



Project no. 609551
Project acronym: SyncFree
Project title: *Large-scale computation without synchronisation*

European Seventh Framework Programme ICT call 10

Deliverable reference number and title: D.6.3
Year 2 Dissemination Report
Due date of deliverable: September 30, 2015
Actual submission date: February 12, 2016
Start date of project: October 1, 2013
Duration: 36 months
Name and organisation of lead editor
for this deliverable: INRIA
Revision: 2
Dissemination level: PU

Contents

1	Dissemination and use of the knowledge	1
1.1	Overview of dissemination activities	1
1.1.1	Publications	1
1.1.2	Events sponsored by SyncFree	2
1.1.3	Presentations	2
1.1.4	Other dissemination	3
1.2	Dissemination activities per WP	4
1.2.1	WP1	4
1.2.2	WP2	4
1.2.3	WP3	4
1.2.4	WP4	5
1.2.5	WP5	6
1.2.6	WP6	6
2	Improvements to the project Web Page	6
	Appendix I: Publications in the second year	8
	Appendix II: Presentations/Talks in the second year	11

About this document

This is a revised version of D6.3 scheduled at M24. The changes concern the second part of the document which, in the original version, proposed a very brief overview of the exploitation activities. The exploitation strategy is now object of a separate deliverable (D6.5). In this new version, Section 2 deals with the improvements applied to the project web page as suggested by the experts during the second review meeting.

1 Dissemination and use of the knowledge

Similarly to the first year, the SyncFree consortium has done active dissemination across many areas of academia and industry. This includes numerous quality conferences, workshops, and other public events. The goal of this document is to provide an exhaustive documentation of these activities.

1.1 Overview of dissemination activities

Continuing the trend of the first year, the project has forward an ambitious dissemination plan, continuously involving the SyncFree members.

Thanks to the pro-activity of its researchers, SyncFree performed in publications, in sponsored events and in presentations. These activities have been integrated with several other more informal activities aimed at the maximization of knowledge sharing.

1.1.1 Publications

The list of publications covering the second year is reported in Appendix I. Publications remain the main medium for dissemination. Our members have worked hard in order to ensure the SyncFree outcomes are published in top-ranked events. The project targets two kinds of conferences: top-level academic conferences and industrial developer conferences.

Note that, in the project area, conferences are the most prestigious venues, and journals remain generally secondary. The consortium has also taken into consideration magazines aimed at a wide audience. Top-level publications entail peer-review processes with high rejection rates, implying that an immediate dissemination is not always possible. During the second year, SyncFree has achieved an excellent research record publishing a paper in a relevant journal (ACM SIGOPS Operating Systems Review) and another one in a prestigious magazine (EU Research) as well as several papers in well-known conferences and scientific events, including MIDDLEWARE 2015, SYSTOR 2015, SRDS 2015, EuroSys 2015, NETYS'15, CloudCom 2015, PaPoC 2015, W-PSDS 2015, DAIS 2015, PPDP 2015, 4th ACM SIGPLAN Erlang Workshop, W-PSDS 2014, Middleware 2014 and OPODIS 2014.

1.1.2 Events sponsored by SyncFree

As reported in D6.2, the project has sponsored in its first year several relevant international events. Likewise, in its second year SyncFree has sponsored several international workshops, to gather both industry and academic experts in the field of consistency in distributed systems, together with the SyncFree researchers to discuss and share knowledge:

- *3rd Workshop on Planetary-Scale Distributed Systems (W-PSDS 2015)*. This workshop is co-located with SRDS 2015 (top ranked conference in the area of distributed computing), 28 September - 1 October, 2015, in Montreal (Canada). The workshop comes from a simple empirical observation about distributed systems that have gained a more prominent role in the everyday life of many people. Additionally, the scale of these systems has also been increasing over the past decade, and systems that interconnect users at a planetary scale are now common place. Users expect these systems to be robust, highly available, safe, and efficient. This leads to the emergence of significant challenges for the scientific community to find solutions that provide all these characteristics. The goal of the workshop is to bring researchers and practitioners from the large-scale distributed systems communities to discuss the current state of the art, emerging challenges and trends, as well as novel solutions, implementation and deployment of large scale, and in particular of planetary-scale, distributed systems and applications.
- *Workshop on Principles and Practice of Consistency for Distributed Data (Pa-PoC 2015)*. It is co-located with EuroSys 2015 (another top level conference in the field of systems), 21-24 April, 2015, in Bordeaux (France). It is the direct successor of the EuroSys 2014 Workshop on Principles and Practice of Eventual Consistency, which brought together approximately 40 researchers and practitioners in the areas of distributed systems, programming languages, databases and concurrent programming. This workshop aims to investigate the principles and practice of weak consistency models for large-scale, distributed shared data systems. It brings together theoreticians and practitioners from different horizons: system development, distributed algorithms, concurrency, fault tolerance, databases, language and verification, including both academia and industry.

