

# Transactional Cloud Applications: Status Quo, Challenges, and Opportunities



# Who are we?



**Rodrigo  
Laigner**  
(University of  
Copenhagen)



**George  
Christodoulou**  
(TU Delft)



**Kyriakos  
Psarakis**  
(TU Delft)



**Asterios  
Katsifodimos**  
(TU Delft)



**Yongluan  
Zhou**  
(University of  
Copenhagen)



Slides, and pointers:

<https://delftdata.github.io/tutorial-sigmod25/>



## Disclaimer: not your typical tutorial

This is **not** your typical tutorial

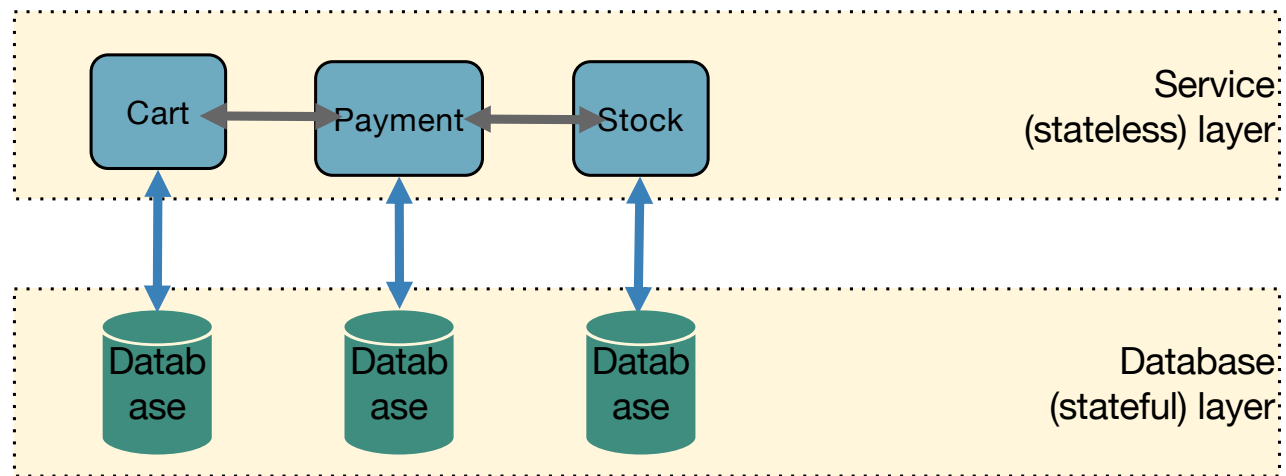
The list of presented systems is **not** exhaustive

This is **not** a survey (although it's a good start)

Presenters may be **opinionated** (at times)

# | Tutorial Scope

## Cloud services (canonical example)



Check & update **stock**, verify **payment**, then checkout the shopping **cart**.

# Use Cases for Transactional Cloud Applications

**Target Use-cases:** *Low-latency Cloud applications requiring transactional consistency, containing complex business logic.*

Booking/reservation, systems, trading

Ad serving & bookkeeping

Fraud detection & payments

Inventory management

## Target Industries

Banking, e-commerce, trading platforms, retail, etc.

