Future Generation Computer Systems 87 (2018) 198–212

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

A resilient and distributed near real-time traffic forecasting application for Fog computing environments

Juan Luis Pérez *, Alberto Gutierrez-Torre, Josep Ll. Berral, David Carrera

Universitat Politècnica de Catalunya (UPC) - BarcelonaTech, Barcelona Supercomputing Center, Spain

HIGHLIGHTS

- A data Distribution algorithm for FCD collection in city-wide Fog deployments.
- A distributed traffic learning and forecasting model, particularly designed for Fog deployments in the city.
- Validation of the two previous elements through the simulation of different network stability conditions, using as a source real FCD data collected in Barcelona for one week period in 2014.

ARTICLE INFO

Article history:
Received 17 November 2017
Received in revised form 21 March 2018
Accepted 9 May 2018
Available online 18 May 2018

Keywords:
Fog computing
Resource management
Internet of Things
IoT
Big data
Analytics
Cloud computing

ABSTRACT

In this paper we propose an architecture for a city-wide traffic modeling and prediction service based on the Fog Computing paradigm. The work assumes an scenario in which a number of distributed antennas receive data generated by vehicles across the city. In the Fog nodes data is collected, processed in local and intermediate nodes, and finally forwarded to a central Cloud location for further analysis. We propose a combination of a data distribution algorithm, resilient to back-haul connectivity issues, and a traffic modeling approach based on deep learning techniques to provide distributed traffic forecasting capabilities. In our experiments, we leverage real traffic logs from one week of Floating Car Data (FCD) generated in the city of Barcelona by a road-assistance service fleet comprising thousands of vehicles. FCD was processed across several simulated conditions, ranging from scenarios in which no connectivity failures occurred in the Fog nodes, to situations with long and frequent connectivity outage periods. For each scenario, the resilience and accuracy of both the data distribution algorithm, and the learning methods were analyzed. Results show that the data distribution process running in the Fog nodes is resilient to back-haul connectivity issues and is able to deliver data to the Cloud location even in presence of severe connectivity problems. Additionally, the proposed traffic modeling and forecasting method exhibits better behavior when run distributed in the Fog instead of centralized in the Cloud, especially when connectivity issues occur that force data to be delivered out of order to the Cloud.

1. Introduction

Several technologies relevant to the expansion of the Internet of Things (IoT) have emerged in the last years, including network functions virtualization (NFV), fifth generation (5G) wireless systems, and Fog computing. The combination of these technologies opens a new range of potential applications in the context of Smart Cities. There is a fast growth in the number of projects planning to deliver new services to citizens, based on the deployment of a large number of Fog nodes near the edge, in the streets of modern cities, bridging the gap between devices and Cloud-based services. Fog nodes can host lightweight services on near real-time, like for instance the collection and processing of streams of data. This technology is a foundational enabler for the future development of advanced services such as for instance traffic monitoring and planning through the combination of street sensors data (e.g. vehicle tracking, air quality measurements) and meteorological information. Although Fog nodes offer a constrained computing capacity compared to their Cloud counterparts, they still have capabilities to process data in near real-time to provide localized service to users, minimizing the communication requirements with the Cloud, or ensuring application resilience against back-haul connectivity outages between the Fog and Cloud layers.

Modern cities demand new approaches to deliver localized services to their citizens, and at the same time, network operators

---

* Corresponding author.
E-mail addresses: juan.luis.perez@bsc.es (J.L. Pérez), alberto.gutierrez@bsc.es (A. Gutierrez-Torre), josep.berral@bsc.es (J.L. Berral), david.carrera@bsc.es (D. Carrera).

https://doi.org/10.1016/j.future.2018.05.013

look for new advanced services that can take advantage of the new hyper-connected society that is expected for the coming 5G era. The incredibly high bandwidth that 5G networks will offer to their users will restrict the possibility to define new services that run only in centralized Cloud-based locations. Therefore, the development of decentralized architectures that leverage the Fog computing paradigm (computing between the edge and the Cloud) is a mandatory requirement for an efficient deployment of 5G technologies over the next couple of years. In this context of highly connected cities with distributed Fog-based computational capabilities, applications will require a superior ability to adapt to the continuous changes that occur within the dynamics of a modern city: the only way to provide the required flexibility will be through the use of advanced Artificial Intelligence techniques that help systems *learn* and model the behavior of crowds in near real-time. It is only under these conditions, with the combination of 5G networks, the Fog computing paradigm and the exploitation of AI techniques, that it will be possible to develop the complex services that cities demand.

In this paper, we present a distributed architecture for a traffic modeling and prediction service, designed for a city-wide scenario based on the Fog computing paradigm. In this context, we assume that a set of advanced antennas (e.g. 5G stations [1] enabled with Fog computing capabilities, acting as a Fog node) are distributed across the city, and they are used to receive telemetry and location data as generated by vehicles. Each vehicle sends data to the nearest antenna and its associated Fog node. Therefore, data is collected and locally processed in Fog nodes (either located at the Edge or in-between the Edge and the Cloud as intermediate nodes), and then forwarded to a central Cloud location for further analysis as well as data warehousing purposes. The proposed architecture combines a real-time data distribution algorithm with enhanced resilience against back-haul connectivity issues, and a traffic modeling technique based on the use of Conditional Restricted Boltzmann Machines (CRBM) to learn traffic patterns. In combination, these two techniques provide resilient and completely decentralized city-wide traffic forecasting capabilities.

