities scattered all over the world and dealing with a huge amount of data. In this context, the peer-to-peer communication paradigm imposed itself as the prevalent model to cope with the requirements of large scale distributed systems. Peer-to-peer systems rely on a symmetric communication model where peers are potentially both client and servers. Many peer-to-peer overlay networks, organizing nodes in a logical network on top of a physical network, have been proposed in order to deal with large-scale and dynamic behavior. Following this trend, we intend to step away from general-purpose overlay networks that have been proposed so far and build domain-specific overlays customized for the targeted application and/or functionality. Among the core functionalities that we are targeting here are efficient and real-time search, notification and content dissemination in the context of social Web.

[5] EU FP6 Grid4All (2006-2009). `http://www.grid4all.eu/`.

[4] GoogleWave: *Communicate and collaborate in real time*. `http://wave.google.com/`.

Internet-based applications comprise a large number of applications deployed over the Internet. Such applications however share some common characteristics. First of all, a basic assumption is that participating entities are potentially able to communicate with every other entity using internet protocol (IP). This has a large impact on the possible structure of an overlay network. However, the characteristics of the underlying network in terms of delay and bandwidth might have to be taken into account in order to fit our targeted application. A peer-to-peer overlay network logically connects peers on top of IP. Two main classes of such overlays dominate: structured and unstructured. The differences relate to the choice of the neighbors in the overlay and, the presence of an underlying naming structure. Overlay networks represent the main approach to build large-scale distributed systems. An overlay network forms a logical structure connecting participating entities on top of the physical network, be it IP or a wireless network. Such an overlay might form a structured overlay network [27, 29, 30] following a specific topology or an unstructured network [10, 33] where participating entities are connected in a random or pseudo random fashion. In between lie weakly structured peer-to-peer overlays where nodes are linked depending on a proximity measure providing more flexibility than structured overlays and better performance than fully unstructured ones. Proximity-aware overlays connect participating entities so that they are connected to close neighbors according to a given proximity metric reflecting some degree of affinity between peers. In our context the interest in the same documents.

The actual application that we are targeting in this area is related to resource management (text documents) with the most important property that is the real-time possibility to edit and modify the content of a document by each peer. Among the different aspects we need to deal with, we have, collaborative storage, resource discovery, allocation, content distribution and indexing. Core functionalities of such applications are search, notification and dissemination.

The proposed approach has to deal with the placement of replicas of a same document. Due to real-time high responsiveness requirements, copies of documents must be placed as near as possible to interested peers. To this end, we need some indexing mechanism that will allow the different peers to discover documents they are interested in. This is somehow application specific as the documents are not static. Indeed, they are permanently modified by the different peers. This will lead us to propose an adequate querying mechanism. On the other hand, an obtained document can be changed by the acquiring peer. Consequently, we have to propose a dissemination mechanism in order to inform as quickly as possible the other peers of the changes in order to respect the real-time constraint.

For such a purpose, epidemic algorithms, also called gossip-based algorithms [9, 3] will be consistently used. In the context of distributed systems, epidemic protocols are mainly used to create overlay networks and to ensure a reliable information dissemination. The principle underlying the technique, in analogy with the spread of a rumor among humans via gossiping, is that participating entities continuously exchange information about the system in order to spread it gradually and reliably. Epidemic algorithms have proven efficient to build and maintain large-scale distributed systems in the context of many applications such as broadcasting [3], monitoring, resource management, search, and more generally in building unstructured peer-to-peer networks [15].

3.1.2 Replication, consistency maintenance and collaborative editing

To maintain consistency between replicas, a general solution, used for instance in database clusters, is known as State Machine Replication (SMR) [17]. In SMR, all replicas execute all operations in some global order on all sites, thanks to a consensus algorithm [18]. This approach does not scale well, as it puts a synchronisation bottleneck inside every application. Furthermore, it is not well adapted to our collaborative and real-time usage scenarios.

To support collaboration and disconnected computing, INRIA-Regal designed the Telex optimistic platform [2]. Each replica executes independently. Via epidemic communication, sites exchange their *actions*, augmented with concurrency and conflict *constraints*. The system ensures eventual *commitment*, i.e., all sites converge to a state that satisfies application invariants, according to *proposals* that the different sites exchange with one another. The application invariants, which Telex will guarantee, may be completely arbitrary. Telex commitment creates a total order of non-commuting operations, to which all replicas eventually converge. However, it performs commitment in the background, off the critical path of applications, which execute in real time, without delay due to network latency or disconnection.

Many collaborative and real-time web systems [11] do not make use of total ordering nor consensus. In contrast, each site executes edit operations (insert and delete into a linear buffer) in a local order, which implies that operations can be made to commute. Historically, this was done by Operational Transformation (OT) [31], i.e., locally shifting positional parameters to take concurrent inserts and deletes into account. OT is a generic approach, as integration is parameterized by transformations that are adapted to the document type. Many OT algorithms make assumptions that are not scalable, do not support dynamic groups, and/or are not suited to real-time interaction. For instance, some depend on a central site [32] for concurrency detection; the Jupiter [19] algorithm used in Google Wave [11] assumes a total-order broadcast to ensure that replicas integrate the same operations in the same order. A peer-to-peer approach is better suited to real time, as replicas do not depend on the responsiveness of the central site, and there is no single point of failure. MOT2 [6] was designed for peer-to-peer communication and does not depend on non-scalable assumptions. However, to detect concurrency requires sharing the full communication history between replicas, which is costly and decreases real-time performance.

Recently, Regal and Score have proposed a novel approach, the Commutative Replicated Data Type (CRDT) [36, 25]. This is a data type where all concurrent operations commute by design; positions are replaced by global identifiers. As epidemic propagation ensures that all replicas eventually execute all operations, eventual convergence is guaranteed; no complex concurrency control is required. CRDT algorithms have much weaker requirements which make them efficient in environments subject to churn and network failures. They do not require detection of concurrency between operations in order to ensure consistency, and they make fewer assumptions regarding operation propagation. However, existing CRDT algorithms suffer from meta-data overhead that is excessive for real-time interactions, where the communication granularity can be as small such as a single character.

An inherent limitation of commutativity-based approaches, either OT or CRDT, is that it is impossible to reject a remote update. In other words, they assume that concurrent updates are never conflicting. This makes them ill suited to maintaining high-level invariants. Conflict-detection systems built above commutativity are generally ad-hoc and clumsy.

3.1.3 Security and privacy

Two of the main technical challenges in the context of STREAMS are access and usage control. Access control grants access to data only to authorized persons. Usage control [23, 12] refers to what happens to data after it has been released to authorized persons, i.e how the user may, must and must not use it. Users receive data together with certain obligations they need to satisfy such as do not distribute data to other users.

Several solutions for peer-to-peer security have been proposed [35]. Some of these solutions address the anonymity issue where participants should be hidden from the eavesdropper during communication. Anonymity can be classified into requester anonymity where the identity of the initiator of the message is hidden and storage anonymity where the destination of the message is hidden [20]. Other solutions tackle attacks such as routing attacks (Chord, CAN, Pastry), storage and retrieval attacks, denial of service attacks and free-riders. The basic solution to avoid attacks is to employ trusted entities or servers to control activities in the system. Alternatively, public keys can be used

to identify nodes and hence malicious nodes can be avoided. In our project we do not handle any of these challenges concerning peer-to-peer security. We rather address the issue of an easy access control mechanism suitable for collaborative environments and the problem of data privacy violation due to non compliance to obligations and to data disclosure to malicious peers.
It should be noted that the protection settings available in all online social networks do not enable users to easily define their access control requirements. They are either restrictive or loose. Moreover, research effort on security in social networks is directed at the moment on privacy-preserving issues to allow statistical analysis on social data without compromising members' privacy. Access control design in social networks is a new research area and very little approaches have been proposed. A detailed survey on access control proposals can be found in [5].
The most mature works in access control are given in [5, 4] where the authors have presented an access control model with policies are specified in terms of constraints on the type and trust level of relationships existing between users. Relevant features of this model are the use of certificates for granting relationships' authenticity, and the client-side enforcement of access control according to a rule-based approach. In this approach, a user requesting to access an object must demonstrate that it has the rights of doing that by means of a proof. To establish this proof, the system needs a central node which stores and manages certificates specified by users. It is clear that their access control model in not suited in peer-to-peer context. Moreover, there is no statement how to revoke access rights and how to modify dynamically the trust levels.
To our knowledge, the only approach that addresses data privacy violation is PriServ [14], a DHT privacy service that combines the Hippocratic database principles with the trust notions. Hippocratic databases enforce purpose-based privacy while reputation techniques guarantee trust notions. However, this approach focuses on a database solution and moreover it does not deal with the discovery of the malicious users and of the update of the trust values for the different users.
In [34] a classification of trust models is given. There exist two main classes of trust models: credential-based and reputation based. In the former category if a user has a credential satisfying policies, it is a trusted user. In the reputation model trust in a person is the result of his past actions. Reputation models can be further divided into two further sub-models: one considers individual reputation while the other considers additionally social relationships among agents. Our solution for data privacy violation should be general and should be able to be combined with any existing reputation mechanism.

3.2 Scientific and Technical Objectives, Progress Beyond the State of the Art

STREAMS project will study the challenges of deploying social real-time web application on peer-to-peer network.

Project Objectives

STREAMS project aims to advance the state of the art on peer-to-peer networks for social and real-time applications.
In real-time social software, communication between participants is performed by following the social connection between these participants. We aim to design a peer-to-peer infrastructure for supporting real-time social web applications and therefore will have to propose a novel network topology that is socially-aware, where two peers socially near are close in the topology. Moreover, propagation of user changes to other interested users must be performed as fast as possible giving the user a feeling of instantaneous communication. A suitable social gossiping protocol must be defined on top of the chosen network topology, this protocol being suitable for real-time communication.

Scalability is generally considered as an inherent characteristic of peer-to-peer systems. It is traditionally achieved using replication technics. Unfortunately, the current state of the art in peer-to-peer networks does not address replication of continuously updated content due to real-time user changes. STREAMS project aims to propose generic algorithms that work for different granularities for the transmitted changes. Suitable algorithms for maintaining consistency over structured data types in real-time collaboration will be also designed. Support for detection and handling of conflicts should be offered. Besides real-time collaboration, the algorithms designed in this task should be suitable for offline mode of collaboration.
There exists a tension between sharing data with friends in a social network deployed in an open peer-to-peer network and ensuring privacy. One of the most challenging issues in social applications is how to balance collaboration with access control to shared objects. Interaction is aimed at making shared objects available to all who need them, whereas access control seeks to ensure this availability only to users with proper authorisation. Besides access control solutions suitable for the real-time collaboration, the project will analyse violation of data privacy due to data disclosure to malicious peers which misuse data. For this we propose a novel audit-based compliance control where obligations are checked a-posteriori and not enforced a-priori, as well as suitable reputation mechanisms. Proposed solutions for access and usage control will be combined.
All proposed solutions in the context of STREAMS project will be theoretically and practically evaluated for their suitability for real-time peer-to-peer environments The theoretical evaluation will be performed in terms of time and space complexities of the algorithms proposed, while the practical evaluation will be made on traces of human produced data modifications provided by our industrial partner. Moreover, as a proof of concept and prototypes developed, a demonstrator that integrates all the proposed solutions on top of which we can connect the XWiki collaborative platform will be deployed.

Breakthrough and Originality

A huge work in industry and academia is or is being dedicated to *cloud computing* for real-time social web applications. We are convinced that *cloud computing* is only a foot-step, the real challenges concern the evolution of those applications towards open peer-to-peer architectures.

Success criteria

STREAMS is basic research project. Therefore, our first success criteria is communication in top-level international conferences and journals in distributed computing, collaborative work, peer-to-peer systems, security and trust. The audience and the impact of this project within our scientific communities and within customers of our industrial partner are also great indicators of success. Finally, the degree of integration of our final demonstrator based on the XWiki collaborative platform is another measure of success.

4 Scientific and Technical Objectives, Project Description

4.1 Scientific programme, Project Structure

STREAMS is a basic research project that aims to explore and address issues related to peer-to-peer real-time social web applications. This project will involve state-of-the art survey, theoretical studies, algorithms design, validation through formal analysis and simulation studies and experimentations by means of performance evaluation on real world usage traces.

This project is divided into five tasks, each task addressing one particular aspect of the proposed research programme.
As imposed by the ANR submission system, Task 1 is dedicated to co-ordination of the project and therefore it lasts the whole duration of the project. It will also support all dissemination activities (web-site, conferences, etc.).

The diagram presented in Figure 1 gives an overview of dependencies and relationships between the technical tasks of the project.

