The Data-centre Whisperer: Relative Attribute Usage Estimation for Cloud Servers

Ruairí de Fréin,
Institiúid Teicneolaíochta Bhaile Átha Cliath, Dublin Institute of Technology, Ireland

***Abstract*—We show that the relative usage of the different attributes of a cloud server can be estimated under time-varying loads. We demonstrate the effectiveness of these estimators by determining how user requests for video –from a video server– affects its usage. Relative Attribute Usage (RAU) estimators are designed by (1) formulating a generative model for the server attributes; (2) using the fact that the load signal has compact support compared to non-idealities in the server's behaviour in the time-frequency domain; and (3) using power-weighting to refine the estimates. The resulting estimators have low complexity. This motivates their candidacy when attribute usage estimates are required for run-time outage diagnosis routines, a task which is commonly referred to as "data-centre whispering". We demonstrate the application of these estimators on a Cloud-testbed in three practical scenarios, when the server is under a (1) periodic, (2) step-increasing and (3) flash-crowd load.**

***Index Terms*—Power-weighted estimators, Blind Source Separation, Video-on-Demand.**

I. Introduction

Understanding the current and future performance of Cloud Computing systems is challenging [1]. (1) The flexibility of programmable computing architectures, promised by Software Defined Networking (SDN) [2]; (2) the desire for dynamic resource [3] and service allocation to improve the efficiency of existing computing infrastructures (by increasing the efficiency of server usage [4]); and finally, (3) the reality of a heterogeneous computing substrate –as data-centres slowly evolve to meet the needs of SDN, reuse existing machines, and incorporate new technologies (cf. the discussion in [1] about the composition of an operational data-centre)– provide a challenging network management scenario for the Network Manager (NM) of the present day.

One of the chief attributes of the *self-regulating* or *self-optimising* system envisaged by SDN is that not only can services and resources migrate (different machines can be used for different tasks at different times), but the load on these systems can change at any time, in response to an arbitrary train of user requests [5]. Therefore, it is important to be able to estimate how this dynamically changing environment is affected by different usage patterns [6].

We discuss the problem of probing a server with a train of user requests for a Video-on-Demand (VoD) service [7] in a Cloud-like set-up; we determine how different attributes of the server respond under different usage loads [8]. This work is relevant to the Blind Source Separation (BSS) community as we demonstrate how work on Time-Frequency (TF) disjointness in [9], [10], for example, can be applied in a new domain with good effect. We call this process *Relative Attribute Usage* (RAU) estimation. RAU estimation learns a typical signature for the performance of a machine/server, in response to user requests, which arise due to time-varying system load. We develop a model for RAU estimation and derive RAU estimators. We posit that they provide the means for the NM to passively or actively probe the secrets of an operational data-centre, in *run-time* in a lightweight manner. They provide the intelligence required to form an initial outage diagnosis. They facilitate the "data-centre whispering" functionality increasingly demanded of NMs, who manage data-centre environments consisting of millions of machines.

Contributions: (1) We pose the challenging RAU problem as a supervised deconvolution problem, which is reminiscent of the BSS literature. (2) Using this simple model for RAU given a user request train, we derive Power Weighted Estimators (PWEs) for RAU estimation. (3) We examine the interplay between our PWEs for the RAU, and a number of different load traces.

Organisation: In Section II we describe the generative model for the attributes collected from the kernel of a UNIX server which captures the evolution of the number of active user requests and the imperfect behaviour of the server. In Section III we derive Power Weighted (PW) RAU estimates and outline the assumptions underpinning them. We examine the statistics of three traces in Section IV: a periodic, step-increase and a flash-crowd trace, and then we estimate RAU of approximately 220 attributes for each trace.

II. Background & Generative Model

We ask the question: "how are the different attributes of a server used when Video-on-Demand is requested?" We attempt to estimate the video-server attribute usage relative to the load on the server. When a client requests VoD from a server whose resources are shared with potentially multiple other users, the effect of the request can be measured by calling a System Activity Report (SAR), which is a widely available [L]UNIX command (cf. http://linux.die.net/man/1/sar). The SAR returns a vector of measurements $\mathbf{x}(i) = [x_1(i), \ldots, x_N(i)]^T \in \Re^N$ –at the discrete time index i– that quantify the performance of the N different attributes of the server.

