SOM-based Behavioral Analysis for Virtualized Network Functions

Giacomo Lanciano[2,1] Antonio Ritacco[1] Tommaso Cucinotta[1] Marco Vannucci[1]
Antonino Artale[3] Luca Basili[3] Enrica Sposato[3] Joao Barata[4]

[1]Scuola Superiore Sant'Anna, Italy [2]Scuola Normale Superiore, Italy [3]Vodafone, Italy [4]Vodafone, Portugal

Introduction

Network operators are facing increasingly demanding requirements coming from the new connectivity scenarios of the future Internet. A recent trend is the wide and massive deployment of **cloud computing** technologies, for realizing flexible infrastructures management strategies able to cope with the new challenges in the area.

IP convergence has facilitated the migration of networking services from the traditional deployment of physical appliances, towards the novel **Network Function Virtualization (NFV)** approach [1]. Virtualized Network Functions (VNF) are essentially software applications that can be deployed on a private, virtualized infrastructure of the operator.

NFV data centers require an efficient distributed **monitoring infrastructure** that gathers continuously system-level metrics from all the physical hosts and VMs deployed in the system. Such data are fed to data center automation functions and human operators, to enable necessary operations.

Anomaly detection is among the principal concerns of operators, since the capability of detecting suspect performance degradation is fundamental to the purpose of establishing automated proactive strategies to minimize the risk of SLA violations.

