

# Democratizing Scalable Cloud Applications

*An Approach Using Stateful Functions on Streaming Dataflows*



Kyriakos Psarakis



George Christodoulou



Marios Fragkoulis


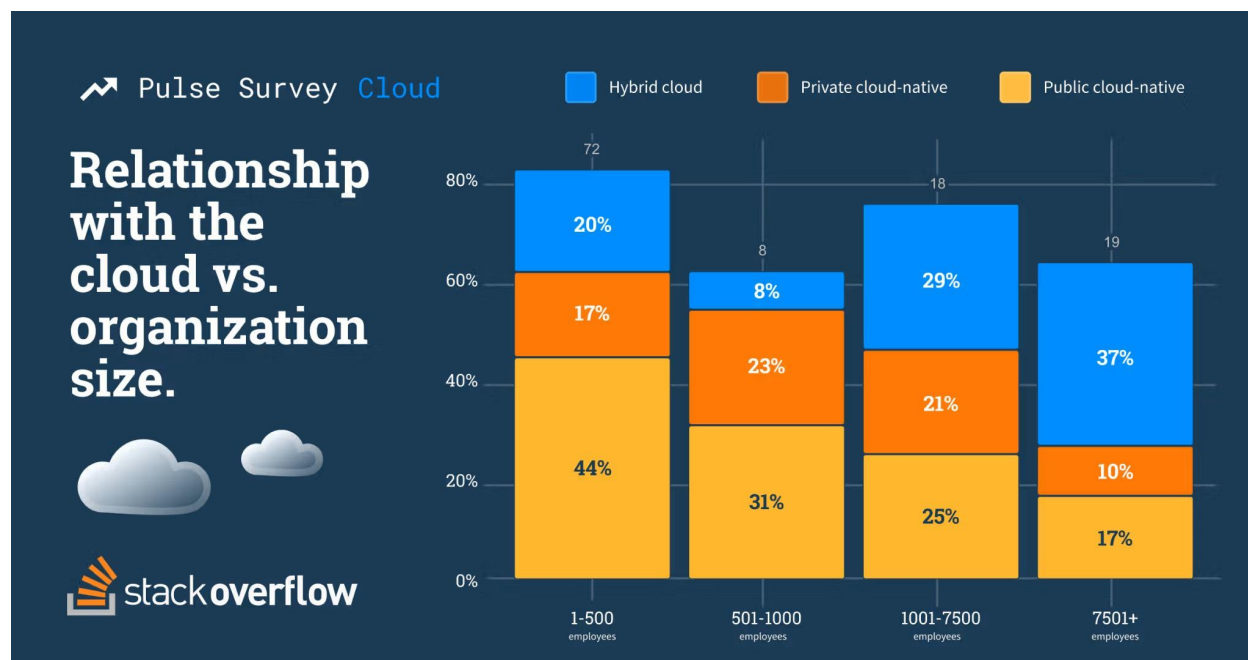


Asterios Katsifodimos



# Cloud Transition

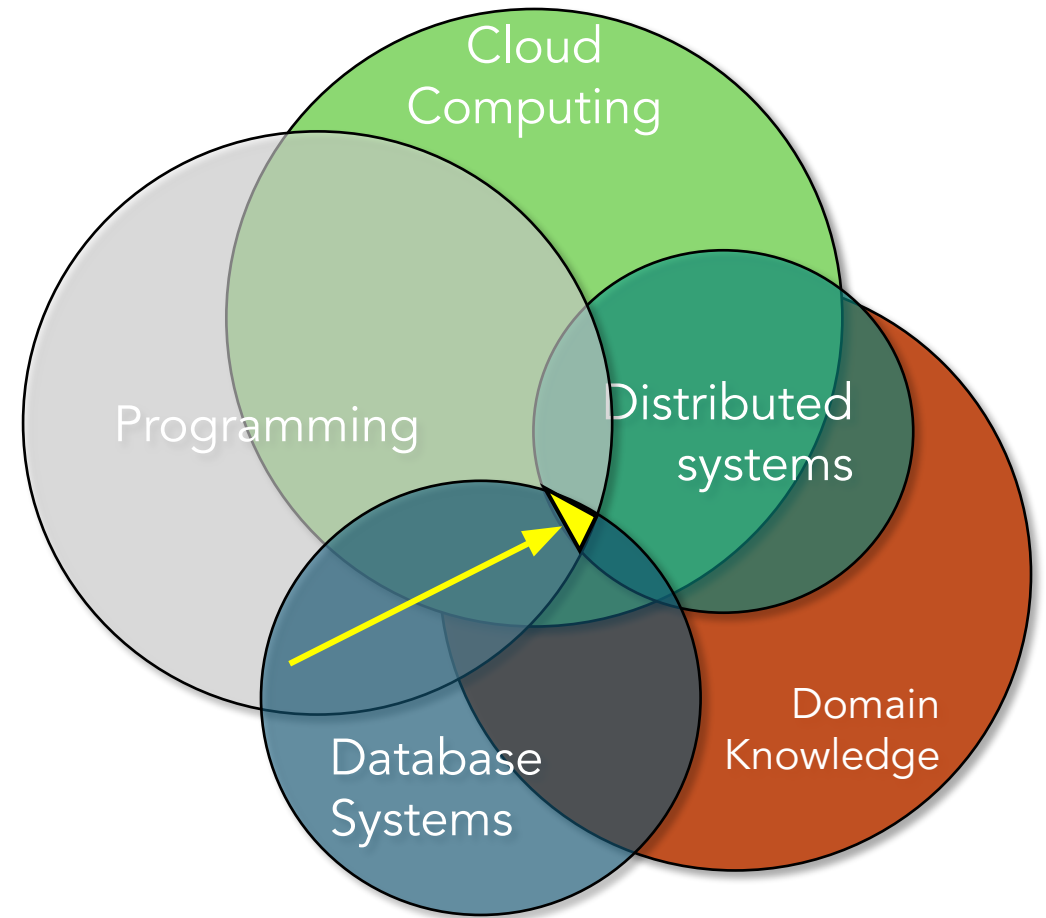
- Cloud adoption has reached levels above 60% in 2021
