On On Serverless Stateful Functions

Distributed Transactions on Serverless Stateful Functions Martijn De Heus, Kyriakos Psarakis, Marios Fragkoulis, Asterios Katsifodimos. To appear in the proceedings of the 15th ACM International Conference on Distributed and Event-based Systems (DEBS) 2021.



- Transactions
- Microservices & Cloud computing
- Stateful Functions & Coordinator Functions
- Evaluation
- Conclusion & discussion



Transactions



Transactions

A transactions are a <u>set</u> of operations on data with 4 (**ACID**) guarantees:

- **A**tomicity
- **C**onsistency
- Isolation
 - o (<u>serializability</u>)
- **D**urability



Traditional Architecture





Distributed transactions

- Two-Phase Commit ensures atomicity
- Two-Phase Locking ensures serializability

Combining Two-Phase Locking and Two-Phase Commit can be used to implement **distributed** transactions



Microservice systems



Microservices





Microservices

Benefits

- Scaling components separately
- Deploying / updating components separately
- Deploying components on specific hardware
- Scaling development in organisations



Two Generals' Problem



The network is **unreliable**, it is never certain whether a message will arrive

If **Service A** never receives message at 3, it does not know whether **Service B** has performed action at 2.



Processing guarantees



AT_MOST_ONCE : only try to send once and do not care about the results, action 2 may not be performed

AT_LEAST_ONCE: keep retrying until a response is received, action 2 may be performed multiple times



Saga orchestration





Saga orchestration





Saga orchestration



Transactions but with completely no <u>isolation</u> <u>Atomicity</u> with of fault-tolerance is also hard to achieve



Cloud computing



Cloud computing

Make computing accessible for everyone

- Eliminate up-front investment requirements and offer compute resources on demand
- Simplify distributed computing
 - Deployment and scalability
 - State management
 - Processing guarantees (exactly-once)
 - Distributed transactions



Cloud computing

Make computing accessible for everyone

- Eliminate up-front investment requirements and offer compute resources on demand
 - Simplify distributed computing
 - Deployment and scalability
 - State management
 - Processing guarantees (exactly-once)
 - Distributed transactions



How can we improve on this?

Make computing accessible for everyone

- Eliminate up-front investment requirements and offer compute resources on demand
- Simplify distributed computing
 - Deployment and scalability
 - State management
 - Processing guarantees (exactly-once)
 - Distributed transactions



Serverless computing

Function-as-a-Service (FaaS) models as the first iteration of serverless

- Eliminate up-front investment requirements and offer compute resources on demand
- Simplify distributed computing
 - Deployment and scalability
 - State management
 - Processing guarantees (exactly-once)
 - Distributed transactions



Function-as-a-Service





Serverless computing

Stateful-Function-as-a-Service (SFaaS)

models as the second iteration of serverless

- Eliminate up-front investment requirements and offer compute resources on demand
- Simplify distributed computing
 - Deployment and scalability
 - State management
 - Processing guarantees (exactly-once)
 - Distributed transactions





Functions-as-a-Service





Stateful Function-as-a-Service



Couples *messaging* and *state management* to ensure <u>exactly-once</u> guarantees across both







Function instances encapsulate specific state based on their **address** (combination of *type* and *id*)

Function instances can perform four side effects:

- Updates to their encapsulated state
- Messages to other function instances
- Delayed messages to other function instances
- Messages to egresses

Function instances can be **invoked** by:

- Messages from other function instances
- Messages from ingresses



```
@functions.bind("example/greeter")
1
    def greet(context, greet_request: GreetRequest):
2
        state = context.state('seen_count').unpack(SeenCount)
3
        if not state:
4
            state = SeenCount()
                                                            state update
5
            state.seen = 1
6
7
        else:
8
            state.seen += 1
        context.state('seen_count').pack(state)
9
                                                                                egress message
10
        response = compute_greeting(greet_request.name, state.seen)
11
12
        egress_message = kafka_egress_record(topic="greetings", key=greet_request.name, value=response)
13
        context.pack_and_send_egress("example/greets", egress_message)
14
```

state access





ŤUDelft

Coordinator functions for distributed transactions



Serverless computing

- Eliminate up-front investment requirements and offer compute resources on demand
- ✓ Simplify distributed computing
 - Deployment and scalability
 - ✓ State management
 - Processing guarantees (exactly-once)
 - Distributed transactions



