

Towards autonomous vehicular clouds

Stephan Olariu^{1,*}, Mohamed Eltoweissy², Mohamed Younis³

¹Department of Computer Science, Old Dominion University, Norfolk, Virginia, USA; ²Pacific Northwest National Laboratory, Richland, Washington DC, USA; ³Department of Computer Science and Electrical Engineering, University of Maryland, College Park, MD 20742, USA

Abstract

The dawn of the 21st century has seen a growing interest in vehicular networking and its myriad potential applications. The initial view of practitioners and researchers was that radio-equipped vehicles could keep the drivers informed about potential safety risks and increase their awareness of road conditions. The view then expanded to include access to the Internet and associated services. This position paper proposes and promotes a novel and more comprehensive vision namely, that advances in vehicular networks, embedded devices and cloud computing will enable the formation of autonomous clouds of vehicular computing, communication, sensing, power and physical resources. Hence, we coin the term, autonomous vehicular clouds (AVCs). A key feature distinguishing AVCs from conventional cloud computing is that mobile AVC resources can be pooled dynamically to serve authorized users and to enable autonomy in real-time service sharing and management on terrestrial, aerial, or aquatic pathways or theaters of operations. In addition to general-purpose AVCs, we also envision the emergence of specialized AVCs such as mobile analytics laboratories. Furthermore, we envision that the integration of AVCs with ubiquitous smart infrastructures including intelligent transportation systems, smart cities and smart electric power grids will have an enormous societal impact enabling ubiquitous utility cyber-physical services at the right place, right time and with right-sized resources.

Keywords: autonomous systems, cloud computing, cyber-physical systems, resource management, vehicular networks

Received on 12 January 2011

Copyright © 2011 ICST and Olariu *et al.* This article is published under the terms of ICST's copyright sharing agreement, which allows authors to retain full rights in the content of their publications.

doi: 10.4108/icst.trans.mca.2011.e2

1. The vehicular model

The past 20 years have seen an unmistakable trend to make the vehicles on our roads smarter and the driving experience safer and more enjoyable. A typical car or truck today is likely to contain at least some of the following devices: an on-board computer, a GPS device, a radio transceiver, a short-range rear collision radar device, a camera, supplemented in high-end models with a variety of sophisticated sensing devices. Some high-end vehicles already offer the convenience of an event data recorder

(EDR) that collects transactional data from most of the vehicle sub-assemblies. It is not widely known that some GM vehicles as old as model year 1994 were equipped with an EDR-like device able to store retrievable data. In general, the EDRs are intended to be tamper-proof, very much like the well-known black boxes on board commercial and military aircraft. Among other things, the EDRs are designed to optimize the 'up time' of vehicles by sophisticated self-checks and by scheduling vehicles for maintenance on a per-need basis as opposed to a fixed calendar date. In 2006 the National Highway Traffic Safety Administration (NHTSA) declared its intent to standardize the various EDR devices provided, on a voluntary basis, by car and truck manufacturers [1]. As it turns out, the EDRs are already having a societal impact, as they are finding their way into the courtrooms where insurance companies use EDR logs in their litigations.

*Corresponding author. Email: olariu@cs.odu.edu

As technology is moving closer and closer to packing sophisticated resources in individual vehicles, many manufacturers are turning their attention to making the vehicles on our roads more fuel- and energy-efficient than ever. It is sufficient to recall that the past decade has seen the emergence of hybrid vehicles from the automotive engineer's drawing board into production, to the point where today half a dozen car and truck manufacturers offer hybrid vehicles on the American market.

In addition to their sophisticated array of sensing and computation capabilities, the availability of virtually unlimited power supply and growing Internet presence will make our vehicles perfect candidates for housing powerful on-board computers augmented with huge storage devices that, collectively, may act as networked computing centers on wheels.

