RESEARCH STATEMENT

ALEXANDRE DA SILVA VEITH

My research interests include distributed systems, big data analytics, Machine Learning (ML), Reinforcement Learning (RL), FoG/edge computing, and Internet of Things (IoT) issues. A common thread in my research is to understand the theory and the design of solutions for improving application performance metrics. I have employed to mathematical tools and performance models from the areas of algorithms, graph theory, probability, RL, and queueing theory. Broadly speaking, my research belongs to the field of *Distributed Systems*.

1. Background and Current Work

Society is increasingly instrumented with sensors integrated to mobile phones, IoT, and in systems for monitoring operational infrastructure, transportation and precision agriculture. These sources provide continuous data streams gathered and analysed by Data Stream Processing (DSP) systems to extract information for decision making in near real-time.

A DSP is often structured as a directed graph or dataflow whose vertices are operators that execute a function over the incoming data and edges that define how data flows between the operators. A dataflow has one or multiple sources (*i.e.*, sensors, gateways or actuators); operators that perform transformations on the data (*e.g.*, filtering, and aggregation); and sinks (*i.e.*, queries that consume or store the data). Traditionally, DSP systems were conceived to run on clusters of homogeneous resources or on the cloud. In a traditional cloud deployment, the whole application is placed on the cloud to benefit from virtually unlimited resources. However, processing all the data on the cloud can introduce latency due to data transfer. *Edge computing* is an attractive solution for performing certain stream processing operations, as many *edge devices* have non-trivial compute capacity and are often geographically closer to where the data is generated.

The task of assinging/reassingning DSP application operators to infrastructures with heterogeneous resources (*e.g.*, edge-cloud infrastructure) considering the resource constraints, is generally referred to as *operator placement* and it has proven to be NP-hard. To tackle this problem, we need to understand the general principles of designing large systems, which give performance guarantees considering the infrastructure limitations (*i.e.*, memory, CPU, bandwidth, energy) and user requirements (*i.e.*, budget, response time, energy and bandwidth consumption). In my Ph. D., I attempt to devise approaches and algorithms to deal with the operator placement problem.

2. Strategies for Big Data Analytics through Lambda Architecture

Initially, I investigated real-life architectures for DSP systems to understand the issues related to the operator placement problem. I focused on Lambda Architecture in big data analytics which combines *batch-processing* and *DSP*. On one hand, batch-processing reduces the data transfer, increases the end-to-end latency, and processes high volumes of data providing an accurate view upon the data. On the other hand, DSP provides real-time analysis. During this investigation, we identified a lack of solutions to handle Lambda Architecture on heterogeneous infratructure. To tackle this problem, we proposed an orchestration system with feedback loop to the application execution statistics [4], and provided analysis and decision strategies to reduce the migration, aggregation, and replication overheads to maximize the throughput [2, 3]. The results of the research led us to propose more elastic algorithms considering the unpredictable behavior on DSP applications.

3. Algorithms for Elastic Big-Data Stream Analytics

We surveyed the state of the art on DSP engines and mechanisms for exploiting resource elasticity features of cloud computing [1]. Resource elasticity allows for an application or service to scale out/in according to fluctuating demands. By investigating the elasticity in the course of solving the operator placement problem under edge-cloud infrastructures, we looked at the issue as single- and multi-objective optimization problems, where we investigated techniques and methods such as queueing theory, series-parallel decomposable graphs, and reinforcement learning. Our analysis allow to examine the problem in four different aspects:

(1) *Design and model solutions for the operator placement problem with single-objective optimization*: We propose a model based on queueing theory and a set of strategies to place operators onto cloud and edge while considering characteristics of resources and the requirements of applications [5]. We consider DSP applications with multiple sources and sinks distributed across the infrastructure and we model the computation and the communication services as M/M/1 queues for estimating the waiting and service times. The strategies use techniques from series-parallel graphs for decomposing the dataflow graph by identifying key behaviors as forks and joins, and then dynamically split the operators across edge and cloud to minimize the end-to-end latency.

(2) *Design and model solutions for the multi-objective optimization in the operator placement problem*: We model the multi-objective optimization using the Simple Additive Weighting method to create a single aggregate cost metric [6]. The cost metric comprises the end-to-end latency, the WAN traffic, and the messaging cost (messages exchanged between the edge and the cloud) for placing DSP applications across edge and cloud resources. The solution extended the strategies from the aforementioned problem by including the new cost metric and by implementing into R-Pulsar (a data analytics software stack for collecting, processing, and analyzing data at the edge and/or at the cloud), for specifying dynamically how IoT applications will be split across the resources.

(3) *Design a model for the Markov Decision Process (MDP) method accounting for the single-objective optimization solution and applying it to RL algorithms*: After the initial placement, operators may need to be reassigned due to variable load or device failures. The search space for operator reassignment can be enormous depending on the number of operators, streams, computational resources and network links. Moreover, it is important to minimise the cost of migration while improving latency. We model the operator reassignment problem as an MDP and investigate the use of RL and Monte-Carlo Tree Search (MCTS) algorithms to devise reconfiguring plans that improve the end-to-end latency of data stream events[**WORK IN PROCESS.**].