The proposed architecture is validated using real traffic logs from one week of Floating Car Data (FCD) in the city of Barcelona, provided by one of the largest road-assistance companies in Spain, comprising thousands of vehicles from their fleet only in the city of Barcelona. The dataset (further described in Section 5.2) comprises data collected over one week between 10/27/2014 and 11/01/2014 across the Barcelona metropolitan area. Fig. 1 shows a heat-map of the vehicle tracking data, comprising over 890,000 data samples and a fleet of more than 100 cars moving simultaneously around the city at some times.

Using this FCD dataset, we simulated using the provided FCD across several conditions, from scenarios in which no connectivity failures occurred between the Fog nodes and the Cloud, to scenarios with long and frequent connectivity outage periods. For each one of those scenarios, we have analyzed the resilience and accuracy of the data distribution algorithm and the learning methods.

While current frameworks dealing with FCD analytics focus on how to distribute load towards anomaly detection, modeling and trend prediction on Cloud infrastructures and leveraging Map-Reduce mechanisms to handle traffic data, in this work we focus on (1) the scenario where analytics can be partial or completely performed on the Edge instead of on the Cloud; and (2) the proper transmission of data between Fog nodes on the Edge and the Cloud towards delivering data to be aggregated or learned models to be used. The current case of use targets city-wide traffic data, but Edge-Cloud architectures can be used for enhancing Smart City applications, like power monitoring and control of elements in public spaces, connectivity on demand from smart phones towards public services, or sensor data recopilation from smart phones towards retrieving environmental data [2].

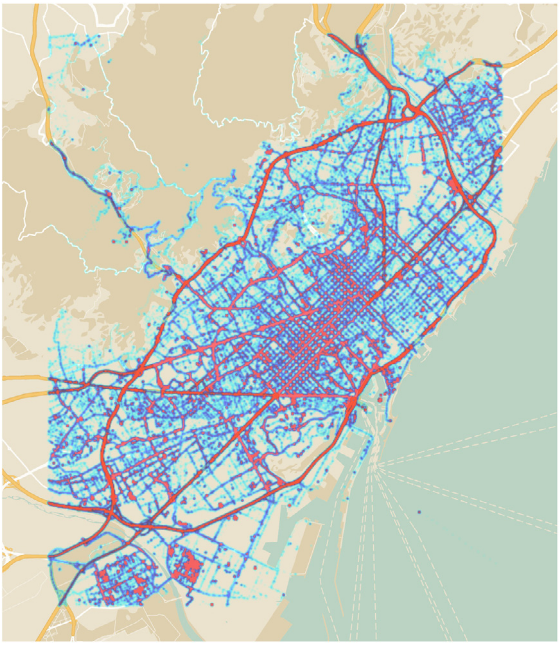

Fig. 1. Barcelona metropolitan area map, combined with a heat-map overlay of the FCD dataset used for the simulations presented in this paper. The dataset contains more than 890,000 data samples of road-assistance cars moving around the city.

Experiments show that the here presented architecture for data distribution running in the Fog nodes is resilient to back-haul connectivity issues, and it is able to deliver data to the Cloud location even in presence of severe connectivity problems. Additionally, the proposed traffic modeling and forecasting method based on CBRMs, not only is able to predict telemetry features at short terms but also exhibits better behavior when modeling local data at Fog nodes instead of a centralized model in the Cloud, useful when connectivity issues force data to be delivered out of order to the Cloud, providing an extra degree of autonomy to Fog nodes.

In summary, the three major contributions of this paper are:

- Data Distribution algorithm for FCD collection in city-wide Fog deployments. The algorithm is designed to be resilient to back-haul connectivity issues, to avoid data to be lost under connectivity outage periods, and to favor distributed data modeling in the Fog. The paper also provides an analysis of the behavior of the algorithm under different scenarios of lost connectivity.
- A distributed traffic learning and forecasting model, particularly designed for Fog deployments in the city, in which data collected by the data distribution algorithm is fed into a distributed set of Conditional Restricted Boltzmann Machines (CRBM). The neural networks learn traffic patterns across the city and can be leveraged to forecast future traffic conditions. The distributed approach is superior to a Cloud-centralized schema in terms of resilience against network connectivity outages.
- Validation of the two previous elements through the simulation of different network stability conditions, using as a source real FCD data collected in Barcelona for one week period in 2014.

The paper is structured as follows: Section 2 introduces the background on distributed architectures and IoT management. Section 3 presents the proposed solution towards the current problem. Section 4 describes in detail the components of the presented approach. Section 5 shows the different evaluation experiments