2. Why vehicular networks?

Wireless technology was available for the past 60 years yet, with few exceptions, it did not find its way into the arena of vehicular communications until very recently. In order to understand the sea change that we have witnessed in the past decade or so, it helps to recall that the US Department of Transportation (US-DOT) estimates that in a single year, congested highways due to various traffic events cost over \$75 billion in lost worker productivity and over 8.4 billion gallons of fuel [2]. The US-DOT also notes that over half of all congestion events are caused by highway incidents rather than by rush-hour traffic in big cities [1]. Further, the NHTSA indicates that congested roads are one of the leading causes of traffic accidents, projecting data extrapolated from January to September 2009 statistics, that, for 2009, an estimated 25576 fatalities are directly attributable to traffic-related incidents [3].

Unfortunately, on most US highways, congestion is a daily event and with rare exceptions, advance notification of imminent congestion is unavailable [4–6]. It is worth mentioning that in the transportation science community several solutions for reducing the effects of congestion were contemplated over the years [7, 8]. One of the proposed solutions involves adding more traffic lanes to our roadways and streets. While at first sight this seems to be a reasonable course of action, a recent study has pointed out that this strategy is futile in the long run as it is likely to lead to more congestion and to increased levels of pollution [8]. On the other hand, it has been argued that given sufficient advance notification, drivers could make educated decisions about taking alternate routes; in turn, this would improve traffic safety by reducing the severity of congestion and, at the same time, save time and fuel [7, 9].

Under present-day technology traffic monitoring and incident reporting systems employ inductive loop detectors (ILDs), video cameras, acoustic tracking systems

and microwave radar sensors [7]. By far the most prevalent among these devices are the ILDs embedded in highways every mile (or half-mile) [10–12]. ILDs measure traffic flow by registering a signal each time a vehicle passes over them. Each ILD (including hardware and controllers) costs around \$8200; in addition, the ILDs are connected by optical fiber that costs \$300000 per mile [9, 13]. Interestingly, official statistics shows that over 50% of the installed ILD base and 30% of the video cameras are defective [7, 12]. Not surprisingly, transportation departments worldwide are looking for less expensive, more reliable and more effective methods for traffic monitoring and incident detection.

To be effective, innovative traffic-event detection systems must enlist the help of the most recent technological advances. This has motivated extending the idea of Mobile Ad-hoc Networks (MANET) to roadway and street communications. The new type of networks, referred to as Vehicular Ad-hoc Networks (VANET) that employ a combination of Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, have been proposed to give drivers advance notification of traffic events. In V2V systems, each vehicle is responsible for inferring the presence of an incident based on reports from other vehicles.

As we just saw, the original impetus for the interest in VANET was provided by the need to inform fellow drivers of actual or imminent road conditions, delays, congestion, hazardous driving conditions and other similar traffic-related concerns. Therefore, most VANET applications focus on traffic status reports, collision avoidance, emergency alerts, cooperative driving and other similar concerns [14–16, 40, 42]. Almost across the board, the community of researchers and practitioners anticipate that advances in VANET, or other emerging vehicle-based computing and communications technology, are poised to have a huge societal impact. Because of this envisioned societal impact, numerous vehicle manufacturers, government agencies and standardization bodies around the world have spawned national and international consortia devoted exclusively to VANET. Examples include Networks-on-Wheels, the Car-2-Car Communication Consortium, the Vehicle Safety Communications Consortium, Honda's Advanced Safety Vehicle Program, among many others. We refer the interested readers to the survey articles [17, 18] where many US and European initiatives and standards are discussed in detail.

