# Tracking State Footprints: How Agents Can Transact

Oto Mraz [1], Rares Şerban [1], Kyriakos Psarakis [2], Burcu Kulahcioglu Ozkan [1], and Asterios Katsifodimos [1]

[1] Delft University of Technology, Netherlands
{o.m.mraz,r.c.serban,b.ozkan,a.katsifodimos}@tudelft.nl
[2] Ververica GmbH, Netherlands
kyriakos.psarakis@ververica.com

Can AI agents transact? We argue that they must: as multi-agent systems (MASs) increasingly write code, deploy infrastructure, modify databases, and call web services, lost updates or stale reads can be catastrophic. Although MASs increasingly execute plans in parallel, current orchestrators do not track the state that agents read and write. As a result, concurrency anomalies manifest even in simple coding tasks. We frame MAS coordination as a data management problem and propose to describe agents by their *state footprint*: the state they read and write across their own local context and state, as well as the state of the orchestrator and external systems. We posit that MASs require guarantees similar to those of databases, but providing them raises new challenges and opportunities: unlike database transactions, agents do not read from a fixed schema or an isolated snapshot, and cannot be replayed deterministically upon failure. They can, however, resolve conflicts semantically instead of aborting, enabling new forms of concurrency control and conflict resolution. Towards agents that can transact, we outline a vision for next-generation agent orchestrators and transactional interfaces for external systems to participate in agentic transactions [1].

## References

[1] Oto Mraz, Rares Şerban, Kyriakos Psarakis, Burcu Kulahcioglu Ozkan, and Asterios Katsifodimos. Tracking State Footprints: How Agents Can Transact. *arXiv preprint arXiv:2610.03140*, 2026. Under revision for VLDB 2026.