

Staying CALM Beyond Deterministic Queries

Baccaert, Tim; Ketsman, Bas

Publication date:
2022

[Link to publication](#)

Citation for published version (APA):
Baccaert, T., & Ketsman, B. (2022). *Staying CALM Beyond Deterministic Queries*. Poster session presented at Dutch-Belgian Database Day, Utrecht, Netherlands.

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Distributed Consistency Beyond Queries

Tim Baccaert and Bas Ketsman

Vrije Universiteit Brussel


System Assumptions

Shared-Nothing

State is local to each node, communication is achieved through message passing

Asynchronous


Messages may arrive out-of-order, and with arbitrary (but finite) delay.



Programs and Coordination


Connection

Is there a path from A to B?




Deadlock Detection^[3]

Is there a cycle in the waits-for graph?





Garbage Collection^[3]

Is a memory object reachable from R?



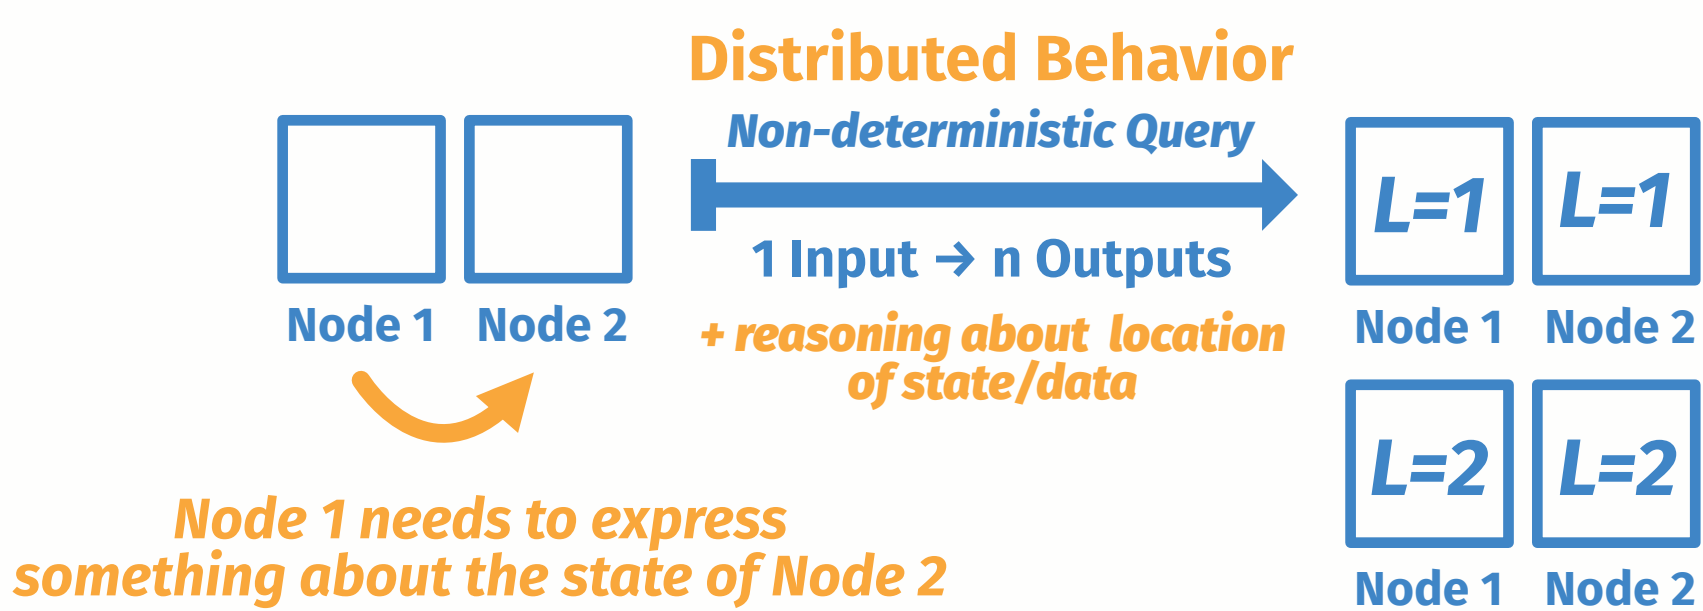
Leader Election

Do all nodes pick the same leader?



Distributed Behaviors

To express Leader Election semantically, we need **non-determinism** and **data placement**.



Modeling Systems With Behaviors

*This list is non-exhaustive, more examples in the full paper.