## ACID in the world of services

### **Atomicity**

All three services execute, despite system/user errors

### **Consistency**

FK constraints: shopping cart contains only products that exist in stock

### **Isolation**

No stock updates visible without having payment cleared

No payment without stock updates reflected in stock service

### **Durability**

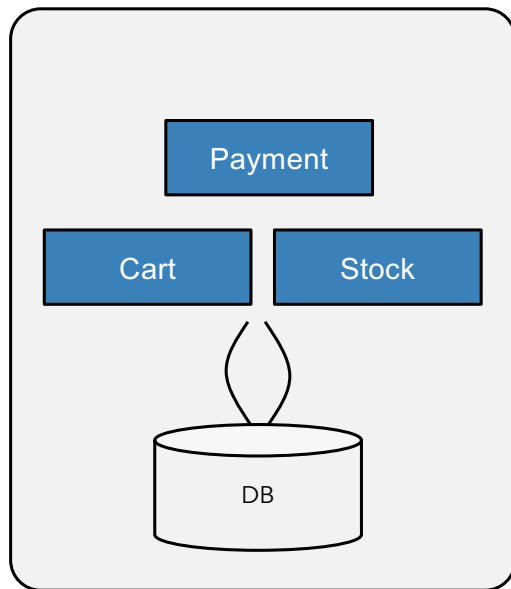
Data is safe, possibly replicated

**| How people develop  
microservices nowadays**



# Service Architectures Primer: Monolith

## Monolith



## YES,

Programmers only deal  
with app logic

Database ensures ACID


## BUT

Cannot migrate easily to  
the Cloud

May not scale on  
commodity hardware

# Service Architectures Primer: $\mu$ Service

"Textbook" Microservice



YES,

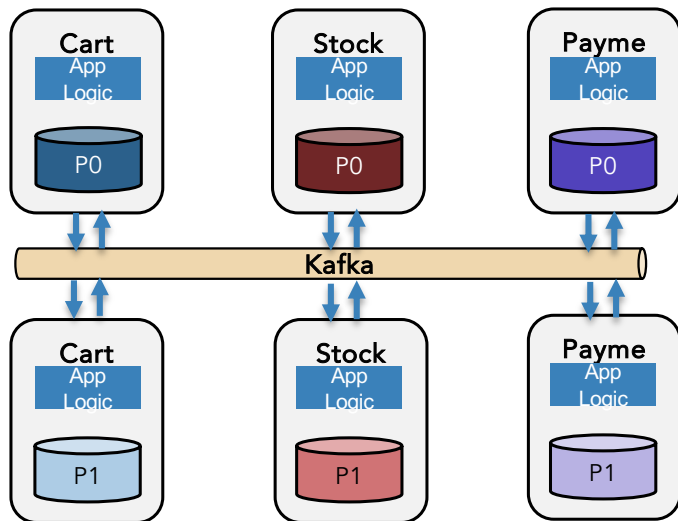
Scales better  
Code modularity  
Independent deployment/scaling

BUT

Requires orchestration  
Retries/idempotence hard  
Consistency/Transactions in app code (SAGAs, 2 Phase Commit/Open XA)

# Service Architectures Primer: event-driven $\mu$ Service

Partitioned, event-driven architecture



YES,

High Performance

Fault-tolerant

Replayable & Debuggable

BUT

Inherits  $\mu$ Service issues

Few people can program and keep it running

Ad-hoc transactions + event-driven = trouble

## **Different 'shades' of consistency guarantees**

Cloud applications largely avoid transactions through databases  
Exceptions apply, e.g., DBOS

Transactions are implemented in user code.  
Ad-hoc (with tons of hacks)  
Two-phase commit  
SAGAs