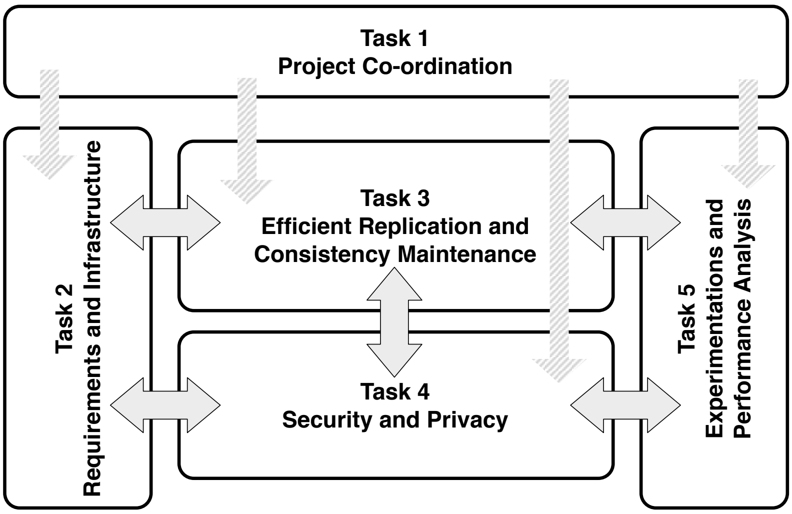

Figure 1: Tasks dependencies

The purpose of Task 2 is twofold. First, it aims to collect requirements from real data usage collected by the industrial partner and to gather basic infrastructure services that are needed by Task 3 (*Efficient Replication and Consistency Maintenance*) and Task 4 (*Security and Privacy*). Based on these requirements, a peer-to-peer infrastructure will be selected and adapted if needed. The collected real usage traces will also serve as input of Task 5 (*Experimentations and Performance Analysis*). This task will take place during the first and half year period.

Task 3 aims primarily to study theoretically replication and consistency approaches to real-time social web applications. Expected outcomes are a suitable comparison of existing approaches (commutative replicated data-types, operational transformation approach, consensus based approaches, etc.), design of new suitable algorithms and their implementations. The main inputs of this task will be the outcomes of Task 2 (*Requirements and Infrastructure*) and also few requirements which will come from the Task 4 (*Security and Privacy*) since access control policies should be replicated and privacy issues might impact the consistency maintenance technics. This task will start in the fourth month of the project and will last four months before the end to let time to integrate all the results.

Task 4 targets the design of new access control methods and privacy aware technics suitable for real-time applications hosted in peer-to-peer social system. The proposed work of this task will consider the requirements from Task 2 (*Requirements and Infrastructure*) before formally defining and verifying the new methods. A tight cooperation with Task 3 (*Efficient Replication and Consistency Maintenance*) will be required as consistency of access control and trust policies is crucial and replication of data might lead to violation of privacy rules. This task will start during the fourth month of the project and will last until almost the end of the project (four months before), the remaining time being required to finalise the integration with other outcomes.

In Task 5 partners will experiment the prototyped implementations of their algorithms. They will perform simulation and performance measurements on real traces to demonstrate the suitability and the viability of their proposition. This task will also serve as *the place of integration* of the proposed solutions of other tasks and therefore will receive as inputs the outcomes of other tasks. We expect that on top of this integration, the industrial partner will be able to deploy an adapted version of his collaborative solution (XWiki). This task will starts after one year and half, and it will last until the end of the project.

4.2 Project management

Task 1: Project co-ordination

Score (leader)	Regal	Asap	Cassis	XWiki	M1 (start)
6 p×m	2 p×m	2 p×m	2 p×m	2 p×m	M36 (end)

Project co-ordination is the responsibility of the project leader, Gérald Oster from Partner 1 Score. He will act as the scientific and administrative supervisor. With the help of steering committee, he will supervise the implementation of the project as a whole and the articulation between and complementarity of the various tasks. He will assist task leaders in their coordination with other to ensure the agenda is respected and the deliverables are produced in time.
A steering committee will be composed of one representative of each partner:

Partner	**Representative**
Partner 1: Score / LORIA, Université Henri Poincaré - Nancy I	Gérald Oster
Partner 2: Regal / LIP6, INRIA Paris - Rocquencourt	Marc Shapiro
Partner 3: Asap / IRISA, INRIA Rennes - Bretagne Atlantique	Achour Mostéfaoui
Partner 4: Cassis / LORIA, INRIA Nancy - Grand Est	Michaël Rusinowitch
Partner 5: XWiki SAS	Fabio Mancinelli

This committee will meet physically or via telephone conference at least on a bi-monthly intervals of the whole duration of the project. This steering committee will be in charge of:

- technically co-ordinating the project;
- supervising the administrative progress of the project;
- reviewing and attesting the technical and scientific soundness of the produced deliverables.

This committee will be competent for everything pertaining to the achievement of the scientific programme, and will make all decisions (to be approved by the ANR) regarding any change in the programme.

Figure 2 gives an overview of the project timeline.
A two-day kickoff meeting will mark the beginning of the project. An open workshop will be organised at the end of the project, to demonstrate and disseminate outcomes of the project to scientific and industrial communities. We will organise a partner meeting and a workshop every year, attended by all the people involved in the project.

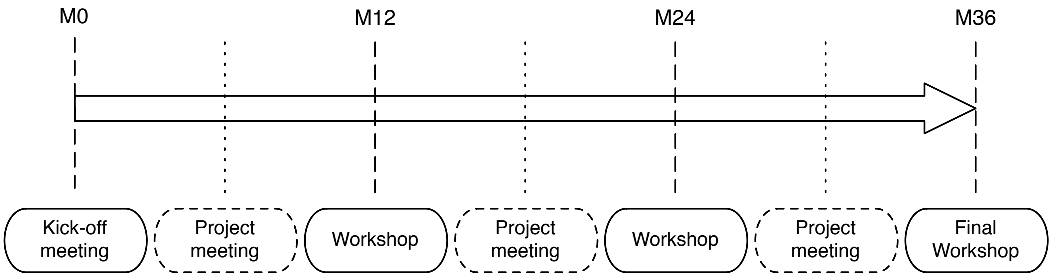

Figure 2: Project timeline and scheduled meetings

Deliverables	Deadlines
D1.1: Kick-off meeting	M0
D1.2: Website	M3 + continuous update
D1.3: Consortium agreement	M12
D1.4: Yearly workshop	M12
D1.5: Yearly workshop	M24
D1.6: Final open workshop	M36
D1.7: Final report	M36
Coordination meetings	M6, M18, M24

4.3 Detailed description of the work organised by task

Task 2: Requirements and Infrastructure					
Asap (leader)	Score	Regal	Cassis	XWiki	M1 (start)
18 p×m	12.2 p×m	0 p×m	6 p×m	9 p×m	M18 (end)

Objectives

The targeted goal of this task is two-fold.

The first objective is to collect the initial requirements. There are two kinds of requirements in this project: *user requirements* and *service requirements*. *User requirements* will mainly come from our industrial partner in the form of execution traces. *Service requirements* will emerge from the theoretical and technical constraints of other tasks.

The second objective is the definition of a peer-to-peer infrastructure for supporting real-time social web applications. The proposed infrastructure must remain open in order to integrate additional services that are not directly addressed in the STREAMS project.

State of the art

We refer to Section 3.1.1 for a comprehensive state of the art on peer-to-peer infrastructure and propagation algorithms.

Chosen methods and anticipated solutions

Concerning user requirements, Partner 5 will instrument his collaborative platforms in order to collect executions traces of the use of its services by its customers. A major role will play traces resulted from the instrumentation of the real-time XWiki system designed in the context of Wiki 3.0 project where Partner 1 and Partner 5 participate. These executions traces will be anonymised before being spread within the STREAMS project. We also expect to extract some requirements from customer feedbacks.
In real-time social software, communication between participants is performed by following the social connection between these participants. For designing a peer-to-peer infrastructure for supporting real-time social web applications it is necessary therefore to find a network topology that is *socially-aware*, where two peers socially near are close in the topology. Moreover, propagation of user changes to other interested users must be performed as fast as possible giving the user a feeling of instantaneous communication. A suitable social gossiping protocol must be defined on top of the chosen network topology, this protocol being suitable for real-time communication. Reliability of the peer-to-peer network has also to be considered, traditionally, additional copies of shared data must be kept in the network without rising privacy violations. Moreover, the mechanism of propagation of changes must take into account security constraints since document changes should not transit through unauthorised users.
For achieving the objective of this task, we will review literature surveys and draw on partners previous experience in designing gossiping protocol and proximity-aware network topology. We will also analyse infrastructures of existing real-time social web applications such as GoogleWave.

Work plan

The first activity is related to gathering of requirements. Industrial partner (Partner 5) will instrument his collaborative solutions (XWiki platform) in order to collect executions traces. In parallel to this activity, other partners involved in this task, will focus their work on the infrastructure. Under the lead of Partner 3 (ASAP team) who is internationally recognised academic expert in the domain of dynamic peer-to-peer protocols and architectures, partners will perform a literature survey, searching for proximity-aware peer-to-peer topologies, social proximity metrics, and gossiping protocols suitable for real-time communication.

Deliverables	Deadlines
D2.1: Requirements for peer-to-peer real-time social applications	M12
D2.2: Peer-to-peer socially-aware network architecture and real-time gossiping protocols	M18

Success criteria

The outcomes of this task should be: (i) Library of requirements and real execution traces that illustrate social activities such as real-time communication and collaborative editing; (ii) a peer-to-peer network architecture based on a socially-aware topology; (iii) social gossiping protocols suitable for real-time communication.

Risks

Information collected in execution traces might be not satisfactory for extracting the requirements necessary for this task. This can be due to the difficulty of instrumenting the XWiki platform but also due to the anonymizing process that might suppress valuable information for the project. However,

we expect that this risk is minimised due to the great open-source community behind XWiki. But if such difficulty happens, we will make use of our other good industrial contacts at Google, Facebook, SAP and other French SMEs that are deploying social networks and real-time web applications.
We see another potential risk in the combination of the requirements of peer-to-peer protocols that for reasons of network reliability assume a replicated shared data and the requirements of privacy control models that try to prevent disclosure of data to malicious peers.

Task 3: Efficient Replication and Consistency Maintenance

Score (leader)	Regal	Asap	Cassis	XWiki	M4 (start)
22 p×m	6 p×m	14 p×m	8 p×m	2 p×m	M32 (end)

Objectives

The objective of this task is to design efficient consistency maintenance algorithms suitable for the real-time collaboration over peer-to-peer environments. These algorithms should be as generic as possible and should work for different granularities for the transmitted changes. Suitable algorithms for maintaining consistency over structured data types in real-time collaboration should be also designed. Support for detection and handling of conflicts should be offered. Besides real-time collaboration, the algorithms designed in this task should be suitable for offline mode of collaboration.

State of the art

In decentralized collaborative systems, data are replicated to provide data availability, reliability, scalability and fault tolerance. In these systems, optimistic replication approaches are used to maintain consistency of the different copies of the same data. These approaches tolerate divergence between copies after they have been modified and ensure that they will eventually converge. Due to social nature of these applications, the convergence state must be acceptable by users, and therefore should preserve their intentions, viz., updates should be cumulative and not overwrite one another.
As explained in Section 3.1.2 there are two general families of consistency protocols, those based on consensus to order non-commuting operations, and those where all operations commute. The former comprises State-Machine Replication (SMR). This is general and safe, since it is equivalent to a local, sequential execution; however, consensus constitutes a serious bottleneck.
The latter family comprises Operational Transformation (OT) and Commutative Replicated Data Types (CRDTs). CRDTs appear to be superior to OT for scenarios where distributed users can concurrently and in real time, such as co-operative editing, wikis or version control [36, 25]. However, the overhead of existing CRDT algorithms is too high for real-time editing.
The Telex middleware [2] attempts to provide the best of each family. Commutative and non-commutative pairs of operations may co-exist. The application specifies invariants that the middleware must enforce; the middleware schedules and commits operations according to the invariants, synchronising as much as required by the application, but no more. Telex executes optimistically, i.e., it performs real-time, local updates without any network-inflicted delays. Telex supports arbitrary invariants; currently, however, it is efficient only for a limited subset of invariants. Telex was designed for disconnected operation, and its effectiveness in real-time interactions has not been assessed.

