Dynamically Switching Data Centre Management Strategies

Gaston Keller, Michael Tighe, Graham Foster, Hanan Lutfiyya and Michael Bauer
Contact: gkeller2@csd.uwo.ca

Western
UNIVERSITY · CANADA

Abstract

Data centres nowadays must *minimize power consumption (goal #1)* while satisfying their clients' resource needs *(goal #2: minimize Service Level Agreement (SLA) violations)*. Pursuing both goals at the same time is challenging, as the goals are often in conflict. **We propose achieving both goals by dynamically switching between two management strategies, each with a single goal, based on current data centre state. Dynamic strategy switching offers improved results over a single management strategy.**

Problem & Solution

We define a management strategy for a virtualized data centre to consist of a set of policies that guide the following three operations: (i) **VM Placement**: placing newly arrived VMs, (ii) **VM Relocation**: migrating (moving) VMs off of highly utilized machines to avoid performance degradation, and (iii) **VM Consolidation**: migrating VMs off of underutilized machines so that they may be shut down to save power.

Developing a management strategy to both minimize power consumption and SLA violations is challenging, as the goals are often in conflict. We propose achieving both goals by dynamically switching between two management strategies -- each one pursuing one of the goals -- based on current data centre state. We propose three methods for dynamically switching strategies:

1. SP-DSS: Defines thresholds for the metrics *SLA violations* and *power efficiency*, and switches to the strategy whose threshold has been exceeded;

2. Goal-DSS: Calculates how well the system is performing towards the goals of *0% SLA violations* and *optimal power efficiency*, and switches to the strategy that would improve achievement of the most distant goal;

3. Util-DSS: switches strategies based on the rate of change of overall data centre utilization (i.e., rate of change of data centre load).