# | Transactions

## Two-Phase Commit in Services


Coordination for committing or rolling back a transaction

Two phases: Prepare phase and Commit phase

Serializable guarantees

Reduced throughput

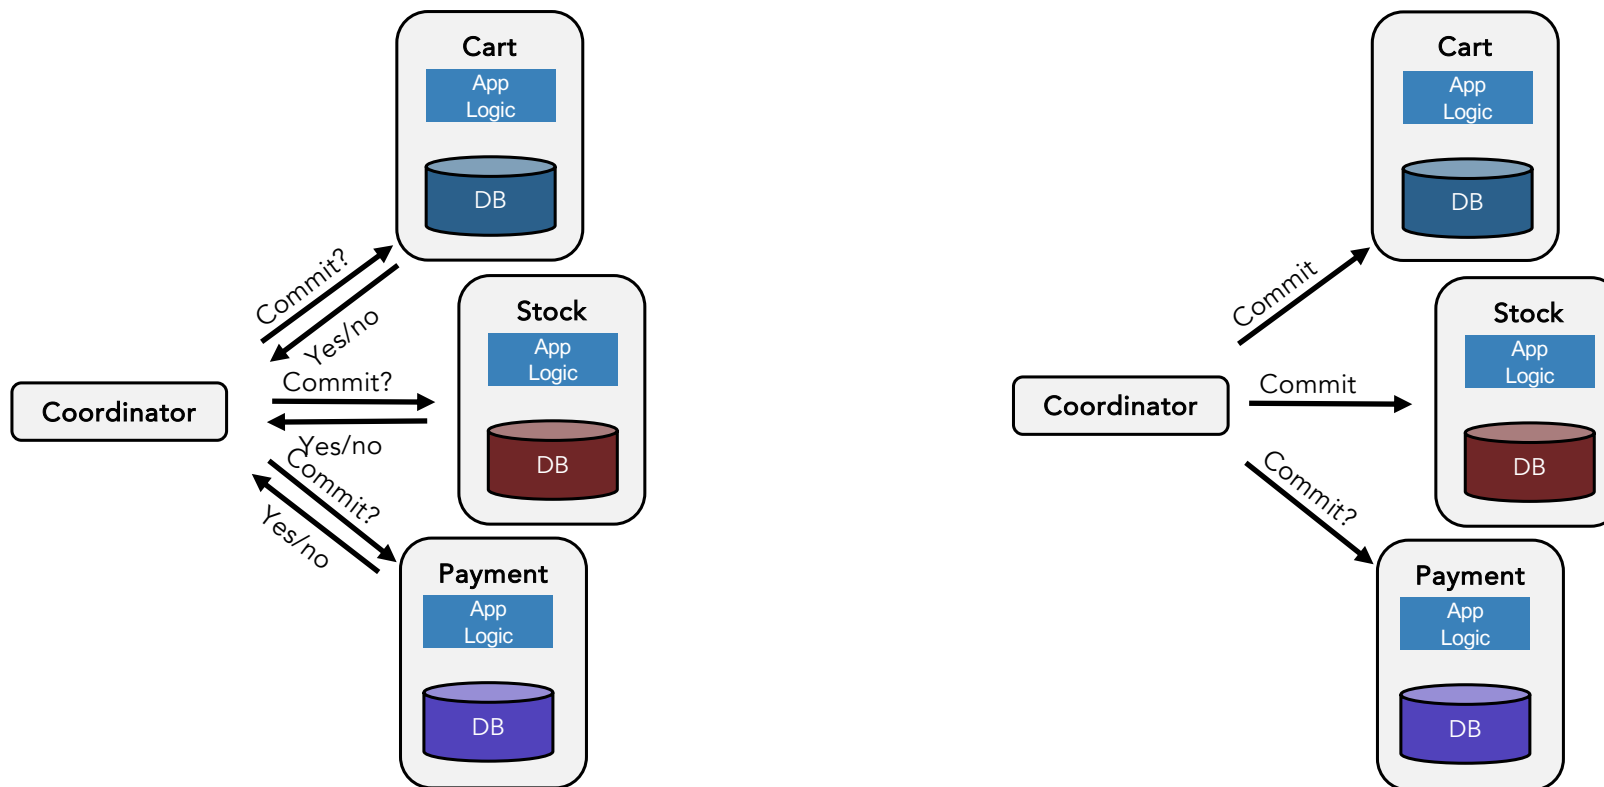
Single point of failure



vs



## Two-Phase Commit example



## **SAGAs in Services**

Talk about the sagas, compensations, etc.

The result is eventually consistent, under many assumptions.

The world implements SAGAs without knowing it.




## SAGAs in Services [1]

Distributed transaction as multiple single local transactions

Compensatory transactions in case of failure

Durable and distributed logs of all messages (usually Kafka topic)