- A part of the cloud landscape is scalable cloud applications



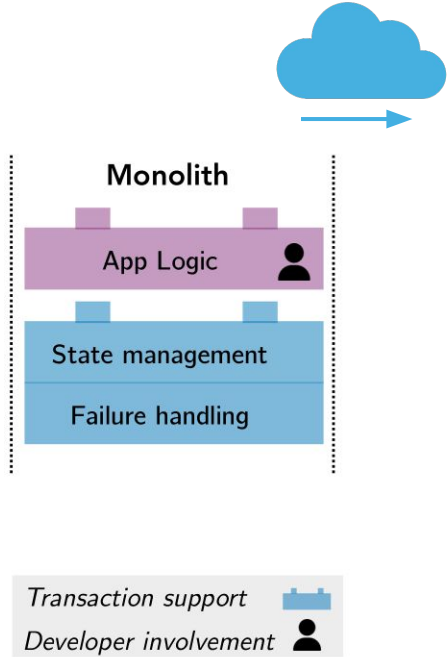
Shopping Cart Cloud Deployment

# Challenges

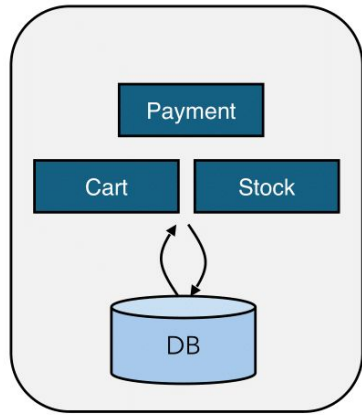
- Deep expertise in multiple domains
- Resurfaces issues solved by database systems:
  - Transactions
  - Fault tolerance