- Id** Nodes have a unique identity.
- Id+All** Each node knows the unique identity of all other nodes.
- Id+All+Ord** Each node agrees on an order over its stored data.
- Id+Ind** Each node knows part of a distributed index (i.e., the data partitioning strategy)

The CALM Theorem: A Coordination Spectrum

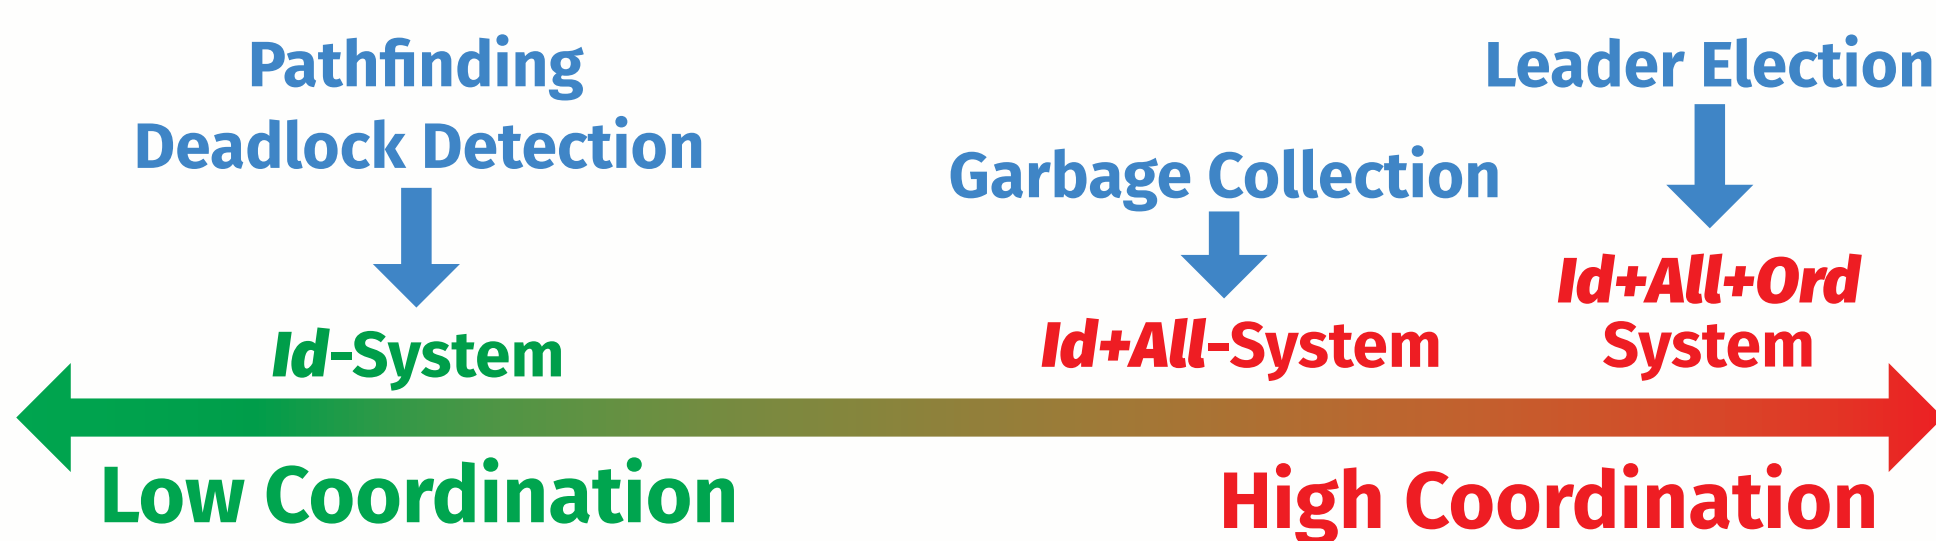
The CALM Theorem^[1,2] (Revisited)




Monotonic Query^[1] **Id-System**^[1] No Coordination^[1]

Semi or Disjoint Monotonic Query^[2] **Id+Ind-System** or **Id+Ind*-System**^[2] Light Coordination^[2]

C-Monotonic Behavior **C-System** Depends on C



Use the system model to define coordination: "A stronger system model implies higher inherent coordination."



[1] Tom J. Ameloot, Frank Neven, and Jan Van den Bussche. 2013. Relational Transducers for Declarative Networking. *J. ACM*. 60, 2, 1-38.
 [2] Tom J. Ameloot, Bas Ketsman, Frank Neven, and Daniel Zinn. 2016. Weaker Forms of Monotonicity for Declarative Networking: A More Fine-Grained Answer to the CALM-Conjecture. *ACM Trans. Database Syst.* 40, 4, 1-45.
 [3] Joseph M. Hellerstein and Peter Alvaro. 2020. Keeping CALM: When Distributed Consistency is Easy. *Commun. ACM*. 63, 9, 72-81.