Proposed Approach

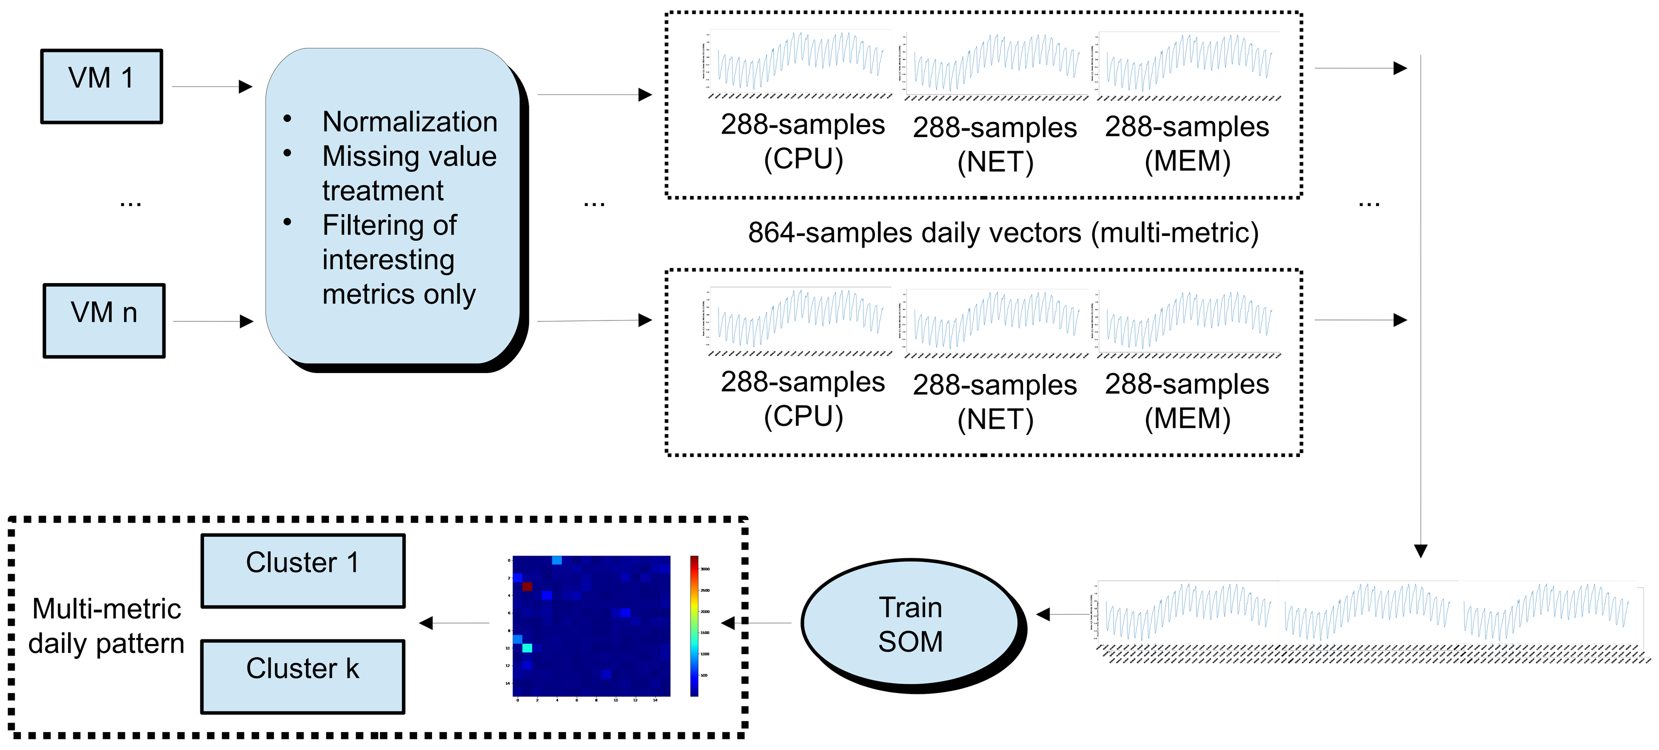

Figure 1. Overview of the SOM-based clustering workflow.

The proposed approach performs a pattern analysis of the VMs behavior, focusing on the resource consumption metrics, related to the underlying infrastructure, as well as the application-level metrics.

A **Self-Organizing Map (SOM)** is used to cluster VM metrics patterns, leveraging on its ability to preserve the topology in the projection. The overall workflow – summarized in Figure 1 – consists in:

1. The raw data are preprocessed to address possible data-quality issues and to retain only the information related to the relevant metrics.
2. The input samples to the SOM are then constructed by diving the time horizon under analysis according to a predefined period (e.g., one day) and merging the individual metrics data related to the same period in a single vector, for each VM separately.
3. Then, such data are fed to the SOM that outputs for each of them the best matching neuron, providing a clustering.

The output of such analysis can be used by a data center operator to **visually inspect** the behaviors captured by the trained neurons, to spot possible suspect/anomalous ones and check which VMs are associated with them. Furthermore, a VM can be observed trough its movement among its best matching units (BMUs) during the time horizon under analysis, so that any changes to a *"far-away"* BMU could be used to **trigger an alarm**.

