Contents lists available at ScienceDirect

Vehicular Communications

www.elsevier.com/locate/vehcom

Designing mobile content delivery networks for the internet of vehicles

Cristiano M. Silva [a,*], Fabricio A. Silva [b], Thiago R. Oliveira [a], Wagner Meira Jr. [c], Jose Marcos S. Nogueira [c]

[a] Departamento de Tecnologia, Universidade Federal de São João Del-Rei (DTECH/UFSJ), Ouro Branco, MG, Brazil
[b] Departamento de Informática, Universidade Federal de Viçosa (DIN/UFV), Florestal, MG, Brazil
[c] Departamento de Ciência da Computação, Universidade Federal de Minas Gerais (DCC/UFMG), Belo Horizonte, MG, Brazil

ARTICLE INFO

Article history:
Received 16 June 2016
Received in revised form 7 November 2016
Accepted 11 November 2016
Available online xxxx

Keywords:
Vehicle-to-roadside communication
Communication architecture
Content delivery

ABSTRACT

Content delivery is a key functionality for developing the Internet of Vehicles. In such networks, vehicles act as sensors of the urban mobility by constantly exchanging messages with another vehicles, the cellular network, and also the infrastructure (roadside units). However, the task of delivering content in such dynamic network is far from trivial. In this work, we investigate the development of Content Delivery Networks (CDN) in the context of vehicular networks. Roadside units support the communication by replicating and delivering contents to vehicles within their range of coverage. Initially, we devise a strategy for measuring the performance of the content delivery in vehicular networks. Then, we use the proposed metric for designing a deployment strategy allowing us to identify the better locations for deploying the roadside units in order to properly support the dissemination of a variety of contents, each content requiring specific levels of performance. We compare our deployment strategy to the intuitive strategy of allocating roadside units at the densest locations of the road network. The results demonstrate that our strategy requires less roadside units than the baseline for non-massive deployments in order to achieve similar levels of performance, incurring in less costs when setting up the content-delivery infrastructure.

1. Introduction

The Internet of Vehicles (IoV) is an emerging technology that investigates the inter-vehicle communication. Differently from the traditional Internet, the Internet of Vehicles imposes novel challenges due to specific features of this network. The high mobility of vehicles leads to constant changes in the network topology. In addition, the communication channel must deal with several vehicles attempting to send and receive data simultaneously, possibly at high speeds [1]. Furthermore, vehicles do not (necessarily) demand continuous coverage, but they may engage in "infofueling" [2] as they opportunistically drive through roadside units. Moreover, the network is assumed to operate in large scale scenarios with different traffic levels.

IoV applications have been evolving from simple small traffic announcements to high quality streaming and real time data transmissions, demanding novel strategies for delivering contents [3]. While the vehicular communication may take place in an ad hoc basis, the performance of the vehicular network is (usually) highly improved by deploying a set of roadside units supporting the vehicular communication [4]. On the other hand, the deployment of a large scale infrastructure may demand huge investments, which may postpone the adoption of vehicular networks [5,6]. Thus, when planning the roadside infrastructure, we typically intend to minimize the number of roadside units, while maximizing the network performance.

In this work, we investigate the design of Mobile Content Delivery Vehicular Networks (MCDVN). Traditional Content Delivery Networks (CDN) rely on the replication of contents in servers located along the network, while requests are load-balanced among the servers, providing high content availability. In the context of vehicular networks, the role of servers is assumed by roadside units deployed along the road network. Since a given content may be meaningful only to a given region of interest, we assume that each content type is related to a target region where it must be made available. Furthermore, given the wide range of envisioned IoV applications, we also consider each content to require distinct

* Corresponding author.
E-mail addresses: cristiano@ufsj.edu.br (C.M. Silva), fabricio.asilva@ufv.br (F.A. Silva), thiagool@ufsj.edu.br (T.R. Oliveira), meira@dcc.ufmg.br (W. Meira Jr.), jmarcos@dcc.ufmg.br (J.M.S. Nogueira).

http://dx.doi.org/10.1016/j.vehcom.2016.11.003
2214-2096/© 2016 Published by Elsevier Inc.

performance levels from the network (defined in terms of V2I contact probability and V2I contact duration).

To assist the network designer in planning Mobile Content Delivery Vehicular Networks, we propose the Sigma Deployment (Σ), a strategy for measuring the performance of the content delivery in infrastructure-based vehicular networks. Sigma Deployment allows the network designer to check whether the vehicular network provides the proper support for the dissemination of a subset of contents demanding distinguished levels of performance from the network, taking into account the set of locations where each content must be made available.

After formally presenting and discussing the Sigma Deployment, we propose an Integer Linear Program formulation for the problem, and we also present a greedy heuristic (Sigma-g) for solving it. Sigma-g considers the set of contents to be made available, the locations where each content must appear, and also the network performance demanded by each tuple formed by {content type, location}. Sigma-g outputs the number of required roadside units and also the location of each roadside unit in order to reach the target deployment. We evaluate Sigma-g using the realistic mobility trace of Cologne, Germany. The results demonstrate Sigma-g requiring up to 80% less roadside units then the baseline for non-massive deployments.