(4) *Include the energy consumption, edge-cloud pricing policies and the DSP application reconfiguration overhead to the multi-objective optimization solution on RL algorithms*: We extended our multi-objective optimization solution including realistic price models, energy consumption, and the overhead for reconfiguring the DSP application. The price models embrace the cost models of two major actors, Amazon IoT Core, and Microsoft Azure IoT Hub, in a real-life edge-cloud scenario. The reconfiguration overhead consists of the total downtime for migrating the operator code and its state. The migrations happen in parallel due to the pause-and-resume approach then the total downtime is equal to the longest migration time [7].

4. Distributed Systems - A Research Agenda

During my Ph.D., I noticed that the overhead of single- and multi-objective optimization solutions for the operator placement problem in heterogeneous infrastructure is not negligable. This is mainly due to the inherent nature of distributed systems and the large combinatorial space created by vast infrastucture and complex applications. Models and approaches from queueing theory, probability and artificial intelligence, can be used to devise elegant and simple solutions. Considering the fields mentioned above, I proposed solutions covering the operator placement on edge-cloud infrastructure, by evaluating them under simulated environments and by considering applications with a significant number of operators, as well as, various operator behaviors regarding selectivity, data reduction, and state size. However, I identified issues with I did not deal during the Ph.D. mainly because of time constraints.

In the near future, I intend to investigate and understand real-life and emerging data stream processing frameworks for FoG/Edge/IoT environments in order to model their behavior and impact in heterogeneous and constrained devices. Simultaneously, I am interested in the principles of Machine Learning for designing systems and improving performance metrics with supervised and unsupervised learning. My research will focus on how these techniques can be built in a scalable manner while maintaining performance guarantees and respecting the environment constraints.

In the future, though performance and scalability will remain essential, I also intend to look at issues such as fault tolerance and new methods for application elasticity (*e.g.*, operator replication/fission, data shadowing, backpressure identification and resolution, and etc.), which will become more relevant. I also believe that as systems become increasingly large and inter-dependent, simplicity in design and a precise model will be major factors. Indeed, many of our proposed solutions involve the design of models covering infrastructure and application aspects, and applying them to a real environment can contribute to addressing the above challenges.

My research fundamentally will focus on investigating big data and distributed computing scenarios on emerging technologies, for instance, FoG, edge, and IoT. As research statement I consider: can we achieve ultra-low response time over the Internet, assuming that the various network components give certain performance guarantees? I also intend to devote the other part of my work on practical systems by proposing architecture and infrastructure models, which have immediate relevance and impact in the industry. Similarly, I intend to collaborate with industry peers in understanding and developing solutions for practical problems. My experience in research working jointly with colleagues and my prior record of participation with industry will help me achieve this goal. I am excited at the prospect of learning, contributing, giving shape and making an impact in this upcoming and challenging field.

References

[1] Marcos Dias de Assunção, **Alexandre da Silva Veith**, Rajkumar Buyya Distributed data stream processing and edge computing: A survey on resource elasticity and future directions. *Journal of Network and Computer Applications*, 1 - 17, 2018.

[2] Julio C. S. Anjos, Kassiano Matteussi, Paulo R. R. De Souza Jr, **Alexandre da Silva Veith**, Gilles Fedak, Jorge Luis Victoria Barbosa, Claudio R. Geyer. Enabling Strategies for Big Data Analytics in Hybrid Infrastructures *International Conference on High Performance Computing Simulation (HPCS)*, 869 - 876, 2018.

[3] Paulo R. R. De Souza Jr, Kassiano Matteussi, **Alexandre da Silva Veith**, Julio C. S. Anjos, Jobe D. D. dos Santos, Claudio R. Geyer. Aten: A Dispatcher for Big Data Applications in Heterogeneous Systems. *International Conference on High Performance Computing Simulation (HPCS)*, 585-592, 2018.

[4] **Alexandre da Silva Veith**, Julio C. S. Anjos, Edison Pignaton de Freitas, Thomas J. Lampoltshammer, Claudio R. Geyer. Strategies for Big Data Analytics through Lambda Architectures in Volatile Environments. 4th *IFAC Symposium on Telematics Applications TA 2016*, 114 - 119, 2016.

[5] **Alexandre da Silva Veith**, Marcos Dias de Assunção, Laurent Lefèvre. Latency-Aware Placement of Data Stream Analytics on Edge Computing. *Service-Oriented Computing*, 215 - 229, 2018.

[6] Eduard Gibert Renart, **Alexandre da Silva Veith**, Daniel Balouek-Thomert, Marcos Dias de Assunção, Laurent Lefèvre. Distributed Operator Placement for IoT Data Analytics Across Edge and Cloud Resources. *IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CCGrid 2019)*, 459 - 468, 2019.

[7] **Alexandre da Silva Veith**, Felipe Rodrigo de Souza, Marcos Dias de Assunção, and Laurent Lefèvre, Julio C. S. Anjos. Multi-Objective Reinforcement Learning for Reconfiguring Data Stream Analytics on Edge Computing. To appear on *48th International Conference on Parallel Processing (ICPP 2019)*, 2019.