Chosen methods and anticipated solutions

CRDTs appear to be superior to the accepted state of the art in co-operative tasks, represented by OT algorithms such as Jupiter. Telex has the advantage of adapting to the consistency requirements

expressed by the application. However, neither has been used so far in real-time social collaboration scenarios. Therefore, the approach of this task will be to compare approaches such as CRDTs and Telex with the state of the art of consistency maintenance for real-time social web applications. We will study the limitations, either theoretical or practical of our current systems. An obvious approach to improving existing CRDTs is to batch and compress, in order to amortise the overhead. We believe this incurs no loss of functionality.
Currently CRDT and OT algorithms assume a linear structure of the shared document and a single granularity for the changes performed on the document. Most commonly the shared document is composed of a sequence of characters and changes can be performed at the character level. In this task we will investigate the suitability of different granularities such as character, word, line or string for the real-time communication. A fine grained granularity is more suited for the real-time aspect of collaboration where users are immediately aware of changes of other users, but it might represent a high overhead for the communication. Consistency maintenance algorithms for different granularities suited for the real-time collaboration will be designed and evaluated.
Besides real-time collaboration we want to offer support for offline work where users can work in isolation and publish their changes at a later time. A user should be able to choose the working mode that suits him better and easily commute from one work mode to the other. Switching from offline mode of collaboration to real-time collaboration might pose some issues as the integration of large changes generated in isolation might disturb the real-time collaboration of users currently working in this mode.
One common limitation of OT and CRDTs is that they are not well suited to maintaining high-level invariants and to conflict detection. For instance, concurrent edits of an XML tree cannot be guaranteed to maintain a correct XML structure; similarly for other structured data types such as schemas or ontologies. Solutions used so far are *ad hoc*. Within STREAMS we will conduct a scientific study of approaches to maintaining such high-level invariants.
Telex has many advanced aspects that are well adapted to STREAMS. However, it represents a substantial piece of software. The work to be performed is to develop Telex applications for social web tasks, and adapt it to a dynamic, real-time environment.
The design of algorithms in this task should take into account requirements of task 4 where operations refer not only to the content of the document but also to access rights and to obligations.
The algorithms designed in this task will be evaluated in task 5.

Work plan

We will start first with a survey of existing OT and CRDT algorithms to test their suitability for the real-time aspect of the collaboration. Novel efficient real-time algorithms will be built that can work for different granularities of the transmitted changes.
Further, Telex will be adapted to work for a dynamic and real-time environment and to offer support for both real-time and offline collaboration. This software will be used for testing the algorithms designed in this task.
In order to offer support for conflicts, i.e. to allow coexistence between commutativity and non-commutativity of operations, agreement protocols and the real-time CRDT algorithms previously designed in this task will be combined. Non-commuting operations will be applied in a total order. Applications allowing the presence of conflicts will be designed and tested in Telex.
Finally, novel algorithms will be designed for structured data types such as XML documents.

Deliverables	Deadlines
D3.1: Design of real-time consistency maintenance algorithms and adaptation of Telex for real-time	M20
D3.2: Study of conflicts and design of novel algorithms for structured data	M32

Success criteria

The algorithms designed in this task prove to be very efficient for the real-time collaboration according to the evaluation in task 5.

Risks

CRDT algorithms that seem very promising for the real-time collaboration prove to have a high overhead.
Dealing with conflicts in real-time environments might not be suitable.
We can design efficient algorithms for the real-time peer-to-peer collaboration, but they do not fit user intentions.

Task 4: Security and Privacy

Cassis (leader)	Score	Regal	Asap	XWiki	M4 (start)
29 p×m	19.6 p×m	0 p×m	0 p×m	2 p×m	M32 (end)

Objectives

Data centralisation in the hands of a single company is an inherent threat to privacy. Our vision is to move away from centralised authority-based collaboration towards a distributed trust network where each user can decide with whom to share their data. In this context and under the constraint of a real-time collaboration, the objective of this task is to propose a flexible security approach. Two security aspects will be addressed: data access and usage control.

State of the art

The major part of latency in access control-based social networks comes from using a shared data-structure containing access rights that is stored on a central server. So controlling access consists in locking this data structure and verifying the access is valid.
As mentioned in section 3.1.3 there is no suitable access control mechanism for real-time collaboration in social networks. To overcome the latency problem, team CASSIS proposed in [13] a flexible access control model based on replicating the access data-structure on every site. Thus, a user will own two copies: the shared social data and the access data-structure. It is clear that this replication enables users to gain performance since when they want to manipulate (read or update) the shared document, this manipulation will be granted or denied by controlling only the local copy of the access data-structure. As real-time social web applications have to allow for dynamic change of access rights, it is possible to achieve this goal when duplicating access rights. To do a single administrator control model was proposed where only one user can update the shared access data-structure.
As mentioned in section 3.1.3 the only approach that addresses data privacy violation is PriServ [14]. However, the approach is very database oriented and it is limited to relational tables. We are interested in collaborative sharing and editing of documents frequently used in social networking such as textual, XML, audio and video documents. Moreover, the approach does not propose neither

a mechanism of discovering the malicious users that do not respect the obligations required for using the data nor an approach for updating the trust values associated to users.

Chosen methods and anticipated solutions

Concerning the access control issue, the approach previously proposed for a single administrator will be generalized for multiple administrators where several users can edit the shared access data-structure. Unlike single administrator model, when dynamic access changes are initiated by several users we have to deal with the resolution of conflict changes.
The shared data's updates and the access data-structure's updates are applied in different orders at different user sites. The absence of safe coordination between these different updates may cause security holes (*i.e.* permitting illegal updates or rejecting legal updates on the shared data). For example, consider scenario where two users U_1 and U_2 working concurrently on a shared friend list. U_1 creates a friend profile node with full access, and later gives it read-only permission. U_2 receives the new node and start to edit it. Next, before receiving the read-only permission, U_2 broadcasts his modification. The problem is now to define a correct state of convergence in this situation. Inspired by the *optimistic security* concept introduced in [24], we propose an optimistic approach that tolerates momentary violation of access rights but then ensures the copies to be restored in valid states with respect to the stabilized access control policy. A PhD student from team CASSIS will work on this solution.
Concerning the issue of violation of data privacy due to data disclosure to malicious peers which misuse data we propose an audit-based compliance control where obligations are checked a-posteriori and not enforced a-priori. The audit-based compliance control proposed in [7] that requires an auditing authority with the task and the ability to observe the critical actions of users is not suited for a peer-to-peer environment as the one required in the context of STREAMS project. We assume that the environment supporting the real-time collaboration allows users to keep a secure log of their actions. This log will keep user modifications made on the shared data as well as the obligations the receiver has to fulfill when data is shared. A certain shared document will have associated a log showing user contributions on this document. Assuming the presence of such an environment that logs user actions is not unreasonable as it is available on each collaborative system for traceability reasons. This log can be automatically analysed by the system for discovering malicious users. Each user is assigned a level of trust in the network. A trust level is an assessment of the probability that a peer will not cheat. A peer with a high trust level is considered as honest and a peer with a low trust level as malicious. Trust levels are updated (incremented/decremented) in order to reflect peers behavior. A peer that was spotted as a cheater according to the log will have its corresponding trust value decremented. Any reputation approach can be used in association with this mechanism for discovery of malicious users. A PhD student from team SCORE will work on this topic.
A formalization of the two approaches will be performed as well as prototype applications testing the two approaches. Testing will be performed using peer-to-peer simulators and the overhead of access control and usage control will be measured under the assumption of real-time constraints.
Finally, we plan to combine the solutions we obtained for access and usage control. The combined solution will be tested in task 5.

Work plan

Having the input from task 2 we start by defining the collaboration model and protocols for both access and usage control including the types of access rights and of obligations users have to respect after they have access to data. Team CASSIS will work then in parallel with team SCORE on the solutions for access and usage control respectively. Interactions with task 3 will take place as consistency and consensus algorithms have to take into account both operations referring to document content and to user rights and obligations. The two solutions for access and usage control will be formalized and prototypes will be implemented and tested. Finally the two approaches will be combined and tested in task 5.

Deliverables	**Deadlines**
D4.1: Access control specification, design and prototype	M24
D4.2: Usage control specification, design and prototype	M24
D4.3: Combination of access and usage control solutions	M32

Success criteria

The outcome of this task shall be an access and usage control mechanism that is enough flexible for users to be used in a collaborative environment and that requires just a small overhead for the real-time communication. The access control solution proposed should be proved to be a safe mechanism guaranteeing that no unauthorized user is granted access to data. The usage control mechanism should discover the malicious persons in the collaboration network. The pilot application tested with real clients in task 5 will give us a feedback if the proposed mechanisms are successful.

Risks

The access and usage control mechanism could have a too high overhead for the real-time communication. Moreover, it could be a burden for users to specify the access control rights and usage control obligations.

Task 5: Experimentations and Performance Analysis

Regal (leader)	Score	Asap	Cassis	XWiki	M18 (start)
8 p×m	14 p×m	8 p×m	7.8 p×m	15 p×m	M36 (end)

Objectives

In this task we will perform both theoretical and practical evaluation of the solutions proposed in the other tasks. Theoretical evaluation in terms of time and space complexities will be done. And, the efficiency of algorithms for groupware and social software will be evaluated on history of human produced data modifications. For instance, CRDT approaches have been evaluated on Wikipedia articles history. These evaluations will be performed regarding the peer-to-peer nature of our context but also regarding the real-time (low granularity of modifications, high responsiveness of editing) constraints of the targeted applications. The final objective is to integrate our solutions within a demonstrator which will be connected to the XWiki collaborative platform.

State of the art

Some of the existing CRDT algorithms for linear structures such as Treedoc [25] and Logoot [36] have been implemented and their complexity in terms of time and space has been experimentally evaluated on Wikipedia histories or Subversion histories. We will evaluate the existing and newly produced algorithm against real-time (i.e. with a low granularity) and distributed histories.
There exist several ways of validating our algorithms. We can make use of peer-to-peer simulators such as PlanetSim [26] and PeerSim [16]. Or, we can design prototypes and deploy them on large-scale distributed systems. For instance, performance and viability of Uniwiki [21], a peer-to-peer wiki built

on top of a structured overlay, have been evaluated by deploying it on Grid'5000 platform. Grid'5000 platform is a large-scale distributed environment which gathers 5,000 CPUs distributed geographically over 9 sites in France. It is an experimental platform that can be easily controlled, reconfigured and monitored to run experiments. PlanetLab [1] is another open platform for developing, deploying and accessing planetary-scale services, that can be used to run such experiments. It is worth to mention that Partner 3 (ASAP) is currently building an experimental network that is composed of an open cluster and laptops of volunteers. As this network will be constituted by very different computers connected through very different types of network with very different characteristics, it will be a perfect test bed for running experimental evaluation in the same conditions as in a real peer-to-peer network. Independently of the experimental approach which is applied (simulated executions or real execution on an experimental platform), it is necessary to have valuable data to serve as inputs of the experiments.

Chosen methods and anticipated solutions

Evaluation will be performed in two steps: a theoretical evaluation followed by a practical evaluation using human produced traces in different real-time collaborative systems.

Work plan

Optimistic access and usage control solutions proposed in task 4 (*Security and Privacy*) will be first evaluated theoretically, then they will be evaluated by means of peer-to-peer simulators. We will evaluate their safety with formal verification technics but also measure their performance overheads. We plan to perform validations of the proposed algorithms in task 3 (*Efficient Replication and Consistency Maintenance*) on Grid'5000 platform to test their scalability and performance in large dynamic peer-to-peer systems. As soon as the experimental network of partner 3 (ASAP) is available, we will run our experiments also this platform since the conditions will be more realistic than the one available on Grid'5000. For these experimentations, we will use traces collected by our partners but also traces collected from other platform such as GoogleWave or any distributed version control systems (Git, Mercurial, etc.)
At the end of the project, we expect to be able to integrate all the proposed solutions (socially-aware peer-to-peer network, efficient replication algorithms, optimistic access control mechanisms coupled with privacy aware models) in one demonstrator on top on which we interconnect the real-time version of the collaborative platform of our industrial partner.

Deliverables	**Deadlines**
D5.1: Theoretical evaluation results of proposed solutions	M24
D5.2: Practical experimentation results of proposed solutions	M32
D5.3: Practical experimentation results of the integrated solution	M36

Success criteria

The success of this task is the measure of the scalability, efficiency and safety of the proposed solutions in the different tasks of the project. All the conducted experimentations (theoretical or practical) will give valuable knowledge that will directly contribute to the state of the art. Consequently, we aim to publish these results in significant venues.
The final anticipated integration of all proposed solutions will establish the success criteria of the whole project. If such integration succeeds, then this will open the way to the decentralisation of real-time social web applications.

