

# Master Project Proposal: Tree-based Decentralized and Robust Causal Dissemination

Advisor: João Carlos Antunes Leitão  
NOVA Laboratory for Computer Science and Informatics (NOVA LINCS)  
and  
Departamento de Informática  
Faculdade de Ciências e Tecnologia  
Universidade NOVA de Lisboa

First Semester, 2018/2019

## 1 Context & Motivation

Causal consistency has become an extremely relevant consistency model in practical geo-replicated storage systems. This relevance comes from the fact that this is the strongest consistency model that can be implemented in an available way in light of the CAP theorem [3]. The increase popularity of the edge-computing paradigm has also created some pressure to build new storage systems with causal guarantees that can operate both in the cloud and at the edge of the system (i.e., on smaller private/regional data centers or in infrastructures on National ISPs and so forth). One way to tackle this is to devise novel causal dissemination schemes that can operate efficiently on this challenging and heterogeneous environment.

In this thesis we plan to evolve solutions for reliable data dissemination based on trees for peer-to-peer systems [7, 5] that can offer additional causality guarantees while operating efficiently on an hybrid cloud-edge setting.

## 2 Thesis Goals

The goal is to design, implement, and experimentally evaluate novel causal-based dissemination solutions based on trees, that should be robust and efficient. The design of these solutions will be guided by requirement of novel replications schemes based on the dissemination of small metadata information that can guide the replications protocol [2, 6]. In summary this thesis aims at:

- Proposal a novel design (or novel designs) for tree based causal dissemination solution.

- Implement and document prototypes of all developed solutions.

- Conduct an experimental evaluation of the prototypes.

## 3 Notes & Observations

The thesis work will be advised by João Leitão (jc.leitao@fct.unl.pt). Please send an e-mail if you are interested in this master thesis proposal.

This thesis project will be pursued in close collaboration with PhD Student Pedro Fouto and MSc student Gonçalo Tomás.

This master thesis proposal is related with the COSMOS Project (Causal Consistency on the Network Edge) financed by the Fundação para a Ciência e Tecnologia.

Other relevant bibliography include the Kronos [4], and ChainReaction [1] systems.

## References

- [1] S. Almeida, J. a. Leitão, and L. Rodrigues. Chainreaction: A causal+ consistent datastore based on chain replication. In *Proceedings of the 8th ACM European Conference on Computer Systems*, EuroSys '13, pages 85–98, New York, NY, USA, 2013. ACM.
- [2] M. Bravo, L. Rodrigues, and P. Van Roy. Saturn: A distributed metadata service for causal consistency. In *Proceedings of the Twelfth European Conference on Computer Systems*, EuroSys '17, pages 111–126, New York, NY, USA, 2017. ACM.
- [3] E. Brewer. Pushing the cap: Strategies for consistency and availability. *Computer*, 45(2):23–29, Feb. 2012.
- [4] R. Escriva, A. Dubey, B. Wong, and E. G. Sirer. Kronos: The design and implementation of an event ordering service. In *Proceedings of the Ninth European Conference on Computer Systems*, EuroSys '14, pages 3:1–3:14, New York, NY, USA, 2014. ACM.
- [5] M. Ferreira, J. Leitao, and L. Rodrigues. Thicket: A protocol for building and maintaining multiple trees in a p2p overlay. In *2010 29th IEEE Symposium on Reliable Distributed Systems*, pages 293–302, Oct 2010.
- [6] P. Fouto, J. Leitão, and N. Preguiça. Consistência causal em sistemas geo-distribuídos com replicação parcial. In *Proceedings of the 9th Portuguese Symposium on Informatics*, Coimbra, Portugal, Sept. 2018.
- [7] J. Leitao, J. Pereira, and L. Rodrigues. Epidemic broadcast trees. In *2007 26th IEEE International Symposium on Reliable Distributed Systems (SRDS 2007)*, pages 301–310, Oct 2007.