Big Data Infrastructures: Economics of Resources, Energy, Data, and Applications (Keynote on 2nd Data Storm Summer School 2015)

Luís Veiga
INESC-ID Lisboa
Instituto Superior Técnico
Universidade de Lisboa

July 2015

Abstract

Big data processing infrastructures encompass the full spectrum of cloud-based services: ranging from Infrastructure-as-a-Service (IaaS), to Platform-as-a-Service (Paas) and Data-as-a-Service (DaaS) both for storage as well as processing. Our recent work proposes new resource management and task scheduling mechanisms and algorithms, covering IaaS, PaaS, and DaaS. These are driven by economic inspired notions (such as *yield*, *partial-utility* and *quality-of-data*) that express the efficiency and effectiveness of applying resources to the tenants' workload, taking into account the clients' and workloads' perceived usefulness to a given allocation.

At the IaaS level, system virtual machines (VMs), the cloud providers goals are to maximize revenue, by minimizing the operational expenditure, typically through consolidation of several VMs per physical server. We present a novel cost model and new scheduling algorithms for system-level VMs, along with their implementation in a state-of-the-art cloud simulation framework. Simulations with synthetic and real-world traces, show that the utility-based scheduling allows more VMs to be allocated, thus allowing extra revenue per resource allocated, and shorter waiting times for clients, when comparing with a utility-oblivious redistribution of resources. Scheduling algorithms are also extended to take into account energy efficiency. For this assessment the typical data center topologies are compared against community networks and community cloud scenarios, which are backed by peer-to-peer architectures.

Regarding the PaaS level, we propose new resource management and application monitoring mechanisms for the VMs used by Platform-as-a-Service (PaaS). For PaaS providers, a Java VM was extended with an integrated resource management API, heap resizing policies for yield maximization and concurrent checkpoint for migration of the execution state. Overall, these new mechanisms impose small penalties, measured in the execution of typical benchmarks, while allowing the use of application-tailored policies.

Regarding the DaaS level, we address the storage and processing layers, so that the underlying infrastructure and resource management should fit data and application behavior effectively and efficiently. In Big-data noSQL storage, such as HBase, the main requirement is high-availability, but classical consistency models necessarily degrade performance. Still, an important class of applications can tolerate and benefit from relaxed consistency by bounding inconsistent access in application-specific manner. We explore the spectrum between pessimistic and eventual consistency and provide a novel consistency model with Quality-of-Service tailored to geo-replication (Quality-of-Service for Consistency of Data).

In Big-data processing, specifically in Java frameworks such as Hadoop / Map-Reduce and workflows/dataflows of MR jobs, traditional WMS enforce strict temporal synchronization. In continuous processing, resources are wasted due to the small impact that input data often might have on the final output. thus, no data- and application-aware reasoning is performed in order to evaluate impact of new executions on the final dataflow output. This fails at delivering high resource efficiency for long-lasting applications while keeping costs low. We propose a novel dataflow model with Quality-of-Data (QoD) for continuous and incremental processing. The WMS framework enforcing the QoD model trades off results accuracy with resource savings, achieving resource efficiency, controlled performance, and task prioritization, targeting data-intensive dataflows, where processing steps communicate data via a noSQL database.

http://www.gsd.inesc-id.pt/~lveiga/papers/DataStorm-2015-Big-Data-Infrastructures.pdf

References:

[1] José Simão and Luís Veiga, Partial Utility-driven Scheduling for Flexible SLA and Pricing Arbitration in Cloud, article in IEEE Transactions on Cloud Computing, Nov. 2014, IEEE, [DOI Article link]

[2] José Simão and Luís Veiga, Flexible SLAs in the Cloud with Partial Utility-driven Scheduling, presented at IEEE 5th International Conference on Cloud Computing Technology and Science (CloudCom 2013), Dec. 2013, IEEE, (Best-Paper Award Runner-up)

[3] Leila Sharifi et al., Energy Efficiency Dilemma: P2P-cloud vs. mega-datacenter, presented at IEEE 6th International Conference on Cloud Computing Technology and Science (CloudCom 2014), Dec. 2014, IEEE, (Best-Paper Candidate)

[4] José Simão et al., A2-VM: A Cooperative Java VM with Support for Resource-Awareness and Cluster-Wide Thread Scheduling, presented at 19th International Conference on COOPERATIVE INFORMATION SYSTEMS (CoopIS 2011), Sep. 2011, LNCS, Springer,

[5] José Simão and Luís Veiga, QoE-JVM: An Adaptive and Resource-Aware Java Runtime for Cloud Computing, presented at 2nd International Symposium on Secure Virtual Infrastructures (DOA-SVI 2012), OTM Conferences 2012, Sep. 2012, Springer, LNCS,

[6] José Simão and Luís Veiga, Adaptability Driven by Quality Of Execution in High Level Virtual Machines for Shared Environments, article in International Journal of Computer Systems Science and Engineering pp. 59-72,Nov. 2013, CRL,

[7] José Simão and Luís Veiga, A Progress and Profile-driven Cloud-VM for Improved Resource-Efficiency and Fairness in e-Science Environments, presented at 28th ACM Symposium On Applied Computing (SAC 2013), Mar. 2013, ACM,

[8] José Simão et al., A Checkpointing-enabled and Resource-Aware Java VM for Efficient and Robust e-Science Applications in Grid Environments, article in Concurrency and Computation: Practice and Experience pp. 1421-1442, Sep. 2012, Wiley,

[9] Sérgio Esteves et al., Quality-of-Service for Consistency of Data Geo-Replication in Cloud Computing, presented at International European Conference on Parallel and Distributed Computing (Euro-Par 2012), Aug. 2012, Springer, LNCS,

[10] Álvaro Garcia Recuero et al., Quality-of-Data for Consistency Levels in Geo-replicated Cloud Data Stores, presented at IEEE CloudCom 2013, Dec. 2013, IEEE,

[11] Sérgio Esteves et al., FluX: a quality-driven dataflow model for data intensive computing, article in Journal of Internet Services and Applications (JISA) pp. 1-23, Apr. 2013,Springer

[12] Sérgio Esteves and Luís Veiga, WaaS: Workflow-as-a-Service for the Cloud with Scheduling of Continuous and Data-intensive Workflows, article in Computer Journal, ISSN 0010-4620, Dec. 2014, Oxford University Press

[13] Sérgio Esteves et al., Incremental Dataflow Execution, Resource Efficiency and Probabilistic Guarantees with Fuzzy Boolean Nets, article in Journal of Parallel and Distributed Computing (JPDC), to appear, Elsevier,

Acknowledgements: PhD and MSc students: Sérgio Esteves, Leila Sharifi (graduating Ph.D. 2016), José Simão (graduated 2015)