



HAL
open science

Internet of Vehicles: From intelligent grid to autonomous cars and vehicular fogs

Eun-Kyu Lee, Mario Gerla, Giovanni Pau, Uichin Lee, Jae-Han Lim

► **To cite this version:**

Eun-Kyu Lee, Mario Gerla, Giovanni Pau, Uichin Lee, Jae-Han Lim. Internet of Vehicles: From intelligent grid to autonomous cars and vehicular fogs. *International Journal of Distributed Sensor Networks*, 2016, 12 (9), pp.1550147716665500. 10.1177/1550147716665500 . hal-01398101

HAL Id: hal-01398101

<https://hal.sorbonne-universite.fr/hal-01398101v1>

Submitted on 16 Nov 2016


HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Internet of Vehicles: From intelligent grid to autonomous cars and vehicular fogs

International Journal of Distributed
Sensor Networks
2016, Vol. 12(9)
© The Author(s) 2016
DOI: 10.1177/1550147716665500
ijdsn.sagepub.com


Eun-Kyu Lee^{1,2}, Mario Gerla¹, Giovanni Pau³, Uichin Lee⁴ and
Jae-Han Lim⁵

Abstract

Recent advances in communications, controls, and embedded systems have changed the perception of a car. A vehicle has been the extension of the man's ambulatory system, docile to the driver's commands. It is now a formidable sensor platform, absorbing information from the environment (and from other cars) and feeding it to drivers and infrastructure to assist in safe navigation, pollution control, and traffic management. The next step in this evolution is just around the corner: the Internet of Autonomous Vehicles. Pioneered by the Google car, the Internet of Vehicles will be a distributed transport fabric capable of making its own decisions about driving customers to their destinations. Like other important instantiations of the Internet of Things (e.g. the smart building), the Internet of Vehicles will have communications, storage, intelligence, and learning capabilities to anticipate the customers' intentions. The concept that will help transition to the Internet of Vehicles is the vehicular fog, the equivalent of instantaneous Internet cloud for vehicles, providing all the services required by the autonomous vehicles. In this article, we discuss the evolution from intelligent vehicle grid to autonomous, Internet-connected vehicles, and vehicular fog.

Keyword

Internet of vehicles, vehicular cloud, fog computing, vehicular network, content discovery

Date received: 15 April 2016; accepted: 17 July 2016

Academic Editor: Michele Amoretti

Introduction

The urban fleet of vehicles is rapidly evolving from a collection of sensor platforms that provide information to drivers and upload filtered sensor data (e.g. global positioning system (GPS) location and road conditions) to the Internet cloud to a network of autonomous vehicles (AUVs) that exchange their sensor inputs with each other in order to optimize a well-defined utility function. This function, in the case of autonomous cars, is a prompt delivery of the passengers to destinations with maximum safety and comfort and minimum impact on the environment. In other words, one is witnessing in the vehicle fleet the same evolution from sensor web (i.e. sensors are accessible from the Internet to get their data) to the Internet of Things (IoT, the components

with embedded sensors are networked with each other and make intelligent use of the sensors). In the intelligent home, the IoT formed by the myriad of sensors and actuators covering the house internally and

¹University of California–Los Angeles, Los Angeles, CA, USA

²Incheon National University, Incheon, Korea

³LIP6, Université Pierre et Marie Curie (UPMC), Sorbonne Universités, Paris, France

⁴Korea Advanced Institute of Science and Technology, Daejeon, Korea

⁵Electronics and Telecommunications Research Institute, Daejeon, Korea
Eun-Kyu Lee and Uichin Lee are principle corresponding authors.

Corresponding author:

Jae-Han Lim, Electronics and Telecommunications Research Institute,
Daejeon 34129, Korea.

Email: ljhar@etri.re.kr



Creative Commons CC-BY: This article is distributed under the terms of the Creative Commons Attribution 3.0 License

(<http://www.creativecommons.org/licenses/by/3.0/>) which permits any use, reproduction and distribution of the work without

further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<http://www.uk.sagepub.com/aboutus/openaccess.htm>).

externally can manage all the utilities in the most economical way, with maximum comfort to residents and with virtually no human intervention. Similarly, in the modern smart grid, the IoT formed by all components, large and small, can manage power loads in a safe and efficient manner, with the operators now playing the role of observers.

In the vehicular network, like in all the other IoTs, when the human control is removed, the AUVs must efficiently cooperate to maintain smooth traffic flow in roads and highways. Visionaries predict that the vehicles will behave much better than drivers, handling more traffic with lower delays, less pollution, and certainly better driver and passenger comfort. However, the complexity of the distributed control of hundreds of thousands of cars cannot be taken lightly. If a natural catastrophe suddenly happens, say an earthquake, the vehicles must be able to coordinate the evacuation of critical areas in a rapid and orderly manner. This requires the ability to communicate efficiently with each other and also to discover where the needed resources are (e.g. ambulances, police vehicles, information about escape routes, and images about damage that must be avoided). Moreover, the communications must be secure to prevent malicious attacks that in the case of AUVs could be literally deadly since there is no standby control and split-second chance of intervention by the driver (who may be surfing the web).

This efficient communications and distributed processing environment can be provided by a new network and computing paradigm specifically designed for vehicles—the vehicular fog. This mobile cloud provides several essential services, from routing to content search, spectrum sharing, dissemination, attack protection, and so on to AUV applications via standard, open interfaces. This article discusses the evolution from intelligent vehicle grid to vehicular fog and autonomous, Internet-connected vehicles. In particular, we highlight the advantages of the Internet of Autonomous Vehicles and at the same time expose its challenges stemming from networking for content distribution to possible attacks.

The rest of the article is organized as follows: section “Vehicular ad hoc network contents” discusses the characteristics observed in emerging vehicle applications. Section “Instantiating the Internet of Vehicles” introduces our vision of trends toward an intelligent vehicle grid and impact on the AUV, followed by detailed description of the vehicular fog—functions, computing, networking, and resources in section “Vehicular fog.” In section “Vehicular fog and AUV challenges,” we discuss the research challenges in the vehicular fog when applying it to an autonomous driving application on the road. Section “Attribute-based content discovery in vehicular fog” examines a content discovery challenge in depth. Finally, we conclude the article in section “Conclusion.”

Vehicular ad hoc network contents

A vehicular ad hoc network (VANET) has enjoyed a variety of applications, from safety and comfort to entertainment and commercial services. This section discusses three important characteristics of application of content items in VANET and discusses their impacts on an emerging application, a connected, AUV.

First, vehicles in emerging applications share a huge amount of sensor data (contents) and collaborate to complete a common task. On-board sensors record a myriad of physical phenomena, and vehicle applications collect such sensor records, even from neighboring vehicles, to produce value-added services.¹ In the CarSpeak application, a vehicle accesses sensors on neighboring vehicles in the same manner as it accesses its own.² The vehicle then runs an autonomous driving application using the sensor collection. The MobEyes enables a vehicle to video-record all surrounding events including car accidents while driving.³ Thereafter, if indeed an accident is reported, mobile agents (e.g. police) search the vehicular network for witnesses as part of their investigations. Collaboration in the sharing and processing of sensor data will be one of the *strong assets of the AUVs*. The continuous sharing of position data is essential to guarantee stability of the autonomous fleet. The crowdsourcing of road conditions (poor pavement conditions, obstacles, accidents, etc.) using the collection of available sensors will allow smooth driving even in perilous conditions.