Differently from all previous approaches, Sigma Deployment is a content-oriented metric for evaluating the performance of the infrastructure supporting the dissemination of several types of contents within vehicular networks. As far as we are concerned, this is the first work proposing a content-oriented deployment strategy accounting for distinguished levels of performance for each specific content.

This document is organized as follows: in Section 2, we present the related work concerning content delivery for vehicular networks and infrastructure deployment strategies. In Section 3, we discuss Mobile Content Delivery Vehicular Networks and the Sigma Deployment. In Section 4, we present a set of experiments for evaluating Sigma-g. The conclusion is presented in Section 5.

2. Related work

The fast increasing demand for Internet services (particularly in the context of the Web during the 1990s) has driven the research community for investigating strategies for dealing with network congestion and server overloading [7]. Content Delivery Networks arise as a feasible strategy for tackling such issues [8]. The basic idea behind CDNs is to allocate content replicas at different-and-strategically-located servers, allowing the redirection of requests for the most appropriate machine.

In this work we investigate the application of Content Delivery Networks for vehicular networks. By considering roadside units as surrogate servers, we propose a strategy for measuring the effectiveness of the infrastructure deployment, and then we evaluate the overall QoS of the vehicular network in order to disseminate several types of contents with distinguished levels of performance, according to the content type, and the physical location where the content must be made available.

2.1. Content delivery in vehicular networks

Some studies from the literature have tackled the content delivery problem in the context of vehicular networks (a detailed description can be found in [9]). In a summarized way, in [10], the authors propose a pro-active data replication solution to select vehicles as replicas. The solution described in [11], on the other hand, replicate content in roadside units based on real mobility trace analysis. Similarly, the work [12] investigates an optimization model to replicate content in roadside units by taking into account the demand and popularity of each content.

Kumar and Kim [13] focus on replicating content for video streaming applications in vehicular delay tolerant networks. Zhang and Cao [14] take advantage of vehicles platoons for content dissemination. There are also solutions envisioned to replicate contents, while keeping them within a given region of interest [15].

Our proposal differs from these approaches in the sense that they assume roadside units providing full coverage of the road network. Since massive deployment of roadside units is likely to demand very high investments, we turn our attention for identifying the minimum number of roadside units required for delivering several content types, where each content type requires a different level of performance, accordingly to the location of the content in the network.

2.2. Deployment of roadside units

The deployment of roadside units is also a problem that has been tackled by researchers in the last years. However, several challenges still remain given the diverse requirements concerning the vehicular networks. Alpha Coverage [16] provides worst-case guarantees on the interconnection gap while using significantly fewer roadside units. Zheng et al. [17] present the evaluation of a deployment strategy considering the contact opportunity. Lee and Kim [18] propose a greedy heuristic to place the infrastructure aiming to improve the vehicles connectivity while reducing the disconnections. Xie et al. [19] address the placement of roadside units into a grid road network assuming knowledge of source and sink of each vehicle. Liya et al. [20] propose a randomized algorithm that calculates an approximate distance for deploying roadside units by approaching the optimal distance step by step from the initial distance $d_0 = 2R_0$, where R_0 is the transmission range. Differently from these approaches, our solution considers the contact probability as well as the contact duration as metrics for roadside units deployment.

Chi et al. [21] propose three optimal algorithms to allocate the roadside units: Greedy, Dynamic and Hybrid algorithms. Liu et al. [22] propose a new roadside units' deployment strategy for file downloading in vehicular networks. Cheng et al. [23] propose a geometry-based coverage strategy to handle the deployment problem over urban scenarios. Patil and Gokhale [24] propose a Voronoi diagram-based algorithm for the effective placement of roadside units using the packet delay and the loss as criteria. Authors seek to provide a collaborative mechanism for dynamic resources management in vehicular networks, and to provide QoS assurances to applications.

Optimization models for the deployment are presented in [25–27]. Aslam et al. [25] use Binary Integer Programming to solve the allocation of infrastructure: they eliminate minor roads and model major roads as a grid. Wu et al. [26] model the deployment of roadside units as a Integer Linear Programming considering both strategies of communication such that the aggregate throughput in the network is maximized. Differently, we handle road networks of arbitrary topology and do not eliminate minor roads.

There are also strategies based on the maximum coverage problem. Trullols et al. [28] use a realistic data set and propose modeling the placement as a Knapsack Problem (KP) and also as a Maximum Coverage Problem (MCP-g). Cataldi and Harri [29] propose the allocation of roadside units considering the Maximum Coverage Problem over a benefit function.

Assuming previous knowledge of vehicles trajectories is a common strategy in deployment studies. However, when we intend to maximize the number of distinct vehicles contacting the infrastructure, we may rely just on the migration ratios of vehicles between distinct locations of the road network. Since the migration

Please cite this article in press as: C.M. Silva et al., Designing mobile content delivery networks for the internet of vehicles, Veh. Commun. (2016), http://dx.doi.org/10.1016/j.vehcom.2016.11.003