Eventual consistency

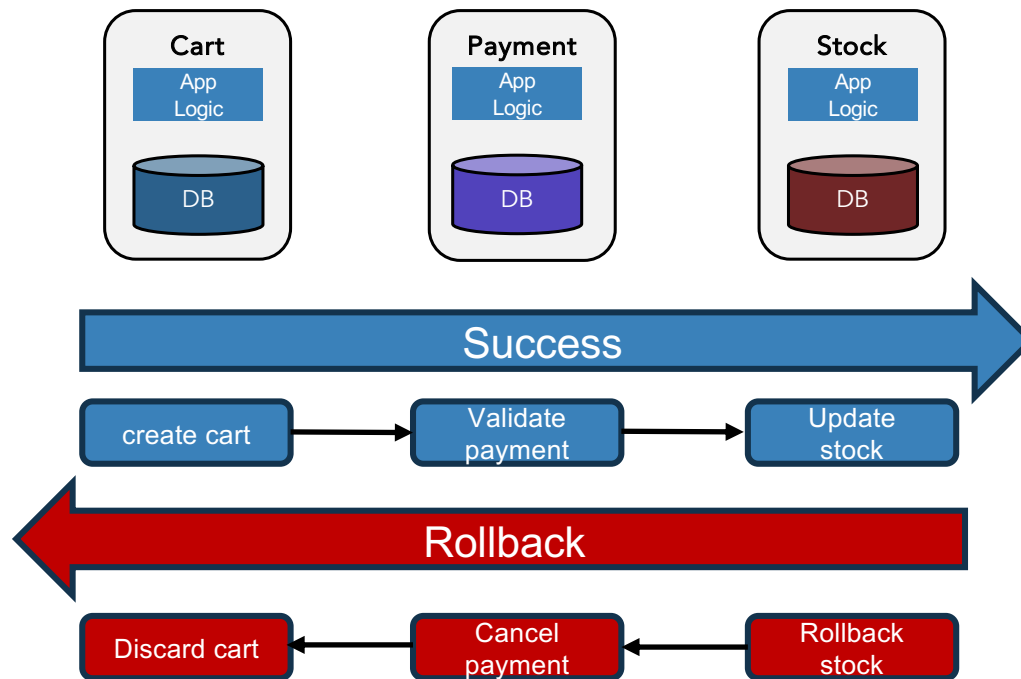


VS



[1] H. Garcia-Molina, K. Salem, "**Sagas**", [SIGMOD '87]


## SAGAs in Services [1]



[1] H. Garcia-Molina, K. Salem, "**Sagas**", [SIGMOD '87]

## Yes, but

YES,



BUT



© \_yes\_but

# Ad-hoc Transactions; all over in Microservices

## Ad Hoc Transactions: What They Are and Why We Should Care

Chuzhe Tang<sup>1,2</sup>, Zhaoguo Wang<sup>1,2</sup>, Xiaodong Zhang<sup>1,2</sup>, Qianmian Yu<sup>1,2</sup>

Binyu Zang<sup>1,2</sup>, Haibing Guan<sup>3</sup>, Haibo Chen<sup>1,2</sup>

<sup>1</sup>Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University

<sup>2</sup>Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China

<sup>3</sup>Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University  
zhaoguowang@sjtu.edu.cn

## Developer's Responsibility or Database's Responsibility? Rethinking Concurrency Control in Databases

Chaoyi Cheng<sup>\*†</sup>

The Ohio State University

Spyros Blanas

The Ohio State University

Mingzhe Han<sup>\*</sup>

The Ohio State University

Michael D. Bond

The Ohio State University

Nuo Xu<sup>‡</sup>

The Ohio State University

Yang Wang

The Ohio State University

FINDING 1. *Every studied application uses ad hoc transactions. Among the 91 ad hoc transactions in total, 71 cases are considered critical to the web applications.*

FINDING 3. *There are 7 different lock implementations and 2 validation implementations among the 8 applications we studied. Except for Broadleaf, developers consistently use the same lock/validation implementation in individual applications.*

**C4.** Due to the complex interplay between microservices' behaviors, asynchronous events are generated to trigger computations. However, avoiding anomalies related to the unintended interleaving of events across microservices is a challenging task.

**C6.** Developers lack viable and efficient abstractions for transactional queuing in microservice architectures. As a result, anomalies arise due to lack of isolation and ad-hoc fault-handling, leading to challenges on ensuring application correctness.

**C8.** Due to the distributed nature of microservice architectures, eventual consistency is often taken as the de facto consistency model by practitioners. This choice introduces a series of challenges on reasoning about distributed states and invariants.

# | Challenges with microservices & the like



# >90% of programmers' time spent in machine/network failures

(a.k.a. "plumbing")