Next, vehicle applications are only interested in content itself, not its provenance—named content-centric dissemination.⁴ When checking the traffic jam of vehicular traffic in the Internet, people visit favorite service pages providing ample information. In contrast, vehicle applications flood query messages to a local area, not to a specific vehicle, and accept responses regardless of the identity of content providers. This pattern occurs because the sources of information (vehicles) are mobile and geographically scattered. Content-centric dissemination will also play a major role in the *management and control of the autonomous car fleet* for two reasons. First, the AUV will travel at high speed and short distance from neighbors (on highways) and must have very up-to-date information of surrounding vehicles up to several kilometers in order to maintain a stable course. Thus, in the content-centric dissemination style, the vehicle periodically solicits position, speed, and direction from the rest of the fleet. Second, in case of accident ahead, the vehicle must alert the driver (who may have been occupied in other matters) of the urgency so that the driver has the option of manual intervention. In this case, to prepare the driver for takeover, the vehicle retrieves photos and possibly video of the accident scene for the cameras of the vehicles facing the accident. Content-centric dissemination allows access to the best

cameras with the needed data, without prior knowledge of the cars that offer the data.

Finally, vehicles consume a great amount of contents while at the same time producing the contents—that is, vehicles function as rich data “prosumers.”⁵ Such contents show several common properties of location relevance—local validity and local interest. *Local interest* represents that nearby vehicles are the bulk of potential content consumers. This concept is further extended so as to distinguish the scope of consumers. For instance, all the vehicles in the vicinity want to receive safety messages, while only a fraction are interested in commercial advertisements. *Local validity* indicates that vehicle-generated content has its own spatiotemporal scope of validity to consumers. For instance, a speed-warning message near a sharp corner is only valid to vehicles approaching the corner, say within 500 m. Likewise, road congestion information may invalidate after 30 min. The location relevance implies the scalability of the data collection/storage/processing applications, since old data is discarded. It also implies that the data should be kept on the vehicles rather than uploaded to the Internet, leading to enormous spectrum savings. Thus, this property will be key to the *scalability of the AUV* concept, given the huge amount of data collected by AUV sensors.

Instantiating the Internet of Vehicles

Vehicles with embedded sensors generate copious amounts of data every second. At the same time, roads are instrumented with smart dust components,⁶ radio-frequency identification (RFID) tags,⁷ and embedded microcontrollers. These *things* constitute a *vehicle grid*, that is, an intelligent transportation infrastructure analogous to the energy grid for intelligent power generation and distribution. The next trend we want to report is the emergence of the vehicular fog. It instantiates the Internet of Vehicles by inter-networking all things that sense and move in the grid and by coordinating them to provide a computing environment. That is, the vehicular fog emerges as a computing and networking model for the systematic implementation of protocols and services required for vehicle applications in the grid.

One of the major benefits of the vehicular fog will be *autonomous driving*. Recall that the AUV must be capable of sensing its surroundings and of self-driving without human inputs.⁸ To this end, it uses a myriad of on-board sensors, ranging from radio detection and ranging (RADAR), GPS, video cameras to controller area network (CAN) bus sensors that monitor vehicle’s internal operation status. An advanced autonomous driving system processes all the sensory data, constructs the traffic map, identifies appropriate paths and

avoids obstacles on such paths, and makes driving safe and comfortable. Google and Daimler-Benz recently demonstrated autonomous driving system prototypes on real roads (Google driverless car, http://en.wikipedia.org/wiki/Google_driverless_car; Daimler-Benz Intelligent Drive, http://techcenter.mercedes-benz.com/en/intelligent_drive/detail.html). Academia has also demonstrated meaningful results by running autonomous driving tests.^{9,10} In the future, as addressed in Kumar et al.,² accessing sensors of neighboring vehicles will significantly improve the accuracy and safety of the driving. The vehicular fog will provide the ideal system environment for the coordinated deployment of the sensor aggregation, fusion, and database sharing applications required by the future AUVs.

Vehicular fog

The concept of a vehicle and a software system in it are evolving toward an intelligent agent performing local collaborations with neighboring vehicles by sharing contents. We claim that a vehicular fog is the core system environment that makes this evolution possible. This section discusses two integral functions of the fog, that is, networking and computing, which is followed by a case scenario of autonomous driving.

Information-centric networking

In addition to the characteristics of vehicle contents discussed in the previous section, the vehicular fog is distinguished from other IoT instantiations with two unique properties of mobility and vehicle-to-vehicle (V2V) communications. In the smart grid, for instance, most energy components are stationary and controlled in a centralized, hierarchical manner. This enormously helps scalability from room to building to city. However, vehicles move around fast and thus cannot be hierarchically partitioned and controlled. The mobility also makes their network connectivity unstable. Instead of relying entirely on communication infrastructure, vehicles form an ad hoc network and communicate each other directly. Thus, vehicle interactions via V2V communications are critical. In this mobile, ad hoc network setting, nodes (vehicles) join and leave a network frequently, and it is not trivial to assign Internet protocol (IP) addresses to such mobile nodes. We note that the VANET protocol still assumes using IP address to represent a host. It is also not easy to discover the IP address of the publisher of a specific content in the network since the content of interest may not be consistently bound to a unique IP address. The vehicular fog must take into consideration this addressing problem in its network design.

Information-centric networking (ICN) has been emerging as a potential solution to solve the addressing

limitation discussed above.¹¹ ICN was initially conceptualized as a general form of communication architecture to achieve efficient content distribution on the Internet. ICN focuses on *what (content)* instead of *where (host)* to fulfill primary demands from both content consumers and publishers. Consumers are interested in content regardless of the originator and publishers strive to efficiently distribute content to consumers. To this end, ICN uses content names instead of IP addresses so that contents are decoupled from publishers. Some of the recently proposed ICN architectures in the Internet context include Data-Oriented Network Architecture (DONA), Named Data Networking (NDN), Publish-Subscribe Internet Routing Paradigm (PSIRP), and Network of Information (NetInf).¹²

NDN, of these architectures, has been recently extended to vehicular networks.^{13–15} NDN defines two types of packets: Interest from consumers and data (i.e. content) from publishers.¹⁶ Content name in these packets is used for routing. A consumer requests content by broadcasting an Interest with its name toward potential publishers. When a publisher receives the Interest and has data matching the Interest, it replies with the data back to consumer using the Interest path in reverse. NDN allows routers on the path to cache the content so that they can reply the cached content to consumers once they receive the matching Interest. This way, NDN achieves an effective content distribution that the vehicular fog critically requires to support its content oriented applications.