# Towards the Ideal Cloud Runtime

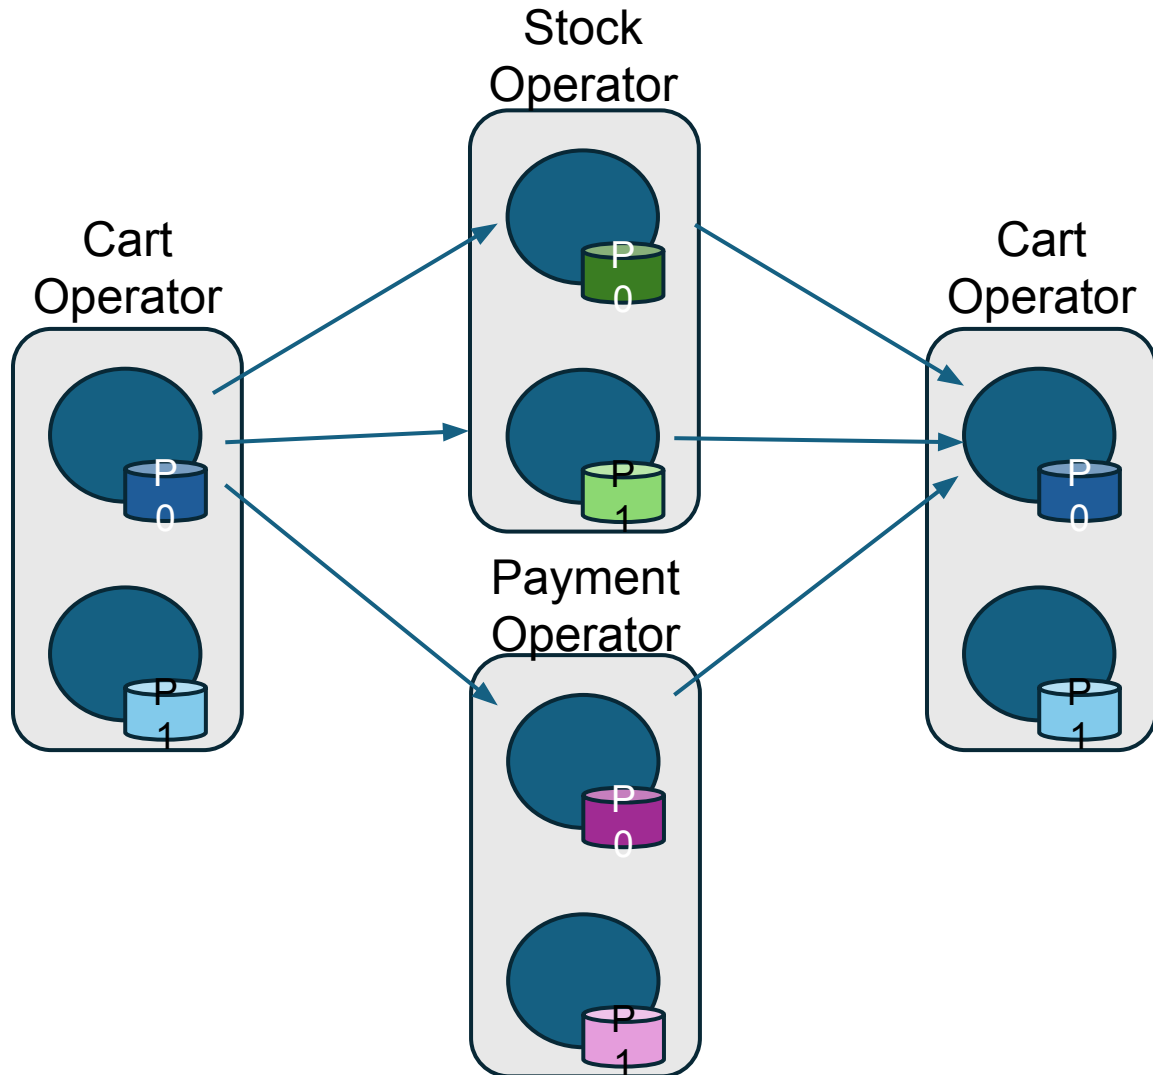


# The Dataflow Realization



a) Monolith

# The Dataflow Realization



Benefits of the dataflow design:

- Solution native to the problem
- Strong guarantees
- High performance
- Coarse grained fault tolerance
- Clean API

# Styx's Low-Level API

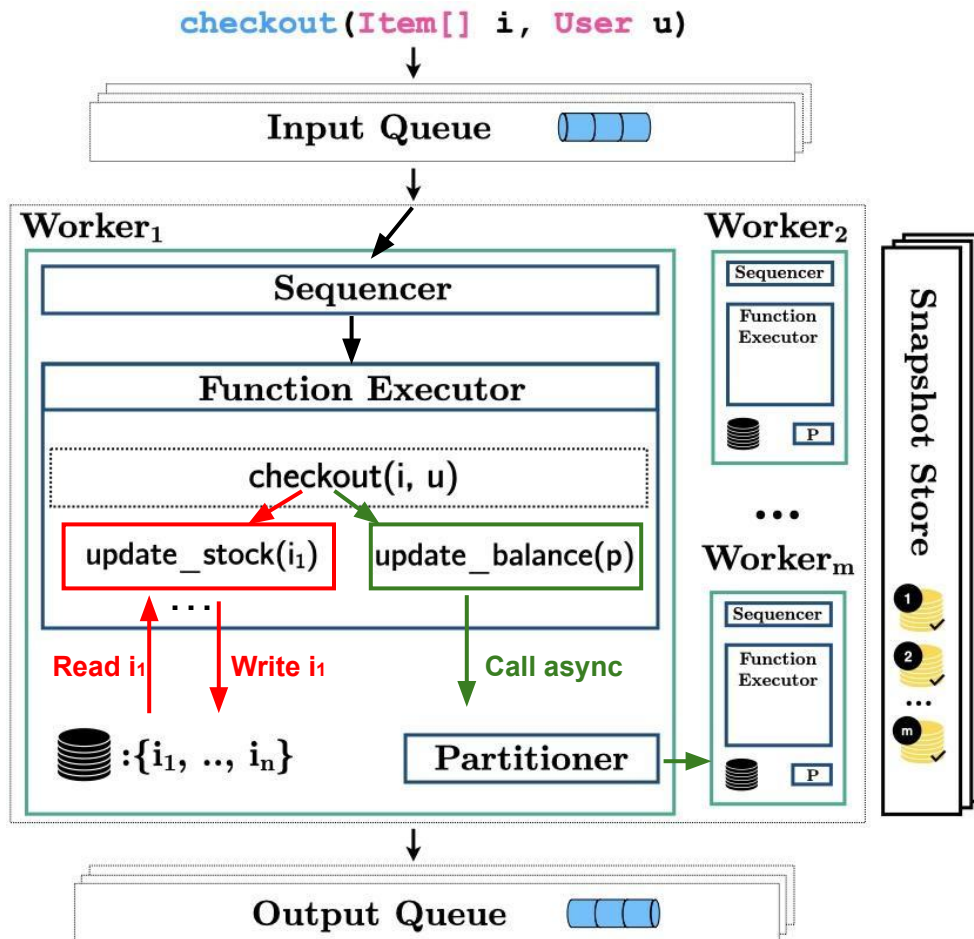
- Styx works with Stateflow<sup>[1]</sup>
- Styx low-level API does not expose:
  - Failure handling
  - Transactional semantics
  - Scalability

```
1 ...
2 # Check if there is enough stock
3 if stock_amount <= 0:
4     raise NotEnoughStock(f'No stock left for item: {context.key}')
5 ...
```

```
1 from styx import Operator
2 from shopping_cart.operators import stock, payment
3
4 cart = Operator('cart', n_partitions=4)
5
6 ...
7
8 @cart.register
9 def checkout(context):
10     order_id = context.key
11     items, user_id, total_price, paid = context.state.get()
12
13     for item_id, qty in items:
14         context.call_async(operator=stock,
15                             function_name='decrement_stock',
16                             key=item_id,
17                             params=(qty, ))
18     context.call_async(operator=payment,
19                         function_name='pay',
20                         key=user_id,
21                         params=(total_price, ))
22
23     paid = True
24     context.state.put((items, user_id, total_price, paid))
25
26     return "Reservation Successful"
```

[1] K. Psarakis, W. Zorgdrager, M. Fragkoulis, G. Salvaneschi, and A. Katsifodimos. Stateful Entities: Object-Oriented Cloud Applications as Distributed Dataflows. In EDBT, 2024.