Experimental Results

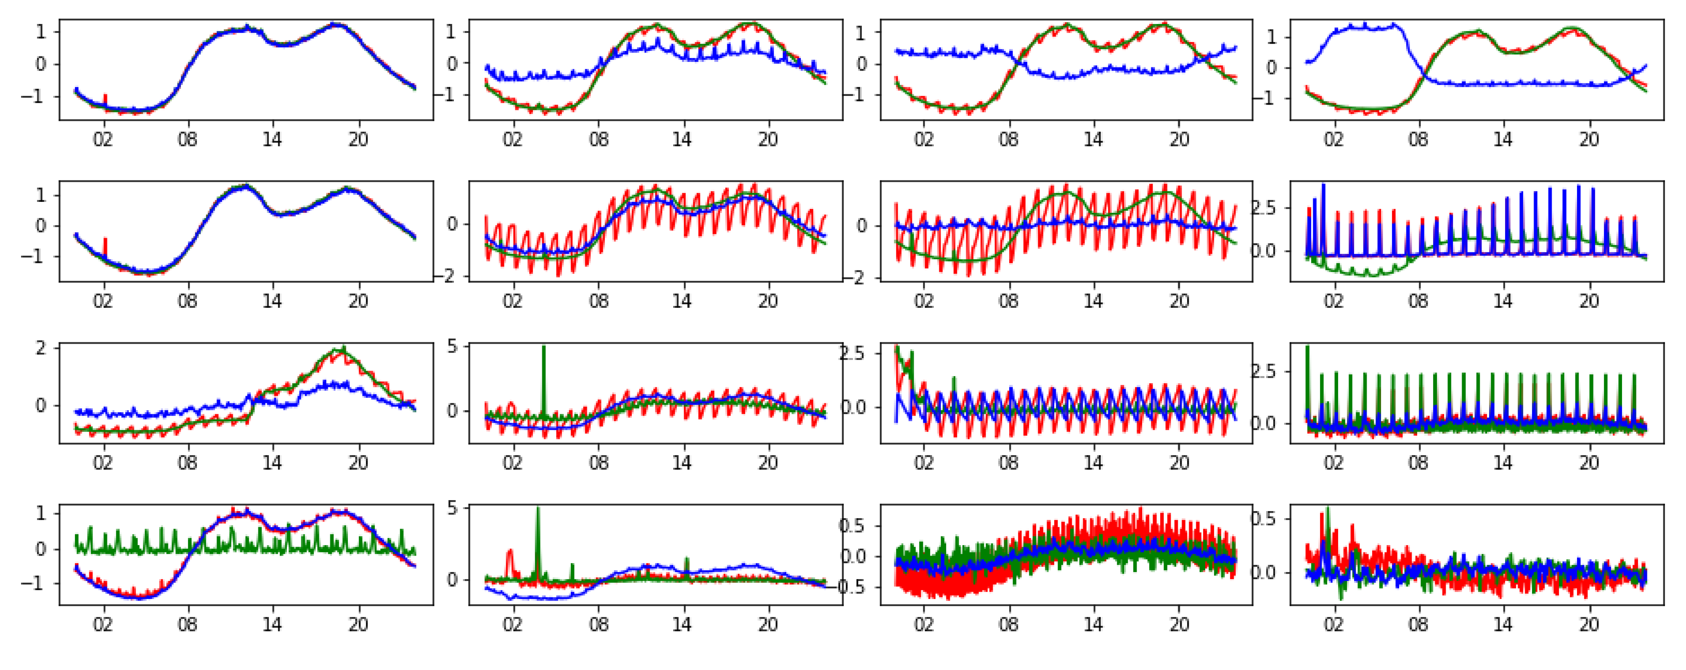

Figure 2. Example of INFRA resource consumption clusters identified with the multi-metric SOM analysis. The red, green and blue curves in each plot correspond to the `cpu|usage_average`, `net|usage_average` and `cpu|capacity_contentionPct` metrics, respectively.

An example output of the proposed approach is the set of clusters highlighted in Figure 2. Each subplot represents the weights in the trained SOM, that jointly identify the characterizing daily behavior for the VMs that fell into that cluster. Some of the insights revealed by the analysis are:

- One of the most recurrent patterns is the one identified by the top-left neuron, occurring in 35.6% of the input samples.
- `cpu|capacity_contentionPct` exhibits a suspect behavior, since its shape follows closely the daily traffic pattern of the involved VMs, whereas in a healthy scenario it is expected to have this metric to stay flat at zero, or at most to undergo a slight increase during peak hours only.
- A significantly different pattern is the one that can be observed in the top-right neuron of Figure 2, representing the 7.84% of the observed patterns. It exhibits a higher CPU contention during night, when the VM has lower traffic, than during the day.