Computing in a vehicular network

In a vehicular network, as discussed, most of the contents picked up by vehicles hold location relevance—that is, most of our queries are about the world surrounding us and our neighboring vehicles are the best probes. For instance, a driver dispatches a query to find the cause of a sudden traffic jam (say, a minor accident on a block ahead). This type of information is created, stored, and distributed within a vehicular fog. It is too costly to upload every small content to the Internet, and too time-consuming to search and download interesting contents from the Internet cloud. Besides, frequent Internet connections quickly deplete sharable communication resources and thus affect performance of other vehicle applications. In the vehicular fog, the data of interest may be scattered across many vehicles and will require in-loco data aggregation and query resolution. The vehicular fog, thus, must be able to provide a computing environment to support localized data processing.

Recent research proposals on mobile fog computing (MFC) resolve the problem above using a self-organized ad hoc computing model. Vehicles in the

vicinity opportunistically form a local group (vehicular fog) for a cooperative computing in which vehicle contents and services are produced, maintained, and consumed. The MFC model leverages the increasing processing and storage capacity of the vehicles and mobile devices. It constructs a distributed computing environment using the collection of vehicles' computing resources, which primarily aims at extending the capability of vehicle interactions.

In the MFC research context, Bonomi et al.¹⁷ catch two features in a vehicular network—dynamic connectivity and interactions: cars to cars, cars to access points, and access points to access points. Authors introduce the concept of fog computing. Fog computing is a highly virtualized platform that provides computing, storage, and networking services between end devices and micro data centers located at the edge of network.¹⁸ This connection can be extended to traditional Internet clouds. Vehicular cloud computing, another computing model for a vehicular fog, has also been studied in academia.^{19–21} Gerla²² introduces its concept and discusses how it differs from mobile cloud computing that promotes direct interactions between mobile nodes of limited resources and a conventional Internet cloud.²³ Nodes access and offload expensive operations to the Internet cloud providing unlimited computing resources. They also store/download contents to/from the Internet.

Computing resources in the fog

An ad hoc virtual network (vehicular fog network) is created for collaborations among network members (vehicles) to produce advanced vehicular services that individual alone cannot make. Unlike Internet, computing platforms such as Amazon Web Service that are created and maintained by a cloud provider, a vehicular fog is temporarily created by interconnecting resources available in the vehicles and roadside units (RSUs). Such networked resources together function as a common virtual platform on which the efficiency of collaboration is maximized. MFC and ICN together contribute to creating the fog and to running the virtual platform efficiently.

Resources in the vehicular fog are distinguished from the ones in the conventional cloud. Each vehicle has three categories of resources: data storage, sensors, and computing as shown in Figure 1. The storage stores vehicle contents generated from applications and sensors as well as traditional multimedia files. It supports data sharing between fog members by accepting a search query and replying with matched contents. Sensors are able to self-actuate as well as to detect events of physical world. Following the Internet of Things model, each sensor is connected to the Internet,

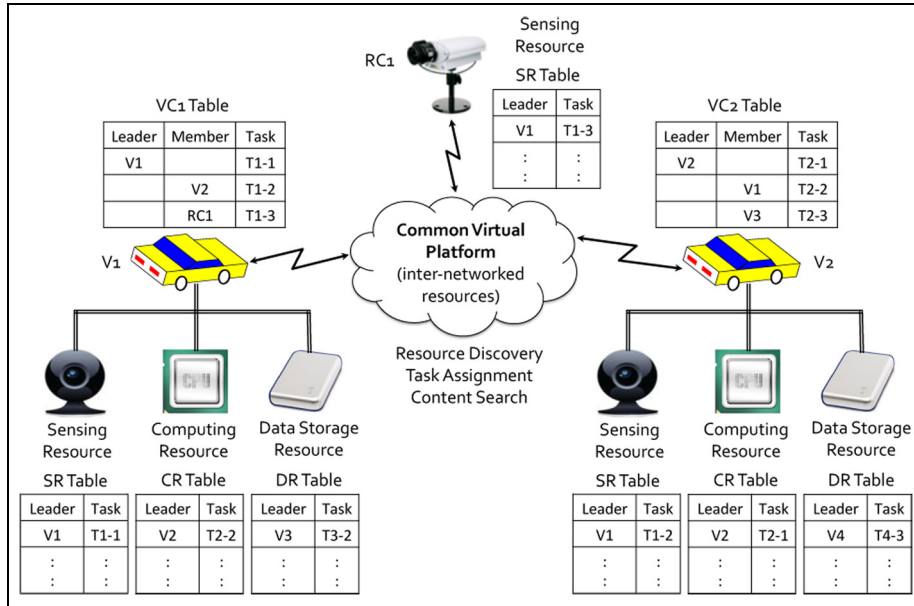


Figure 1. Fog resources include data storage, sensors, and computing. They are shared to create a common virtual platform.

so that it can be read and controlled by external systems.

In the vehicular fog, the resources are inter-networked via purely peer-to-peer connections. Each vehicle negotiates the level of resource sharing directly with each other. For efficiency, one vehicle in the fog may be elected as a broker based on some selection metrics (e.g. connectivity). Then, it mediates the resource sharing process as well as other fog operations. An RSU, joining the fog as a stationary member in Figure 2, can be a good candidate for the negotiator role. We also envision the deployment of resource-constrained RSUs such as cameras. They may not have enough storage and computing power, but still have reliable connections to vehicles. If this is the case, they can store and manage data indexes for effective content discovery.




Figure 2. Resources in the fog are inter-networked in a purely decentralized manner. We borrow the V2V/V2I communication architecture from VANET.

Autonomous driving on the road: case scenario

Given the collection of resources from vehicles and RSUs and their potential interconnections, we illustrate how the fog establishes a virtual computing platform and to enable fog-type collaboration in it. We use a simple autonomous driving scenario as shown in Figure 2.

To discover fog resources. Suppose that a vehicle V_1 (a fog leader) self-organizes as a vehicular computing fog to complete an autonomous driving application. The application requires images of next three road segments in order to improve the accuracy of context awareness, yet resources in V_1 only covers one road segment. The

fog leader sends out a rreq to recruit vehicles and RSUs in the right positions that can provide the right sensing resources such as a camera.

To form a mobile fog. Upon receiving rreps containing resource information from nearby vehicles and RSUs, the leader selects two fog members (say, a vehicle V_2 and a road camera RC_1) and forms a new fog.

To assign tasks and collect results. The fog leader, then, assigns tasks of taking a picture of the next two block scenes and of returning the data back to it.

To publish and share contents. After collecting images from fog members, the leader processes the collection to create new content that is published to the entire network. V_1 consumes the content for its autonomous driving application. At the same time, the leader asks other vehicles (V_4 in Figure 2) to store and keep the content in their storage for the purpose of potential reuse of the contents around the fog. When, sometime later, the following vehicles V_6 and V_7 run their autonomous driving applications, they request the contents by broadcast an Interest message with the content name. Finding a match, V_4 can transmit the matched contents to V_6 and V_7 directly without contacting V_1 .

