Vehicular Communications 4 (2016) 47–56

Contents lists available at ScienceDirect

Vehicular Communications

www.elsevier.com/locate/vehcom

Virtual machine migration and management for vehicular clouds

Tarek K. Refaat [a], Burak Kantarci [b], Hussein T. Mouftah [a,*]

[a] *School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, ON, Canada*
[b] *Department of Electrical and Computer Engineering, Clarkson University, Potsdam, NY, USA*

ARTICLE INFO

Keywords:
Cloud computing
Data center management
Vehicular ad hoc networks
Vehicular clouds
Virtual machine migration

ABSTRACT

Vehicular Cloud Computing is a growing research field which consolidates the benefit of cloud computing into vehicular ad hoc networks. However, few studies address vehicles as potential Virtual Machine hosts. Due to the rapidly changing environment of a vehicular cloud, a host can easily change or leave coverage. As such, Virtual Machine Management and Migration schemes are necessary to ensure cloud subscribers have a satisfactory level of access to the resources. This paper proposes several Vehicular Virtual Machine Migration (VVMM) schemes: VVMM-U (Uniform), VVMM-LW (Least Workload), VVMM-MA (Mobility Aware) and MDWLAM (Mobility and Destination Workload Aware Migration). Their performance is evaluated with respect to a set of metrics through simulations with varying levels of vehicular traffic congestion, Virtual Machine sizes and levels of load restriction. The most advanced scheme (MDWLAM), takes into account, the workload and mobility of the original host as well as those of the potential destinations. By doing so a valid destination will both have time to receive the workload and migrate the new load when necessary. The behavior of various algorithms is compared and the MDWLAM has been shown to demonstrate best performance, exhibiting migration drop rates that are negligibly small.

1. Introduction

The incorporation of cloud computing into various aspects of technology is a growing trend, both in terms of depth and breadth of spectrum. The most common application users come across is typically the cloud storage. More complex applications include cloud-based software or processing [1]. With the progress in the research of the realization of the Internet of Things (IoT), cloud computing becomes more appealing as devices with network interfaces increase [2]. Vehicles are an example of such devices where interconnectivity is heavily studied. Typically, the field of Vehicular Ad Hoc Networks (VANETs) studies the communication from vehicle to vehicle (V2V), vehicle to infrastructure (V2I) and infrastructure to vehicle (I2V), with the purpose being to exchange safety or traffic information [3]. However, a new paradigm considers utilizing vehicles for cloud computing, effectively creating a Vehicular Cloud (VC) [4–7].

As vehicle designs incorporate an increasing amount of electronics, resources such as storage space and processing power, make the VC more realizable every day [8,9]. A vehicular cloud can make use of many of these often underutilized resources, operating like a typical cloud data center [10,11]. The unique setting of a vehicular environment facilitates services including Entertainment-as-a-Service (ENaaS), Traffic-Information-as-a-Service, Network-as-a-Service, and Storage-as-a-Service, among others [12–16].

In order to maximize flexibility, when allocating hardware resources to cloud users, Virtual Machines (VMs) are used. A VM allows a user to set up an environment and run their unique processes, without worrying about the physical resources involved. The resources are handled by hypervisors that, in turn, isolate VMs, virtually, although several may share the same physical resource, rendering each VM (and its corresponding user) oblivious of its neighbors. Moving the VM from one physical location to another is simple and potentially seamless. Such VM Migrations (VMMs) can be triggered to handle hot-spots or for resource load balancing [12].

There are several inherent obstacles that VMM schemes must address, and for the traditional cloud, there are several approaches available [13,17]. Aside from the typical issues of handling migration triggers, the VC has unique obstacles to address. The main aspect that is almost exclusive to the VC is the dynamicity and mobility of the resources, namely the vehicles.

It is worthwhile noting that while studies propose vehicles as sources of data for the VC, there is a limited number of studies addressing the use of vehicles as actual VM hosts [18]. This study focuses on vehicles as VM hosts and specifically, the VMM algo-

* Corresponding author at: School of Electrical Engineering and Computer Science, University of Ottawa, 800 King Edward Avenue, Ottawa, ON, K1N 6N5, Canada. Tel.: +1 613 562 5800x2173.
E-mail addresses: tarek.k.r@ieee.org (T.K. Refaat), bkantarc@clarkson.edu (B. Kantarci), mouftah@uottawa.ca (H.T. Mouftah).