If the number of active VoD users of the server at time i is denoted $a(i) \in \mathcal{Z}$ –a real, integer-valued, non-negative time-series– the measurements $\mathbf{x}(i)$ can be written as a function of the load on the server, and non-idealities in its behaviour, using an instantaneous mixing model, reminiscent of [9], [10],

$$\mathbf{x}(i) = \mathbf{H}(i)\mathbf{s}(i). \tag{1}$$

The roles of the mixing matrix $\mathbf{H}(i)$ and the perturbation signals $\mathbf{s}(i)$ are described, starting with a few special cases.

TCP Socket Count: Let the N-th attribute, $x_N(i)$, denote the number of active TCP sockets on the server at time index i. An additional VoD request at time $i+1$ increments the TCP socket count, which is our proxy for the load signal $x_N(i) = a(i)$, by 1, therefore, $x_N(i+1) = x_N(i)+1$. Row 1 of Fig. 1 illustrates one user requesting video, then a second user, followed by a third. The first request starts at t_1 and the video finishes at t_2. The effect of this session on the TCP socket count is represented by the signal $u(i-t_1) - u(i-t_2)$. The effect of three users requesting video at different times is illustrated by the step functions activated at the time points, t_3, t_4, t_5 and t_6, e.g. $x_N(i) = (u(i-t_3) - u(i-t_6)) + (u(i-t_4) - u(i-t_6)) + (u(i-t_5) - u(i-t_6))$. The aggregate affect of these user requests is illustrated in Row 2 of Fig. 1, $x_N(i) = a(i)$.

Arbitrary Attributes: The effect of an additional user session on the rate of context switching (attribute $x_1(i)$), or the CPU usage (attribute $x_2(i)$) of the server, for example, is not as clear-cut. A simple model for the n-th attribute is

$$x_n(i) = \alpha_n a(i) + s_n(i). \tag{2}$$

The term α_n represents the true RAU of the n-th attribute, relative to the TCP socket count. It is illustrated in Row 1 of Fig. 1 by scaling the step-functions by α_n. The performance of servers is not ideal. We add the term, $s_n(i)$, to capture perturbations in the server's performance from the ideal performance state for the n-th attribute in Row 3 of Fig. 1. The perturbation signal is composed of the sum of a perturbation signal component, $\epsilon_n(i,k)$, for each active user request,

$$s_n(i) = \sum_{k=1}^{a(i)} \epsilon_n(i,k). \tag{3}$$

The duration and variance of these perturbation components, $\epsilon_n(i,1)$, $\epsilon_n(i,2)$ and $\epsilon_n(i,3)$, are illustrated in Row 3 of Fig. 1.

Assumptions: This model makes certain fundamental assumptions which we justify in the discussion below and by pointing to the effectiveness of the resulting estimators in comprehensive numerical experiments.

Linearity: We have assumed that the server is relatively lightly loaded, and thus, it can be assumed to be in an approximately linear mode of operation.

The linearity assumption is not always true for all loads. The server may become saturated by requests, and so the use of a linear RAU model may then be unreasonable.

Uncorrelated Perturbations: Non-ideality in the server's performance is uncorrelated for each attribute and user.

We expect the performance of the server attributes to exhibit greater variation as the load on the server increases. This assumption is supported by empirical evidence. One way for this to occur is for the contribution of the perturbation components for each active user with respect to the n-th attribute to be uncorrelated.