Other events sponsored by SyncFree are:

- *Workshop on Highly-Scalable Distributed Systems*, January 2015, Paris
- *Workshop on DB consistency in the cloud*, September 2014, Paris.

1.1.3 Presentations

The members of the consortium have been extremely active at academic or industry events giving a number of presentations. The full list is available in Appendix II, and includes:

- **Keynotes.** During the second year, SyncFree members have presented three keynotes (Workshop on Large-Scale Distributed Systems and Middleware 2014, W-PSDS 2014 and MASCOTS 2014). They deal, respectively, with CRDTs in Riak 2.0, Causal Consistency for client side applications, and the trade-off between performance and programmability in large scale distributed data store.
- **Tutorials.** Marc Shapiro and Nuno Preguiça have presented a tutorial about consistency, *All About Consistency: getting it right*, at CodeMesh 2015. The tutorial is aimed at system/application developers and focuses on CRDTs in practice.
- **Conference/Workshop talks.** Researchers involved in SyncFree have assured a significant presence throughout the whole year in many worldwide recognised events, including international conferences and workshops. Technical presentations have addressed the most relevant aspects of the project research and have been presented in both industry and academic venues.
- **Posters.** Besides, a poster about coordination-free cloud storage has been proposed at the 25th ACM Symposium on Operating Systems Principles (SOSP '15).
- **Other talks or discussions.** As an ideal complement to the presentations cited above, a number of technical presentations have been done also in smaller or more informal events, such as summer schools, seminars, local workshops and talks in universities.

1.1.4 Other dissemination

- **Open source.** Aiming at the release of high-quality open-source libraries along with tutorials and example programs doing large-scale calculations, results from the SyncFree project are openly available on-line on GitHub. Those repositories are playing a double role, facilitating capabilities for internal co-operation as well as maximizing the potential visibility of results.
- **Project web page** (<https://syncfree.lip6.fr/>). The project web page has been improved and extended during the second year (see Section 2) in the attempt to consider the project outcomes in a wider context suitable to an external audience.
- **Social Networks, Blogs and other media channels.** A website has been launched (<http://lasp-lang.org/>) for LASP. Furthermore, aiming at an active discussion on the SyncFree results, the consortium members are very active also on social networks (e.g. Twitter) and in well-known blogs. Some relevant posts in blogs are listed below:
 - Lasp: A language for distributed, coordination-free programming. The Morning Paper, blog by Adrian Colyer (Aug. 17, 2015) <http://blog.acolyer.org/2015/08/17/>

- Lasp.
Another Word for It, blog by Patrick Durusau (Aug. 1, 2015)
<http://tm.durusau.net/?cat=2055>
- Lasp: A language for distributed, coordination-free programming.
The Trendy Things, blog (Aug. 17, 2015)
<http://thetrendythings.com>

1.2 Dissemination activities per WP

According to their purpose and scope, the different WPs have contributed, directly or indirectly, to the dissemination activities. A brief overview is provided in the next subsections.

1.2.1 WP1

As a part of WP1, the report "D1.2 Formal Mathematical Requirements" has been delivered, providing a more precise description of the use-cases described in D1.1 in terms of invariants and mathematical definitions, as well as formal description. Therefore, the WP (ended at M18) has provided the consolidation of use cases through the identification and the formalization of key requirements. Papers describing the main outcomes are expected in the next period.

1.2.2 WP2

The results and dissemination efforts of WP2 fall in two categories:

- the design of novel data types with CRDT characteristics
- and research on highly-scalable data stores and protocols supporting CRDTs.

A number of papers have been published or accepted for publication based on the results of WP2. They have been or will be presented by their authors at the respective conferences and workshops.

Papers on implementations of CRDTs have been published at NETSYS'15, PaPoC'15 and APLAS 2015, as well as papers on the design of CRDT data stores at Middleware 2014 and PaPoC'15.

