SLA for Sequential Serverless Chains

Mohamed Elsakhawy and Michael Bauer

Seventh International Workshop on Serverless Computing (WoSC7) 2021
https://www.serverlesscomputing.org/wosc7/papers/p6
Dec 7th 2021

Western

The Problem

- Serverless Execution performance
  - Complex to assess/define guarantees
    - FaaS unique abstraction
    - Dependency on the provider's ecosystem
  - Has a financial impact
- Serverless Chains
  - Composites of Serverless Functions
  - Defining performance guarantees for serverless chains
    - More Complex
    - And more crucial

FaaS SLA Framework (*FaaS2F*)

- A Framework to define performance guarantees
  - Leverges resource (CPU, Memory, Tx and Rx) utilization derivative values (i.e. fingerprints)
  - Classifies a function's execution into SLA abiding/non-abiding
  - Machine Learning classifier

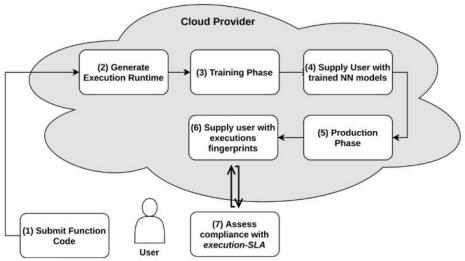

Figure 2: FaaS2F High-level Diagram

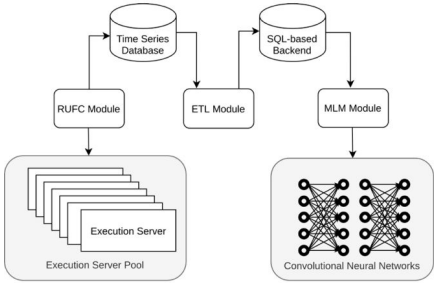

Figure 3: FaaS2F Architecture

Western

Sequential Serverless Chains

We focus on Sequential Serverless Chains

- Fixed size
- Variable size

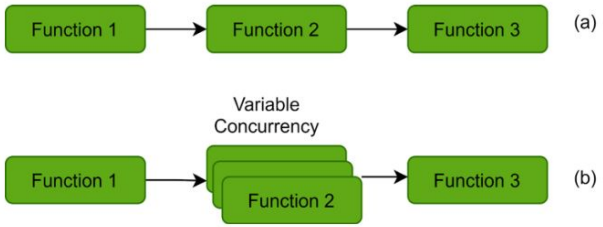

Figure 5: Fixed-size (a) and Variable-size (b) chains

Western

SLA for Sequential Serverless Chains

- We utilize stacked resource utilization fingerprints of the constructing functions.
- FaaS2F is trained on the stacked fingerprints

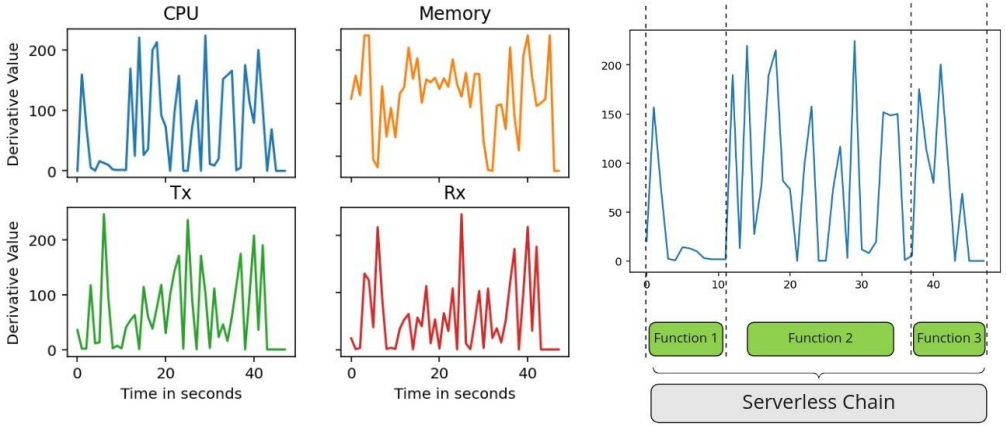

Experimental Validation

- Two Fixed-size Chains
  - Chain 1
    - 3-function ML pipeline
  - Chain 2
    - 6-function Image/video app.
- One Variable-size Chain
  - Chain 3
    - Map-reduce