To maintain the fog. In the meantime, the leader may receive a fog leave message from a fog member. Then, it selects a replacement among nodes that sent rrep in the resource discovery phase and have sufficient resources to complete the task assigned to the member that just left. The leader reassigns the task and updates the fog table.

To release the fog. When the fog leader decides not to use the fog any more, it sends a fog release message to all the fog members V_2 and RC_1 .

Vehicular fog and AUV challenges

The evolution from manually operated to AUV will pose several new challenges. Some of these challenges come from the massive deployment of sensors on the AUV and the huge amount of data that the AUV picks up from the environment. Other challenges result from the fact that the AUV “drives itself autonomously” while the driver may be busy with background activities and not capable of intervening immediately in case of emergencies. After all, a much-advertised AUV benefit is the ability of the driver to engage in other activities as if he or she were on a train—“with wheels.” In this section, we review these challenges and their impacts on vehicular protocols and applications and more generally on the vehicular fog architecture design.

Naming and discovering contents in instantaneous fogs

The previous section shows that the “narrow waist” network layer is NDN (ICN, more generally) that finds content using naming hierarchy. In fact, due to node mobility, one cannot assume that there is a geographically consistent name hierarchy such that the prefix location gives a hint about the location of the target content. In the vehicular fog, however, most queries will be location relevant and content is found by exploiting geographic relevance. For instance, we wish to find a

video clip of a museum in a certain area of the city, or witnesses in a car accident, or information about pavement conditions on a given route (e.g. potholes and bumps), an ambulance near a train station, or a photo or video of a congested street we are supposed to drive through. This “environment monitoring” service will become popular when there are lots of AUVs on the road, equipped with all sorts of sensors, from vibration sensors to video cameras and GPS, and capable to capture every detail of the environment. Today, Google cars roam the city and map topology, and combine the actual pictures of the buildings. Visionaries believe that AUVs will map the entire “world” more so than regular cars, and they will maintain the index to this “mapped world.” Finding the desired content in this large volume of environment data stored on the AUVs will be a challenge for the networking service in the fog. Section “Attribute-based content discovery in vehicular fog” examines this issue in detail.

Content sharing via V2V communications

Beacons and alarms. One important service built within the vehicular fog is “Beaconing and Alarms.” Recall that the AUV sensors (from optical to Lidar) do most of the work in an attempt to keep the vehicle and its passengers out of trouble. Sensors alone, however, are not sufficient to maintain stable operations in high speeds and extremely reduced inter-vehicle spacing. This is particularly true in vehicle platoons (Figure 3). In this case, it was found that communications from front to rear trucks are necessary to avoid the onset of oscillations. Likewise, V2V communications are necessary to avoid the formation of shock waves in a long column of AUVs when a slow down or accident occurs in front.²⁴

Intersection collisions will not be so critical when most cars are autonomous, since the AUVs (unlike human drivers) abide by the signals and speed limits and approach intersections with caution. Nevertheless, V2V communications will still be required among lead cars facing 4-stop intersections in order to implement



Figure 3. An example driving of vehicle platoon.

the “smart traffic light.”²⁵ The electronic light schedules groups of cars across the intersection just as a real traffic light would do, dramatically reducing delays.

File and media downloading. Efficient downloading of multimedia to drivers and passengers (e.g. TV shows, movies, and games) will be a critical marketing strategy for the autonomous driving. Previous research in this area has shown that in the crowded wireless access spectrum, the download of popular content from the web is best done using bit torrent techniques via V2V support.²⁶ Downloading from WiFi access points or long-term evolution (LTE) alone will not work.

Content distribution to AUVs is also motivated by safety considerations. For instance, drivers in the middle of a convoy traveling bumper to bumper at 60 miles/h will be reassured, when they are able to capture the video of the lead car. It will give them the impression of “being in control” without having to work on the commands. Even more important will be the immediate delivery of the video, or image, of an accident scene to AUV drivers to alert them of the severity of a problem ahead and let them judge if they should take on the control.

A possible scenario of media file propagation is as follows: the beacons inform the AUV’s upstream of the presence of an accident in location (x, y), say. A particular AUV determines that the accident can impact its drive and submit an “interest” (in NDN terminology) to the location in question. The first video camera facing the accident responds by returning the video, following the Pending Interest Table (PIT) pointer trail in reverse. Other vehicles can join the multicast tree as well. Clearly, this broadcast can be supported only by V2V communications. LTE would introduce too much latency and would not scale.²⁷

Collaborations among connected vehicles

Intelligent transport. The introduction of the autonomous driving will greatly enhance intelligent transport. The AUVs will be able to use the existing highway network much more efficiently than manually operated cars because they can be packed in compact platoons and convoys. They can also make efficient use of preferred (or pay-per-service) lanes, by maintaining a “train on wheel” configuration on such lanes, and by allowing efficient in-and-out lane switches using a combination of sensors and V2V communications in a much safer way than human could (given the high speeds involved). The AUVs can also become aware of other mobiles sharing the road such as pedestrians and bicycles. Vehicles can track them with their sophisticated sensors/Lidars and share the information of “bike ahead”

with vehicles behind and one of two lanes across through V2V communications.

Recovery from infrastructure failure. The AUVs depend on the infrastructure (e.g. WiFi access points, RSUs, and LTE) for several non-safety functions such as advanced sensor data processing and intelligent transport. In the case of a major infrastructure failure caused by an earthquake, say, some of these functions must be taken over by human drivers. However, there is a gray period, between when a massive infrastructure failure occurs and when the human takes over of navigation, during which the AUV systems must deal with the problems on their own. This is a very critical window because the AUVs only know about their immediate neighbors. After the disaster, they have lost knowledge of the neighbors beyond the reach of their sensors, which was provided by an Internet transaction server (ITS) server. To avoid a second disaster, caused by the AUVs going out of control, it is important to maintain a V2V-supported propagation of traffic conditions and congestion state on adjacent roads. This background “crowdsourcing” of traffic will allow the AUV systems to make intelligent routing decisions (to avoid obstacles or blocked roads in case of earth quakes) so that the human drivers can progressively takeover with confidence.

Connecting to the Internet cloud: vehicular traffic management

Tasks in vehicular traffic management and route optimization include “measuring the vehicular traffic” in real time and “informing vehicles of the new routes.” To measure the traffic, the Department of Transportation in the last decade instrumented the highways with sensors under the pavements and video cameras, which is a costly solution. Then, the information about the “best route” was conveyed to drivers with billboards, radio announcement, and, more recently, the Internet. Unfortunately, sending the same instruction to all the vehicles had the effect of creating “route flapping” problems and route instabilities. Everybody rushes to the newly announced route.