WP2 has provided efforts to maximize the dissemination of results through significant talks in the context of important events. An example is the presentation by the WP leader at the UPMARC Workshop which is organized by the Uppsala Programming for Multicore Architectures Research Center and is an academic/industrial workshop on the analysis of programs running on systems that do not provide sequential consistency.

1.2.3 WP3

During this period, people involved in WP3 have been able to publish and present in conferences the following results from Tasks 3.1 and 3.2 (some joint with WPs 2 and 4):

- A new model for synchronizing state-based CRDTs, by propagating deltas instead of full state was presented at NETYS 2015. This is an extended version of the work presented in PaPEC 2014 workshop. A C++ implementation is available at <https://github.com/CBaquero/delta-enabled-crdts>.
- CRDTs that maintain numeric invariants will be presented in SRDS 2015. A generic mechanism for transferring information exactly-once has been presented in W-PSDS.
- A new consistency model to maintains invariants across CRDTs, explicit consistency, and the design of a system that enforces such consistency model has been presented in a vision paper at LADIS 2014 and selected for publication at OSR in Jan 2015 and later as a complete paper at EuroSys 2015. This work has also been presented as part of a lecture in the Portugal—UT Austin Summer school in distributed computing.
- We have presented two techniques for improving the execution of transactions have been presented at PaPoC'15 (presenting an algorithm to minimize conflicts in transactions) and DAIS'15 (proposing a new technique for maintaining information for tracking causality).

Two works have already been accepted for publication and will be presented during the next period:

- Partial replication of CRDTs at CloudCom 2015.
- A solution for accessing causally consistent information in clients that supports the need for data center hand-over will be presented in Middleware 2015.

A number of other works are under submission and we will continue to actively try to disseminate project's results in scientific venues.

1.2.4 WP4

Dissemination activities for WP4 during Y2 focused on two main new results, the Lasp programming model and the CISE proof system. For Lasp, we published three papers in international workshops and one in an international conference (PPDP 2015). In addition, we launched a website, lasp-lang.org, and information on Lasp was posted on several widely disseminated technical blogs. We gave talks at four European research institutions (KTH, IST, UCL, UPMC) and three professional software development conferences (Erlang Factory SF Bay, GOTO Chicago, Strange Loop). We also published an article in InfoQ, a well-known Web magazine on professional software development. We officially visited one company, 360T Treasury Systems, to discuss system architecture based on the Lasp model. For CISE, we published a paper at POPL 2016. We gave talks at IMDEA, Technion, and Amazon, at the CurryOn conference (colocated with ECOOP 2015), and three workshops, Formal Reasoning in Distributed Algorithms (FRIDA 2015), Imperial Concurrency Workshop 2015, and Dagstuhl Seminar 15191. In Y3 we expect to continue developing and disseminating both Lasp and CISE, and we expect also to disseminate the work on verification of Antidote programs.

1.2.5 WP5

WP5 is still in a preliminary phase where evaluations and experiments are being planned. Indeed, dissemination activities are expected in later phases of the project.

1.2.6 WP6

The project leader (Marc Shapiro) has been extremely active throughout the whole year ensuring dissemination. Those activities have been developed according to two convergent main axis aimed at a horizontal dissemination (the project as a whole) and at a vertical dissemination (distinguished outcomes from the project). He has directly contributed to papers that have been presented at top-level conferences (such as EUROSYS 2015, SRDS 2015, SYSTOR 2015 and MIDDLEWARE 2015) and, indirectly, to many others. Furthermore he published an interview-style paper in EU Research. During the second year, he has presented a keynote and a tutorial, as well as a significant number of presentations and talks in relevant venues and events.

2 Improvements to the project Web Page

The SyncFree website (<https://syncfree.lip6.fr>) has been improved in order to facilitate the dissemination of the results achieved as well as to better address an external audience. The SyncFree homepage is showed in Figure 1.

The changes with respect to the previous version can be summarized as follows:

The content for internal purposes forward to the intranet , i.e. content not interesting/suitable for an external audience. It should help the visitor to focus on the content really relevant to understand SyncFree and its outcomes.

The existing content has been reorganized to facilitate the understanding of the project scope, purpose and achievements. Key contributions, such as main publications and talks, have been highlighted to provide a further guide to the visitors.

Specific content describing the main outcomes of the project has been added both with the links to the software repositories (where applicable). We trust this overview is relevant for the exploitation and we plan to progressively extend those sub-pages providing further information and details.

