Predicting Web Service Levels During VM Live Migrations

Helmut Hlavacs and Thomas Treutner
Research Group Entertainment Computing, University of Vienna, Austria

Introduction

- **VM live migration important for energy efficiency**
- Enables us to establish energy efficient target distribution of VMs
- Supposedly no perceivable service downtime while live migrating
- Live migration is resource intensive (iterative page copying)
- **Experiments: Influence on service levels while migrating?**
- **Modelling: Predict service levels based on utilization?**

Scenario

- Virtualized data center, static consolidation (P2V)
- Provisioning for peak load, still bad energy efficiency
- E.g., 9-5-cycles, very low utilization at night
- **Live migration enables dynamic consolidation**
- **But: Seldom used, fear of possible side effects**
- **⇒ Identify and quantify effects on (web) service levels**
- **⇒ Find most influencing utilization metrics**

Experiment

- **Two servers, a single VM, migrating forth and back**
- VM disk image on central node (Gbit, open-iscsi)
- qemu-kvm VM: Linux, Apache2, PHP5, MediaWiki
- SQL VM and load generation on an extra nodes
- **Logging utilization of servers and VM, more than 100 variables**
- **Rise load from 50 to 600 concurrent virtual users and back**
- **Migrate every 15min, track response time of last 5min**
- Maximum allowed response time: 1s

