Predicting VM Behavior for DRS & DPM Cost/Benefit Analysis

Chengwei Wang, Canturk Isci*, Chirag Bhatt*, Ganesha Shanmuganathan, Anne Holler, Minwen Ji

* Mentor

Background / Introduction

- Distributed Power Management (DPM) uses prediction to make decision on tradeoff between cost and benefit after power-off
- Distributed Resource Scheduler (DRS) uses prediction to make decision on tradeoff between cost and benefit after vmotion
- Our contributions: 1. Define Fundamental Concepts for VM Behavior Prediction 2. Extract Workload Traces from Real Customer Dump Files for Prediction Analysis and Implementation. 3. Use Traces to Implement Current Predictor (Baseline), Perfect Predictor (ORACLE) & Improved Novel Predictor (GlitchFree)

Fundamental Concepts

• What's VM Behavior Prediction ?

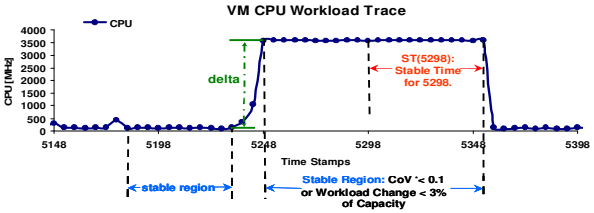

- At any sampling point, foreseeing its stable time based on former trace
- Foreseeing gap between beginning value of next stable region & current workload value, delta_i (d(i))

• What's EWMA Prediction? "Exponentially Weighted Moving Average." Prediction Method

$$EWMA_ST(k) = a*OP + a*(1-a)*SR(1) + a*(1-a)^2*SR(2) + a*(1-a)^i*SR(i)$$

- EWMA_ST (k) : EWMA Prediction of Stable Time at Sample k.
- OP: Open Period ; SR(i): Previous ith Stable Region; a: Weight (0.5)

• How to Evaluate a Predictor? Goodness Measure

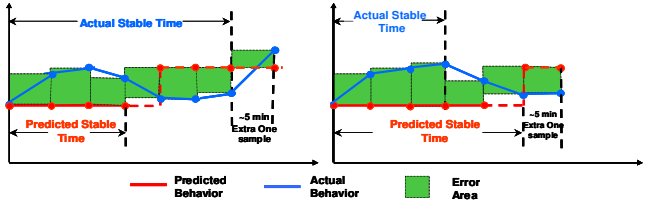

- Goodness Measure (GM) = Total Error Area / Number of samples
- Length = max (actual_stable_time, predicted_stable_ time) + 5 mins

• What's a "Glitch" ?Glitch is the bursting in VM Trace

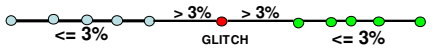

VM Behavior Predictors & Evaluation

• Result First -- ORACLE vs Baseline vs GlitchFree

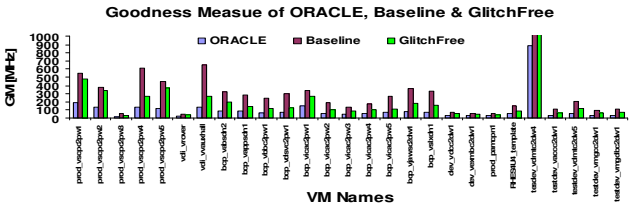

- GlitchFree outperforms Baseline in all VMs, on average 1.7X improvement
- ORACLE predictor is the bound of prediction accracy – potential of improvement

• Novel Predictor – GlitchFree ("Forward" EMWA without Glitch Pollution)