Risks

Since the proposed experimental evaluations rely on the real execution traces collected by our industrial partner, there exist a very small risk: information available in the traces might not be appropriate. Indeed, at the beginning of the project, partner 5 can only provide traces collected from his non-real-time collaboration platform. These traces will be useful since they contain already a lot of informations (kind and size of user modifications, concurrency between these modifications, etc.), but they were not produced in *real-time*. Therefore an extrapolation of the provided data is needed. Fortunately, we expect to get real-time traces from our partner when their platform will support this type of collaboration. This is the targeted objective of the Wiki3.0 in which partner 1 and partner 5 are involved. Consequently, there exist a risk that these real-time traces will be available lately within the project.

4.4 Tasks Schedule, deliverables and milestones

Figure 3 recalls the proposed project timeline.

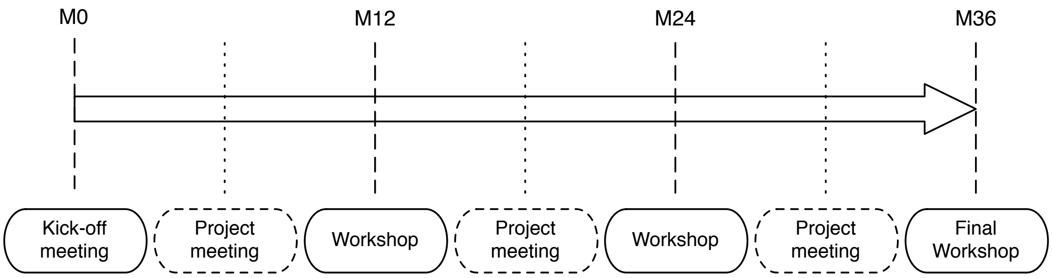

Figure 3: Project timeline and scheduled meetings

A two-day kickoff meeting will mark the beginning of the project. An open workshop will be organised at the end of the project, to demonstrate and disseminate outcomes of the project to scientific and industrial communities. We will organise a partner meeting and a workshop every year, attended by all the people involved in the project.

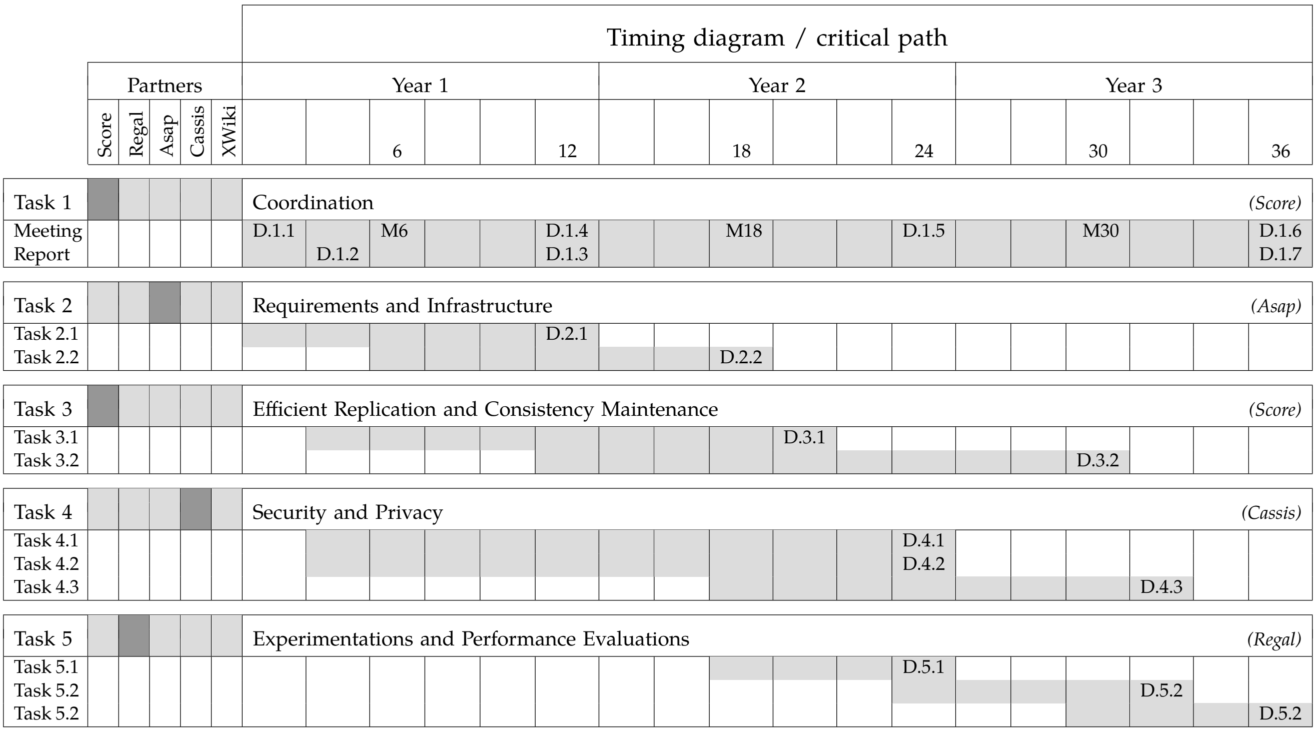

Figure 4: Project deliverables schedule

Full list of deliverables and their deadlines

D1.1: Kick-off meeting	M0
D1.2: Website	M3 + continuous updates
D1.3: Consortium agreement	M12
D1.4: Yearly workshop	M12
D1.5: Yearly workshop	M24
D1.6: Final open workshop	M36
D1.7: Final report	M36
D2.1: Requirements for peer-to-peer real-time social applications	M12
D2.2: Peer-to-peer socially-aware network architecture and real-time gossiping protocols	M18
D3.1: Design of real-time consistency maintenance algorithms and adaptation of Telex for real-time	M20
D3.2: Study of conflicts and design of novel algorithms for structured data	M32
D4.1: Access control specification, design and prototype	M24
D4.2: Usage control specification, design and prototype	M24
D4.3: Combination of access and usage control solutions	M32
D5.1: Theoretical evaluation results of proposed solutions	M24
D5.2: Practical experimentation results of proposed solutions	M32
D5.3: Practical experimentation results of the integrated solution	M36

5 Dissemination and Exploitation of Results, Management of Intellectual Property

5.1 Dissemination of results

Dissemination and transfer of knowledge are both internal and external activities to the project consortium. Within the former, it is a process of improving knowledge among partners. As for external dissemination, it will be focused on scientific communication.
All consortium partners will contribute to the dissemination activities by means of:

- Writing of research papers and participation to peer-reviewed international conferences and journals.
- Setting-up and maintenance of various dissemination tools, e.g. web-sites.
- Organising and participating to the final open workshop where outcomes of the project will be demonstrated.

All documentations are expected to be public. They will first circulate inside the project and will be made public as soon as the involved partners and work package leaders have declared their consent. An internet website will be developed from the very start of the project, whose main objective is to diffuse the STREAMS's objectives and results as wider as possible, throughout the community and over. All deliverables will be published on this website.
The industrial partner (partner 5 XWiki) expects to use the project's outcomes in order to improve its open-source collaborative solutions.

5.2 Consortium agreement

To protect interests and intellectual property rights of each partner, a Consortium Agreement will be signed in the first twelve months of the project. This agreement will settle any remaining issues related to the consortium organisation and responsibilities of all partners. It will define the intellectual property of previous results and software brought by each partner in the project, as well as the one the intellectual property and exploitation of outcomes of the project. The Consortium Agreement will precisely specify under which terms each partner will make available to the consortium its pre-existing knowledge and software required for the project. The consortium agreement will mainly follow the principles above:

- Existing knowledge and software: Each partner will retain its intellectual property of its pre-existing knowledge and software that he will make available to the consortium.
- Exploitation of outcomes: Results (knowledge and software) are the property of their authors. Artifacts include source codes, data sets, benchmarks, execution traces, and so on. Intellectual property of artifacts developed in the context of the project by a partner will belong to this partner. Intellectual property of artifacts developed in the context of the project in collaboration with other partners will be shared between these partners. Partners will be encouraged to make all artifacts developed in the context of the project widely available under a non-restrictive open source license such as BSD, CeCill. Sensitive data such as execution traces and usage logs will be anonymised before being made public.
- Publication: Each partner will be free to publish its own results, knowledge or artifacts without requiring the permission of other partners. Partners will inform other partners in advance about future submissions and publications related to the project. Obviously, partners may not violate the confidentiality or intellectual property of the other partners.

6 Consortium description

6.1 Partners description & relevance, complementarity

Partner 1: SCORE - LORIA / Université Henri Poincaré Nancy I

LORIA, the Lorraine Laboratory of Information Technology Research and its Applications, is a mixed research unit - UMR 7503 - shared by several institutions: CNRS (National Centre of Scientific Research), INPL (National Polytechnic Institute of Lorraine), INRIA (National Institute for Research in Computer Science and Control), Henri Poincaré University Nancy 1 and Nancy 2 University. LORIA is a laboratory of more than 450 individuals including more than 150 researchers and teaching-researchers, a third of doctorate students and post-doctorates and a third of engineers, technicians and administrative staffs. LORIA's research teams are involved in more than 40 industrial contracts underway for more than 3 millions euros, and participate in more than 125 cooperation projects with more than 32 different countries.

Score team is a joint research group of Henri Poincaré University Nancy 1, Nancy 2 University and INRIA Nancy - Grand Est. Score team investigates cooperative, distributed, and process-aware Web Information Systems. Its research are organised in two main streams:

- Process Engineering which is interested in process-aware information systems that manage and execute operational processes involving people, applications, and information sources on the basis of process models.
- Collaborative Distributed Systems which is concerned with the development of collaborative systems but with a scientific focus on data consistency in peer to peer architectures. Interactions between these two axes are mainly governed by shared issues, especially on awareness, coordination, and privacy and security management.

Members of Score team have a high expertise in consistency maintenance approaches for optimistic replicated data such as operational transformation approach for instance. These approaches have been applied in the context of several research and industrial projects such as the ANR XWiki Concerto (peer-to-peer wiki) and the European FP6 QualiPSo project (collaborative software development platform).
Score acts as the leader of this project.

Partner 2: REGAL - LIP6 / INRIA Paris - Rocquencourt

INRIA, the French National Institute for Research in Computer and Control Sciences is an world leader in fundamental and applied research, in the areas of information and communication science and technology. The Institute plays a major role in technology transfer, by research training, scientific and technical information, development, providing expert advice and participating in international programs.
INRIA has eight research centres (Paris-Rocquencourt, Rennes, Sophia Antipolis, Grenoble-Lyon, Nancy, Bordeaux, Lille and Saclay). Its workforce numbers 3,700, of whom 2,900 are scientists, organised in 152 research project-teams. Many INRIA researchers are also professors, whose approximately 1,000 doctoral students work on theses as part of INRIA research project-teams.
Regal is a joint research group of LIP6 and INRIA Paris - Rocquencourt.
Regal investigates large-scale distributed systems, and especially peer-to-peer architectures. An important focus is large-scale replication in very dynamic settings. They investigate adaptive algorithms, in order to react to changes in the environment and in the application.

Some our their research areas are:

- Data management in large scale configurations: to deploy and locate data, and to manage consistency.
- System monitoring and failure detection: we investigate failure detectors with provable properties in dynamic environments.
- Replication: replication of data improves availability and responsiveness but updates raise the issue of consistency. We have several interests in this area: fault-tolerant replication techniques, optimistic techniques (which allow local updates but cause replication divergence) and adaptive replication. Our research aims to compare and evaluate existing algorithms and to combine their best features into new, high-performance protocols.
- Dynamically-adaptive operating systems

Partner 3: ASAP - IRISA / INRIA Rennes - Bretagne Atlantique

The ASAP (As Scalable As Possible) team is located both in IRISA/INRIA Rennes and INRIA Saclay, lead by Anne-Marie Kermarrec. Its research activities range from theoretical bounds to practical protocols and implementations for large-scale distributed dynamic systems to cope with the recent and tremendous evolution of distributed systems. Effectively we observed huge evolutions: (i) Scale shift in terms of system size, geographical spread, volume of data, and (ii) Dynamic behaviour due to versatility, mobility, connectivity patterns.
These characteristics lead to a large amount of uncertainty. Mastering such uncertainty is actually our goal. We aim at providing a wide range of applications (from content delivery networks to sensors networks, from backup systems to voice over IP, from publish-subscribe systems to genomic databases). We focus our research on two main areas: information management and dissemination. We believe such services are basic building blocks of many distributed applications in two networking contexts: Internet and wireless sensors. These two classes of applications, although exhibiting very different behaviours and constraints, clearly require scalable solutions. ASAP team regroups 14 researchers (including 1 professor, 1 senior researcher, 2 associate professors and 2 researchers).