Recently, the introduction of on-board navigators has changed all that. The navigator service agency can learn instantaneous traffic flows and patterns from the mobile fog, and can deliver differentiated route instructions to vehicles thus avoiding route flapping. In the envisioned “mobile fog enabled” traffic management, on-board vehicle navigators periodically send time, GPS coordinates and final destination information to a navigation server in the Internet. The server estimates road segment loads and delays, and constructs the traffic load map as well as the traffic pattern matrix. It

then computes optimal incremental routes and returns such routes to vehicles upon request. An important benefit of the individual interaction between navigator server and on-board navigator (as opposed to traffic billboard announcements) is the fact that the former allows to balance the load among multiple route options. Moreover, the on-board navigator may choose, within some limits, between different route recommendations depending on the driver profile (aggressive or conservative driver) and type of vehicle (say, combustion or electric engine). Simulation results confirm the convergence to the optimal, minimum delay route solution at quasi-steady state.²⁸ This application is a good example of synergy between vehicular fog and Internet cloud. In particular, the sensing of segment traffic congestion is done in the vehicle fog (by means of reporting time and GPS position successive snapshots), as well as the route “actuation,” through instructions received by the on-board navigators from the navigator server. The navigator server, implemented in the Internet cloud, does the rest. Namely, it computes the traffic pattern, from the destination ID carried by each on-board navigator message. It computes optimal incremental routes and dispatches such routes to the on-board navigators.

Enhancing performance of underlying fog network

Congested wireless medium and cognitive radios. The dedicated short-range communication (DSRC) spectrum, in principle, can support V2V traffic, or at least the traffic for beacons and emergency services. However, visionaries anticipate that the DSRC 75 MHz spectrum will be quickly exhausted by basic safety applications. In such cases, previous studies have shown that the V2V requirements must be supported by the WiFi spectrum in a dynamic spectrum sharing mode, competing with residential users in an urban environment.²⁹ The cognitive radio functions must be supported by a multi-radio AUV platform. Their capability can be enhanced by AUV crowdsourcing of the occupancy of the 802.11b/g channels ahead. Collective tracking of available channels using sophisticated on-board radios will allow careful mapping of the available spectrum.

System issues and virtualization. Virtualization is one of the most fundamental features in the Internet cloud. It also plays an important role in the vehicular fog, especially in the support of AUVs. Because of the rich assortment of sensors on-board, the AUV fleet may be required to perform “data mining” like tasks such as recognizing a fugitive in a certain geographic area. The AUVs can do some initial filtering and correlation of images of interest. But, for final processing, these data must be uploaded to a virtual image of the pattern

recognition process in the Internet cloud. Another important function of virtualization is the customization of the sensor platform to different applications, often executed also for the privacy of the drivers. For example, the car manufacturer can access all CAN bus sensors and cameras, while a neighbor vehicle may access only the outward pointing camera.

Security and privacy

A major incentive for participants in the vehicular fog is to protect data and allow users to decide what information could be exposed and what information should be kept private. Moreover, functions, data, and trust validations of mobile applications can be delegated to a vehicular fog, if mobile devices and mobile users become temporarily disconnected. The fog also provides protection from devices that have been penetrated by the adversary, or exhibit uncontrolled, disruptive behavior.

In addition to the common security requirements like confidentiality, integrity, privacy, and authentication, the AUV is very vulnerable to vicious attacks that may, say, disable the steering or the brake system. The latter attacks are of concern with normal cars with a human driver in control. They are extremely dangerous for AUVs because there is no driver on instant standby. For this reason, the protection from attacks both external (from access points or from conventional vehicles) as well as internal (from other AUVs) must be designed with stricter standards. Accessing to the cars’ internal mechanism and possibly to on-board diagnostics (OBD) and CAN bus must be allowed when the AUV is out of control because of either internal malfunctioning or malicious attack. In this sense, proper enforcement of access control emerges as a first-line protection strategy in the vehicular fog. A simple management using password and role assignment is not enough.³⁰

In addition, botnet research has been paid special attention as threat using botnets becomes reality in the IoT³¹ and its consequence in the vehicular fog will be more disastrous. Denial of service (DoS) is also of importance because most communications including V2V rely on wireless medium. Radio-frequency (RF) jamming can create large communication-blind areas, failing to deliver warning messages in a timely manner required for critical safety applications.³²

Attribute-based content discovery in vehicular fog

Using attribute in the content discovery

Section “Naming and discovering contents in instantaneous fogs” shows that the networking service (i.e. NDN) finds content using naming hierarchy, which is

limited in the vehicular fog as names are inconsistent with geographical relevance. A solution approach to resolve the limitation is to exploit geolocation information along with names in a content discovery process.

To investigate this approach further, we adopt a general term *attribute*.³³ An attribute is a word describing a characteristic of an object, and its example includes name, geolocation, date of generation, expiration, type, and so on. In general, each object has a set of attributes that help describe and identify it. The concept of the attribute has been studied in many research areas, especially in security (e.g. identification,³⁴ encryption,^{35,36} access control,^{33,37,38} and signature^{39,40}). There, a network node is identified and qualified by a set of attributes (e.g. “UCLA,” “CS,” and “Professor”), not by an IP address. This way, a node does not need to memorize the identity of potential communication partners, handling scalability issue in a modern, large network.

Attributes can be used in the content discovery, which is comparable with the index and keywords in the Internet search. A search engine collects web data using crawlers, does indexing, and creates a search database. Then, users send a query with several keywords to the engine to find contents of interest. NDN, designed initially for the Internet, executes the content search using one attribute (i.e. content name) at the network layer.

The content discovery in the vehicular fog is distinguished from the Internet search in that it works in a decentralized manner without a central search engine. Moreover, the vehicle content does not accumulate within the fog; it easily invalidates or becomes unreachable due to mobility. By accommodating these differences, the content discovery process exploits attributes as follows. When producing a content, its owner tags the content with a set of attributes and publishes both content and corresponding attributes to the fog. According to the fog’s configuration, they may be stored/maintained in more than one fog members together or separately. When a data consumer finds a content, it prepares a query with an arbitrary set of attributes representing the content of interest and then broadcasts the query to the fog. Any fog member having content whose attributes match the query responds directly to the consumer. Optionally, the consumer may retry the query with an adjusted attribute set in a way to increase the matching probability when his first trial fails to find one.

Evaluating performance of content discovery

A vehicular fog is formed by randomly encountered vehicles and RSUs and operates in a distributed manner. On the one hand, contents may be stored in more than one fog members and thus the consumer may receive multiple results. On the other hand, the query

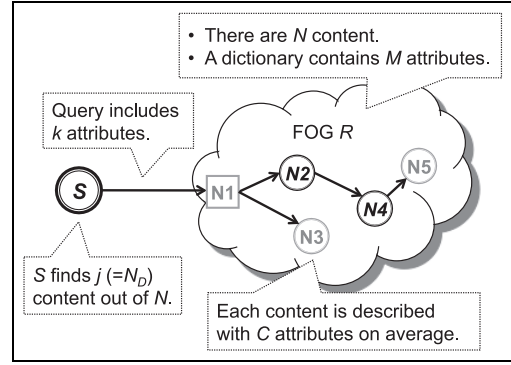


Figure 4. A vehicular fog scenario for experiment of attribute-based content discovery.

may not match and the consumer fails to find required content. This section is to enumerate how well a query finds contents in the fog and to assess performance of the attribute-based content discovery.