The past few years have witnessed a rapid convergence of Intelligent Transportation Systems (ITS) and VANET leading to the emergence of Intelligent Vehicular Networks with the expectation to revolutionize the way we drive by creating a safe, secure and robust ubiquitous computing environment that will eventually pervade our highways and city streets. In support of traffic-related communications, the US Federal Communications Commission (FCC) has allocated 75 MHz of spectrum in the

5.850–5.925 GHz band specially allocated by the FCC for Dedicated Short-Range Communications (DSRC) [5, 19]. It was recently noticed that the DSRC spectrum set aside by the FCC, by far exceeds the needs of traffic-related safety applications. This observation has motivated the emergence of a host of other applications that can take advantage of the allocated spectrum. Not surprisingly, we see more and more third-party providers offering non-safety-related applications ranging from location-specific services, to on-the-road peer-to-peer communications, to Internet access, to on-line gaming and other forms of mobile entertainment. In due time, we will see the emergence of commercial applications targeted at the traveling public and distributed via the excess bandwidth in DSRC. As a pleasant side benefit, the unsightly billboards that flank our highways will disappear and will be replaced by in-vehicle advertising that the driver can filter according to their wants and needs.

3. Looking into the crystal ball

Recently, there has been a good deal of commercial and research interests in utilizing broadband communications and wireless technologies to provide Internet connectivity to the public on the road [15, 20–23]. One such system that has received quite a bit of attention in the recent literature is known as Drive-thru Internet [24, 25] and relies on dedicated road-side access points (AP) on roadways and city streets to enable the driving public to connect to the Internet, at least while they are within the coverage of the infrastructure. An important feature of Drive-thru Internet systems is the multi-access sharing of the same AP's bandwidth by the vehicles that are simultaneously under its coverage [26]. Since, as a rule, the vehicles move very fast, and the coverage limitations are dictated by present-day technology, the amount of data that a passing vehicle can download from any given AP is rather limited, a state of affairs that is promoting vehicular interchanges and the blossoming of vehicular peer-to-peer connections that are key to mitigating the effects of the limited Internet coverage conferred by Drive-thru Internet and other similar schemes [14, 27–29].

We also fully expect third-party infrastructure providers to deploy various forms of road-side infrastructure as well as advanced in-vehicle resources such as embedded powerful computing and storage devices, cognitive radios and cognitive radio networks, and multi-modal programmable sensor nodes. As a result, in the near future, vehicles equipped with computing, communication and sensing capabilities will be organized into ubiquitous and pervasive networks with virtually unlimited Internet access while on the move. This will revolutionize the driving experience making it safer and more enjoyable [14, 20, 21].

The huge array of on-board capabilities is likely to remain underutilized by safety applications alone. The

realization of this fact has already motivated the investigation of offering value-added services including on-line gaming, mobile infotainment, along with various location-specific services. We conjecture that the potential is even far beyond that. Specifically, we propose to 'take vehicular networks to the clouds' so that our prized transportation means can take their natural place integrated with our productivity, comfort, safety and economic prosperity.

4. Cloud computing

The notion cloud computing started from the realization of the fact that instead of investing in infrastructure, businesses may find it useful to rent the infrastructure and sometimes the needed software to run their applications [30, 31]. This powerful idea has been suggested, at least in part, by ubiquitous and relatively low-cost high-speed Internet, virtualization and advances in parallel and distributed computing and distributed databases. One of the key benefits of cloud computing is that it provides scalable access to computing resources and information technology (IT) services [23].

Cloud computing is a paradigm shift adopted by a large number of infrastructure providers who have a large installed infrastructure that often goes underutilized. Hand in hand with cloud computing go 'cloud IT services' where not only computational resources and storage are rented, but specialized services are also provided on demand. In this context, a user may purchase the amount of services they need at the moment. As their IT needs grow and as their services and customer base expand, the users will be in the market for more and more cloud services and more diversified computational and storage resources. In general there are three delivery models for cloud services: Software as a Service (SaaS): customers rent software hosted by the vendor; Platform as a Service (PaaS): customers rent infrastructure and programming tools hosted by the vendor to create their own applications; and Infrastructure as a Service (IaaS): customers rent processing, storage, networking and other fundamental computing resources for all purposes [17, 31–33].