Integration of specific content to address a large audience . Multimedia content has been added to provide a straightforward path towards SyncFree. That content is integrated with specific sub-pages that should exhaustively answer FAQ on CRDTs and on SyncFree.

Generic improvements including, among others, visit counter, more visible contact details and improvements to the website look.

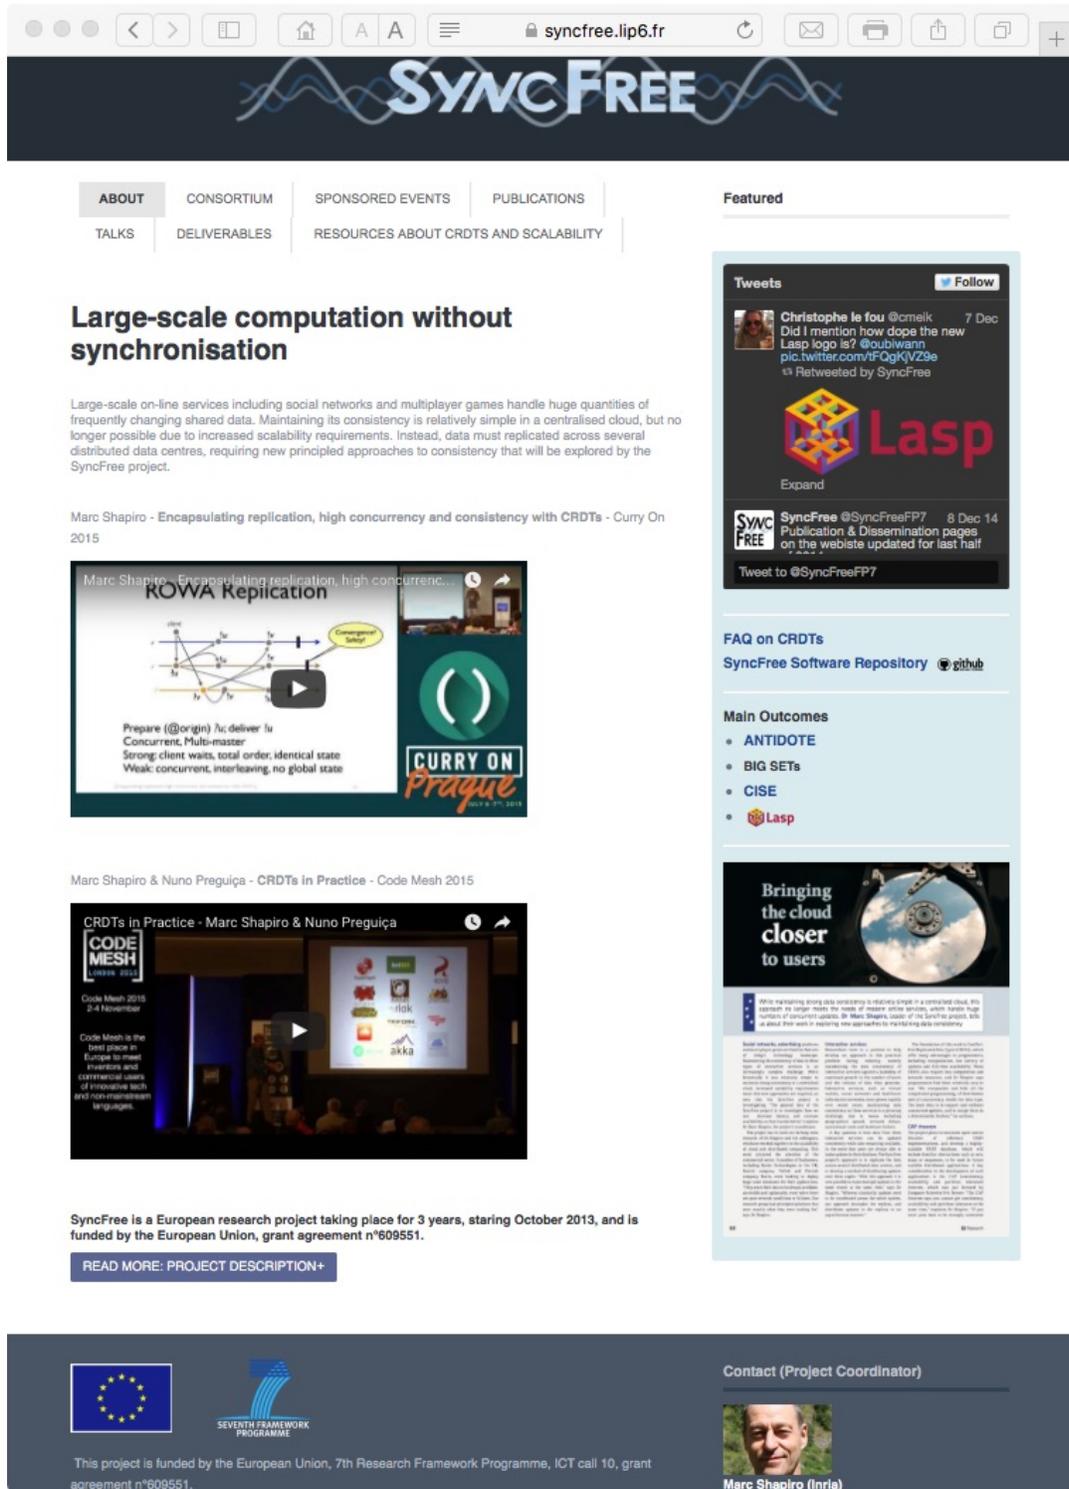


Figure 1: SyncFree website.

Appendix I: Publications in the second year

Journal/Magazine Papers

- Valter Balegas, Sergio Duarte, Carla Ferreira, Rodrigo Rodrigues, Nuno Preguiça, Mahsa Najafzadeh and Marc Shapiro. *Towards Fast Invariant Preservation in Geo-replicated Systems*. ACM SIGOPS Operating Systems Review, Special Issue on Repeatability and Sharing of Experimental Artifacts. Volume 49 Issue 1, January 2015. Pages 121-125.
- Marc Shapiro, *Bringing the cloud closer to users*, EU Research, September 2015, pp.68-69.

Conference Publications

- Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, Marc Shapiro, *Cause Im Strong Enough: Reasoning about Consistency Choices in Distributed Systems*, POPL 2016.
- Ricardo Gonçalves and Paulo Sérgio Almeida and Carlos Baquero Moreno and Vitor Fonte, *Concise Server-Wide Causality Management for Eventually Consistent Data Stores*, Proceedings of the 15th International Conference on Distributed Applications and Interoperable Systems (DAIS'15).