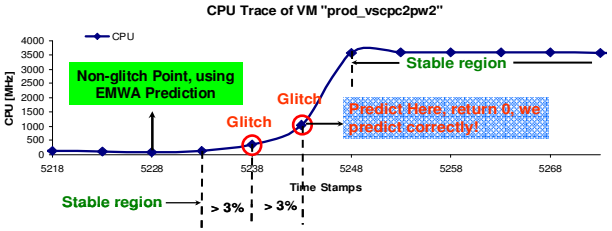

• Perfect Predictor – ORACLE

```
oracle_d(i)  = max(0, (d(i)+d(i+1))/2)
oracle_ST(i) = CPU_ACTUAL_ST(i)
```

→ To minimize error

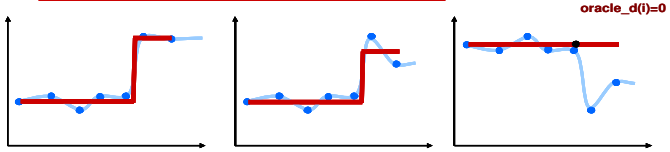

• Current Predictor – Baseline ("Backward" EMWA)

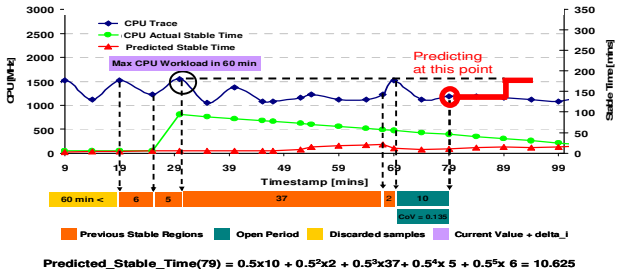

Predicted_Stable_Time(79) = 0.5x10 + 0.5²x2 + 0.5³x37 + 0.5⁴x 5 + 0.5⁵x 6 = 10.625

Statistic Analysis for Predictor Implementation & Improvement

• We implement & analyze predictors on Real Customer Workload Traces .

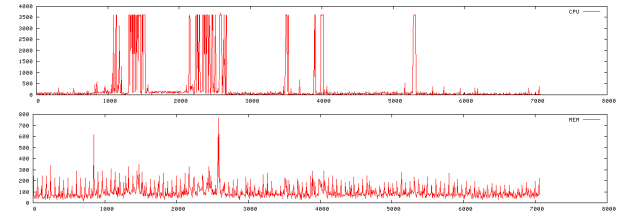

- Extract VM Traces from customer drm dump files
- 2000+ Customers' VM traces from Natixis, Boeing, USPS, Fernuni, Lockheed.
- Check @ http://engweb.vmware.com/~cisci/CHENGWEI/public_html/Trace_Plots/

• Why GlitchFree Outperforms Baseline ?

- Reason 1 : High Probability of Glitches

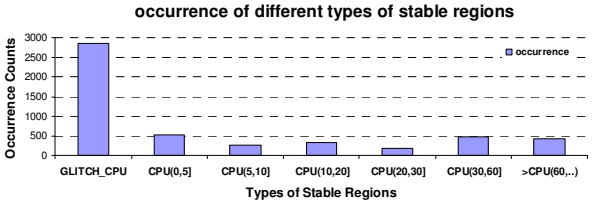

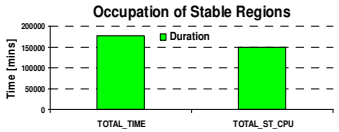

- With high probability, Baseline will encounter a glitch, therefore it gets polluted history information.
- GlitchFree doesn't predict at glitches & it always choose the max no-glitch workload in pre 60 mins as delta_i

- Reason 2 : Backward Prediction Manner

- At each predicting point, Baseline goes back to find open period and stable regions. Sometimes, they are not the actual ones
- GlitchFree calculates open period and previous stable regions from the first sample point, which is a forward manner with precise history information

Future Work / Reference

• Probability Based GlitchFree Predictor (in progress)

Use on-line detection of stable region probabilities for making decisions on future stable region which is minused by open period to get the predicted stable time

• Related Publication

Isci et al. "Long-term Workload Phases: Duration Predictions and Applications to DVFS" MICRO'05