5. Autonomous vehicular clouds

The huge fleet of energy-sufficient vehicles that crisscross our roadways, airways, and waterways, most of them with a permanent Internet presence, featuring substantial on-board computational, storage and sensing capabilities can be thought of as a huge farm of computers on the move. These attributes make vehicles ideal candidates for nodes in a cloud as described above. Indeed, the owner of a vehicle may decide to rent out their in-vehicle capabilities on demand, or a per-instance, or a per-day, per-week or per-month basis, just as owners of large com-

puting facilities find it economically appealing to rent out excess capacity to seek pecuniary advantages.

More significantly, we postulate that vehicles will autonomously self-organize into clouds utilizing their corporate resources on demand and largely in real time in resolving critical problems that may occur unexpectedly. The new vehicular clouds will also contribute to unraveling some technical challenges of the increasingly complex transportation systems with their emergent behavior and uncertainty.

We believe it is only a matter of time before the huge vehicular fleets on our roadways, streets and parking lots will be recognized as an abundant and underutilized computational resource that can be tapped into for the purpose of providing third-party or community services. However, what distinguishes vehicles from standard nodes in a conventional cloud is autonomy and mobility. Indeed, large numbers of vehicles spend substantial amounts of time on the road and may be involved in dynamically changing situations; we argue that in such situations, the vehicles have the potential to cooperatively solve problems that would take a centralized system an inordinate amount of time, rendering the solution useless.

In [34] Olariu *et al.* argued convincingly that it is only a matter of time before the huge vehicular fleets on our roadways, streets and parking lots will be recognized as an abundant and under-utilized computational resource that can be tapped into for the purpose of providing third-party or community services. However, what distinguishes vehicles from standard nodes in a conventional cloud is autonomy and mobility. Indeed, large numbers of vehicles spend substantial amounts of time on the road and may be involved in dynamically changing situations; Olariu *et al.* [34] argued that in such situations, the vehicles have the potential to cooperatively solve problems that would take a centralized system an inordinate amount of time, rendering the solution useless.

In this work we continue to explore various aspects of the vehicular cloud (VC) concept proposed by Olariu *et al.* in their seminal paper [34].

With this in mind, we have coined the term autonomous vehicular cloud (AVC) to refer to: *a group of largely autonomous vehicles whose corporate computing, sensing, communication and physical resources can be coordinated and dynamically allocated to authorized users.*

In our view, the AVC concept is the next natural step in meeting the computational and situational awareness needs not only of the driving public but also of a much larger segment of the population. A primary goal of the AVC is to provide on-demand solutions to events that have occurred but cannot be met reasonably with pre-assigned assets or in a proactive fashion.

It is important to delineate the structural, functional and behavioral characteristics of AVCs. As a step in this direction, in this position paper, we identify autonomous

cooperation among vehicular resources as a distinguishing characteristic of AVCs. Another important characteristic of AVCs is the ability to offer a seamless integration and decentralized management of cyber-physical resources; an AVC can dynamically adapt its managed vehicular resources allocated to applications according to the applications' changing requirements and environmental and systems conditions.

As far as a simple taxonomy goes, AVCs can be public, private or various hybrids thereof. The public AVC will provide (typically short-term) services on the Internet, whereas a private AVC is proprietary and provides (typically long-term) services to a limited set of users and would belong to specific vehicle fleets such as FedEx, UPS, Costco or Wal-Mart. As an example of a hybrid AVC, one may consider inter-AVC cooperation as discussed in Section 6.

It is not too far-fetched to imagine, in the not-so-distant future, a large-scale federation of AVCs will be established *ad hoc* in support of mitigating a large-scale emergency. One of these large-scale emergencies could be a planned evacuation in the face of a potentially deadly hurricane or tsunami that is expected to make landfall in a coastal region [35, 36]. Yet another such emergency would be a natural or man-made disaster apt to destroy the existing infrastructure and to play havoc with cellular communications. In such a scenario, a federation of AVCs could provide a short-term replacement for the infrastructure and also provide a decision-support system.

