

Computing over widely-replicated data in a hybrid cloud

Mots clés :

- **Directeur de thèse** : Marc Shapiro
- **Co-encadrant(s)** :
- **Unité de recherche** : Laboratoire d'informatique de Paris 6
- **Ecole doctorale** : École Doctorale Informatique, Télécommunications, Électronique de Paris
- **Domaine scientifique principal**: Divers

Résumé du projet de recherche (Langue 1)

The Regal group of INRIA and LIP6 studies distributed systems and operating systems. Our aim is to design systems that are better by some pragmatic, objective metric: e.g., scalability, response time, throughput, fault tolerance, etc. This requires studying algorithms, understanding their bottlenecks, and improving their design. For instance, we recently showed how to completely remove the consistency bottleneck (in some restricted cases) with the concept of a Conflict-Free Replicated Data Type (CRDT) [SSS 2011]. Cloud computing platforms are evolving towards a hybrid, so-called “fog” model. On the one hand, modern protocols support strong consistency and transactions at the scale of a single or geo-replicated data centres. On the other hand, data is moving outside of data centres, using computing and storage resources near the edge. This enables a wide range of options: an application at the edge can be highly responsive and available, but consistency and fault-tolerance are hard; conversely, a data centre has greater, more elastic computing resources and can more guarantee strong consistency and integrity.

Résumé du projet de recherche (Langue 2)

We propose to study the possibilities and trade-offs joining these two worlds, from a number of perspectives: system, language, protocols, security, etc. Here are some possible topics: -* Replication and consistency: exploring the spectrum between strong consistency at the centre and eventual consistency at the edge; leveraging application semantics; understanding the responsiveness-availability vs. consistency trade-offs; how this translates for the application semantics (“isolation levels”). -* Synchronisation-free, scalable mechanisms for maintaining atomicity, consistent snapshots, causal ordering, concurrency detection, transactions and partial replication. -* Computation models for widely replicated big data, beyond MapReduce. Designing data types for replication and sharding. Incrementally propagate updates that happen at the edge (e.g., along a data flow graph) to all replicas and to downstream results. -* Securing replicated data: ensuring integrity, confidentiality, and information flow of widely replicated data; securing clients against one another.

Informations complémentaires (Langue 1)

-* Mandatory: One or two three-month internships during the the PhD. -* Publication in international conferences. -* Participation in EU projects.

Informations complémentaires (Langue 2)

Applicants should have an excellent academic record, be strongly interested in of distributed algorithms and systems, and have good programming and experimental skills. Previous experience with distributed programming in Java is an advantage. Please provide a CV, the list of Masters or PhD courses and your marks, an essay relevant to the topic (1 to 4 pages), and at least two references (whom we will contact ourselves for a recommendation).