Partner 4: CASSIS - LORIA / INRIA Nancy - Grand Est

The aim of the INRIA CASSIS project is the design and the development of tools for checking the safety of systems with an infinite number of states. Our analysis of systems is based on a symbolic representation of the sets of states as formal languages or logical formulas. Safety is obtained by automated proofs, symbolic exploration of models, or tests generation. The applications are embedded softwares on smart cards, for example, security protocols and distributed systems. Moreover, INRIA CASSIS project is involved in the design and the validation of new access control models for distributed systems characterized by interactivity, scalability and high security levels. CASSIS is a partner in the European FP7 project AVANTSSAR (2007-2011) which aims to propose a rigorous technology for the formal specification and *Automated VAlidatioN of Trust and Security of Service-oriented ARchitectures*.

Partner 5: XWiki SAS

XWiki SAS is an OpenSource company created in 2004 that is specialized in the development of collaborative solutions for the enterprise. Its main product is named after the company and provides an open platform for the development of collaborative applications for editing and sharing information in the context of the enterprise. The XWiki platform has been used to address different types of markets and needs, such as collaborative watch, project management and the implementation of advanced intranets and extranets. The platform is also corroborated by a wide OpenSource community that contributes to its development and that uses it in different domain. XWiki SAS employs almost 30 people, some of them in its Romanian branch. Among the XWiki SAS's customers we mention the AFP (Agence France Presse), the Aelia Group, and EADS. XWiki SAS' solutions have also been used for deploying large scale solutions in the context of the Curriki Project (a non-profit organization that distributes educational OpenSource material in order to improve education worldwide), and EMC, worldwide leader in storage solutions. XWiki SAS has already participated and is currently leading other research and development projects such as XWiki Concerto (peer-to-peer wiki), SCRIBO (Semiautomatic and Collaborative Retrieval of Information Based on Ontology) and the latest Wiki3.0 which aims at including real-time and social features in the XWiki platform.

Complementarity and added value of each partner

Score and Regal have a strong expertise in consistency maintenance algorithms for replicated data. Regal research is more focused on large scale distributed system while Score research is more focused on computer-supported cooperative work. Recently, both teams proposed new commutative replicated datatypes which reveled being very promising approaches. Additionally, each of these two teams have a specific expertise in other optimistic replication approach: Score is one of the international expert in operational transformation (OT) approach while Regal is the inventor of the Actions-Constraints Framework (ACF) which permits to express invariant to be preserved during reconciliation.
Asap brings to the consortium its expertise on large-scale distributed dynamic system. They have international expertise on providing abstractions and algorithms to build server-less large-scale distributed applications involving a large set of volatile, geographically distant, potentially mobile and/or resource limited computing entities.
Cassis is expert in formal methods dedicated to security protocols. Score and Cassis already collaborated in the past to formally model operational transformation approach and apply automated verification technics to check the correctness of the existing OT algorithms.
XWiki builds open-source collaboration solutions and is one of the world leader in enterprise wiki systems. Traces collected from usage of their systems by their international customers is a valuable data source for the project.
All partners of the consortium already successfully collaborate in several research projects. For instance, the four academic partners collaborated in the INRIA ARC Recall (2006-2007) on massive collaboration over peer-to-peer wiki systems. Score and XWiki participated in ANR XWiki Concerto (2006-2009) and are cooperating in the context of the Wiki3.0 project (2009-2011).

6.2 Relevant experience of the project coordinator

Gérald Oster is an Associate Professor at University Nancy 1, France. He received his Ph.D. degree in 2005 from the Department of Computer Science at Nancy-Université, France and he was a postdoctoral researcher for one year at ETH Zurich, Switzerland. His domain of research is distributed collaborative editing systems with a focus on optimistic replication in peer-to-peer systems. He participated to several research projects such as RNTL LibreSource (2002-2004), INRIA ARC Recall (2006-2007), ANR XWiki Concerto (2006-2009), EU FP6 QualiPSo (2006-2010) and Wiki3.0 (2009-2011).
He is reviewer for the following international journals and conference in distributed collaborative systems: *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, *ACM Conference on Computer-Supported Cooperative Work (CSCW)*, *International Conference on Collaborative Computing (CollaborateCom)*, *International Conference on Ubiquitous Computing (UbiComp)*, *International Conference on Web Information*

Systems Engineering (WISE), *International Conference on Cooperative Information Systems (CoopIS)*, *International Conference on Advanced Information Systems Engineering (CAiSE)*, *International Conference on Cooperative Design, Visualization and Engineering (CDVE)*, *International Workshop on Collaborative Peer-to-Peer Systems (COPS)*.
He is or was member of the program committee of the following international conferences and workshops: *International Conference on Cooperative Information Systems (CoopIS)* and *International Workshop on Collaborative Editing Systems (IWCES or CEW)*.
In 2005, he received the LORIA medal for his participation in the RNTL LibreSource (2002-2004) in which part of his PhD research results was transferred to industry.

Selected publications related to the project

[MD09] G. Oster, P. Molli, S. Dumitriu and R. Mondéjar, UniWiki: A Collaborative P2P System for Distributed Wiki Applications. In *Proceedings of the 18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises, WETICE 2009*, pp.87-92, June 2009.

[IOM07] C.-L. Ignat, G. Oster, P. Molli, M. Cart, J. Ferrié, A.-M. Kermarrec, P. Sutra, M. Shapiro, L. Benmouffok, J.-M. Busca and R. Guerraoui, A Comparison of Optimistic Approaches to Collaborative Editing of Wiki Pages. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2007*, pp.474-483, November 2007.

[OMU06] G. Oster, P. Molli, P. Urso and A. Imine, Tombstone Transformation Functions for Ensuring Consistency in Collaborative Editing Systems. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2006*, pp.38-48, November 2006.

[OUM06] G. Oster, P. Urso, P. Molli and A. Imine, Data Consistency for P2P Collaborative Editing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2006*, pp.259-268, November 2006.

[IROM06] A. Imine, M. Rusinowitch, G. Oster and P. Molli, Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2), pp.167-183, 2006.

7 Scientific justification for the mobilisation of the resources

We do not claim any budget for a large equipment funding. Most experiments can be performed on standard desktops or laptops. However, in Task 5, we plan to make use of Grid'5000[6] for experiments at large scale.
Likewise, the project does not require any subcontracting nor internal expenses.
Our travel requests include the project meetings, which all participants must attend, and partner-to-partner visits that will occur during the project.
To ensure a tight and successful collaboration, there will be frequent exchanges and visits between partner. We plan that the PhD students involved in the project make visit to all partners.
Finally, our travel budget includes attendance in international conferences, where we will present the results of STREAMS to international research audiences.

[6] Grid'5000. `http://www.grid5000.fr/`.

Overall summary of the financial plan

	Score	*Regal*	*Asap*	*Cassis*	*XWiki*	*Total*
Permanent personnel	36 m.m	8 m.m	24 m.m	24 m.m	30 m.m	**122 m.m**
Non-permanent personnel funded by ANR	33 m.m	8 m.m	18 m.m	24 m.m	0 m.m	**83 m.m**
Non-permanent personnel non-funded by ANR	4.8 m.m	0 m.m	0 m.m	4.8 m.m	0 m.m	**9.6 m.m**
Total project cost	**472,218 €**	**251,691 €**	**304,715 €**	**455,899 €**	**263,243 €**	**1,747,767 €**
Total requested budget:	**171,797 €**	**58,040 €**	**100,048 €**	**97,531 €**	**118,459 €**	**545,875 €**

7.1 Partner 1: SCORE - LORIA / Université Henri Poincaré Nancy I

Equipment

Not applicable.

Personnel costs

Permanent personnel

Name	*Status*	*%*	*Months*
Gérald Oster	MCF	50	18
Claudia Ignat	CR1	25	9
Pascal Urso	MCF	20	7.2
Pascal Molli	MCF/HdR	5	1.8

Non-permanent personnel not funded by ANR

Name	*Status*	*%*	*Months*
Hien Thi Thu Truong	PhD	20	4.8

Non-permanent personnel funded by ANR

Name	*Status*	*%*	*Months*
CDD1	PhD	100	24
CDD2	Post-doc	100	9

Description of CDD1 - Position: PhD fellowship
The recruited PhD. will be mainly involved in task 3 (*Efficient Replication and Consistency Maintenance*). First, he will participate in the gathering of requirements related to replication in task 2 (*Requirements and Infrastructure*). Then he will work on the literature survey, the design and the implementations of novel consistency maintenance solutions which will be based either on CRDTs and/or Operational Transformation approaches. Finally, he will implement and participate to the experimentations performed in task 5(*Experimentations and Performance Analysis*).

All these tasks are planned to be achieved during a PhD thesis. Score team only requests the ANR funding of two years. The third year will be financed by the team on its own budget.
A PhD. thesis proposal is available in Section 8.3.1.
Description of CDD2 - Position: Post-doctoral fellowship
The recruited postdoctoral fellow will be mainly involved in task 3 (Efficient Replication and Consistency Maintenance) on the subtask of designing efficient algorithms for structured data in real-time peer-to-peer environments. He will first study existing algorithms for real-time communication based either on CRDTs or operational transformation mechanism and then propose a novel reconciliation approach for structured data. The proposed approach will be extended to deal with conflicts. The implemented algorithms will be tested in task 5 (Experimentations and Performance Analysis).
Score team requests 9 month of funding for the postdoctoral position. An additional period of 9 months will be funded by the Wiki 3.0 project. In fact the candidate will be first recruited for working 9 months in Wiki 3.0 project on reconciliation algorithms for real-time communication and afterwards for working 9 months in the STREAMS project on extending these algorithms for peer-to-peer environments and for dealing with conflicts.
The research subject for the postdoc position is available in Section 8.3.3.

Subcontracting

Not applicable.

Travel

We request the approximate funding for about three national meetings attended by four people (a total of 12). This amount includes the planned meetings to which almost all participants have to attend. To this budget, we add the organisation the estimated cost of the final open workshop (room rental and associated costs).
We plan and therefore request budget for approximately nine registrations at various international conferences and six registrations for national conferences, where we will present the research results obtained in this project.
Finally, to ensure a tight and successful collaboration with partners, PhD students and post-doctoral personnels working on the project will make short-period visits to other academic partners. Therefore, we request the budget for about two periods of two months.

Travel and short-period visits

Kind	*Number*
National meetings	12
National conferences	6
International conferences	9
Short-period visits to partners	2 × 2-months

Expenses for inward billing (Costs justified by internal procedure of invoicing)

Not applicable.

Other working costs

Statistically, 30% of the computer inventory is renewed every year. Therefore during this three years project, we expected that approximately one computer per participant will be purchased.

Small equipment

Kind	*Number*
Workstation/laptop	5

7.2 Partner 2: REGAL - LIP6 / INRIA Paris - Rocquencourt

Equipment

Not applicable.

Personnel costs

Permanent personnel

Name	*Status*	*%*	*Months*
Marc Shapiro	DR1/HdR	20	8

Non-permanent personnel funded by ANR

Name	*Status*	*%*	*Months*
CDD1	Engineer	100	8

Description of CDD1 - Position: Engineer
The recruited engineer will mainly work on task 5 with some minor implication in task 3. The targeted objectives will be the experimentations of the proposed replication and consistency algorithms in the Telex system, a replication infrastructure for collaborative and nomadic applications which supports optimistic and partial replication. A small part of the work will be dedicated to adapt Telex to the real-time context of the project. Then, the main task will consist in implementing the proposed solutions and run experimentation using real traces provided by the industrial partner.

Subcontracting

Not applicable.

Travel

We request the approximate funding for about three national meetings attended by two people (a total of 6). This amount includes these planned meetings to which almost all participants have to attend, but also the visits to other partners that will occur during the project. We expect about 6 participations to top-level international conferences in distributed systems and peer-to-peer computing.

Travel

Kind	*Number*
National meetings	6
International conferences	6

Expenses for inward billing (Costs justified by internal procedure of invoicing)

Not applicable.

Other working costs

We plan to buy two hefty computers for use in the project.

Small equipment

Kind	*Number*
Workstation/laptop	2

7.3 Partner 3: ASAP - IRISA / INRIA Rennes - Bretagne Atlantique

Equipment

To be able to test and evaluate peer-to-peer protocols, we are installing an open cluster that will constitute the core of an experimental network (also composed of laptops of the STREAMS members and the computer of some external volunteers). As this network will be composed by very different computer architectures and connected by different types of network (ADSL, RENATER, etc.), it will present, in a manageable scale, the same conditions as real peer-to-peer networks. Part of this cluster is already financed by another ANR project (Shaman) and by regional funds for a total amount of to 75,000 euros. The requested participation for the STREAMS project amounts to 6,000 euros, in order to buy ten blade computers, a switch and a KVM switch.