6. Application scenarios

The main goal of this section is to illustrate the power of the AVC concept. We touch upon several important scenarios illustrating various aspects AVCs that are extremely important and that, under present-day technology are unlikely to see a satisfactory resolution. The outlined scenarios are representative of two application categories: (i) traffic management and (ii) asset management. Other categories include situational awareness, for example tipping and cueing for threat analysis and mitigation, and nomadic social and business spaces, for example mobile marketplace (Sue may want to sell perfumes on her way back from work).

6.1. Traffic management scenarios

The goal of this subsection is to discuss a number of traffic management scenarios where AVCs can be useful. While only a small number of scenarios are presented here, it is easy to extrapolate from them and to contemplate many variations of these scenarios where AVCs are likely to make a difference.

Scenario 1: synchronizing traffic lights after clearing an accident. Consider for example, a city block where a

traffic-related event (e.g. an accident) has occurred and where, as a consequence, a large number of vehicles are co-located. Once the traffic event has been cleared, relying on the existing scheduling of the traffic lights will not help dissipate the huge traffic backlog in an efficient way. We envision a solution to this problem where the vehicles themselves will pool their computational resources together creating the effect of a powerful super-computer that will recommend to a higher authority a way of rescheduling the traffic lights that will serve the purpose of de-congesting the afflicted area as fast as possible. We note that, in general, the solution cannot involve a handful of traffic lights but may require rescheduling the traffic lights in a large geographic area.

As mentioned before, the ability of vehicles to pool their resources, in a dynamic way, in support of the common good will have a huge societal impact alleviating, among others recurring congestion events that plague our cities around the morning or afternoon rush hour. Also, and very importantly, while congestion is a daily phenomenon, proactively solving the problem is infeasible because of the dynamic nature of the problem, and of the huge computational effort its resolution requires. The problem is best solved if and when it occurs in an on-demand fashion dedicating the right amount of resources rather than conservatively pre-allocating abundant resources based on the worst case, which is becoming increasingly infeasible. The key concept that allows the problem to be solved efficiently and economically is the engagement of the necessary resources from the available vehicles participating in the traffic event and their involvement in finding a solution autonomously without waiting for an authority to react to the complicated situation on the ground.

Scenario II: autonomous mitigation of recurring congestion. In face of traffic congestion some drivers often pursue detours and alternate routes that often involve local roads. Making the decision behind the steering wheel is often challenging. The driver does not know whether the congestion is about to ease or is worsening. In addition, when many vehicles decide to execute the same travel plan, local roads become flooded with traffic that exceeds its capacity and sometimes deadlocks take place. Contemporary ITS and traffic advisory schemes are both slow to report traffic problems and usually do not provide any mitigation plan. An AVC-based solution will be the most appropriate and effective choice. Basically, vehicles in the vicinity will be able to query the plan of each other and estimate the impact on local roads. In addition, an accurate assessment of the cause of the congestion and traffic flow can be made by contacting vehicles close to where the bottleneck is. In addition, appropriate safety precautions can be applied to cope with the incident, e.g. poor air quality due to the smoke of a burned vehicle.

Interestingly, this approach can be applied not to drivers on the road but also to those who are about to leave home. Delayed start and telecommuting may be considered as an alternative in order to increase productivity and avoid wasted energy and time.

Scenario III: sharing on-road safety messages. The trend in the car manufacturing industry is to equip new vehicles with major sensing capabilities in order to achieve efficient and safe operation. For example, Honda is already installing cameras on their Civic models in Japan. The cameras track the lines on the road and help the driver stay in lane. A vehicle would thus be a mobile sensor node and an AVC can be envisioned as a huge wireless sensor network with very dynamic membership. It would be beneficial for a vehicle to query the sensors of other vehicle in the vicinity in order to increase the fidelity of its own sensed data, get an assessment of the road conditions and the existence of potential hazard ahead. For example when the tire pressure sensor on a vehicle reports the loss of air, vehicles that are coming behind on the same lane should suspect the existence of nails on the road and may consider changing the lane. The same happens when a vehicle changes lane frequently and significantly exceeds the speed limit; vehicles that come behind, and which cannot see this vehicle, can suspect the presence of aggressive drivers on the road and consider staying away from the lanes and/or keeping a distance from the potentially dangerous driver. The same applies when detecting holes, unmarked speed breakers, black ice, etc. Contemporary VANET design cannot pull together the required solution and foster the level of coordination needed for providing these safety measures.