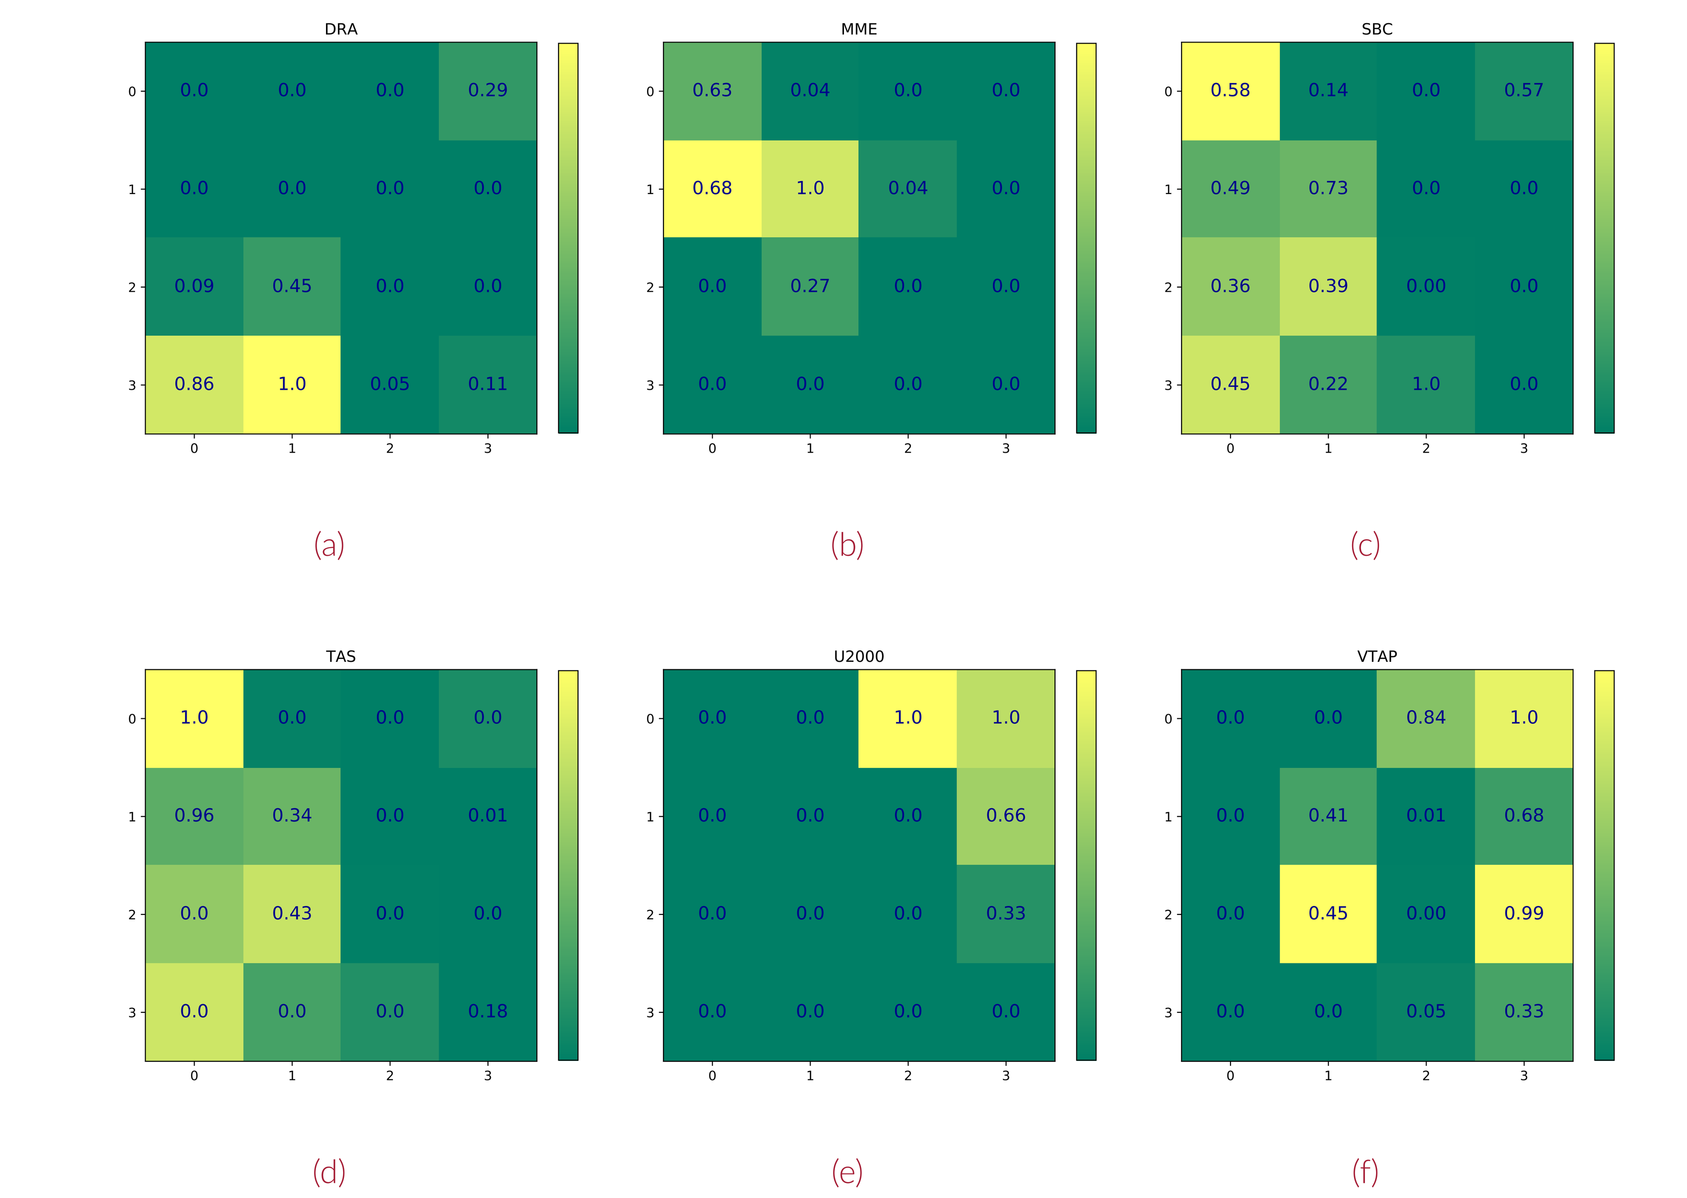

Figure 3. SOM clusters and corresponding per-VNF hitmaps identified by the multi-metric SOM-based analysis. For confidentiality reasons, the total number of hits in the hitmap cells has been rescaled to 1, to avoid disclosing the actual figures.

Starting from the output reported in Figure 2, it possible to study how different VNFs behave in terms of their resource consumption patterns. In Figure 3, the hitmaps report how many daily patterns of each given VNF map onto the SOM neurons, highlighting that:

- Both **SBC** and **TAS** VNFs mostly follow the usual "nightly/daily" pattern (i.e., neuron $(0, 0)$), characterized by a low workload over the night and a high workload over the day, with peaks around noon and 6pm.
- **DRA** VNF samples are captured by neurons $(3, 0)$ and $(3, 1)$ and exhibit the classical "nightly/daily" pattern for the `cpu|capacity_contentionPct` metric and periodic peaks (every 30 minutes) for the others.
- A consistent number of **VTAP** VMs are captured by neuron $(2, 3)$, characterized by hourly peaks.