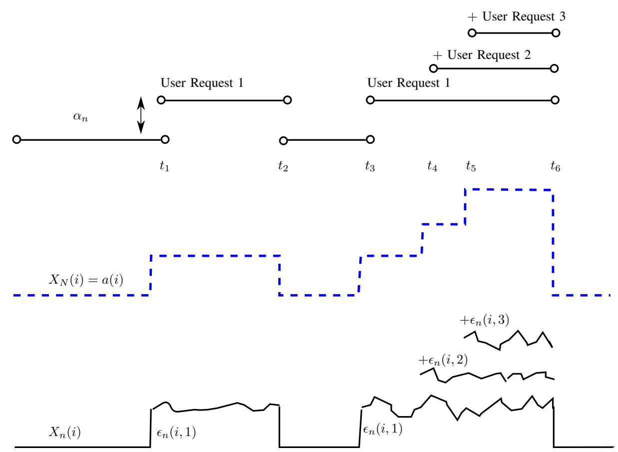

Fig. 1. Attribute Generative Model: Heaviside step functions are stacked in row 1 to illustrate the activation of different video sessions. Row 2 illustrates the TCP socket count attribute, $x_N(i) = a(i)$, which arises in response to these user requests. The imperfection in server behaviour with respect to the n-th attribute is illustrated using the perturbation signals $\epsilon_n(i,k)$ in row 3. The discrete time index i evolves from left to right.

Constant Activation: Attributes are activated for the duration of the session at a constant level, e.g $\alpha_n(u(i-t_i)-u(i-t_j)))$.

Certain attributes may only be affected for a short duration of the activated VoD session. Moreover, the effect of a request may cause the attribute usage to ramp-up or down. We assume that the relative usage is constant as a first approximation. This model may benefit from further refinement and is a topic of our on-going work. It is important to note that the SAR returns the attribute measurements in the form of an average of the measurements over a configurable interval, and thus, the observed attribute values are typically smoothed-out.

Finally, we posit that the server under examination is only serving one type of workload so that the load placed on the server corresponds to one service. This is a reasonable assumption: workload isolation is a major topic of research in the networking community [1]. Isolation of different workloads may allow the NM to perform network planning for workloads which are more predictable –a school of thought in the field of computer networking holds that multiplexing different workloads on a set of common machines or links should be avoided as it may have unexpected consequences on the performance of each of the services, for example, unexpected network delays due to poor resource allocation, or the saturation of different server attributes.

We complete this section by defining the mixing matrix:

$$\mathbf{H}(i) = [\mathbf{I_N} | a(i)\boldsymbol{\alpha}]. \tag{4}$$

It consists of an $N \times N$ identity matrix $\mathbf{I_N}$ and the vector of RAU scalars $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \ldots, \alpha_N]^T$, one for each of the attributes. Finally, the perturbation signal vector is defined as $\mathbf{s}(i) = [s_1(1), s_2(2), \ldots, s_{N-1}(i), 0, 1]^T$. Recall that the TCP socket count corresponds to the N-th attribute, and so, we define its perturbation signal to be zero. The $(N+1)$-th entry of $\mathbf{s}$ serves to "add-in" the scaled load contribution to the corresponding attribution measurement.

III. Relative Attribute Usage Estimators

Goal: The goal of this paper is to estimate the true attribute usage relative to the load on the server, e.g. $\alpha_n, \forall n$ in the presence of server perturbation signals $s_n(i)$ and a time-varying load $a(i)$. In effect we treat the load signal as the *source signal* we desire and we treat the fluctuations in the server's performance from the ideal server behaviour as an *interfering source signal*, which corrupts our estimate of "by how much" the load uses different attributes. The problem is that the source and the interfering signals are *mixed* by the multiplexing/mixing action of the network load-balancer –the network entity that assigns workload to different servers– and the inherent non-ideality of real-world servers. To obtain RAU estimators, we refine the generative model above, with results in the BSS literature [11].

To begin we posit that the load on the server has compact support in the discrete TF domain. We rely on the Discrete Fourier Transform (DFT) to transform each attribute time series into the discrete TF domain, [10], [12], using a sliding window, $w(i)$, positioned at time τ for a local-in-time estimate

$$T^w\{x_n(i)\}(\omega, \tau) : x_n(i) \mapsto X_n(\omega, \tau) \quad (5)$$

where ω is the discrete frequency index. Superior linear transforms may suggest themselves once the TCP socket count signal has been observed for some previous observation period. It follows that the DFT of the perturbation signal, $s_n(i) \in \mathbb{R}$, and the load signal, $a(i) \in \mathbb{R}$, are $S_n(\omega, \tau) \in \mathbb{C}$ and $A(\omega, \tau) \in \mathbb{C}$ respectively. By appealing to the properties of the DFT, the first $N-1$ attribute signals may be written as the sum of two TF signals. The TCP socket count consists of one term because we assume that it is exact.