- Manuel Bravo and Paolo Romano and Luis Rodrigues and Peter Van Roy, *Reducing the Vulnerability Window in Distributed Transaction Protocols*, PaPoC 2015.
- Iwan Briquemont, Manuel Bravo, Zhongmiao Li, and Peter Van Roy. *Conflict-free Partially Replicated Data Types*. In Proceedings of the 7th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2015). IEEE, Nov 2015.
- David Navalho, Sergio Duarte, Nuno Preguiça. *A Study of CRDTs that do Computations*. In Proc. Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC'15), 2015. ACM.
- Paulo Sergio Almeida, Ali Shoker, and Carlos Baquero. *Efficient State-based CRDTs by Delta-Mutation*. In the proceedings of the International Conference of Networked sYSTEMS (NETYS'15) , Agadir, Morocco, May 2015.
- Vinh Tao, Marc Shapiro, Vianney Rancurel. *Merging semantics for conflict updates in geo-distributed file systems*. Proceedings of the 8th ACM International Systems and Storage Conference (SYSTOR '15).
- C. Meiklejohn. *Highly Distributed Computations Without Synchronization*. InfoQ, Feb. 17, 2015.
- C. Meiklejohn and P. Van Roy. *The Implementation and Use of a Generic Dataflow Behaviour in Erlang*. 4th ACM SIGPLAN Erlang Workshop, Vancouver, BC, Sep. 4, 2015.