Scenario IV: dynamic management of HOV lanes. As pointed out by the US-DOT in its 2008 report and guidelines, *‘the primary purpose of an HOV lane is to increase the total number of people moved through a congested corridor by offering two kinds of incentives: a savings in travel time and a reliable and predictable travel time. Because HOV lanes carry vehicles with a higher number of occupants, they may move significantly more people during congested periods, even when the number of vehicles that use the HOV lane is lower than on the adjoining general-purpose lanes. In general, carpoolers, vanpoolers, and transit users are the primary beneficiaries of HOV lanes’* [26].

It is, thus, plainly obvious that the main goal of HOV lanes is to promote traffic fluidity and to prevent traffic slowdowns and congestion. Due to insufficient cyber-physical means (appropriate signs being one key shortcoming) most cities in the USA only use HOV lanes at rush hour. However, AVC could make recommendations for setting up HOV lanes dynamically in the best interest of promoting traffic fluidity and of minimizing travel times for people using the designated HOV lanes. AVCs can enable such a dynamic solution by factoring

data from sensors on board the individual vehicles, e.g. occupancy sensors, and local traffic intensity in order to optimally configure the HOV lanes. Such a solution is infeasible under present-day technology.

The same idea applies to the strategy of marking certain streets and thoroughfares as ‘one-way’ in support of improving the fluidity of traffic. Again, currently such an approach is infeasible mostly because of insufficient signaling means. This, however, should not be a problem in AVC since the drivers will be alerted in real time to road occlusions and other dynamic changes.

6.2. Asset management scenarios

The goal of this subsection is to present a number of plausible scenarios that illustrate the potential use of AVCs for managing existing assets.

Scenario V: mobile experimental and analytics laboratory in support of homeland security. Sensor networks are expected to evolve into long-lived, open, ubiquitous, multi-purpose networked systems. Recently, the authors have proposed ANSWER, an autonomous networked sensor system whose mission is to provide in situ users with real-time, secure information that enhances their situational and location awareness [37]. ANSWER finds immediate applications in homeland security. The architectural model of ANSWER is composed of a large number of sensors and of a set of (mobile) aggregation-and-forwarding nodes, possibly AVC nodes, that organize and manage their sensors (if any) and the sensors in their vicinity. As argued in [37], ANSWER can provide secure, QoS-aware information and analysis services to *in situ* mobile users in support of application-level tasks and queries, while hiding network-level details. We anticipate that an AVC can naturally interface in a symbiotic relationship with ANSWER creating a powerful mission-oriented system.

Scenario VI: augmenting the capabilities of small businesses. Consider a small business employing about 250 people and specializing in offering IT support and services. It is not hard to imagine that, even if we allow for car-pooling, there will be up to 150 vehicles parked in the company’s parking lot. Day in and day out, the computational resources in those vehicles are sitting idle. We envision harvesting the corporate computational and storage resources in the vehicles sitting in the parking lot for the purpose of creating a computer cluster and a huge distributed data storage facility that, with proposer security safeguards in place, will turn out to be an important asset that the company cannot afford to waste.