In our experiment scenario (see Figure 4 and Table 1), a data consumer (S , or source node) prepares a query for a content by selecting k attributes. A vehicular fog (R) randomly produces N contents in total, each of which is tagged with C attributes on average. R also maintains a dictionary of M attributes—that is, S selects k attributes for the query and R selects C attributes for the tagging out of the dictionary. Initially, we assume that each attribute is equally likely selected from the dictionary—homogeneous case. See Appendix 1 for our content discovery model. We release this assumption later. When the query reaches the fog R , a matching process compares k attributes in the query (\mathcal{A}_x) with C attributes of every N contents (\mathcal{A}_y). We say that the query *hits* i th content if $\mathcal{A}_x \in \mathcal{A}_{y(i)}$, where $\mathcal{A}_x = \{a_1, \dots, a_k\}$ and $\mathcal{A}_y = \{\mathcal{A}_{y(i)} = \{a'_1, \dots, a'_C\} | 1 \leq i \leq N\}$. The hits are counted and recorded (N_D) during the process. To measure performance, we compute a variable *hitratio*, $h = N_D/N$. The h value represents the probability that the consumer S successfully discovers required contents under unreliable network connectivity. It also implies how quickly S obtains the content because the higher h is, the more probable S receives the content from close neighbors (we roughly assume a content is duplicated equally likely over the fog). Our evaluation examines the impact of various systems parameters on the hit ratio and thus on the ability of content discovery.

For our evaluation, we implement the attribute-based content discovery on QualNet and run simulations. Since our experiments focus only on content discovery, we assume that there is a wireless protocol connecting network nodes that form a vehicular fog together. Some nodes move around randomly representing vehicles and others do not. In a network, N nodes are deployed, representing N contents. Each

Table 1. Symbol and notations used in our content discovery model.



| Symbol | Notation |
|-----------|---|
| S | A data consumer sending a query to the fog |
| k | Number of attributes selected for the query by S |
| R | A vehicular fog |
| N | Total number of content produced in R |
| C | Number of attributes tagged to each content in R |
| M | Total number of attributes in the attribute dictionary of R |
| $N_D = j$ | Number of qualified content in R (having k attributes) |

node selects C attributes on average out of given M attributes randomly or by following specific probability distribution. For instance, a node i stores only one content $Y(i)$ that is tagged by a set of attributes $\mathcal{A}_{y(i)} = \{a'_{i1}, \dots, a'_{iC}\}$. A new node S tries to find a content X that is described by a set of k attributes. Thus, S selects an attribute set $\mathcal{A}_x = \{a_1, \dots, a_k\}$ from the given M attributes and broadcasts the set (query) to the entire network. Upon receiving the query, the node i checks if $\mathcal{A}_x \in \mathcal{A}_{y(i)}$. If this is the case, the node i sends an acknowledgement and corresponding content $Y(i)$ back to S . Throughout experiments, we vary parameters listed in Table 1, run simulations 500 times for each parameter setting, and illustrate averaged performance results.

Experiments and results

Number of attributes tagged to each content in R . The first experiment investigates the impacts of C and k on h . C varies from 10 to 40 and k varies from 1 to 20, while N and M are fixed to 200 and 50, respectively. Figure 5 shows the results. As k grows, h declines exponentially. With $C = 30$, h begins at around 0.6 when $k = 1$. When k increases to 4 and 8, h sharply drops down to 0.12 and 0.01, respectively. Finally, h goes to 0 when $k \geq 12$. This recommends that the sender S must include less than 12 attributes (24% out of 50 attributes) in the query to find at least one content in R . Figure 5 also illustrates how h changes with C . When $k = 1$, h is 0.2, 0.4, 0.6, or 0.8 with $C = 10, 20, 30,$ or 40, respectively. One thing to note is the threshold of k when h goes below 0.005 (i.e. $N_D = 1$). The thresholds are 4, 8, 12, and 20 with $C = 10, 20, 30,$ and 40, respectively.

Distribution of C . The first experiment assumes the homogeneous case. In practice, however, some content may have more numbers of attributes (group X) whereas some other have less attributes (group Y). For instance, some contents contain much information so that they are described with more numbers of attributes. To realize such a heterogeneous case, we

**Figure 5.** Impact of C and k on h .**Figure 6.** Variance of C (σ) influences h . $N = 200$ and $M = 50$.

consider that the number of attributes that a content in R has follows the normal distribution with mean = C and variance = σ . A high σ means that contents in R have very different numbers of attributes. In the second experiment, we vary C and σ while $N = 200$ and $k = 5$. Figure 6 shows the results confirming that as C increases, h grows. It also demonstrates that h grows with increasing σ values. As σ increases, there are more contents in the group X which are highly likely to hit the query. This stability makes the curve of $\sigma = 20$ less exponential than other curves. One interesting observation is that all the curves converge when $C = 43$. This is the case when a content has 86% of attributes (43 out of 50) on average. After this point, the curves are inverted. When $C = 50$, most content has around 50 attributes (i.e. in the group X) with $\sigma = 1$, which results in high value of h . However, with $\sigma = 20$, there are still notable numbers of contents in the group Y that do not hit the query.

Attribute popularity. Next experiment considers attribute popularity. That is, some attributes are more popular




Figure 7. Reference of attributes with Zipf-like distribution.




Figure 8. Attribute popularity influences h . $N = 200$ and $M = 50$.

than others so that more contents in R has them. This implies that attributes are ranked according to their popularity, and each content is more likely to select highly ranked attributes. To realize the concept, we take the Zipf's law that has been studied in web objects: The web page access rate follows Zipf-like distribution, generally considered as representative of its ranked popularity.⁴¹ In the distribution, the access probability of the i th most popular item is represented

$$P(i; \alpha, M) = \frac{1/i^\alpha}{\sum_{j=1}^M (1/j^\alpha)}, (1 \leq i \leq M)$$

where α is an exponent characterizing the distribution ($0 \leq \alpha \leq 1$). Figure 7 shows the Zipf-like popularity distribution given 100 ranked items. In our scenario, M attributes are ranked, and the consumer S and the fog R select k and C attributes with the distribution. The experiment fixes $\sigma = 1$, varies C and α , and then measures h . Figure 8 shows the impact of attribute




Figure 9. Impact of C and k on h . C follows normal distribution with $\sigma = 10$. Attribute popularity follows Zipf-like distribution with $\alpha = 0.9$.

popularity. The curve with $\alpha = 0.1$ shapes very similar to the one with $\sigma = 1$ in Figure 6. This occurs because the access probability with low α follows a uniform distribution (see Figure 7). As the α value increases, the popular attributes are frequently selected by S and R . In this way, the notion of the popularity affects content discovery performance. Compared to $\alpha = 0.1$, h increases seven times on average in the range of $10 < C < 50$ when $\alpha = 0.9$.