Personnel costs

Permanent personnel

Name	*Status*	*%*	*Months*
Achour Mostefaoui	MCF/HdR	22	8
Marin Bertier	MCF	22	8
Michel Raynal	Prof/HdR	22	8

Non-permanent personnel funded by ANR

Name	*Status*	*%*	*Months*
CDD1	Post-doc	100	12
CDD2	Internship	100	6

Description of CDD1 - Position: Post-doctoral fellowship
The postdoctoral fellowship is related to activities performed in task 2. A more detailed topic proposal is available in Section 8.3.4.
Description of CDD2 - Position: Internship fellowship
The recruited internship will participate in the activity of task 3 (*Efficient Replication and Consistency Maintenance*). More precisely, she will work on the concurrency control algorithms dealing with continuous updates of different levels on granularity on the same shared document.

Subcontracting

Not applicable.

Travel

It is expected that most of the results will be communicated to top-level international conferences in distributed computing, collaborative work and peer-to-peer systems. We expect about 10 participations to these conferences during the 3 years of the project.
Moreover, we request the approximate funding for about three national meetings attended by four people (a total of 12). This amount includes these planned meetings to which almost all participants have to attend, but also the visits to other partners that will occur during the project.

Travel

Kind	*Number*
National meetings	12
International conferences	10

Expenses for inward billing (Costs justified by internal procedure of invoicing)

Not applicable.

Other working costs

Two classical workstations will be bought for the two recruited people (post-doc and internship) for a total cost of 5,000 euros.

Small equipment

Kind	*Number*
Workstation/laptop	2

7.4 Partner 4: CASSIS - LORIA / INRIA Nancy - Grand Est

Equipment

Not applicable.

Personnel costs

Permanent personnel

Name	*Status*	*%*	*Months*
Abdessamad Imine	MCF	33	12
Michaël Rusinowitch	DR1/HdR	33	12

Non-permanent personnel not funded by ANR

Name	*Status*	*%*	*Months*
Asma Cherif	PhD	20	4.8

Non-permanent personnel funded by ANR

Name	*Status*	*%*	*Months*
CDD1	PhD	100	24

Description of CDD1 - Position: PhD fellowship
The recruited PhD. will be involved in task 4 (*Security and Privacy*). A PhD. thesis proposal is available in Section 8.3.2.

Subcontracting

Not applicable.

Travel

Travel and short-period visits

Kind	*Number*
National meetings	9
National conferences	3
International conferences	6

We request the approximate funding for about three national meetings attended by three people (a total of 9). This amount includes the planned meetings to which almost all participants have to attend. We plan and therefore request budget for approximately six registrations at various international conferences and three registrations for national conferences, where we will present the research results obtained in this project.

Expenses for inward billing (Costs justified by internal procedure of invoicing)

Not applicable.

Other working costs

Not applicable.

7.5 Partner 5: XWiki SAS

Equipment

Not applicable.

Personnel costs

Permanent personnel

Name	*Status*	*%*	*Months*
Ludovic Dubost	Architect	10	3
Vincent Massol	Architect	10	3
Fabio Mancinelli	R&D Engineer	20	6
Jean-Vincent Drean	R&D Engineer	15	4.5
Sergiu Dumitriu	R&D Engineer	15	4.5
Thomas Eveilleau	R&D Engineer	15	4.5
Jerome Velociter	R&D Engineer	15	4.5

Subcontracting

Not applicable.

Travel

Travel and short-period visits

Kind	*Number*
National meetings	9

We request the approximate funding for about three national meetings attended by three people (a total of 9). This amount includes the planned meetings to which all partners have to attend.

Expenses for inward billing (Costs justified by internal procedure of invoicing)

Not applicable.

Other working costs

We are planning to reserve a server for the length of the project in order to deploy prototypes and experiment with them. The total cost for this amounts to 6624 euros for the three years.

8 Appendices

8.1 References

[1] PlanetLab. *An open platform for developing, deploying and accessing planetary-scale services*. http://www.planet-lab.org/.

[2] Lamia Benmouffok, Jean-Michel Busca, Joan Manuel Marquès, Marc Shapiro, Pierre Sutra, and Georgios Tsoukalas. Telex: A semantic platform for cooperative application development. In *Conf. Française sur les Systèmes d'Exploitation (CFSE)*, Toulouse, France, September 2009.

[3] Kenneth P. Birman, Mark Hayden, Oznur Ozkasap, Zhen Xiao, Mihai Budiu, and Yaron Minsky. Bimodal multicast. *ACM Transactions on Computer Systems*, 17(2):41–88, May 1999.

[4] Barbara Carminati and Elena Ferrari. Access control and privacy in web-based social networks. *International Journal of Web Information Systems - IJWIS*, 4(4):395–415, 2008.

[5] Barbara Carminati, Elena Ferrari, and Andrea Perego. Enforcing access control in web-based social networks. *ACM Transactions on Information and System Security - TISSEC*, 13(1):1–38, 2009.

[6] Michelle Cart and Jean Ferrié. Asynchronous reconciliation based on operational transformation for P2P collaborative environments. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2007*, pages 127–138, White Plains, New York, USA, November 2007. IEEE Computer Society.

[7] Jan G. Cederquist, Ricardo Corin, M. A. C. Dekker, Sandro Etalle, J. I. den Hartog, and Gabriele Lenzini. Audit-based compliance control. *International Journal of Information Security*, 6(2):133–151, 2007.

[8] Monica Chew, Dirk Balfanz, and Ben Laurie. (Under)mining Privacy in Social Networks. In *Proceedings of the Web 2.0 Security and Privacy workshop - W2SP 2008*, Oakland, California, USA, May 2008.

[9] Alan J. Demers, Dan H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing - PODC'87*, pages 1–12, Vancouver, British Columbia, Canada, August 1987.

[10] Patrick Th. Eugster, Sidath B. Handurukande, Rachid Guerraoui, Anne-Marie Kermarrec, and Petr Kouznetsov. Lightweight probabilistic broadcast. *ACM Transaction on Computer Systems*, 21(4), November 2003.

[11] Google Wave. Google Wave: *Communicate and collaborate in real time*. http://wave.google.com/about.html.

[12] Manuel Hilty, David A. Basin, and Alexander Pretschner. On obligations. In *ESORICS*, pages 98–117, 2005.

[13] Abdessamad Imine, Asma Cherif, and Michaël Rusinowitch. A flexible access control model for distributed collaborative editors. In *Secure Data Management*, pages 89–106, 2009.

[14] Mohamed Jawad, Patricia Serrano-Alvarado, and Patrick Valduriez. Protecting data privacy in structured p2p networks. In *Globe*, pages 85–98, 2009.

[15] Márk Jelasity, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten van Steen. The peer sampling service: Experimental evaluation of unstructured gossip-based implementations. volume 52, February 2003.

[16] Márk Jelasity, Alberto Montresor, Gian Paolo Jesi, and Spyros Voulgaris. The PeerSim simulator. `http://peersim.sf.net/`.

[17] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.

[18] Leslie Lamport. Paxos made simple. *ACM SIGACT News (Distributed Computing Column)*, 32(4):18–25, December 2001.

[19] David A. Nichols, Pavel Curtis, Michael Dixon, and John Lamping. High-latency, low-bandwidth windowing in the Jupiter collaboration system. In *Proceedings of the 8th annual ACM symposium on User interface and software technology - UIST '95*, pages 111–120. ACM Press, 1995.

[20] Charles W. O'Donnell and Vinod Vaikuntanathan. Information leak in the chord lookup protocol. *Peer-to-Peer Computing, IEEE International Conference on*, 0:28–35, 2004.

[21] Gérald Oster, Pascal Molli, Sergiu Dumitriu, and Rubèn Mondéjar. Uniwiki: A collaborative p2p system for distributed wiki applications. In IEEE Computer Society, editor, *Proceedings of the 18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises - WETICE 2009*, pages 87–92, Groningen, The Netherlands, June 2009.

[22] Gérald Oster, Pascal Urso, Pascal Molli, and Abdessamad Imine. Data Consistency for P2P Collaborative Editing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2006*, pages 259–267, Banff, Alberta, Canada, November 2006. ACM Press.

[23] Jaehong Park and Ravi Sandhu. The uconabc usage control model. *ACM Transactions on Information and Systems Security*, 7(1):128–174, 2004.

[24] Dean Povey. Optimistic security: A new access control paradigm. In *Proceedings of the 1999 Workshop on New Security Paradigms - NSPW '99*, pages 40–45. ACM Press, 2000.

[25] Nuno Preguiça, Joan Manuel Marquès, Marc Shapiro, and Mihai Lec tia. A commutative replicated data type for cooperative editing. In *Int. Conf. on Distributed Comp. Sys. (ICDCS)*, pages 395–403, Montréal, Canada, June 2009.

[26] Jordi Pujol-Ahulló and Pedro García-López. Planetsim: An extensible simulation tool for peer-to-peer networks and services. In *Proceedings of 9th International Conference on Peer-to-Peer Computing - P2P 2009*, pages 85–86. IEEE Computer Society, September 2009.

[27] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Shenker. A Scalable Content-Addressable Network. In *Proceedings of ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications - SIGCOMM'01*, pages 161–172, New York, NY, USA, August 2001.

[28] ReadWriteWeb. Introduction to the real-time web. `http://www.readwriteweb.com/archives/introduction_to_the_real_time_web.php`.

[29] Antony Rowstron and Peter Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM Intl. Conf. on Distributed Systems Platforms (Middleware)*, pages 329–350, November 2001.

[30] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM'01)*, pages 149–160, San Diego, CA, USA, August 2001. ACM, ACM Press.

[31] Chengzheng Sun and Clarence Ellis. Operational transformation in real-time group editors: issues, algorithms, and achievements. In *Int. Conf. on Computer-Supported Coop. Work (CSCW)*, page 59, Seattle WA, USA, November 1998.

[32] Nicolas Vidot, Michèle Cart, Jean Ferrié, and Maher Suleiman. Copies Convergence in a Distributed Real-Time Collaborative Environment. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2000*, pages 171–180, Philadelphia, Pennsylvania, USA, December 2000. ACM Press.

[33] Spyros Voulgaris, Daniela Gavidia, and Maarten van Steen. CYCLON: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 13(2), 2005.

[34] Quang Hieu Vu, Mihai Lupu, and Beng Chin Ooi. *Peer-to-peer Computing: Principles And Applications*. Springer, 2009.

[35] Dan S. Wallach. A survey of peer-to-peer security issues. In *International Symposium on Software Security*, pages 42–57, 2002.

[36] Stéphane Weiss, Pascal Urso, and Pascal Molli. Logoot : A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks. In *Proceedings of the 29th International Conference on Distributed Computing Systems - ICDCS 2009*, pages 404–412. IEEE Computer Society, June 2009.

8.2 CVs, Resume

8.2.1 Gérald Oster, Associate Professor, Université Nancy 1 / Partner 1. SCORE

Gérald Oster is an Associate Professor at University Nancy 1, France. He received his Ph.D. degree in 2005 from the Department of Computer Science at Nancy-Université, France and he was a postdoctoral researcher for one year at ETH Zurich, Switzerland.
His domain of research is distributed collaborative editing systems with a focus on optimistic replication in peer-to-peer systems. He participated to several research projects such as RNTL LibreSource (2002-2004), INRIA ARC Recall (2006-2007), ANR XWiki Concerto (2006-2009), EU FP6 QualiPSo (2006-2010) and Wiki3.0 (2009-2011).

Selected publications related to the project

[MD09] G. Oster, P. Molli, S. Dumitriu and R. Mondéjar, UniWiki: A Collaborative P2P System for Distributed Wiki Applications. In *Proceedings of the 18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises, WETICE 2009*, pp. 87-92, June 2009.

[IOM07] C.-L. Ignat, G. Oster, P. Molli, M. Cart, J. Ferrié, A.-M. Kermarrec, P. Sutra, M. Shapiro, L. Benmouffok, J.-M. Busca and R. Guerraoui, A Comparison of Optimistic Approaches to Collaborative Editing of Wiki Pages. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2007*, pp. 474-483, November 2007.

[OMU06] G. Oster, P. Molli, P. Urso and A. Imine, Tombstone Transformation Functions for Ensuring Consistency in Collaborative Editing Systems. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2006*, pp. 38-48, November 2006.