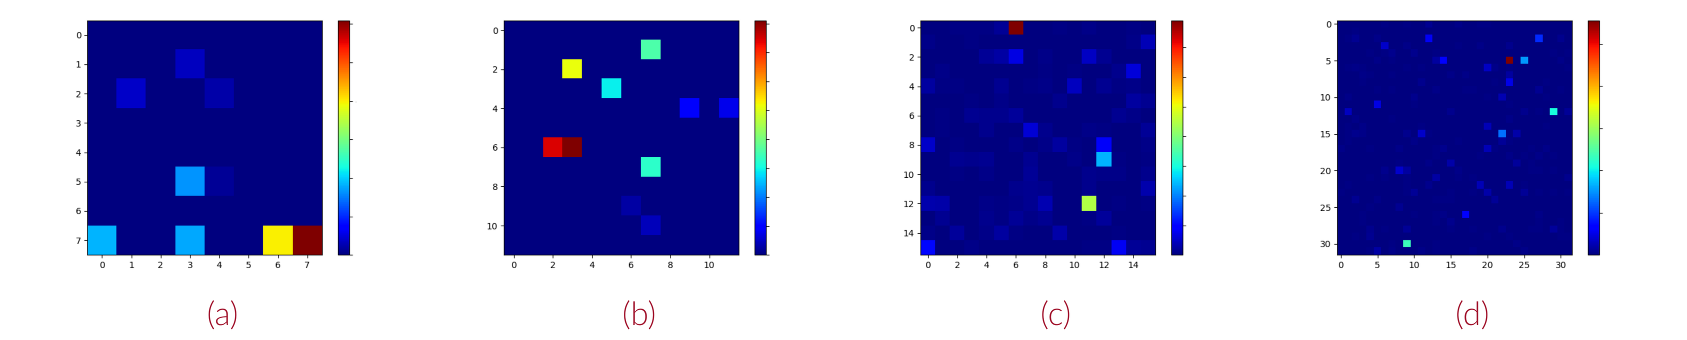

Figure 4. SOMs with low σ values: (a) 8×8, σ: 0.1, lr: 0.2; (b) 12×12, σ: 0.1, lr: 0.2; (c) 16×16, σ: 0.1, lr: 0.9; (d) 32×32, σ: 0.1, lr: 0.8 For confidentiality reasons, the scale has been omitted.

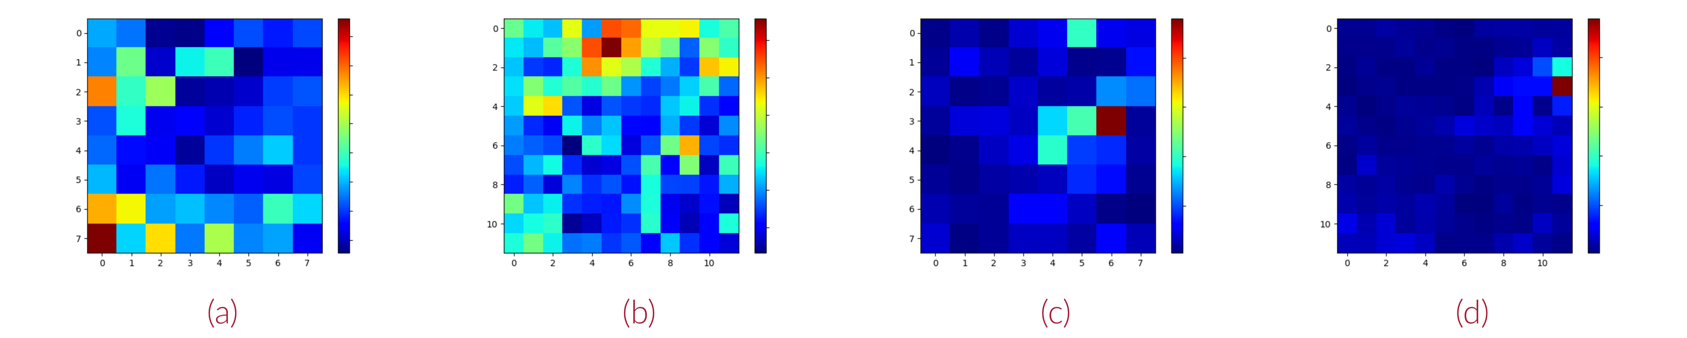

Figure 5. SOMs with high σ values: (a) 8×8, σ: 0.6, lr: 0.2; (b) 12×12, σ: 0.6, lr: 0.2; (c) 8×8, σ: 0.6, lr: 0.9; (d) 12×12, σ: 0.6, lr: 0.9 For confidentiality reasons, the scale has been omitted.

An extensive grid search has been used to identify **optimum SOM hyper-parameters**. In particular, a total of 1600 different configurations have been tested, uncovering some key insights:

- Using a **low** σ (i.e., neighborhood radius) with a **low learning rate** yields the worst results: very few BMUs capture more than 95% of data, resulting in higher quantization errors. Examples of such runs are reported in Figure 4.
- SOMs larger than 12×12 require **high** σ and very **low learning rate** to result in a low quantization errors. Although, in such cases the results tend to become unreadable, since too many neurons specialize on similar patterns. Examples of such runs are reported in Figure 5.

Conclusions

This work presents a SOM-based technique for the classification of behavioral patterns of VMs resource consumption in NFV data centers, currently used at Vodafone network operator. Possible future works:

- Improve automation: while some hyper-parameters can be effectively tuned via a grid search, others need to be tuned manually by operators, depending on the achieved results.
- Enrich the approach by using Deep Learning models for time-series classification, in order to build more effective anomaly detection models.

References

[1] NFV Industry Specif. Group.
Network Functions Virtualisation.
Introductory White Paper, 2012.

giacomo.lanciano@sns.it