Impact of C and k . Having considered the distribution of C and attribute popularity, we repeat the first experiment (Figure 5) with $\sigma = 10$ and $\alpha = 0.9$. Figure 9 shows the experimental results. Compared with the curves in Figure 5, the hit ratio clearly increases with more realistic type of contents and attribute configurations. h reaches 0.07, 0.19, 0.4, and 0.68 with $C = 10, 20, 30,$ and 40 when $k = 4$, where h is increased by 0.2 on average. In particular, the curve with $C = 40$ benefits notably. Even when $k = 20$, h shows above 0.2, representing that the consumer finds more than 20% of content in the fog. As C decreases, two elements (i.e. distribution and popularity) affect the content discovery performance more critically. When $k = 8$, the consumer finds content with $C = 10$ and 20 , respectively. Note that the same configurations show $h = 0$ in Figure 9. The threshold of k making N_D below 1 evaluates this performance, and Figure 9 illustrates that thresholds are 12 and 20 when $C = 10$ and 20 .

Number of attributes in the dictionary. In the last experiment, we fix $C = 20$ and vary M and k . Other variables are same to those in the previous experiment. Results in Figure 10 show how the total number of attributes in the dictionary affects content discovery performance. $M = 40$ means that each content has 20 attributes




Figure 10. Impact of M and k on h . $N = 200$ and $C = 20$.

selected out of total 40 attribute candidates. As M increases, h values decrease quite linearly. This indicates that M has less effect on the performance than C . When $k = 1$, especially, h changes the least—it goes down from 0.88 to 0.62 (28.8%) as M increases from 20 to 50. The graph also illustrates the impact of k with varying M values. When $M = 20$, the difference between the maximum h and the minimum one is 0.34. As M increases up to 50, the difference grows 0.62 (by 84.5%). This denotes that as M increases, k affects performance stronger.

Conclusion

The urban fleet of vehicles is evolving from a collection of sensor platforms to the Internet of Autonomous Vehicles. Like other instantiations of the Internet of Things, the Internet of Vehicles will have communications, storage, intelligence, and learning capabilities to anticipate the customers' intentions. This article claims that the vehicular fog, the equivalent of Internet cloud for vehicles, will be the core system environment that makes the evolution possible and that the autonomous driving will be the major beneficiary in the cloud architecture. We showed a vehicular fog model in detail and discussed the potential design perspective with highlights on AUV for future research. We note that this research extends a previous article.⁴²

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the National Research

Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (2016R1C1B1016084).

References

1. Lee U and Gerla M. A survey of urban vehicular sensing platforms. *Comput Netw* 2010; 54(4): 527–544.
2. Kumar S, Shi L, Gil S, et al. CarSpeak: a content-centric network for autonomous driving. In: *Proceedings of the ACM SIGCOMM*, Helsinki, 13–17 August 2012.
3. Lee U, Magistretti E, Zhou B, et al. MobEyes: smart mobs for urban monitoring with a vehicular sensor network. *IEEE Commun Mag* 2006; 13(6): 52–57.
4. Amadeo M, Campolo C and Molinaro A. CRoWN: content-centric networking in vehicular ad hoc networks. *IEEE Commun Lett* 2012; 16(9): 1380–1383.
5. Bai F and Krishnamachari B. Exploiting the wisdom of the crowd: localized, distributed information-centric VANETs. *IEEE Commun Mag* 2010; 48(5): 138–146.
6. Smart Dust Project, <http://robotics.eecs.berkeley.edu/~pister/SmartDust/>
7. Lee E-K, Yoo YM, Park CG, et al. Installation and evaluation of RFID readers on moving vehicles. In: *Proceedings of the ACM VANET*, Beijing, China, 25 September 2009.
8. Ziegler J, Bender P, Schreiber M, et al. Making Bertha drive: an autonomous journey on a historic route. *IEEE Intell Transp Syst Mag* 2014; 6(2): 8–20.
9. Wei J, Snider J, Kim J, et al. Towards a viable autonomous driving research platform. In: *Proceedings of the IEEE intelligent vehicles symposium (IV)*, Gold Coast, QLD, Australia, 23–26 June 2013.
10. Broggi A, Buzzoni M, Debattisti S, et al. Extensive tests of autonomous driving technologies. *IEEE T Intell Transp* 2013; 14(3): 1403–1415.
11. Amadeo M, Campolo C and Molinaro A. Information-centric networking for connected vehicles: a survey and future perspectives. *IEEE Commun Mag* 2016; 54(2): 98–104.
12. Ahlgren B, Dannewitz C, Imbrenda C, et al. A survey of information-centric networking. *IEEE Commun Mag* 2012; 50(7): 26–36.
13. Wang L, Wakikawa R, Kuntz R, et al. Data naming in vehicle-to-vehicle communications. In: *Proceedings of the IEEE conference on computer communications network*, Orlando, FL, 25–30 March 2012.
14. Yu Y-T, Punihaole T, Gerla M, et al. Content routing in the vehicle cloud. In: *Proceedings of the IEEE MILCOM*, Orlando, FL, 29 October–1 November 2012.
15. Grassi G, Pesavento D, Wang L, et al. Vehicular inter-networking via named data. In: *Proceedings of the ACM HotMobile (poster)*, Jekyll Island, GA, 26–27 February 2013.
16. Jacobson V, Smetters DK, Thornton JD, et al. Networking named content. In: *Proceedings of the ACM CoN-EXT*, Rome, 1–4 December 2009.
17. Bonomi F, Milito R, Zhu J, et al. Fog computing and its role in the internet of things. In: *Proceedings of the ACM workshop on mobile cloud computing (MCC)*, Helsinki, 17 August 2012.