[OUM06] G. Oster, P. Urso, P. Molli and A. Imine, Data Consistency for P2P Collaborative Editing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2006*, pp. 259-268, November 2006.

[IROM06] A. Imine, M. Rusinowitch, G. Oster and P. Molli, Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2), pp. 167-183, 2006.

8.2.2 Claudia Ignat, Chargée de Recherche 1, INRIA / Partner 1. SCORE

Claudia-Lavinia Ignat obtained a B.Sc. in Computer Science from the Technical University of Cluj-Napoca, Romania and a PhD in Computer Science from ETH Zurich, Switzerland. She is currently a researcher at INRIA-Nancy Grand Est in France.
Her research area is collaborative editing with a focus on consistency maintenance over different types of documents such as textual, graphical and XML documents as well as awareness approaches in collaborative environments. She is also currently leading research activities on trust and privacy issues in distributed collaborative editing systems. She participated to several research projects such as INRIA ARC Recall (2006-2007), ANR XWiki Concerto (2006-2009) and Wiki 3.0 (2010-2011).

Selected publications related to the project

[IN08] C.-L. Ignat and M. Norrie, Multi-level editing of hierarchical documents. *Journal of Computer Supported Cooperative Work - JCSCW*, 17(5-6), pp. 423-468, December 2008.

[IPO08] C.-L. Ignat and S. Papdopoulou, G. Oster and M. Norrie, Providing Awareness in Multi-synchronous Collaboration Without Compromising Privacy. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2008*, pp. 659-668, November 2008.

[IO08] C.-L. Ignat and G. Oster, Peer-to-peer Collaboration over XML Documents. In *Proceedings of the 5th International Conference on Cooperative Design, Visualization and Engineering - CDVE 2008*, pp. 66-73, September 2008.

[IOM07] C.-L. Ignat, G. Oster, P. Molli, M. Cart, J. Ferrié, A.-M. Kermarrec, P. Sutra, M. Shapiro, L. Benmouffok, J.-M. Busca and R. Guerraoui, A Comparison of Optimistic Approaches to Collaborative Editing of Wiki Pages. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2007*, pp. 474-483, November 2007.

[IN06] C.-L. Ignat and N. Morrie, Draw-Together : Graphical Editor for Collaborative Drawing. In *Proceedings of the ACM International Conference on Computer Supported Cooperative Work - CSCW 2006*, pp. 269-278, November 2006.

8.2.3 Pascal Urso, Associate Professor, Université Nancy 1 / Partner 1. SCORE

Pascal Urso is an Associate Professor at the Université Henri Poincaré (UHP) since September 2004, and works in the LORIA laboratory – team SCORE. He received the PhD degree in computer science from the Université de Nice Sophia Antipolis in 2002. Prior to its recruitment at UHP, he worked as a post-doctoral fellow at the University of Namur (FUNDP). His research interests include distributed systems, data replication, collaborative systems, peer-to-peer computing and automated theorem proving. Pascal Urso participated to several projects including NoE-Interop (2003-2007), INRIA ARC Recall (2006-2007) and EU FP6 QualiPSo (2006-2010).

Selected publications related to the project

[WUM10] S. Weiss, P. Urso and P. Molli. Logoot-Undo: Distributed Collaborative Editing System on P2P Networks. *IEEE Transactions on Parallel and Distributed Systems*, (to appear).

[WUM09] S. Weiss, P. Urso and P. Molli. Logoot: A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks. In *Proceedings of the 29th IEEE International Conference on Distributed Computing Systems - ICDCS 2009*, pp. 404-412, June 2009

[WUM07] S. Weiss, P. Urso and P. Molli. Wooki: A P2P Wiki-Based Collaborative Writing Tool. In *Proceedings of the 8th International Conference on Web Information Systems Engineering - WISE 2007*, pp. 503-512, December 2007.

[OMU06] G. Oster, P. Molli, P. Urso and A. Imine. Tombstone Transformation Functions for Ensuring Consistency in Collaborative Editing Systems. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2006*, November 2006.

[OUM06] G. Oster, P. Urso, P. Molli and A. Imine. Data Consistency for P2P Collaborative Editing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2006* pp. 259-268, November 2006.

8.2.4 Pascal Molli, Associate Professor, HdR, Université Nancy 1 / Partner 1. SCORE

Pascal Molli graduated from Nancy University (France) and received his Ph.D. in Computer Science from Nancy University in 1996. Since 1997, he is Associate Professor at University of Nancy. He participated in the creation of the INRIA ECOO (Environments for Cooperation) project in 1998 and he was vice-head of the INRIA ECOO Team. From October 2009 to current, He is head of the INRIA SCORE team. Pascal Molli has mainly worked on collaborative distributed systems and focused on problems of consistency of shared data in collaborative environments and awareness models for collaborative editing. He participated to several research projects such as RNTL LibreSource (2002-2004), ANR XWiki Concerto (2006-2009) and European FP6 QualiPSo (2006-2010).

Selected publications related to the project

[WUM10] S. Weiss, P. Urso and P. Molli. Logoot-Undo: Distributed Collaborative Editing System on P2P Networks. *IEEE Transactions on Parallel and Distributed Systems*, (to appear).

[WUM09] S. Weiss, P. Urso and P. Molli. Logoot: A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks. In *Proceedings of the 29th IEEE International Conference on Distributed Computing Systems - ICDCS 2009*, pp. 404-412, June 2009

[MD09] G. Oster, P. Molli, S. Dumitriu and R. Mondéjar, UniWiki: A Collaborative P2P System for Distributed Wiki Applications. In *Proceedings of the 18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises, WETICE 2009*, pp.87-92, June 2009.

[IOM07] C.-L. Ignat, G. Oster, P. Molli, M. Cart, J. Ferrié, A.-M. Kermarrec, P. Sutra, M. Shapiro, L. Benmouffok, J.-M. Busca and R. Guerraoui, A Comparison of Optimistic Approaches to Collaborative Editing of Wiki Pages. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2007*, pp. 474-483, November 2007.

[OMU06] G. Oster, P. Molli, P. Urso and A. Imine, Tombstone Transformation Functions for Ensuring Consistency in Collaborative Editing Systems. In *Proceedings of the International Conference on Collaborative Computing: Networking, Applications and Worksharing - CollaborateCom 2006*, pp. 38-48, November 2006.

[OUM06] G. Oster, P. Urso, P. Molli and A. Imine, Data Consistency for P2P Collaborative Editing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work - CSCW 2006*, pp. 259-268, November 2006.

[IROM06] A. Imine, M. Rusinowitch, G. Oster and P. Molli, Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2), pp. 167-183, 2006.

8.2.5 Marc Shapiro, Research Director, HdR, INRIA / Partner 2. REGAL

Marc Shapiro is a Senior Researcher (Directeur de Recherche) at INRIA. He created the SOR research group (Systèmes d'Objets Répartis, Distributed Object Systems), which he led for ten years. During this period he was Principal Investigator for many research projects, both academic and in collaboration with industry. This includes joint research projects with Chorus Systèmes (published at SOSP), Novell, DEC (published at OSDI), Sun Microsystems, Bull, and CNET.
He spent six years and a half at Microsoft Research Cambridge as Senior Researcher, managing the ten-person Cambridge Distributed Systems Group (Camdis). He recruited several young researchers who moved on to become stars of their field, such as Anne-Marie Kermarrec, Antony Rowstron, Miguel Castro, Youssef Hamidi, and Manuel Costa.
He is now a member of the Regal group and an INRIA employee. Some recently-finished collaborations are: a research grant with Microsoft Research Cambridge, FP6 project Grid4All, ANR project Respire and ARC project Recall.
The Grid4All project was a European FP6 project (2005–2009) with academic partners ICCS, INRIA, KTH, SICS, UPC, UPRC, and industrial partners France Télécom R&D and Antares. Its aim was to democratise access to distributed system technologies and to enable large-scale data sharing for collaborative groups. Marc Shapiro chaired its Scientific Committee and led the data-sharing work package.
In 1997–2000, Marc Shapiro was the Principal Investigator of the European Long-Term Research project PerDiS. PerDiS was influenced by real user requirements for multiple, large, complex object databases; a combination of co-operation and isolation; multiple trust and geographical domains. The PerDiS platform was designed for large-scale data sharing based on sophisticated security and memory management algorithms. The PerDiS persistent memory was used for a suite of co-operative CAD applications for the building industry.

8.2.6 Achour Mostefaoui, Associate Professor, HdR, IRISA / Partner 3. ASAP

Achour Mostefaoui is currently Associate Professor at the Computer Science Department of the University of Rennes, France. He received his Engineer Degree in Computer Science in 1990 from the University of Algiers, and a Ph.D. in Computer Science in 1994 from the University of Rennes, France. His research interests include fault-tolerance and synchronization in distributed systems, group communication, data consistency and distributed checkpointing. Achour Mostefaoui has published more than 70 scientific publications and served as a reviewer for more than 20 major journals and conferences.

8.2.7 Marin Bertier, Associate Professor, HdR, IRISA / Partner 3. ASAP

Marin Bertier is an associate professor at the INSA of Rennes and a member of the ASAP. He received his PhD from the Université Pierre et Marie Curie - Paris VI (France) in 2004. His research interests are in large scale distributed systems, wireless sensor networks, and peer-to-peer systems.

8.2.8 Michel Raynal, Professor, HdR, IRISA / Partner 3. ASAP

Michel Raynal is a professor of computer science at the University of Rennes 1 and a member of ASAP. His research interests include distributed algorithms, distributed computing systems, networks and dependability. He has been Principal Investigator of a number of research grants in these areas.

8.2.9 Abdessamad Imine, Associate Professor, Université Nancy 2 / Partner 4. CASSIS

Dr. Abdessamad Imine has been Associate Professor of Computer Science at the University of Oran, Algeria (1995-2002). He got PhD in Computer Science at the University of Nancy 1 in 2006. Since September 2007, he is Associate professor at the University of Nancy 2 and permanent member in the LORIA-INRIA CASSIS project.
He is working on access control policies for distributed editors, safe updating strategies for XML-based firewall policies and consistency verification of collaborative systems.

Selected publications related to the project

[ICR09] A. Imine, A. Cherif and M. Rusinowitch, A Flexible Access Control Model for Distributed Collaborative Editors. In *Proceedings of the 6th VLDB Workshop, Secure Data Management - SDM 2009*, pp. 89-106, August 2009.

[Imi09] A. Imine, Coordination Model for Real-Time Collaborative Editors. In *Proceedings of the 11th International Conference, COORDINATION 2009* pp. 225-246, June 2009.

[BI09] H. Boucheneb and A. Imine, On Model-Checking Optimistic Replication Algorithms. In *Proceedings of the 29th IFIP WG 6.1 International Conference - FORTE 2009*, pp. 73-89, June 2009.

[IROM06] A. Imine, M. Rusinowitch, G. Oster and P. Molli, Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2), pp. 167-183, 2006.

[IMOR03] A. Imine, P. Molli, G. Oster, and M. Rusinowitch, Proving correctness of transformation functions in real-time groupware. In *Proceedings of the 8th European Conference on Computer-Supported Cooperative Work - ECSCW 2003*, pp. 277-293, September 2003.

8.2.10 Michaël Rusinowitch, Research Director, HdR, INRIA / Partner 4. CASSIS

Michael Rusinowitch received a Thèse d'Etat in Computer Science in 1987. Since 1994 he is Directeur de Recherche at INRIA. His research is mainly concerned with security analysis, theorem-proving, term-rewriting, and their application to software verification. He contributed to the development of automated deduction with constraints, to new proof methods based on induction and rewriting and to the verification of security protocols. He has published his works in more than 70 international conferences and 28 journal papers, and is the author of a book on automated deduction. He has been in program committees of many international conferences and co-chairman of the conference on Rewriting Techniques and Applications, and International Joint Conference on Automated Reasoning in 2004. Former founder and head of the CASSIS group at INRIA Grand - Est since 2003, he has been responsible of several national or international research projects (RNTL, ACI, ARC, PAI, NSF-INRIA, NSF-CNRS). He has participated to many recruiting committees at INRIA, Institut National Polytechnique de Lorraine, and University of Nancy 2. He has also participated to evaluation committees for CNRS, France Telecom R&D, and has been reviewer for Swedish Research Council for Engineering Sciences, Science Foundation Ireland, Swiss National Science Foundation, Conseil de Recherches en Sciences Naturelles et Génie du Canada. He is member of IFIP Working Group 1.6 (Rewriting) and has been Invited Editor of Journal of Symbolic Computation, Theory of Computing Systems, Journal of Automated Reasoning, Information and Computation, Technique et Science Informatiques. He has given invited conferences at RTA 2001, LPAR 2000 and IFIP Theoretical Computer Science (World Computer Congress), 2004.