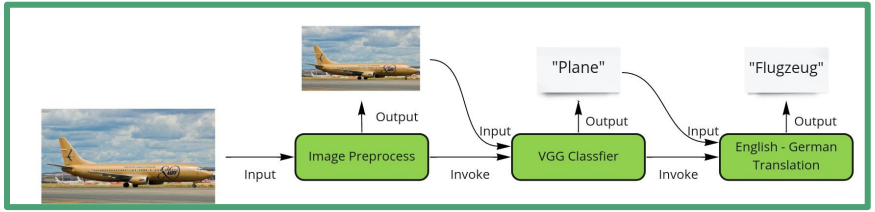

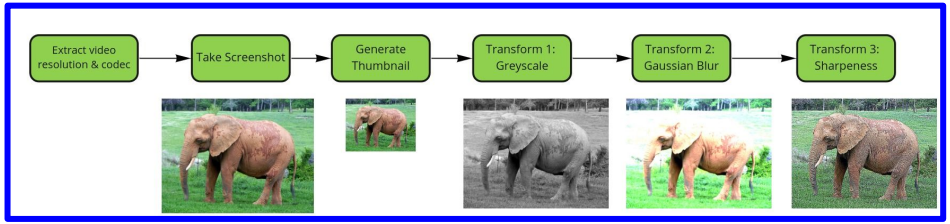

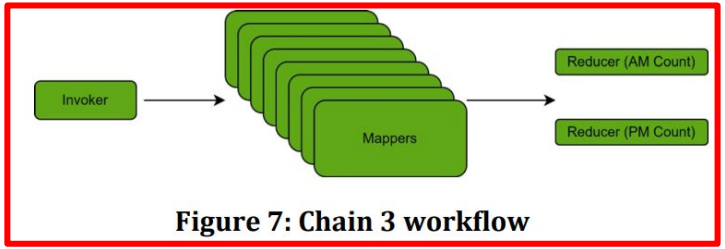

Figure 7: Chain 3 workflow

Western

Classification Accuracy (CNN vs *RandomForests*)

Validation Chain	Resource Utilization Fingerprint			
	CPU	Memory	Tx	Rx
Chain 1	92.8%	89.4%	96.4%	91.8%
Chain 2	98.4%	94%	97.1%	97.6%

Table 1: Chains' Classification Accuracy

Validation Function	Resource Utilization Fingerprint			
	CPU	Memory	Tx	Rx
Chain 1 -Function 1	73.1%	83.4%	96%	95.9%
Chain 1 -Function 2	78.2%	85.7%	97.7%	98.7%
Chain 1 -Function 3	78.8%	83.1%	90%	95.8%

Table 2: Chain-1 Per-Function Classification Accuracy

Validation Function	Resource Utilization Fingerprint			
	CPU	Memory	Tx	Rx
Chain 2 -Function 1	69.4%	85.8%	93%	93.9%
Chain 2 -Function 2	58.2%	56.7%	88.5%	92.5%
Chain 2 -Function 3	87.8%	64.9%	91.4%	87.8%
Chain 2 -Function 4	65.4%	79.2%	88.6%	62%
Chain 2 -Function 5	80.1%	72.4%	91.3%	92%
Chain 2 -Function 6	80.7%	83.5%	90.6%	91%

Table 3: Chain-2 Per-Function Classification Accuracy

Validation Chain	Resource Utilization Fingerprint			
	CPU	Memory	Tx	Rx
Chain 3	79.3%	84.8%	92.7%	82.4%

Table 4: Chain 3 Classification Accuracy

Validation Chain	Resource Utilization Fingerprint			
	CPU	Memory	Tx	Rx
Chain 1	99.5%	98.5%	99.3%	99.6%
Chain 2	98.3%	92%	96.6%	96.3%
Chain 3	91.75%	92.3%	91.2%	91.4%

Table 5: Chains' RFs-based Classification Accuracy

Western

Conclusions and Future work

- FaaS2F can be applied to define SLA for Sequential serverless chains with a high accuracy
- Multiple classifiers can be evaluated
- Expanding our validation chains' categories
- We hope to examine serverless chains encompassing other workflows

Western

Thank you

Questions ?

Western