$$X_n(\omega, \tau) = \alpha_n A(\omega, \tau) + S_n(\omega, \tau), \; 1 \leq n \leq N-1 \quad (6)$$

$$\text{and } X_N(\omega, \tau) = A(\omega, \tau) \text{ for the TCP socket.} \quad (7)$$

The assumption that the load has compact support in TF may be motivated by considering that perturbations in the server's behaviour are likely to occur at higher frequencies than the rate at which new users request video from the server. In addition, the perturbations signal $S_n(\omega, \tau)$ is unlikely to have exactly the same support as the load, $A(\omega, \tau)$. If the load signal is constant for the duration of the observation period, we can remove the effect of the load from the attribute $X_n(\omega, \tau)$ by zero-ing the DC coefficient of the DFT. Therefore, in our evaluation we compare the attributes that arise from time-varying load traces with constant load traces to *isolate* the spectra of the perturbation signals.

The BSS literature [10], [11], [12] often call on the fact that a suitable linear transform $T^w\{\cdot\}(\cdot, \cdot)$ promotes approximate equality in the expression

$$T^w\{a(i)\}(\omega, \tau) \cdot T^w\{s_n(i)\}(\omega, \tau) \approx 0, \quad \forall n, \omega, \tau. \quad (8)$$

The approximate separation of the load and the perturbation signal can be achieved by selecting the set of TF bins θ_a where the load signal is dominant. If the relation in Eqn. 8 holds, then an instantaneous estimate of the RAU, $\alpha_n(\omega, \tau)$ of the n-th feature is obtained by division

$$\alpha_n(\omega, \tau) = \left| \frac{X_n(\omega, \tau)}{X_N(\omega, \tau)} \right| = \left| \frac{X_n(\omega)}{A(\omega, \tau)} \right|, \quad \forall (\omega, \tau) \in \theta_a. \quad (9)$$

We construct a set of estimators that reflect this crucial property. The set θ_a is determined in this paper by choosing the frequency bins corresponding to 90% of the power of $a(i)$. It is clear that the instantaneous RAU estimates will not all be equally useful in our estimate of α_n. We propose a Power Weighting scheme, $C_n(\omega, \tau)$, that distinguishes between different instantaneous estimates, $\alpha_n(\omega, \tau)$, by encoding information about the disjointness associated with each TF bin. This is possible because we know the load, $a(i)$, exactly. The element-wise weight $C_n(\omega, \tau) = |A(\omega, \tau) X_n(\omega, \tau)|$ exhibits the following behaviour in the TF bins θ_a active in our estimators. When $X_n(\omega, \tau) \approx \alpha_n A(\omega, \tau)$ then the associated weight is $C_n(\omega, \tau) = \alpha_n |A(\omega, \tau) A(\omega, \tau)|$, which is large when $A(\omega, \tau)$ is large; and small when $A(\omega, \tau)$ is small. If the case should arise that a TF bin is included in error and the condition $|S_n(\omega, \tau)| \gg |\alpha_n A(\omega, \tau)|$ holds, then the product of the terms in $|S_n(\omega, \tau) A(\omega, \tau)|$ serves to reduce the effect of the associated instantaneous RAU estimate, $\alpha_n(\omega, \tau)$. A Power Weighted objective that expresses (Eqn. 8) is,

$$L_p(\alpha_n^p) := \sum_{(\omega, \tau) \in \theta_a} (C_n(\omega, \tau) \alpha_n^p - \alpha_n(\omega, \tau) C_n(\omega, \tau))^2. \quad (10)$$

Solving $\partial L_p / \partial \alpha_n^p = 0$ for α_n^p yields a PWE for the RAU of the n-th attribute, α_n^p,

$$\alpha_n^p = \frac{\sum_{(\omega, \tau) \in \theta_a} C_n^2(\omega, \tau) \alpha_n(\omega, \tau)}{\sum_{(\omega, \tau) \in \theta_a} C_n^2(\omega, \tau)}. \quad (11)$$