Coordinator functions



This relies on Statefun's exactly_once guarantees and linearizable operations on single function instances



Coordinator functions

Coordinator functions are simply *specialized function instances* to coordinate a <u>serializable transaction</u> or <u>sagas</u>

Coordinator function instances can perform side effects:

- Add messages to other function instances to the saga or transactions
- Add side effects to perform based on the completion of the saga

Coordinator function instances can be **invoked** by:

- Messages from other function instances
- Messages from ingresses











Regular functions

- Functions should be able to **fail** explicitly
- Batching should be adjusted to respect isolation of invocations part of a serializable transaction
- Locking should be introduced
- It should participate in the protocols for sagas and serializable transactions



Explicit failure

```
class NotFoundException(FunctionInvocationException):
1
        def __init__(self, request_id):
2
            super().__init__(request_id)
3
4
    @functions.bind("example/user_function")
5
    async def account_function(context, message: Read):
6
7
        # Get state
                                                               Raise an exception
        user = context.state('user').unpack(User)
8
9
        if not user:
10
            raise NotFoundException(message.request_id)
11
12
        else:
            response = Response(request_id=request_id, status_code=200, state=user)
13
            egress_message = kafka_egress_record(topic="responses", key=request_id, value=response)
14
            context.pack_and_send_egress("ycsb-example/kafka-egress", egress_message)
15
16
                                                                      Exception handler for
    @functions.bind_exception_handler(NotFoundException)
17
                                                                      non-transactional
    async def handle_not_found(context, request_id):
18
        response = Response(request_id=request_id, status_code=404) invocations
19
        egress_message = kafka_egress_record(topic="responses", key=request_id, value=response)
20
        context.pack_and_send_egress("ycsb-example/kafka-egress", egress_message)
21
```

Coordinator functions



TUDelft

Deadlock detection







Based on an extension of YCSB

- Stage 1: insert *x* keys in the system
- Stage 2: perform of workload of read and write operations
 - Extended with transfer operations
 - Uniform distribution

Measured maximum throughput and latency









Statefun is restricted to a number of instances and CPUs other components are scaled so they are not the bottleneck





~20% drop in performance for 100 keys ~10% drop in performance for 5000+ keys









Very poor performance for serializable transaction across 100 keys More comparable performance for 5000+ keys





Rollbacks have some overhead for sagas Sagas still perform better than serializable transactions even at 100% rollbacks





| Keys | Transfer | Deadlocks / |
|------|------------|---------------------|
| | proportion | transfer ops |
| 100 | 0.25 | 9/12014 (0.07%) |
| | 0.5 | 27/24107 (0.11%) |
| | 0.75 | 82/35875 (0.22%) |
| 5000 | 0.25 | 0/60121 |
| | 0.5 | 0/120089 |
| | 0.75 | 0/179794 |





Consistent 90% scalability efficiency for sagas

87% scalability efficiency for serializable transactions at 10% transfers
 75% scalability efficiency for serializable transactions at 100% transfers



Conclusion

- Simple programming model abstracting away all distributed systems concerns from the application developer
- Acceptable performance overhead on non-transactional workloads
- Good scalability for saga-based transactions and acceptable scalability for serializable transactions



Future work

Make Stateful Functions complete

- State access without specific function ID
 - Secondary indices
 - WHERE clause
- Aggregates over state across function instances
- Versioning of state / updating functions
- Research based on use cases developing improved benchmarks



on Serverless Stateful Functions

Martijn de Heus 4367839 Web Information Systems