Scenario VII: efficient tasking of law enforcement officers. Law enforcement officers play a crucial role in keeping the road safe for motorists. Even if a police vehicle is so visible on the road, it serves as a deterrent for aggressive drivers and vehicle safety violators. An AVC can be used as an effective resource-planning tool for

the police squad. Moving vehicles form an AVC and report to the police so that decisions can be made efficiently about deploying troopers in certain spots and/or employing surveillance cameras and aircraft to identify and video tape violators for further assigning fines. That will allow effective usage of officers’ time and enable them to allocate resources for other vital tasks such as criminal investigation and prevention. Implementing this idea through today’s technology is resource-prohibitive and requires major infrastructure investment.

Scenario VIII: dynamic management of parking facilities. Anyone who has attempted to find a convenient parking spot in the downtown area of a big city or close to a university campus where the need for parking by far outstrips the supply would certainly be interested to enlist the help of an automated parking management facility. The problem of managing parking availability is a ubiquitous and a pervasive one, and several solutions were reported recently [38, 39, 41, 43]. However, most of the existing solutions rely on a centralized solution where reports from individual parking garages and parking meters are aggregated at a central (city-wide) location and then disseminated to the public. The difficulty is with the real-time management of parking availability since the information that reaches the public is often stale and outdated. This, in turn, may worsen the situations especially when a large number of drivers are trying to park, say, to attend a down-town event.

We envision that by real-time pooling the information about the availability of parking at various locations inside the city, an AVC consisting of the vehicles that happen to be in a certain neighborhood will be able to maintain real-time information about the availability of parking and direct the drivers to the most promising location where parking is (still) available.

Scenario IX: dynamic asset management in planned evacuations. In cases of predicted disasters, such as hurricanes, massive evacuations are often necessary in order to minimize the impact of the disaster on human lives. However, there are several issues involved in a large-scale evacuation. For example, once an evacuation is underway, finding available resources, such as gasoline, drinking water, medical facilities and shelter, quickly becomes an issue [12]. In its recent report on hurricane evacuations [9], the US-DOT found that emergency evacuation plans often do not even consider availability of such resources. The US-DOT also determined that emergency managers need a method for communicating with evacuees during the evacuation in order to provide updated information. The report suggested that traffic monitoring equipment should be deployed to provide real-time traffic information along evacuation routes.

We now point out natural ways in which AVC can work with the emergency management center overseeing the evacuation in order to provide travel time estimates,

notification of available resources, such as gasoline, food, and shelter, and notification of contra-flow roadways to the evacuees. We anticipate that the vehicles involved in the evacuation will self-organize into one or several inter-operating vehicular clouds that will work hand in hand with the emergency management center. In the course of this interaction, the emergency managers can upload information about open shelters to the central server.

It is important to note that this system would be used to facilitate an evacuation before disaster strikes, so we assume that electricity and network connections are available. In addition to having state authorities send information to the AVCs about evacuations or contra-flow lanes, using role-based communication as described earlier, the AVCs themselves could determine the direction and speed that traffic is flowing. The evacuees entering entrance ramps onto contra-flow roadways (these ramps would likely have been used as exit ramps previously) will be alerted to the direction that traffic is moving. The AVCs could also alert drivers to upcoming entrance ramps that were previously used as exit ramps during non-contra-flow travel. Since the AVC system can easily monitor traffic flow, it could offer recommendations to the emergency center about which roadways are good contra-flow candidates.

Scenario X: AVCs in developing countries. We conjecture that the usefulness and practicality of the AVC concept will become even more apparent in developing countries lacking a sophisticated centralized decision-support infrastructure. We further conjecture that, in such contexts, AVCs will play an essential role in bringing together a huge number of relatively modest computational resources available in the vehicular network into one or several foci of computing and communications that will find and/or recommend solutions to problems arising dynamically and that cannot possibly be resolved with the existing infrastructure. We have seen a similar phenomenon happening with the penetration of cell phones in developing countries where they were adopted rapidly and unhesitatingly by a population that had access to a modest landline telephony system.