Remark: The Power Weighting uses a cross-weight $C_n(\omega, \tau) = |A(\omega, \tau) X_n(\omega, \tau)|$; cross-weighting generally reduces the risk of divisions by approximately zero, which potentially arise with an auto-weighting scheme, $C_n(\omega, \tau) = |A(\omega, \tau) A(\omega, \tau)|$. The compactness of the support of the load means that the standard Maximum Likelihood (ML) approach for estimating the RAU is corrupted by bins in which the load is not present. Maximising $L(\alpha_n)$ is equivalent to maximising the likelihood function L_0 of α_n, e.g. $L_0(\alpha_n) := p(X_n(\omega), X_N(\omega), A(\omega) | \alpha_n)$.

$$L(\alpha_n) := - \sum_{(\omega, \tau) \in \theta_l} |X_n(\omega, \tau) - \alpha_n A(\omega, \tau)|^2. \quad (12)$$

The perturbation signal in the TF bins where $A(\omega, \tau)$ is dominant, $S_n(\omega, \tau)$, is treated as an interfering signal. It is modelled as complex iid white Gaussian noise, with zero mean and variance σ_s^2. Solving $\partial L / \partial \alpha_n = 0$ for α_n yields

$$\alpha_n^\star = \frac{\sum_{\omega \in \theta_l} \mathrm{Re}\{X_n(\omega, \tau) \bar{A}(\omega, \tau)\}}{\sum_{\omega \in \theta_l} |A(\omega, \tau)|^2}. \quad (13)$$

This MLE can produce negative values, and the PWE cannot. Because an increase in the load on the server may cause a increase/decrease in some server attribute, we correct the PWE sign by using the sign of the MLE, $\alpha_n^p = \mathrm{sgn}(\alpha_n^\star) \alpha_n^p$.

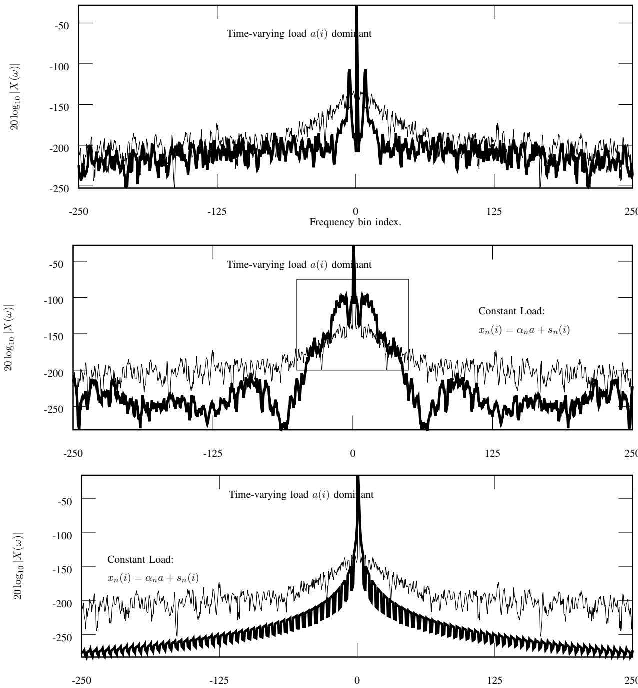

g

Fig. 2. Eqn. 8: Magnitude spectrogram of (1) a time-varying load (heavy line) on the server (top to bottom: periodic, flash-crowd and step-increasing) vs. (2) one of the server attributes under a constant load conditions, e.g. $a = 25$ users (lighter line). The time-varying load dominates the constant load attribute in a small number of low frequency bins (rectangle).

IV. Numerical Evaluation

We evaluate the claim in Eqn. 8 –that the load signal and the perturbations in the server's performance signal are disjoint– using three traces: a periodic, step-increasing and flash-crowd trace. We then evaluate the performance of the RAU PWE.