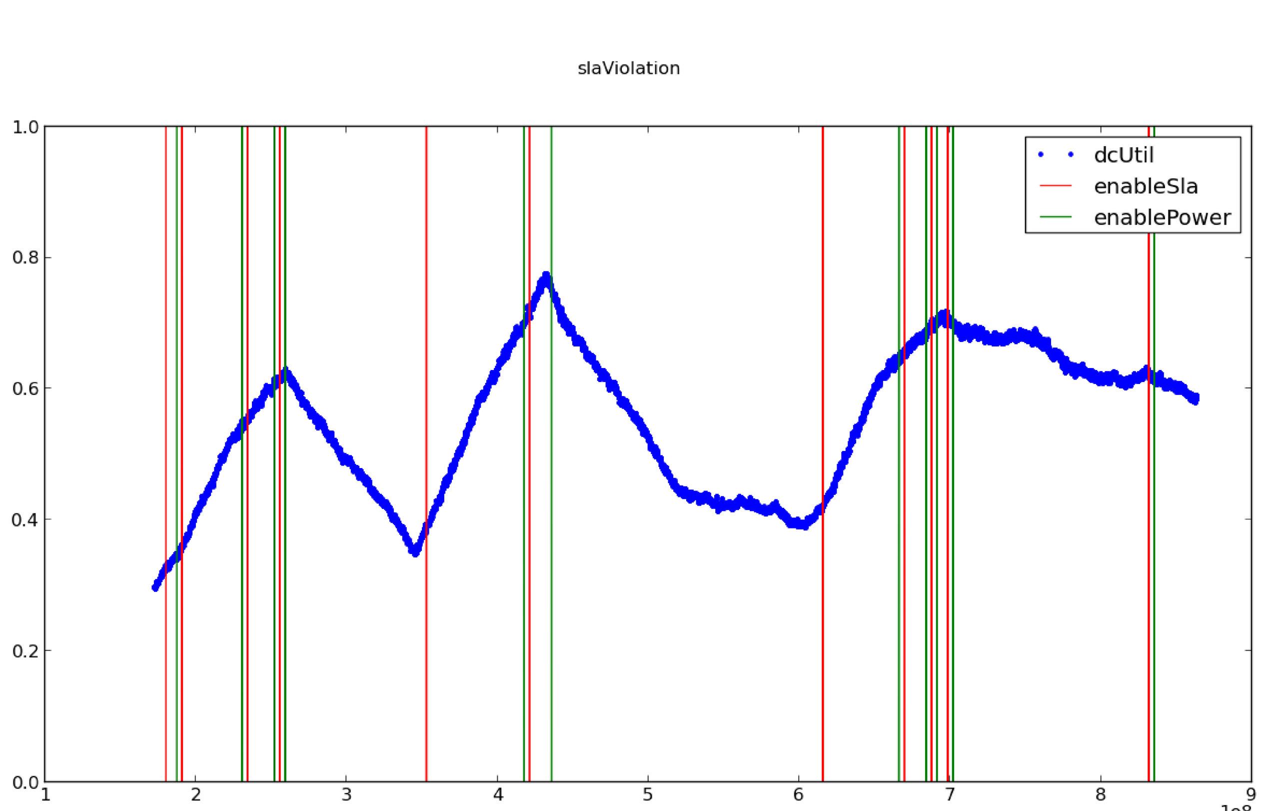

Util-DSS in action during one of the experiments.

Evaluation

We developed two management strategies, *Power* and *SLA*, and a third one, *Hybrid*, combining the VM Placement and VM Relocation policies of SLA with the VM Consolidation policy of Power. The three DSS meta-strategies switched between the Power and SLA strategies.

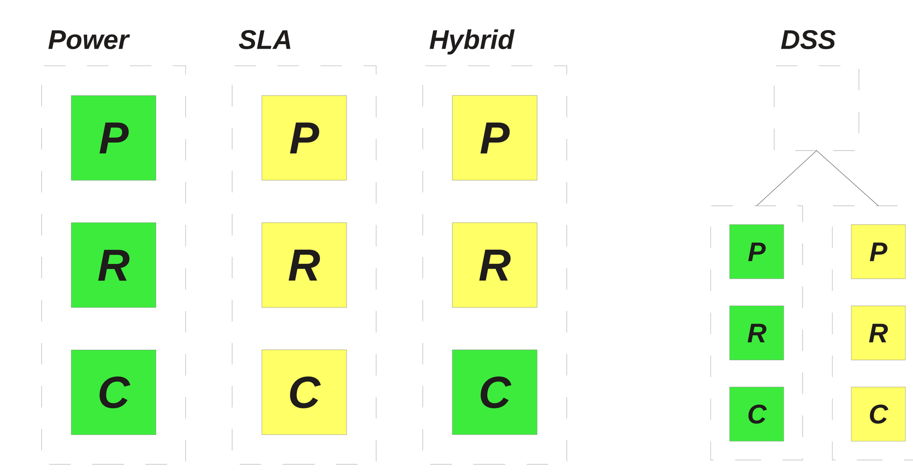

We evaluated all strategies through simulation using DCSim, a data centre simulation tool [1]. The simulated data centre consisted of 200 hosts, whose power consumption was calculated using results from the SPECPower benchmark [2]. The number of VMs within the data centre was varied dynamically between 600 and 1600 to simulate the arrival and departure of VMs. VMs' resource requirements were trace-driven, using 5 different real-world traces: the ClarkNet, EPA, and SDSC traces [3], and two different job types from the Google Cluster Data trace [4].

Hosts	HW Specifications
Small	2 dual-core 3GHz CPUs, 8 GB RAM
Large	2 quad-core 2.5GHz CPUs, 16 GB RAM

VMs	HW Specifications
Small	1 virtual core (minimum 1.5GHz), 512 MB RAM
Medium	1 virtual core (minimum 2.5GHz), 512 MB RAM
Large	2 virtual core (minimum 2.5GHz), 1 GB RAM

Results

Each management strategy was evaluated with the same set of 100 randomly generated workload patterns. Each experiment was repeated only once per workload pattern, as the simulation is deterministic. Results were averaged across all workload patterns and the standard deviation is shown in square brackets.