18. Vaquero LM and Rodero-Merino L. Finding your way in the fog: towards a comprehensive definition of fog computing. *ACM SIGCOMM Comput Commun Rev* 2014; 14(5): 27–32.
19. Whaiduzzaman M, Sookhak M, Gani A, et al. A survey on vehicular cloud computing. *J Netw Comput Appl* 2014; 40: 325–344.
20. Yu R, Zhang Y, Gjessing S, et al. Toward cloud-based vehicular networks with efficient resource management. *IEEE Netw* 2013; 27(5): 48–55.
21. Yan G, Wen D, Olariu S, et al. Security challenges in vehicular cloud computing. *IEEE T Intell Transp* 2013; 14(1): 284–294.
22. Gerla M. Vehicular cloud computing. In: *Proceedings of the IEEE Med-Hoc-Net*, Ayia Napa, 19–22 June 2012.
23. Fernando N, Loke S and Rahayu W. Mobile cloud computing: a survey. *Future Gener Comp Sy* 2013; 29(1): 84–106.
24. Forster M, Frank R, Gerla M, et al. A cooperative advanced driver assistance system to mitigate vehicular traffic shock waves. In: *Proceedings of the IEEE INFOCOM*, Toronto, ON, Canada, 27 April–2 May 2014.
25. U.S. Department of Transportation. Vehicle-to-Infrastructure (V2I) communications for safety, http://www.its.dot.gov/factsheets/v2isafety_factsheet.htm
26. Lee KC, Lee SH, Cheung R, et al. First experience with CarTorrent in a real vehicular ad hoc network testbed. In: *Proceedings of the mobile networking for vehicular environments*, Anchorage, AK, 11 May 2007.
27. Huang J, Qian F, Gerber A, et al. A close examination of performance and power characteristics of 4G LTE networks. In: *Proceedings of the ACM MobiSys*, Lake District, 25–29 June 2012.
28. Kim W and Gerla M. NAVOPT: navigator assisted vehicular route OPTimizer. In: *Proceedings of the international conference on innovative mobile and internet services in ubiquitous computing*, Seoul, South Korea, 30 June–2 July 2011.
29. Kim W, Choi BSC, Oh SY, et al. Cognitive multicast (CoCast) in vehicular networks using OFDM subchannels and network coding. In: *Proceedings of the IEEE ICNC*, Maui, HI, 30 January–2 February 2012.
30. Tesla cars can be hacked to locate and unlock remotely, <http://thehackernews.com/2014/03/tesla-cars-can-be-hacked-to-locate-and.html>
31. Proofpoint uncovers Internet of Things (IoT) cyberattack, <http://investors.proofpoint.com/releasedetail.cfm?releaseid=819799>
32. Pual O, Aguiar A and Gross J. In VANETs we trust? Characterizing RF jamming in vehicular networks. In: *Proceedings of the ACM VANET*, Lake District, 25 June 2012.
33. Hu VC, Ferraiolo D, Kuhn R, et al. *Guide to attribute based access control (ABAC) definition and considerations*. NIST SP no. 800-162, January 2014, <http://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.sp.800-162.pdf>
34. Lee S-H, Wong SH, Chau C-K, et al. InterMR: inter-MANET routing in heterogeneous MANETs. In: *Proceedings of the IEEE MASS*, San Francisco, CA, 8–12 November 2010.
35. Bethencourt J, Sahai A and Waters B. Ciphertext-policy attribute-based encryption. In: *Proceedings of the IEEE symposium on security and privacy*, Oakland, CA, 20–23 May 2007.
36. Akinyele J, Pagano M, Green M, et al. Securing electronic medical records using attribute-based encryption on mobile devices. In: *Proceedings of the ACM workshop on security and privacy in smartphones and mobile devices*, Chicago, IL, 17–21 October 2011.
37. Lee E-K, Gadh R and Gerla M. Resource centric security to protect customer energy information in the smart grid. In: *Proceedings of the IEEE smart grid communications*, Tainan, Taiwan, 5–8 November 2012.
38. Ruj S and Nayak A. A decentralized security framework for data aggregation and access control in smart grids. *IEEE Trans Smart Grid* 2013; 4(1): 196–205.
39. Li J, Au MH, Susilo W, et al. Attribute-based signature and its applications. In: *Proceedings of the ACM ASIACCS*, Beijing, China, 13–16 April 2010.
40. Okamoto T and Takashima K. Decentralized attribute-based signatures. In: K Kurosawa and G Hanaoka (eds) *PKC: LNCS*, vol. 7778. Berlin, Heidelberg: Springer, 2013, pp.125–142.
41. Breslau L, Cao P, Fan L, et al. Web caching and Zipf-like distributions: evidence and implications. In: *Proceedings of the IEEE INFOCOM*, New York, 21–25 March 1999.
42. Gerla M, Lee E-K, Pau G, et al. Internet of vehicles: from intelligent grid to autonomous cars and vehicular clouds. In: *Proceedings of the IEEE world forum on internet of things*, Seoul, South Korea, 6–8 March 2014.

Appendix I

Content discovery model

In a vehicular fog R , there are N content in total and a dictionary of M attributes. Each content has C attributes on average, each of which is randomly selected from the dictionary. Thus, there are NC attributes assigned to all the content in R . We initially suppose that all attributes are selected with equal probability, named homogeneous case. Let r be the average number of appearances per each attribute, that is, $r = \lceil NC/M \rceil$, and assume $r \geq j$. The homogeneous case assumes that each attribute has the same number of appearances.

$$\forall i \in \{1, \dots, M\} : app(a_i) = r$$

A data consumer S prepares a query that contains k attributes selected out of M attributes. Upon receiving the query, the fog R finds matching content (out of N). More generally, we define a function \mathcal{F} to be the number of content (j) in R that have the target k attributes. Let X a random variable that represents the number of qualified content in R given k . We begin with a distribution function $F(j) = P(X \geq j)$. Then, the final goal is to solve the problem function

$$\mathcal{F}(j) = \sum_{j=0}^N P(X = j) \times j \quad (1)$$

where $P(X = j) = P(X \geq j) - P(X \geq j + 1)$.

To compute $P(X \geq j)$, we define a group G as a set of all the content in R and define an event $E_{g_t}(j)$ —an arbitrary subgroup g_t in G contains j content such that all the j content are qualified, that is, they have k attributes, where $1 \leq t \leq {}_N C_j$. Given the universe cases of qualified content $\mathcal{J} = ({}_N P_r)^k$, the event is represented as

$$E_{g_t}(j) = \frac{({}_r P_j)^k \times ({}_{N-j} P_{r-j})^k}{\mathcal{J}} \quad (2)$$

Then, we represent $P(X \geq j)$ using $E_{g_t}(j)$ as

$$P(X \geq j) = P(E_{g_1}(j) \cup E_{g_2}(j) \cup \dots \cup E_{g_{{}_N C_j}}(j)) \quad (3)$$

To compute (3), we assume that an event $E_{g_t}(j)$ is independent of each other. When $j \ll N$, the size of each subgroup g_t in the event $E_{g_t}(j)$ becomes so small that their intersection event gets negligible because independent events would dominate in the entire event space. Now, suppose that j grows to N ($j \sim N$). Then,

the subgroups' size increases, and dependent events would influence $P(X \geq j)$ significantly. However, when j increases, j 's value would dominate in (1), and the value of $P(X \geq j)$ becomes negligible. Therefore, this assumption is reasonable.

$P(X \geq j)$ in (3) can be rewritten as follows. Let $u = {}_N C_j$

$$\begin{aligned} P(X \geq j) &= [P(E_{g_1}(j)) + \dots + P(E_{g_u}(j))] \\ &\quad + (-1)[P(E_{g_1}(j) \cap E_{g_2}(j)) \\ &\quad + \dots + P(E_{g_{{}_N C_{j-1}}}(j) \cap E_{g_u}(j))] + [\dots] + \\ &\quad \dots \\ &= (-1)^2 {}_u C_1 \cdot P(E_{g_1}(j)) \\ &\quad (-1)^{2+1} {}_u C_2 \cdot P(E_{g_1}(j))^2 \\ &\quad + \dots + (-1)^{u+1} {}_u C_u \cdot P(E_{g_1}(j))^u \\ &= \sum_{v=1}^u C_v \cdot P(E_{g_1}(j))^v \cdot (-1)^{v+1} \end{aligned} \quad (4)$$

Finally, we solve the problem function \mathcal{F} by applying (4) to (1).