Experiment Set-up: The testbed set-up consists of a Video-Server (VS) and a Load-Generator (LG) machine, e.g. Dell PowerEdge R715 2U rack servers, each with 64 GB RAM, two 12-core AMD Opteron processors, a 500 GB hard disk, and a 1 Gb network controller. All machines run Ubuntu 12.04 LTS. The server machine runs multiple VLC servers (version 2.1.3). Each VLC server is configured for VoD service. It transcodes the video and audio streams, and streams the videos over the network to VoD clients. Every second, the server machine invokes a SAR. The server is populated with the ten most popular YouTube videos. An experimental run lasts $>15\times10^3$ seconds. The load-balancer generates a Poisson-distributed load on the server $a(i)$. The LG dynamically controls the number of active VoD sessions by spawning and terminating VLC clients. A new client sends a request for a random video to a random VLC server on the VS. A full list of the attributes of the server recorded by SAR and used in the experiments is given the RAU dashboard in Fig. 3. They can be categorised as CPU core utilization, memory and swap space utilization, disk I/O and network statistics. The mean of the Poisson distribution is modulated using a cosine (with amplitude 20 users, period 360s and an offset of the minimum value of the trace), a step-increasing function (which increases the number of users by 5 every 1500s) and a flashcrowd (cf. [13]) in order to drive the time-varying element of the load.

Are the load and perturbation signals disjoint? Firstly, a cosine is chosen to drive the mean of the Poisson-distributed load $a(i)$ because it is time-varying and has compact support. We compute the magnitude spectrum of $a(i)$ and compare this magnitude spectrogram with an attribute taken from an experimental run when the load is constant for all time $a = 25$ user requests, e.g. the load is not time-varying and $s_n(i) \neq 0$. In Fig. 2 row 1, this periodic load clearly dominates the constant-load attribute in the low frequency bins. Then, in row 2 and 3, the flash-crowd and step-increasing load are compared with the constant-load attribute. It still holds that the support of the two signals is approximately disjoint. We note that the flashcrowd is a realistic load signal and examination of its support provides evidence that PWE will be successful.

PWE Dashboard: Fig. 3 illustrates the sign of the ML estimates times the logarithm (base-10) of the RAU estimates. We do not illustrate the MLEs –they are generally very small. These estimates tell us "by how much" a VoD server attribute is used by a single user request –they are VoD's RAU signature. They are computed by taking an 2^{13} sample ($\approx$2-hr) long window to compute our local-in-time DFT of each attribute. This window is shifted successively by 1 sample, and a PWE is computed for each attribute for each shift. The typical PWE is computed by taking the mean of 2000 of these estimates. For a description of each of these attributes, we refer the reader to the manual http://linux.die.net/man/1/sar. Each PWE estimate is illustrated with a dot and its label is given alternating between the left or right hand side of the dots. A significant number of the RAU estimates are non-zero –yielding the *signature* of VoD. There is a clear correlation between the RAU estimates learned across the traces, which illustrates that the estimators are relatively invariant to the statistics of a load signal, for example, different numbers of users being active for longer periods on different experimental runs. In addition, the sign of the estimates is sometimes switched. We posit that this could be due to the effect of a particular load pattern on an attribute affecting the MLE. It is very interesting to see that under different loads the size of the estimators is larger/smaller. We posit that under the periodic load the estimators might be a bit numerically unstable due to the very narrow support of the cosine. If we compare the periodic case with the flash-crowd and the step-increasing load we see that the spectrum of the signals is wider and the PWEs give more stable estimates of the RAU. In terms of the complexity of the estimators, and the deployment of the estimators as a sub-routine in a cloud server, computation of the PWE costs $25F + 7$ FLOPS per estimate, where F is the number of elements in θ_l, which is cheap.