Data Overview

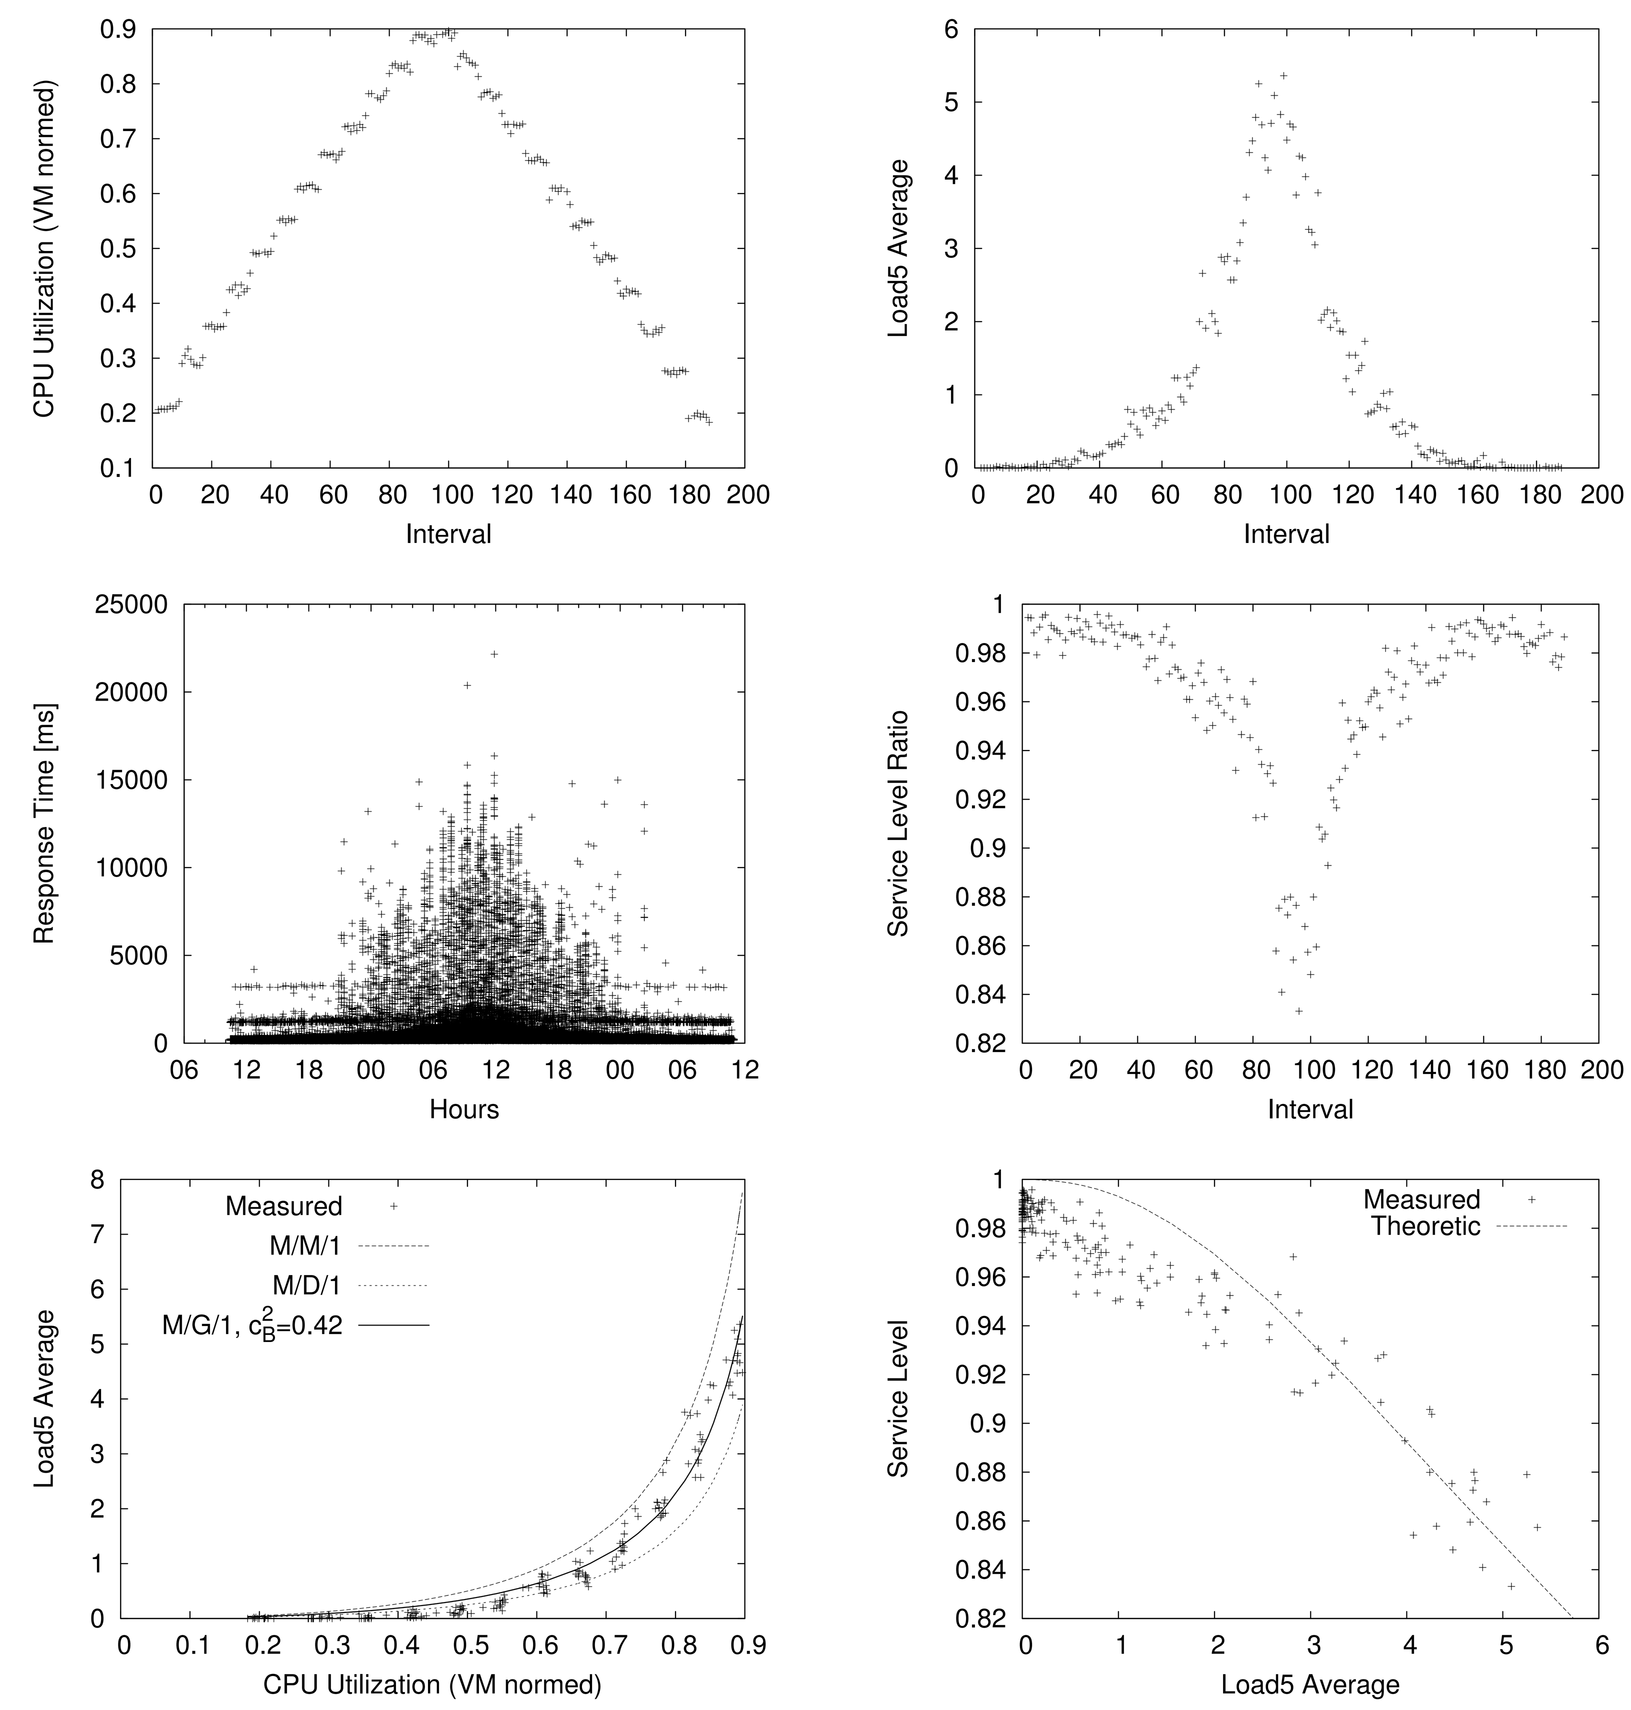

- We can interpret the `UNIX load` as approximation to $\overline{Q}$, the average number of jobs in a Markovian M/M/1 queue, and the VM's CPU utilization as ρ, the system utilization: $\overline{Q}_{M/M/1} = \frac{\rho^2}{1-\rho}$
- `UNIX load` is exponentially averaged by definition and the service times are not necessarily exactly exponentially distributed: $\overline{Q}_{M/G/1} = \frac{\rho^2}{1-\rho} \cdot \frac{(1+c_B^2)}{2}$
- For deterministic service times $c_B^2 = 0$, resulting in $\overline{Q}_{M/D/1} = \frac{\rho^2}{2(1-\rho)}$
- Simple linear regression delivers the coefficient $c_B^2 = 0.42$
- $P(T \le x) = F_T(x) = 1 - e^{-\mu(1-\rho)x}$, the theoretical probability that the response time T is lower than or equal to a limit x for a given service rate μ

Service Levels for Low/Medium/High Workload Scenarios

- **Low** (top right): Slightly increased response times during live migration, seldom response time violations
- **Medium** (bottom left): SLA ratio generally satisfies the 97% limit
- **High** (bottom right): Often and heavy violations, unacceptable low service levels, typically decreased by 20-25% percentage points