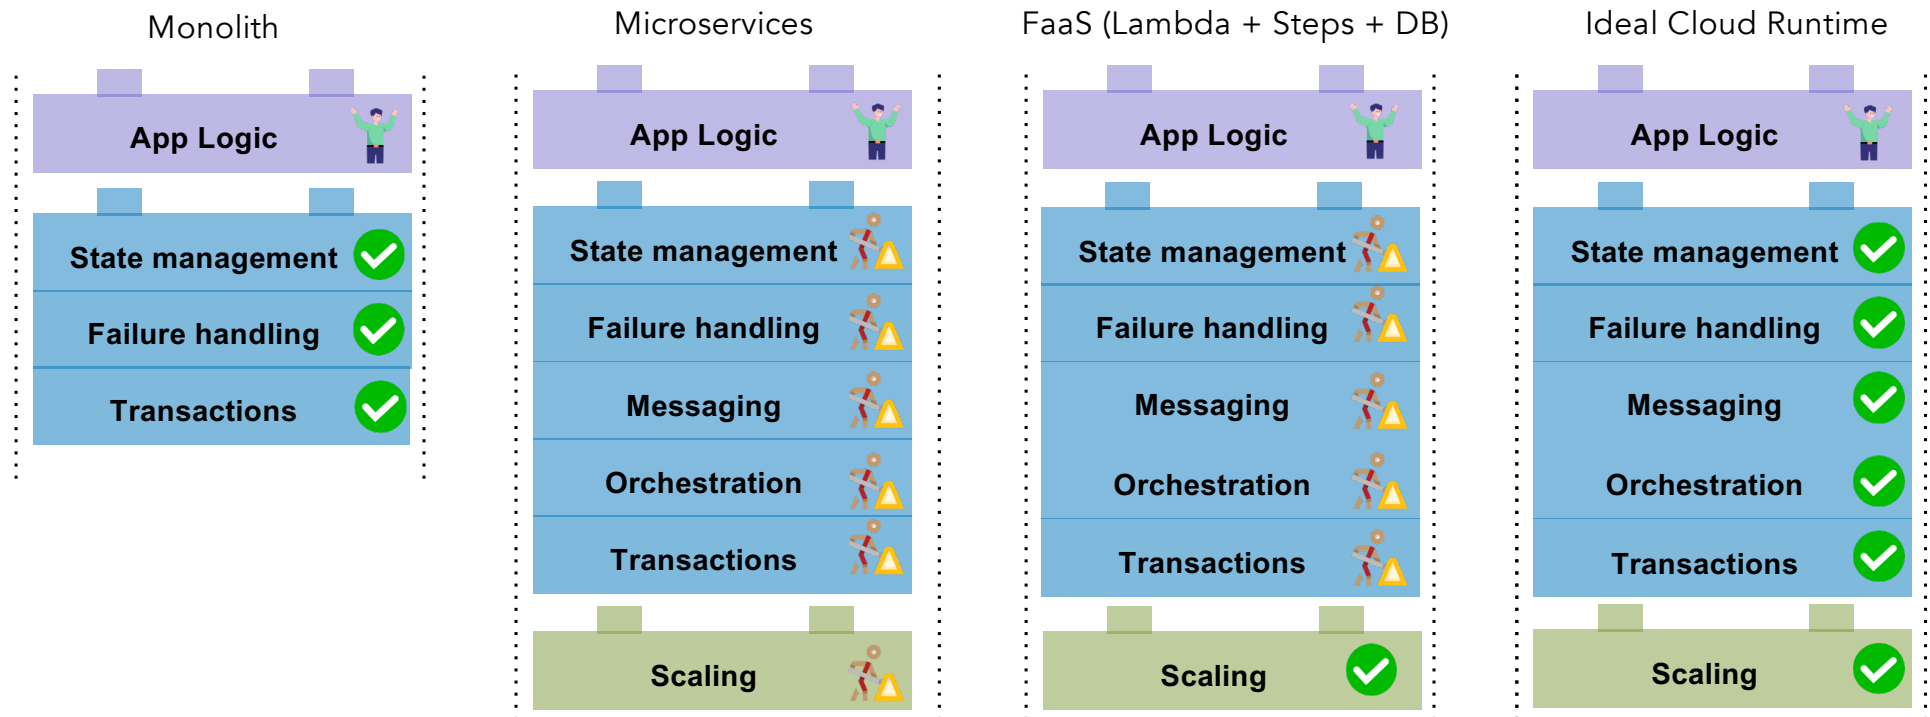
\*Actual code (shrunked) of scalable shopping cart, using Flask and Postgres. Excludes K8s config file hell.



- Retries/Atomicity/Idempotency
- Consistency
- Recovery
- Partitioning
- Scaling
- Maintaining DBs, proxies, containers
- ...

"Useful" application-logic code percentage: 5-10%.

# Requirements for transactional Cloud apps



Transaction support



Sad Developer



Happy Developer

YES,



BUT



📷 \_yes\_but

## The rest of this tutorial

Right after

Programming Models

After the break

Runtimes

Benchmarks

Open Problems