He has been Supervisor or Co-supervisor of 14 PhD. theses (10 of the former students currently have permanent positions in academic research institutions) and of many post-graduate students. Lecturer of international tutorials and master courses.

Selected publications related to the project

[CTR09] N. Chridi, M. Turuani and M. Rusinowitch, Decidable Analysis for a Class of Cryptographic Group Protocols with Unbounded Lists. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium, CSF 2009*, pp. 277-289, July 2009.

[ICR09] A. Imine, A. Cherif and M. Rusinowitch, A Flexible Access Control Model for Distributed Collaborative Editors. In *Proceedings of the 6th VLDB Workshop, Secure Data Management - SDM 2009*, pp. 89-106, August 2009.

[IROM06] A. Imine, M. Rusinowitch, G. Oster and P. Molli, Formal design and verification of operational transformation algorithms for copies convergence. *Theoretical Computer Science*, 351(2), pp. 167-183, 2006.

[CR08] Y. Chevalier and M. Rusinowitch, Hierarchical combination of intruder theories. *Information and Computation*, 206(2-4), pp. 352-377, 2008.

[CKRT08] Y. Chevalier, R. Kuesters, M. Rusinowitch and M. Turuani, Complexity results for security protocols with Diffie-Hellman exponentiation and commuting public key encryption, *ACM Transactions on Computational Logic - TOCL*, 9(4), pp. 1-52, 2008.

8.2.11 Ludovic Dubost / Partner 5. XWIKI

Ludovic Dubost is the XWiki founder and CEO. He studied at the Polytechnique (X90) and Telecom Paris, and he started his career in 1995 at Cap Gemini. In 1996 he joined Netscape where he was the intranet solution architect. He did consulting for bug corporations such as Renault and Airbus for deploying Netscape intranet solutions. In 1998 he joined NetValue, one of the first French start-up that accessed the stock exchange, as chief officer of development and then as CTO. He left NetValue in 2004 after it was bought by Nielsen/NetRatings to found XWiki.
With more than 15 years of experience in web technologies and in software services at Cap Gemini, Netscape and NetValue, Ludovic Dubost is the guarantee of the project methodology and the quality of service of the XWiki SAS company.

8.2.12 Vincent Massol / Partner 5. XWIKI

Vincent Massol is CTO of XWiki SAS, the company who created and is sponsoring the development of XWiki, an open source collaboration platform, based on a second generation wiki.
With 7 years of open source done mostly during his free time (pioneer committer on Apache Maven since 2002, creator of Jakarta Cactus in 1999 and Codehaus Cargo in 2004), Vincent is currently fulfilling his passion for open source during his day time by coordinating the open source XWiki community and doing actual coding for the XWiki project.
Vincent is the author of best-seller “JUnit in Action” from Manning, “Maven: A Developer’s Notebook” from O’Reilly and co-author of “Better Builds with Maven” by DevZuz.

8.2.13 Fabio Mancinelli / Partner 5. XWIKI

Fabio Mancinelli received a PhD in Computer Science in 2006, he is currently the Research Project Manager at XWiki SAS and also collaborates to the development of the XWiki platform. His current activities are the management of all the research projects where XWiki is involved in, and his main interests are collaborative technologies, semantic web and web-oriented architectures.

8.2.14 Sergiu Dumitriu / Partner 5. XWIKI

Free Software advocate, Sergiu Dumitriu spent several years as a contributor to the XWiki open source project, before joining XWiki SAS as a Research and Development engineer. He is currently pursuing a PhD thesis in Computer Science at LORIA since 2008 under a CIFRE contract in partnership with XWiki SAS. His main interests include web technologies, distributed systems and collaborative environments.

8.2.15 Thomas Eveilleau / Partner 5. XWIKI

Thomas Eveilleau is project manager at XWiki since 2006. After a DUT "Services et Réseaux de Communication" at the University of Paris Est and a master in "Digital Communication" at the HETIC (Hautes Etudes des Technologies de l'Information et de la Communication) he joined XWiki as project manager. He handles client projects and also collaborates in the development of solution for research projects with respect to the integration aspects.

8.2.16 Jean-Vincent Drean / Partner 5. XWIKI

Jean-Vincent Drean graduated at EPITECH. He joined XWiki in 2006 as a research and development engineer after working for other integrator companies such as IDEALX and LINBOX. He worked on different client projects and his main activity now is concentrated on the development and the evolution of the main product.

8.2.17 Thomas Mortagne / Partner 5. XWIKI

Thomas Mortagne graduate at EPITECH in 2006. After 1 year as a consultant for ALTEN he started working for XWiki SAS as core programmer on the XWiki Platform. His interests are object oriented development frameworks and web technologies.

8.2.18 Jérôme Velociter / Partner 5. XWIKI

Jerome Velociter graduated from Telecom Lille 1 in 2007, with a specialty in International Business Engineering attended at Institut National des Telecommunications (INT). He joined XWiki SAS early in 2007 as a software engineer, working on XWiki products and customer projects. His Open Source software activities include being an active committer on the XWiki.org project and contributor to the OpenRemote.org project.

8.3 PhD. and Post-doctoral Topics

8.3.1 PhD. about Efficient Consistency and Replication

Replication and Consistency Maintenance for Real-time Editing in Peer-to-Peer Social Applications

Context

Real-time social web applications are now central to our economies and every day life.

> *Real-time information delivery is fast emerging as one of the most important elements of our online experience. No more waiting for the Pony Express to deliver a parcel cross-country, no more waiting for web services to communicate from one polling instance to another. This is information being available to you at nearly the moment it's produced, whether you're watching for it or not.*
>
> **Source:** Introduction to the Real-Time Web, ReadWriteWeb[28]

Currently, all real-time social applications rely on a centralised architecture or on *clouds*. In some way *clouds* infrastructures give a solution to scalability issues but their support for collaborative editing of the same data remain very limited. Moreover privacy concerns on *clouds* are growing. Data that users have to put in the *cloud* are potentially locked out. Our objectives is to deploy real-time social applications on open peer-to-peer networks.

Objectives

The goal of this PhD thesis is to build efficient replication and consistency maintenance algorithms for social real-time web applications. To achieve this objective we propose the following work plan:

- Literature survey on optimistic replication, with a focus on Commutative Replicated Data Types (CRDTs) and on the Operational Transformation (OT) approach.
- Collecting of requirements from inputs of industrial actor such as our XWiki partner but also from sources available on the web. The PhD student should expect to collect executions traces from public platform such as GoogleWave and GoogleBuzz.
- Design of novel consistency algorithms. These algorithms should be suitable regarding the peer-to-peer constraints (no central coordinator, huge amount of participants, large amount of data and updates of these data) but also the real-time aspect of the context (updates might be of different granularity, algorithm must permit a high-responsiveness of the system).
- Theoretical analysis of the proposed algorithms.
- Implementation of the proposed algorithms.
- Experimentation on different testbeds (Grid'5000, PlanetLab, etc.) with traces illustrating real use cases.

8.3.2 PhD. about Security and Privacy

Towards Controlling Access and Usage of Replicated Social Objects

Context
The last few years have witnessed the explosion of online social networks users. They provide computer support to make available an information space where each social network participant can publish and share data, called social object, such texts, videos, annotations, blogs, and, generically resources for a variety of purposes. In order to improve performance and availability of data in such a distributed context, each user may own a local copy of the shared social objects, upon which he may perform updates. Locally executed updates are then transmitted to the other users sharing the same objects. For instance, Google Wave – an online tool for real-time communication and collaboration – enables users to share their documents and to create conversation sessions within these documents..
One of the most challenging problems in online social networks is how to balance the computing goals of collaboration and access control to shared social objects. Indeed, interaction is aimed at making shared objects available to all who need it, whereas access control seeks to ensure this availability only to users with proper authorization. Moreover, it is clear that data replication improves high responsiveness of local updates. However adding standard access control mechanisms to social networks may degrade responsiveness since updates must be granted by some authorization coming from a distant user (such as a central server). Furthermore, due to its dynamic characteristic, social networks have to allow for dynamic change of access permissions since users may join and leave the group in an ad-hoc manner and other shared social objects may be created and destroyed.
Usage control refers to what happens to data after it was granted to authorised persons. In the case of online social networks users are granted access to a social object together with some obligations that they need to satisfy such as do not distributing the object to other users.

Objectives
The goal of this PhD thesis is to build reliable access and usage control schemes that are well suited for online social networks. To do that, we would like to achieve the following objectives:

- Analyzing the current and the future use of access control in social networks. This step enables us to propose new models that will be well adapted to social networks' access control.
- Defining a formal model for specifying and verifying access control policies. Two description languages could be used: Access Control Lists (ACL) and XACML.
- Designing an environment for reconciling replicated social objects with dynamic access control policies.
- Combining the access control approach built in the previous steps with the usage control mechanism proposed in the context of the STREAMS project.
- Testing the combined access and usage control mechanisms using simulators and traces of human collaboration obtained from various social networks.

8.3.3 Post-doctoral Position about Reconciliation of Structured Documents

Reconciliation of structured documents in real-time peer-to-peer environments

Context
Nowadays we are faced with an explosion of real-time and social software such as wikis, blog, micro-blogs, and social networks services (Facebook, LinkedIn, MySpace, etc.). Real-time information delivery is emerging as one of the most important elements of our online experience. One of the representative social software offering the real-time feature is Google Wave, the new communication service developed by Google, that defines the wave as being equal parts conversation and document, where people can communicate and work together with richly formatted text, photos, videos and maps. All these social software are based on a centralised architecture. Besides issues of limited scalability, lack of shared administration costs and limited fault tolerance, data centralisation in the hands of a single vendor is an inherent threat to privacy. Our vision is to move away from centralised authority-based collaboration towards a distributed trust network where each user can decide with whom to share their data. Users define their network of trust containing people that they trust and with whom they wish to collaborate. Peer-to-peer networks support very well this model of collaboration. Structured model such as XML is a general model of representation of collaboration in social software.
In order to provide a high responsiveness, data availability, reliability, scalability and fault tolerance, data is replicated. Optimistic replication approaches are used to maintain consistency of the different copies of the same data. These approaches tolerate divergence between copies after they have been modified and ensure that they will eventually converge.

Objectives
The goal of this research project is to propose an efficient algorithm for the reconciliation of structured documents in real-time peer-to-peer environments. Concurrent edits in the structured documents might generate conflicts and cannot guarantee to maintain a correct structure. The reconciliation mechanism must detect and handle conflicts. The following steps are required:

- Study existing algorithms for maintaining consistency of linear structured documents in real-time collaboration in peer-to-peer networks (e.g. operational transformation and CRDT algorithms)
- Propose a novel efficient approach for the reconciliation of XML documents
- Define the notion of conflict in structured documents and extend the reconciliation approach to deal with conflicts
- Test the reconciliation mechanism in peer-to-peer network simulators and using traces of human produced collaboration

8.3.4 Post-doctoral Position about Peer-to-Peer Infrastructure

Fine-grain and fast updates in peer-to-peer systems

Context
One of the objectives of this project is to design peer-to-peer solutions that offer underlying services required by real-time social web applications and that eliminate the disadvantages of centralised architectures. Unfortunately, all existing real-time social web services rely on a central authority. Moreover, centralisation of the platforms hosting these services makes their scalability and reliability very costly and limit their deployment by customers for their own needs. Decentralised alternatives do not exist due to the state of art of existing collaborative algorithms and peer-to-peer architectures and protocols. Indeed, the current state of the art in peer-to-peer networks does not address replication of continuously updated content due to real-time user changes. The STREAMS project aims to propose generic algorithms that work for different granularities for the transmitted changes. Suitable algorithms for maintaining consistency over structured data types in real-time collaboration need to be designed.

Objectives
The hired post-doc fellow will participate mainly in Task 2 (Requirements and Infrastructure). The goal of this post-doc is, on the one hand, to serve as an intermediate between the industrial partner and the other partners. On the other hand, downward, he has to define the requirements for the experimental plateform and then to design new algorithms for dealing with the nature of the targeted applications

- Study existing peer-to-peer structures and compare them how they comply with real-time collaborative editing of linear structured documents.
- Study the requirements of the industrial partner in order to integrate them as comparison criteria in the above mentioned item.
- Propose new and well-suited data replication techniques among peers along with an adapted gossip-based update service.
- Participate to design the of consistency mechanisms and their integration into the plateform.