- Valter Balegas, Diogo Serra, Sergio Duarte, Carla Ferreira, Marc Shapiro, Rodrigo Rodrigues and Nuno Preguiça. *Extending Eventually Consistent Cloud Databases for Enforcing Numeric Invariants*. 34th International Symposium on Reliable Distributed Systems (SRDS 2015), Montreal, Quebec, Sep. 29, 2015.
- Alejandro Zlatko Tomsic, Marc Shapiro and Tyler Crain, *Scaling geo-distributed databases to the MEC ecosystem*. Workshop on Planetary-Scale Distributed Systems, September 28, 2015. Montreal, Canada.
- Christopher Meiklejohn and Peter Van Roy. *Selective Hearing: An Approach To Distributed, Eventually Consistent Edge Computation*. Workshop on Planetary-Scale Distributed Systems. September 28, 2015. Montreal, Canada.
- Ali Shoker, Paulo Sergio Almeida and Carlos Baquero. *Exactly-Once Quantity Transfer*. Workshop on Planetary-Scale Distributed Systems, September 28, 2015. Montreal, Canada.
- Valter Balegas, Nuno Preguiça, Rodrigo Rodrigues, Sergio Duarte and Carla Ferreira, *Designing Concurrency-aware Geo-Replicated systems*. Workshop on Planetary-Scale Distributed Systems, September 28, 2015. Montreal, Canada.
- Valter Balegas, Sergio Duarte, Carla Ferreira, Rodrigo Rodrigues, Nuno Preguiça, Mahsa Najafzadeh and Marc Shapiro. *Putting Consistency back into Eventual Consistency*. Proceedings of the Tenth European Conference on Computer Systems(EuroSys 2015).
- Ricardo Gonçalves, Paulo Sergio Almeida, Carlos Baquero, Victor Fonte. *Concise Server-Wide Causality Management for Eventually Consistent Data Stores*. 15th IFIP WG 6.1 International Conference, DAIS 2015, Grenoble, France, June 2-4, 2015.
- Christopher Meiklejohn and Peter Van Roy. *Lasp: A Language for Distributed, Coordination-Free Programming*. In 17th International Symposium on Principles and Practice of Declarative Programming (PPDP 2015). July 14-16, 2015, Siena, Italy.
- V. Balegas, N. Preguiça, S. Duarte, C. Ferreira, R. Rodrigues, M. Najafzadeh, M. Shapiro. *The Case for Fast and Invariant-Preserving Geo-Replication*. W. on Planetary-Scale Distributed Systems (W-PSDS). co-located with SRDS, Nara, Japan October 2014.
- M. Saeida-Ardekani, P. Sutra, M. Shapiro. *G-DUR: A Middleware for Assembling, Analyzing, and Improving Transactional Protocols*. Middleware, Bordeaux, France, Dec. 2014.
- Marcos K. Aguilera, Leonardo Querzoni and Marc Shapiro. *Principles of Distributed Systems*. The 18th International Conference on Principles of Distributed Systems. 16-19 December 2014, Cortina d'Ampezzo, Italy.

- Santiago Castiñeira and Annette Bieniusa. *Collaborative offline web applications using Conflict-free Replicated Data Types*. PaPoC 2015.
- Tyler Crain and Marc Shapiro. *Designing a causally consistent protocol for geo-distributed partial replication*. PaPoC 2015.
- Alejandro Z.Tomsic, Tyler Crain and Marc Shapiro. *An empirical perspective on causal consistency*. PaPoC 2015.
- Amadeo Asco and Annette Bieniusa. *Adaptive strength geo-replication strategy*. PaPoC 2015.
- C. Meiklejohn and P. Van Roy. *LASP: Distributed, Eventually Consistent Computations*. PaPoC 2015.
- Marek Zawirski, Nuno Preguiça, Sergio Duarte, Annette Bieniusa, Valter Balegas, Marc Shapiro. *Write Fast, Read in the Past: Causal Consistency for Client-side Applications*, Middleware 2015, Dec. 2015. Vancouver, BC, Canada.
- Deepthi Devaki Akkoorath and Annette Bieniusa. *Transactions on Mergeable Objects*. In Proceedings of the Programming Languages and Systems 13th Asian Symposium, APLAS 2015, December 2015.
- Mathias Weber und Annette Bieniusa. *Access Control for Weakly Consistent Data Stores*. Presentation at the 18. Kolloquium Programmiersprachen und Grundlagen der Programmierung (KPS15), October 2015.

Posters

- Valter Balegas, Sergio Duarte, Carla Ferreira, Rodrigo Rodrigues and Nuno Preguiça. *The quest for coordination-free cloud storage*. 25th ACM Symposium on Operating Systems Principles (SOSP '15), Monterey, CA, Oct. 4-7, 2015.