# Styx Function Flow



Guarantees:

1. Exactly-once processing
2. Exactly-once output
3. Durability in the Snapshot store and I/O queues



# Deterministic Transaction Execution

- Epoch-based deterministic protocol (ensures serializability of the sequence)
- Enables Early-Commit replies (replies to clients before persistence in durable storage)

## 1 Sequencing

$R_1(k_1, k_2) \rightarrow \text{III} \rightarrow T_1$


$R_3(k_2, k_8) \rightarrow \text{III} \rightarrow T_2$

$R_2(k_3, k_8) \rightarrow \text{III} \rightarrow T_3$

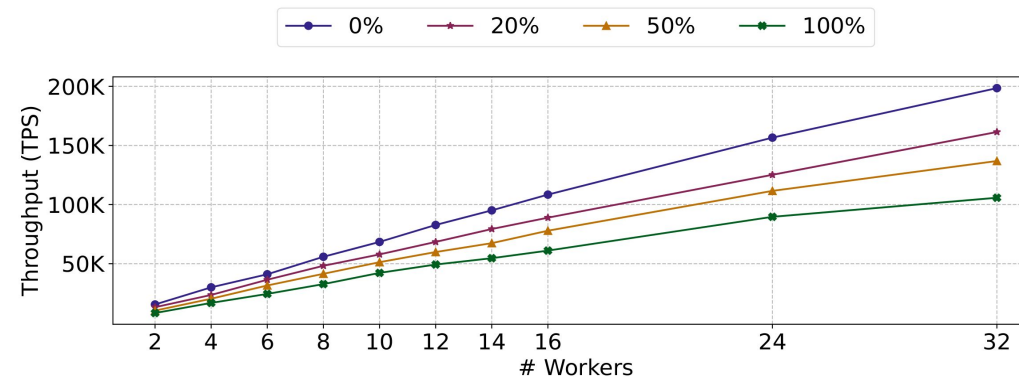
# Experimental Results

Scenario	#keys	Function Calls	Transactions %
YCSB-T	10k	2	100%
Deathstar Movie	2k	9-10	0%
Deathstar Travel	2k	3	0.5%
TPC-C	1m-100m	8 / 20-50	100%

YCSB-T Performance



YCSB-T Scalability (multipartition%)



# Takeaways



Collocation of state and processing is essential for low-latency & high throughput

# Takeaways

# Thank you!



Collocation of state and processing is essential for low-latency & high throughput



Exactly-Once processing guarantees allow for a simple & easy to use API



Deterministic transactions act as a performance enabler



Coarse grained fault tolerance adds minimal overhead

## References:

- [1] K. Psarakis, G. Christodoulou, M. Frankoulis, and A. Katsifodimos. **Transactional Cloud Applications Go with the (Data)Flow**. In CIDR, 2025 (to appear).
- [2] K. Psarakis, W. Zоргdrager, M. Fragkoulis, G. Salvaneschi, and A. Katsifodimos. **Stateful Entities: Object-Oriented Cloud Applications as Distributed Dataflows**. In EDBT, 2024.
- [3] K. Psarakis, G. Siachamis, G. Christodoulou, M. Fragkoulis, and A. Katsifodimos. **Styx: Transactional Stateful Functions on Streaming Dataflows**. In arXiv:2312.06893, 2024.
- [4] M. de Heus, K. Psarakis, M. Fragkoulis, and A. Katsifodimos. **Distributed transactions on serverless stateful functions**. In DEBS 2021.

