



Haifa Research Lab

Coverage-Based Metrics for Cloud Adaptation

Yonit Magid, Rachel Tzoref-Brill, Marcel Zalmanovici

IBM Haifa Research Lab



This project is funded under the European Union's Seventh Framework Programme (FP7) under grant agreement 610802



© 2016 IBM Corporation















CloudWave

- Research project developing the next gen Cloud stack
 - Funded by the EU
- Cloud infrastructure and hosted applications collaborate
 - Increase level of service
 - Lower costs
- Dynamically adapts cloud services to their environment
 - Improving service quality
 - Optimizing resource use
- Enables DevOps-like interfaces with the developer



CloudWave (cont.)











Cloud Adaptations

Migration

 Move one (or more) VMs, not necessarily the one(s) reporting problems

Scale-out

- Add more VM(s)/hardware instances
- Scale-up
 - Add more resources to VM(s) running the application
- Application adaptation
 - App. changes its resource usage, e.g. reducing/delaying services



Cloud Adaptations (cont.)











Combinatorial Model



Using the IBM FOCUS Combinatorial Testing engine traditionally used for testing at the development phase



Combinatorial Model (cont.)

• Model based on HW features. May contain any measurable SW/HW features

VM ₁				VM ₂				VM _n			
CPU	Mem	Net	ID	CPU	Mem	Net	ID	CPU	Mem	Net	ID
Low	Low	Low	Grp_1	Low	Low	Low	Grp_1	Low	Low	Low	Grp_1
Med	Med	Med	Grp_2	Med	Med	Med	Grp ₂	Med	Med	Med	Grp_2
High	High	High		High	High	High		High	High	High	
			Grp _n				Grp _n				Grp_{n}

- Test space (size of the cross-product) is huge
 - E.g. an app with 10 modules (VMs) and 4 HW groups: Each VM: 3 x 3 x 3 x 4. In total: 108¹⁰=2.16e20
- Coverage metric
 - Pairs (can be any tuple) of parameter values (in our example ~8K)



Combinatorial Model (cont.)

Example of measures while running tests

		VM ₁		VM ₂					
CPU	Mem	Net In	Net Out	ID	CPU	Mem	Net In	Net Out	ID
Medium	Medium	High	Low	Gr1	Medium	Low	Medium	High	Gr2
Low	High	Low	Medium	Gr1	Low	High	High	Medium	Gr2
High	Low	Medium	High	Gr1	High	Medium	Low	Low	Gr2
Medium	Low	High	High	Gr1	Low	High	High	High	Gr2
High	High	Low	Low	Gr1	High	Low	High	High	Gr2
Medium	Low	Low	Low	Gr1	High	High	Medium	Low	Gr2
High	High	High	High	Gr1	Low	Low	Low	Medium	Gr2

- Each <u>row</u> is a measurement
 - Every pair of values is a coverage task
- Each <u>test</u> may generate more than one row









Risk Assessment Engine





Risk Assessment Engine - Novelty

- Using Combinatorial testing (CT) outside test design phase
 - Learning new "safe" combinations during production
 - Blurring the boundary between testing and production on the cloud
- Use CT to predict missing coverage of potential adaptations
 - Instead of measuring existing combinatorial coverage
- Using CT for load testing
 - As opposed to its traditional use for functional testing
- Using (load) history to prioritize adaptation actions by risk



Risk Assessment Engine (cont.)

- Goal: estimate the risk of encountering a bug
- How is it calculated?
 - Inputs
 - Model describing the world (dimensions)
 - Test concern (coverage metric)
 - Baseline data-points monitored while running tests (test coverage)
 - Live monitoring data (coverage update)



Risk Assessment Engine (cont.)

- How is it calculated?
 - Calculates
 - $risk_{metric} = \frac{|\{(X,Y)| \ (X,Y) \in win \land (X,Y) \in unseen_combos\}|}{|unseen_combos|}$

 $unseen_combos = \{(X, Y) | Y_{VM_{PID}=tr,g_{PID}}\} \cup unseen_combos$

•
$$risk_{migration} = \frac{\left| \{ (X, Y) \mid Y_{VM_{PID} = trg_{PID}} \land (X, Y) \in unseen_combos \} \right|}{|unseen_combos|}$$

$$\begin{split} v_{1} &= |\{(X,Y) \mid X \in orig_{VM} \land Y \in clone_{VM}\}| \\ \bullet \quad v_{2} &= |\{(X,Y) \mid X \in clone_{VM} \land (X,Y) \in old_model \land (X,Y) \in unseen_combos\}| \\ risk_{scale_out} &= \frac{v_{1} + v_{2}}{v_{1} + v_{2} + |unseen_combos|} \end{split}$$

Risk values are taken as part of the world which tests did not cover

Using a BDD-based symbolic computation for scaling









Preliminary Results

- Tested a TelCo application on top of CloudWave
- Risk metric triggering adaptation



Haifa Research Lab





Haifa Research Lab