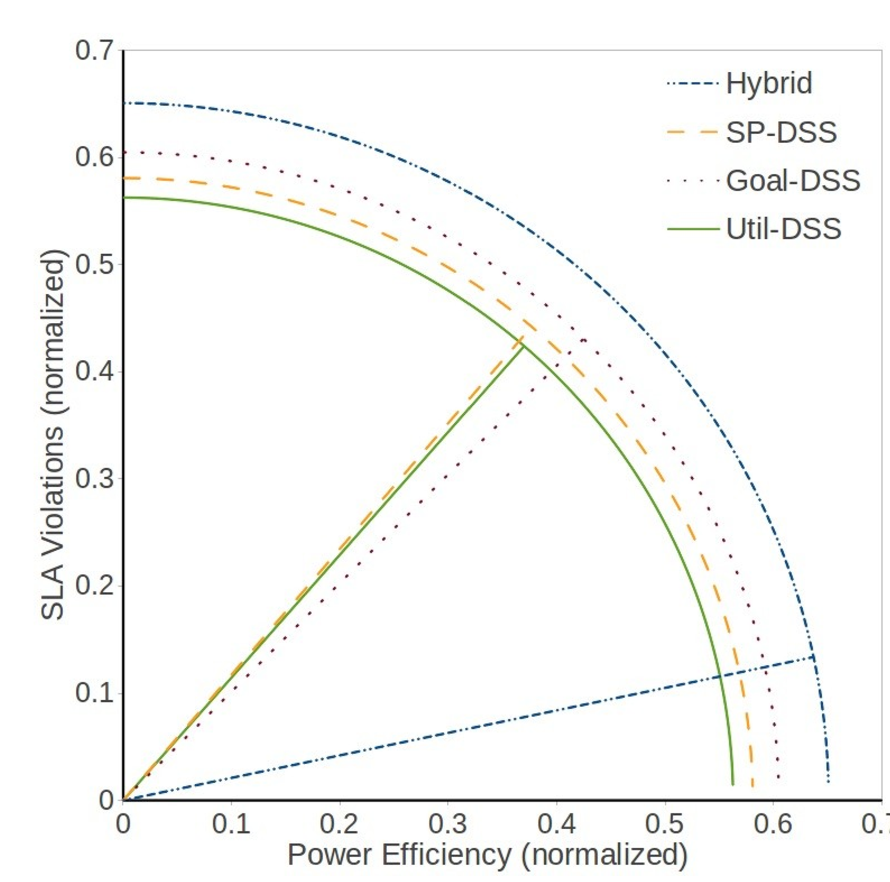

Strategies' scores.

	SLA	Power	Hybrid	SP-DSS	Goal-DSS	Util-DSS
Avg. Active Host Util.	75% [0.4]	88% [0.4]	81% [0.4]	82% [1.0]	81% [1.0]	82% [1.0]
# of Migrations	15818 [2292]	24378 [3311]	14643 [1930]	19187 [3420]	19448 [2754]	19580 [3047]
Power Consumed (kWh)	5488 [703]	4384 [519]	5049 [679]	4766 [548]	4821 [612]	4778 [583]
Power Efficiency	60.6 [2.4]	75.2 [2.0]	65.9 [2.7]	69.7 [1.8]	69.0 [2.4]	69.8 [2.3]
SLA Violation	0.033% [0.01]	0.474% [0.05]	0.092% [0.01]	0.228% [0.05]	0.223% [0.02]	0.220% [0.05]
# of Strategy Switches	N/A	N/A	N/A	20 [5]	54 [5]	30 [10]
s_{norm}	0.0	1.0	0.135 [0.01]	0.440 [0.10]	0.430 [0.04]	0.425 [0.09]
p_{norm}	1.0	0.0	0.636 [0.06]	0.380 [0.09]	0.425 [0.05]	0.373 [0.08]
Score	1.0	1.0	0.651 [0.05]	0.594 [0.05]	0.607 [0.03]	0.576 [0.04]

Results per Strategy.

Conclusions

DSS improves overall performance by about 40% when compared to single-goal strategies (SLA and Power), and by 7-12% when compared to a hybrid strategy designed to pursue both goals simultaneously.

References

[1] M. Tighe, G. Keller, M. Bauer, and H. Lutfiyya, "DCSim: A data centre simulation tool for evaluating dynamic virtualized resource management," in SVM Proceedings, 6th Int. DMTF Academic Alliance Workshop on, Oct. 2012. **[2]** (2012, Oct.) Specpower_ssj2008 benchmark. Standard Performance Evaluation Corporation. [Online]. Available: http://www.spec.org/power_ssj2008/ **[3]** (2012, Oct.) The internet traffic archive. [Online]. Available: http://ita.ee.lbl.gov/ **[4]** (2012, Oct.) Google cluster data. Google Inc. [Online]. Available: http://code.google.com/p/googleclusterdata/