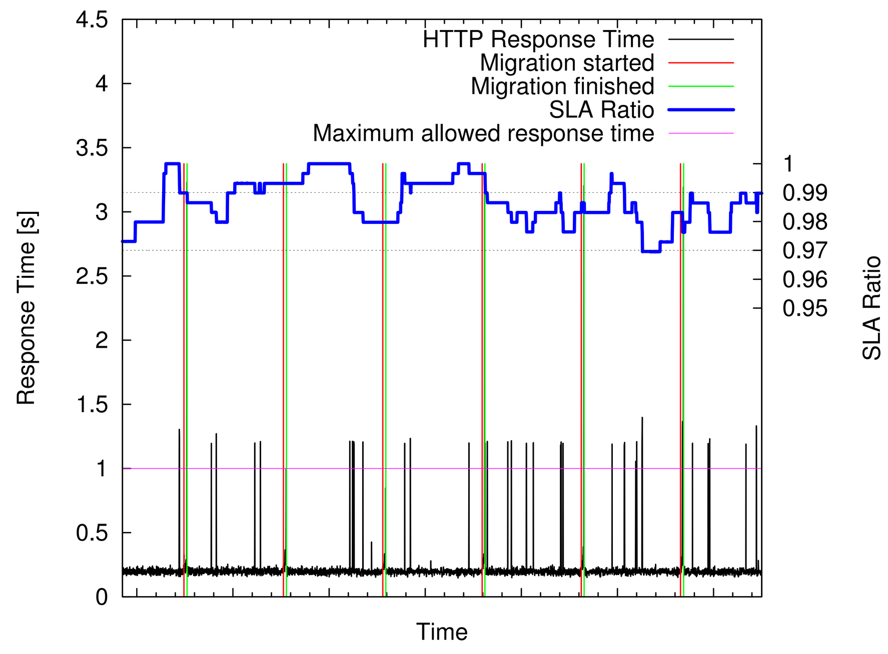

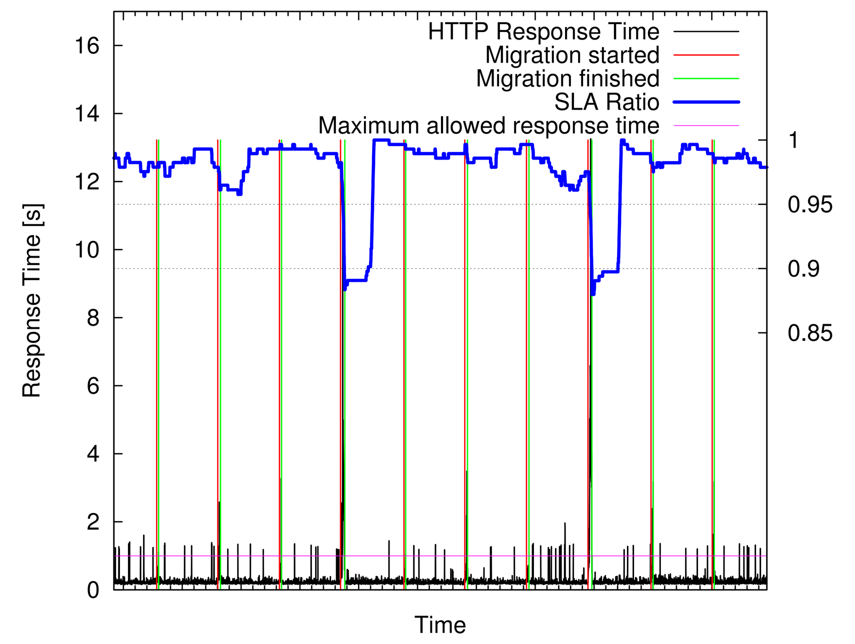

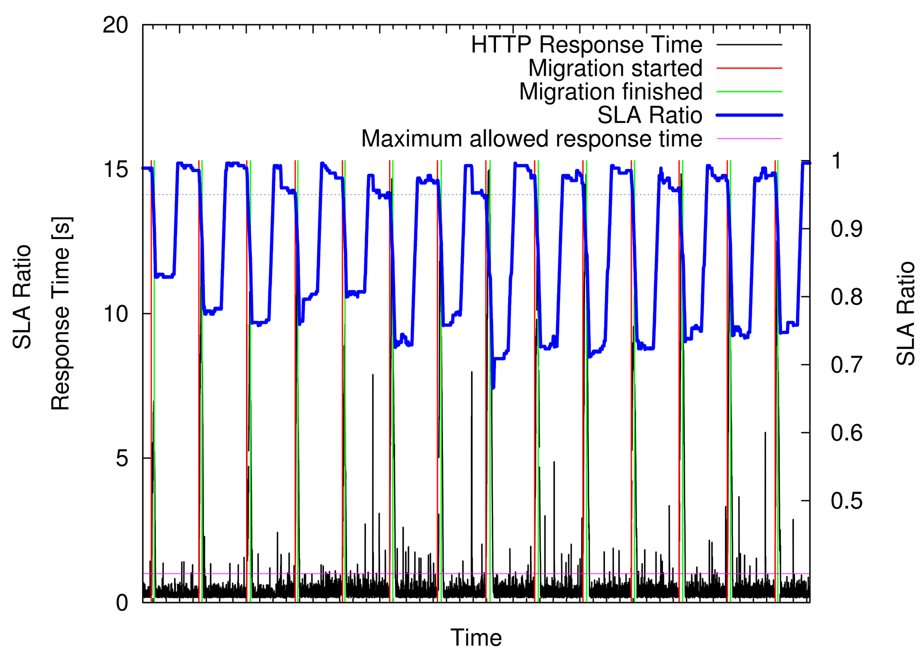

Model Selection

- **Stepwise model selection: Akaike Information Criterion**
- Finds trade-off between number of parameters (model size) and goodness of fit (model quality)
- **For comparison: Exhaustive all-subsets-regression (LEAPS)**
- LEAPS: Find best of all possible models for given range of model size
- Computationally intensive even if number of variables is limited
- R^2_{Adj} increases slightly with increased model complexity. `UNIX load5` contributes R^2_{Adj} of **∼ 90%**

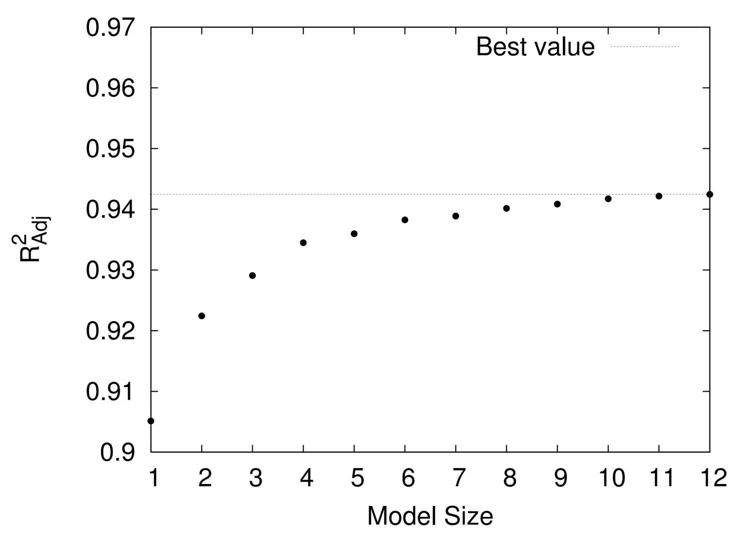

Most Influencing Model Parameters

Variable	Meaning	Estimate	Std. Error	Pr(> |t|)
Intercept		2.395e+00	5.069e-01	3.00e-06
wp01_load5	VM `UNIX load5`	-1.871e-02	2.627e-03	3.67e-12
wp01_swapUsed	VM swap used	-7.656e-07	8.809e-08	< 2e-16
wp01_residentSize_SQ	Suared amount of resident memory used by the `qemu-kvm` process.	-4.652e-14	1.652e-14	0.00506
src_host_cpu_proc_s	Tasks created/s, source host.	-2.475e+00	9.166e-01	0.00716
src_host_cpu_proc_s_SQ		1.091e+00	4.133e-01	0.00856
wp01_cpu_util_vmnorm_SQ		-1.328e-01	3.187e-02	3.63e-05
wp01_cpu_util_vmnorm	CPU util measured inside VM	9.517e-02	2.316e-02	4.64e-05
wp01_load5_SQ	Squared `UNIX load5` of the VM.	-1.140e-03	1.918e-04	5.22e-09
wp01_freeMemRatio_SQ	Squared ratio of free memory inside the VM.	1.976e-02	9.462e-03	0.03727

Conclusions

- **Impact of live migration on SL depends on amount of workload**
- **Tighter SLAs can be fulfilled during low and medium workload**
- **Migrating during high load causes massive decrease of service level**
- **Service level variance during a live migration to 90% predictable using only a single variable**, the `UNIX load5 average`, models with 12 variables can explain 95% of variance
- Systems using live migration as a mechanism to realize a more energy efficient target distribution and have service level targets need to consider the `UNIX load average`, but typical hypervisors do not collect/export this information
- Hypervisors should be extended to export `load` information (cf. free memory)

Future Work

- Influence of additional VMs (idle, utilized, mixed)
- `Linux` and `qemu-kvm`: Kernel Samepage Merging (KSM)
- Database VM migration, currently taboo due to potentially severe influence
- `qemu-kvm` parameters: Bandwidth limits, maximum allowed downtime
- Predict migration delay, energy consumption, service downtime

http://cs.univie.ac.at/research/research-groups/entertainment-computing/ <firstname>.<lastname>@univie.ac.at