http://dx.doi.org/10.1016/j.vehcom.2016.05.001
2214-2096/

List of symbols

T_s	time needed to migrate source workload	s
T_d	time needed to migrate destination workload	s
T	simulated time	s
d_t	total dropped VM workload, at time t in the grid	MB
s_t	total successfully migrated VM load, at time t, throughout the entire grid	MB
v_t	total VM load, successfully migrated to the infrastructure, at time t, throughout the entire grid	MB
l_t	total VM load hosted on all the vehicles currently the grid, at time t	MB
l_m	maximum VM load capacity the grid can offer, based on the expected average number of vehicles at the congestion level simulated	MB/s
ul_t	total upload bandwidth utilized, at time t, by all vehicles throughout the entire grid	MB/s
ul_m	maximum upload bandwidth available in the grid	MB/s
dl_t	total download bandwidth utilized, at time t, by all vehicles throughout the entire grid	MB/s
dl_m	maximum download bandwidth available in the grid	MB/s
P_t	probability, at time t, that a packet is dropped due to interference or connection instability	
P_{max}	maximum value for P_t (set at 0.5)	
Int_t	normalized interference due to concurrent migrations at time t	
CQ_t	normalized connection quality between source and destination, based on distance between them, at time t	
SD_t	distance between source and destination	
SD_m	maximum value for SD_t, the diagonal distance from one corner of the grid to the other	

rithms for such a VC. As the vehicular VM hosts move, the need for a VMM scheme arises. Several Vehicular VMM (VVMM) schemes are proposed, tested and compared against a generic benchmark algorithm. The VC paradigm proposed involves the vehicles (VM hosts) and Roadside Units (RSUs), linked to the cellular infrastructure.

This paper presents the following VVMM algorithms: Uniform (U), Least Workload (LW), Mobility Aware (MA), Mobility and Destination Workload Aware Migration (MDWLAM). The VVMM-U algorithm is the benchmark, whereby migration destinations are selected using a uniform likelihood. The VVMM-LW selects destinations based on those with the least workload. The mobility of the destination is taken into account in VVMM-MA and finally, the destination workload is also considered for MDWLAM. The various algorithms are analyzed using various metrics, varying VM sizes and vehicular traffic congestion levels. The MDWLAM scheme achieves negligibly small migration drop rates and infrastructure migrations that are negligibly small, outperforming the other proposed algorithms.

The rest of the paper is organized as follows. Section 2 summarizes related studies for VC paradigms. Section 3 introduces the VVMM algorithms. The performance of the various algorithms is evaluated in Section 4 and the paper is concluded in Section 5.

2. Related work

The underutilized resources in modern vehicles make a VC quite appealing. As opposed to investing significant capital in a traditional data center, vehicle resources can be leased to and managed by a cloud administrator [19]. The cooperation of vehicles will involve communication issues due to a dynamic environment, network congestion, energy efficiency, and latencies, as well as security and privacy limitations [20]. However, should a VC be realized, there will be a significant benefit due to the existing underutilized resources.

Initially, studies focused on avoiding mobility issues by targeting parked vehicles [21,22]. One study considers avoiding communication issues also by implementing a VC using an airport parking lot. Vehicles parked for relatively long periods of time can be connected to a wired network via ports in their parking spots, effectively creating a traditional data center throughout the parking lot [12]. Such an approach is quite simple to realize as it avoids any wireless communication issues as well as the more challenging issue of mobility.

As research progressed, mobile vehicles for VCs were considered. While the complexity of realizing such a VC increases, it is more appealing. The number of vehicle groups in motion is far greater (and more common) than large parking lots like airports [23]. Such VCs will necessitate wireless V2V, V2I and/or I2V communication. The infrastructure involvement can decrease the complexity of administration as it is a static element of the cloud [24].

A generic model for VANET cloud computing is presented in [6]. The model includes two different types of VC models: permanent and temporary. In the permanent cloud, the majority of vehicles involved are stationary. The mobile vehicles, forming a sub-cloud are the temporary cloud. However, it is important to note that there is always a heavy reliance on the RSU. It was the only gateway for vehicle to vehicle interaction. The authors coin the term VANET-Cloud to describe their system, separated into three layers: Cloud, Communication and Client [6]. Finally, the authors report the future challenges a VANET-Cloud must address including security, data aggregation, energy efficiency and coordination between sub-clouds.

For many studies, vehicles are considered cloud subscribers or data sources for the cloud. Traffic congestion information, video surveillance on public transport, cooperative file downloading and data mining are all examples of potential vehicle utilizations [25]. However, it is important to consider the possibility that vehicles can act as VM hosts.

As VM hosts, VMM schemes are necessary to maintain a quality of service for cloud subscribers. While there are differences between the traditional cloud and the VC, it is important to consider existing VMM schemes for the traditional cloud, such as hot/cold migrations, black-box and gray-box migrations [26–28]. They serve as a starting point for VVMM schemes, albeit with several modifications. The optimization of VMM schemes is presented in [28]. Mainly, the study addresses bandwidth optimization, by determining the exact VMM scheme suitability based on various possible cases. The authors analyze several schemes in detail and summarize the pros and cons of each based on different metrics and overheads [28]. The study also covers optimizing energy efficiency and also storage efficiency for VMM schemes.

One of the few studies addressing the VC with vehicular resources being used for the cloud is presented in [5]. The study describes the job assignment to the various (mainly parked) vehicles, and more importantly incorporating fault-tolerance. This is specifically to address the dynamicity of the resources, as cars may