V. Conclusions

We have designed and demonstrated the application of RAU estimators. We posit that the RAU PWE compute accurate estimates of the usage of the VoD server, irrespective of the load pattern placed on the server. This motivates the usage of

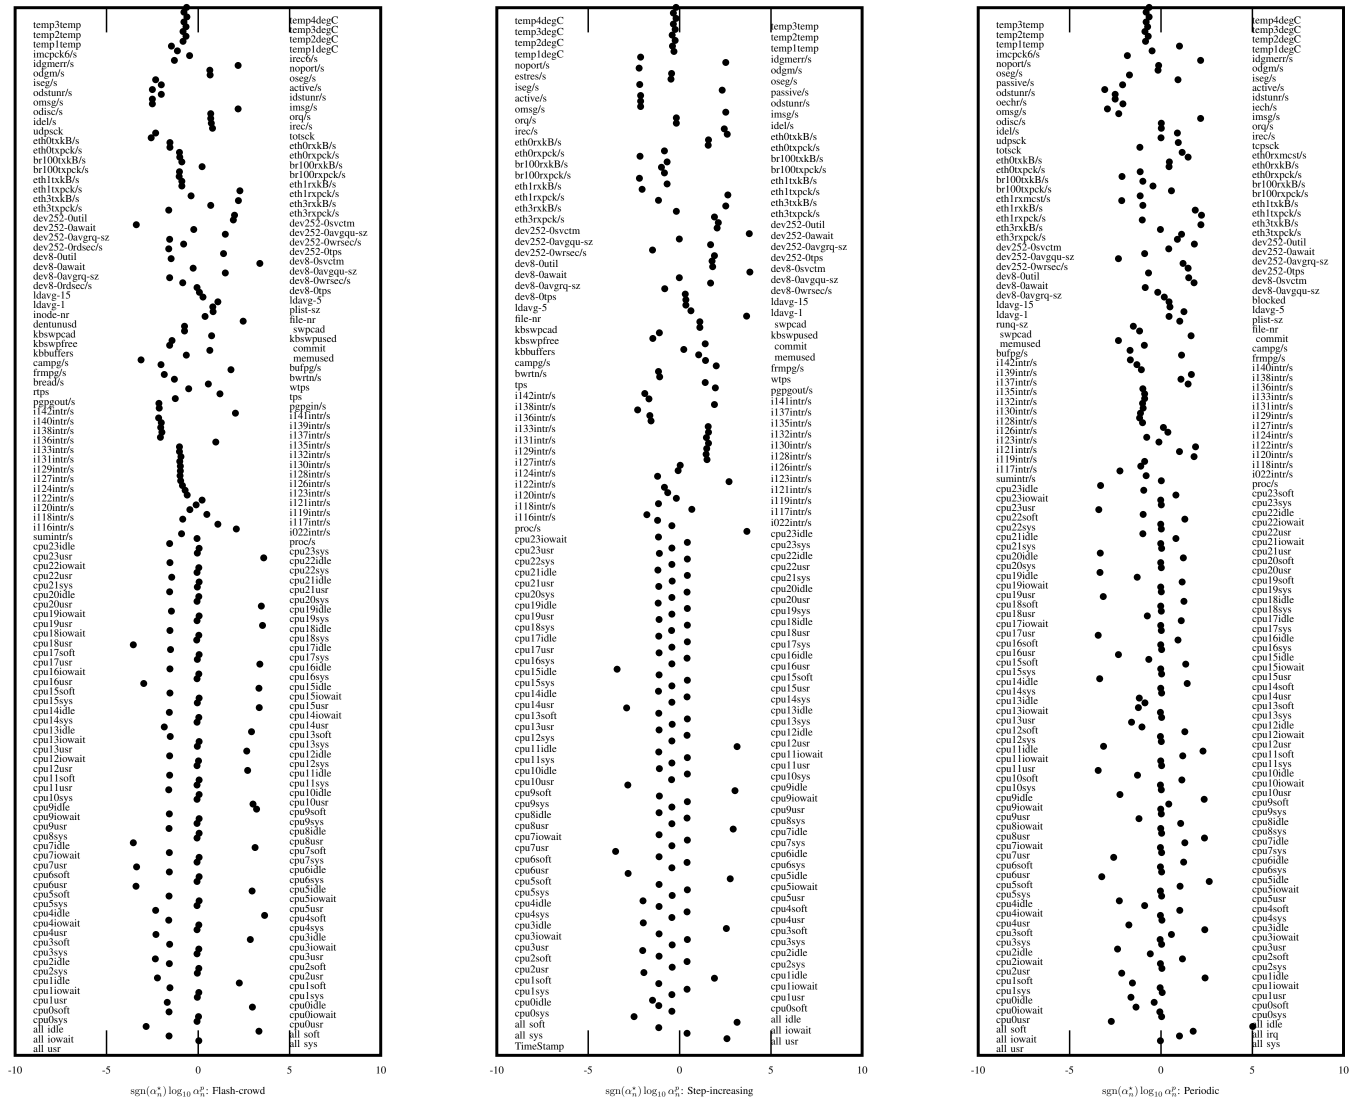

Fig. 3. PWE Dashboard: Each dot represents an RAU estimate for an attribute computed for the flash-crowd, step-increasing and periodic load scenario. Attribute labels are listed successively on the LHS and then RHS of the dots for illustration clarity. There is significant correlation in the signature learned for VoD workload, which is not significantly affected by different load patterns. The sign of the associated MLEs is affected by the workload.

these PWE as a VoD signature, which could be used in run-time data-centre reconfiguration routines, or fault diagnosis.

References

[1] R. de Fréin, J. Pfaff, and T. Paré, "Enterprise data center globality measurement," in *IEEE CIT/IUCC/DASC/PICOM*, 2015, pp. 1861–9.

[2] R. Jain and S. Paul, "Network virtualization and software defined networking for cloud computing: a survey," *Comms Mag., IEEE*, vol. 51, no. 11, pp. 24–31, Nov. 2013.

[3] Z. Xiao, W. Song, and Q. Chen, "Dynamic resource allocation using virtual machines for cloud computing environment," *Par. Dist. Sys., IEEE Trans.*, vol. 24, no. 6, pp. 1107–17, Jun. 2013.

[4] J. Iglesias, M. De Cauwer, D. Mehta, B. O'Sullivan, and L. Murphy, "Increasing task consolidation efficiency by using more accurate resource estimations," *Fut. Gen. Comp. Sys.*, vol. 56, pp. 407 – 20, 2016.

