

Probabilistic Resource Allocation based on a Large Deviations Principle

Paulo Gonçalvesⁱ
INRIA-DANTE, France

EXTENDED ABSTRACT

Users of a Cloud Computing platform can have several numbers of choices regarding server selection (some are compute intensive, some provide better I/O performance, some are superior in networking). Cloud provider such as Amazon offers many different server instances that differ in many aspects with respect to CPU speed, network bandwidth and memory capacity. Each of these instances provides a certain amount of dedicated resource and charges per instance-hour consumed. A Service Provider finds it to be extremely difficult to optimise the best combination of servers to be deployed in a \emph{Cloud} for his business on a certain application. This problem differs from the concept of traditional distributed computing (like Grid), since the numbers of servers are virtually unlimited but bandwidth is limited. The choice of deployment of resources can be dynamically tuned using cloud virtualization, that abstracts the IT resources to allow communication and control on-line. Cost of resources varies significantly depending on server types and Cloud Service Providers.

In most applications, the amount of IT resource that is actually used, is a highly variable quantity that follows the instantaneous activity, and in particular the volume of exchanged traffic when network infrastructures are concerned. Depending on the type of application, the generated workload can be a highly varying process that turns difficult to find an acceptable trade-off between an expensive over-provisioning able to anticipate peak loads and a sub performing resource allocation that does not mobilise enough resources. To bypass this challenge, dynamic bandwidth allocation is an original approach that we chose to investigate in the context of network virtualization. We aim to demonstrate the proof of concept for the case of a Video on Demand (VoD) system by adaptively tuning the provisioned bandwidth to the current application workload. In this work we have resorted

to probabilistic provisioning of resource management; however in some situations it can be used to anticipate resource requirements that can serve as inputs for dynamic resource allocation.

Our approach attempts to capture some properties that describe user behaviors or workload generating mechanism of the system and fits them to a mathematical model satisfying particular properties. We leverage these properties to derive a probabilistic assumption on the mean workload of the system at different time resolutions. Embedding the notion of time scale is very important since time scale is by essence intrinsic to dynamicity. In this study we build our system using epidemic models where Markovian models are widely used and happen to satisfy to the specific property mentioned above.

Epidemic information dissemination has been an active area of research in distributed systems, such as Peer-to-Peer (P2P) or VoD systems. In Eugster et al. (Eugster et al., 2004), it has been already demonstrated that the epidemic algorithms can be used as an effective solution for information dissemination in the P2P systems as deployed on Internet or ad-hoc networks. The authors of Bonald et al. (Bonald et al., 2008) studied random epidemic strategies like the random peer, latest useful chunk algorithm to achieve optimal information dissemination. However the most relevant work to our study is derived in Caron et al. (Caron et al., 2011) where the authors proposed an approach to predict workload for cloud clients. They used auto-scaling algorithm for resource provisioning and validated the result with real-world Cloud client application traces. Our approach encompasses both constructive Markovian model to reproduce epidemic information dissemination and workload provisioning aspects. However, we stress that its originality stems from the analysis of the Large Deviation property of the proposed Markovian model. The resulting characterization can be viewed as a multi-resolution extension of the classical

steady-state distribution for the observable mean value of the random process over different aggregated time scales.

After constructing the Markovian mathematical model, we propose two possible and generic ways to exploit these information in the context of probabilistic resource provisioning. They can serve as the input of resource management functionalities of the Cloud environment. It is evident that we can not define elasticity without the notion of a time scale. The Large Deviation Principle (LDP) is capable of automatically integrating the time resolution in automatic description of the system. It is to be noticed that Markovian processes do satisfy the LDP, but so do some other models as well. Hence, our proposed probabilistic approach is very generic and can adapt to address any provisioning issues, provided the resource volatility can be resiliently represented by a stochastic process for which the LDP holds true.

application domain for his methodological research and dealt with statistical characterisation and modelling of traffic, for protocol quality assessment and control. Recently, he returned to wavelet theory and more generally to multi-scale analysis of dynamic graph signals.

ⁱ This is a joint work with Thomas Begin (Université Lyon 1), Patrick Loiseau (EURECOM) and Shubhabrata Roy (ENS Lyon)

BIBLIOGRAPHY

- T. Bonald, L. Massoulié, F. Mathieu, D. Perino, and A. Twigg. Epidemic live streaming: Optimal performance trade-offs. *ACM SIGMETRICS Performance Evaluation Review - SIG-METRICS '08*, 36(1):325–336, June 2008.
- E. Caron, F. Desprez, and A. Muresan. Pattern matching based forecast of non-periodic repetitive behavior for cloud clients. *Journal of Grid Computing*, 9(1):49–64, March 2011.
- P. T. Eugster, R. Guerraoui, A. Kermarrec, and L. Massoulié. Epidemic information dissemination in distributed systems. *IEEE Computer Society*, 37(5):60–67, May 2004.

BRIEF BIOGRAPHY

Paulo Gonçalves (PhD, HDR) is associate researcher at the *Institut National de Recherche en Informatique et Automatique* (INRIA). His PhD work was on the analysis of non-stationary processes with emphasis on time-frequency and time-scale representations of signals. After his post-doc at Rice University (1994-96) on closely related topics, he focused his research activity on fractal analysis and more specifically on scaling laws estimation from wavelet decompositions. Then, P. Gonçalves broaden up the scope of his investigation to more general wavelet-based statistical inference and also on graph based semi-supervised classification. In the course, he privileged network metrology as the main