Appendix II: Presentations/Talks in the second year

Keynotes

- *From Paper to Product: Putting CRDTs into Riak 2.0.* Russel Brown at 8th Workshop on Large-Scale Distributed Systems and Middleware, Oct. 2014.
- *SwiftCloud: Fault-Tolerant Geo-Replication Integrated all the Way to the Client Machine.* Nuno Preguiça at Workshop on Planetary-Scale Distributed Systems 2014, Oct. 2014.
- *A principled approach to the performance-vs.-programmability trade-off in large-scale distributed data store.* MASCOTS Conference 2014, Sept. 2014, Marc Shapiro.

Tutorials

- Tutorial on consistency at CodeMesh 2015, by Marc Shapiro and Nuno Preguiça.

Conferences/Workshops

- *Reasoning about consistency choices in Distributed Systems* by Carla Ferreira at Workshop on Formal Reasoning in Distributed Algorithms (FRIDA 2015)
- *Distributed, Eventually Consistent Computations.* Christopher Meiklejohn at Strange Loop 2015.
- *Scalable Consistency for Replicated Data.* Annette Bieniusa at BOBKONF'15
- *Coordination-Free Computations* (invited talk). Christopher Meiklejohn at GOTO Chicago 2015
- *Distributed Eventually Consistent Computations.* Christopher Meiklejohn at Erlang Factory 2015.
- *Scalable Consistency for Replicated Data.* Annette Bieniusa at UPMARC Workshop on Memory Models (MM'15)
- *Decomposing consistency.* Invited talk by Marc Shapiro at Workshop on Chemistry of Concurrent and Distributed Programming II.
- *A principled approach to the performance-vs.-programmability trade-off in large-scale distributed data stores.* Invited speaker Mascots 2014 (Paris).
- *Designing for Partition Tolerance with CRDTs.* Carlos Baquero at RICON 2014, Oct. 2014.
- *Eventually Consistent Computations with CRDTs.* Christopher Meiklejohn at RICON 2014, Oct. 2014.
- *SyncFree: Large-Scale Computation Without Synchronization.* Annette Bieniusa and Christopher Meiklejohn at RICON 2014, Oct. 2014.

- *From Paper to Product: Putting CRDTs into Riak 2.0*. At LADIS workshop Cambridge (Large-Scale Distributed Systems) by Marc Shapiro
- *The Case for Fast and Invariant-Preserving Geo-Replication*. At W. on Planetary-Scale Distributed Systems (W-PSDS) by Marc Shapiro
- LADIS Cambridge. Several presentations on SyncFree. Russell Brown invited speaker.
- Invited talk CurryOn (ECOOP), Prague, *Encapsulating replication, high concurrency and consistency with CRDTs* by Marc Shapiro.
- Invited speaker LADIS 2015, *Practical experience with CRDTs in real applications*.
- Invited talk CodeMesh 2015, presentation on CRDTs by Marc Shapiro and Nuno Preguiça.
- *Write Fast Read in the Past*, Journe Virtualisation et Cloud, by Marc Shapiro.

General

- Nuno Preguiça. Rodrigo Rodrigues. *Defining and mixing consistency guarantees*. Talk at Portugal—UT Austin Summer school in distributed computing, Sep 6, 2015.
- *Scalable consistency - all the way to the edge!*. Annette Bieniusa. Dagstuhl Seminar "Distributed Cloud Computing" (15072)
- *Scalable consistency in distributed systems*. Annette Bieniusa. Dagstuhl Seminar "Concurrent computing in the many-core era" (15021)
- *Scalable consistency in distributed systems*. Annette Bieniusa. CS Department Talk Series - University of Kaiserslautern
- *Decomposing consistency: It's the invariants, stupid*. At Universidade Nova de Lisboa by Marc Shapiro
- Invited talk by Marc Shapiro at Royal Holloway London, *Decomposing consistency - It's the Invariants, Stupid!*.
- Marc Shapiro invited to Dagstuhl Seminar, *Distributed Cloud Computing*.
- EuroTM final workshop Amsterdam, Marc Shapiro speaker at panel "The Future of TM"
- *Lasp: A Language for Eventually Consistent Distributed Programming with CRDTs* by Peter Van Roy. Three talks: KTH, Stockholm (March 5, 2015), IST, Lisbon (March 13, 2015), UCL, Louvain-la-Neuve (May 12, 2015).
- *La programmation repartie sans coordination* by Peter Van Roy. Journée en l'honneur de Christian Queinnec. LIP6/UPMC, Paris, June 8, 2015.