[5] A. Pras, J. Schonwalder, Mark Burgess, O. Festor, G.M. Perez, R. Stadler, and B. Stiller, "Key research challenges in network management," *Comms Mag., IEEE*, vol. 45, no. 10, pp. 104–10, 2007.

[6] Ruairí de Fréin, C. Olariu, Y. Song, R. Brennan, P. McDonagh, A. Hava, C. Thorpe, J. Murphy, L. Murphy, and P. French, "Integration of QoS Metrics, Rules and Semantic Uplift for Advanced IPTV Monitoring," *J. Net. Sys. Man.*, pp. 1–36, 2014.

[7] A. Begen, T. Akgul, and M. Baugher, "Watching video over the web Part 2: Applications, standardization, and open issues," *IEEE Internet Comp.*, vol. 15, no. 3, pp. 59–63, 2011.

[8] R. de Fréin, "Take off a load: Load-adjusted video quality prediction and measurement," in *CIT/IUCC/DASC/PICOM, IEEE*, 2015, pp. 1886–94.

[9] S. Rickard and O. Yilmaz, "On the approximate w-disjoint orthogonality of speech," in *IEEE ICASSP*, May 2002, vol. 1, pp. 529 – 32.

[10] A. Jourjine, S. Rickard, and O. Yilmaz, "Blind separation of disjoint orthogonal signals: demixing n sources from 2 mixtures," in *IEEE ICASSP*, 2000, vol. 5, pp. 2985–88.

[11] O. Yilmaz and S. Rickard, "Blind separation of speech mixtures via time-frequency masking," *Sig. Proc., IEEE Trans.*, vol. 52, no. 7, pp. 1830–47, Jul. 2004.

[12] Ruairí de Fréin and S. Rickard, "The Synchronized Short-Time-Fourier-Transform: Properties and Definitions for Multichannel Source Separation," *Sig. Proc., IEEE Trans.*, vol. 59, no. 1, pp. 91–103, 2011.

[13] I. Ari, B. Hong, E.L. Miller, S.A. Brandt, and D.D.E Long, "Managing flash crowds on the Internet," in *Proc. 11th Int. Symp. Modeling, Ana., Simulat. of Comp. and Tel. Sys. (MASCOTS)*, Dec. 2003, pp. 246–9.