7. AVC research issues

The application scenarios discussed above require better V2V and V2I collaboration in order to reach critical and mutually beneficial decisions, effective and unconventional management to cope with the highly dynamic nature of the computing, communication, sensing and physical resources, and well-defined operation structures that enable autonomy and authority in adjusting local settings with the potential of making wide impact.

Currently, our group is initiating several research projects in AVC engineering. We now present some of the

research issues along the three systems engineering dimensions, namely structure, function and behavior (operation and policy).

7.1. Architectural challenges

- Elastic mobile architecture
 - The AVC networking and associated protocol architecture must be developed to accommodate changing application demands and resource availability on the move.
- Resilient AVC architecture in the wild
 - AVC basic structural and composed building blocks must be designed and engineered to withstand structural stresses induced by the inherent instability in the operating environment. Research is needed on architectures enabling vehicle visualization and migration of virtual vehicles.
- Service-oriented network architecture
 - Contemporary layered network architectures, for example the TCP/IP stack, have proven limited in face of evolving applications and technologies. We envision the adoption of service-oriented component-based network architectures with intrinsic monitoring and learning capabilities.

7.2. Functional challenges

- Enabling AVC autonomy
 - Research is needed on developing a trustworthy base, negotiation and strategy formulation methodology (game theory, etc.), efficient communication protocols, data processing and decision-support systems, etc.
- Managing highly dynamic cloud membership
 - There is a critical need to efficiently manage mobility, resource heterogeneity (including sensing, computation and communication), trust and vehicle membership (change in interest, change in location, resource denial and/or failure).
- Cyber-physical control
 - AVCs can be defined by their aggregated cyber and physical resources. Their aggregation, coordination and control are non-trivial research issues.
- Cooperation between AVCs
 - To motivate the need for AVCs to cooperate, imagine that in adjacent areas of a municipality there is a sporting event as well as a rock concert downtown. Both these events are very likely to draw a huge crowd. Now, assuming that due to bad planning by the municipality, the two events end at the same time, creating two distinct zones of massive conges-

tion. In such a scenario, each congestion event will trigger the formation of an ad hoc AVC. These two AVCs will have to coordinate and to solve the congestion problem collectively, since they cannot proceed to selfishly reschedule traffic lights in a way that benefits each of them individually.

7.3. Operational and policy challenges

- Trust and trust assurance
 - In order for the vision outlined above to become reality, the problems of assuring emergent trust and security in AVC communication and information need to be addressed. The establishment of trust relationships between the various players is a key component of trustworthy computation and communication. We argue that since typically most, if not all the vehicles involved, must have met before, the task of establishing proactively, a basic trust relationship between vehicles is possible and may be even desirable (think in terms of vehicles that meet day after day in a parking garage).
 - Also, in order to be effective at cooperative problem solving, an AVC may need to have delegated authority to take local action in lieu of a central authority. Referring to our motivating scenarios, it is clearly useless for the vehicles involved in a traffic jam to produce a workable schedule of the traffic lights that will best promote the rapid dissipation of congestion if they do not have the authority to implement such a schedule. Clearly, the resolution of this problem resides in some form of a trust relationship that needs to be forged between the municipal or county authority and the AVC.
- Contract-driven versus *ad hoc* AVC
 - We anticipate that AVCs will be largely contract-driven, where the owner of the vehicle or fleet consents to renting out some form of excess computational or storage capacity. At the same time, mobility concerns dictate that in addition to the contract-based form of AVC, it should be possible to form an AVC in an *ad hoc* manner as necessitated by dynamically changing situations like those discussed in the scenarios above.
- Effective operational policies
 - In order for the AVCs to operate and inter-operate seamlessly, issues related to authority establishment and management, decision support and control structure, the establishment of accountability metrics, assessment and intervention strategies, rules and regulations, standardization, etc. must all be addressed. Dealing with these will require a broad participation and must involve local, state or even federal decision makers.

- AVC utility computing
 - There is a need for economic models and metrics to determine reasonable pricing and billing for AVC services.

8. Architectures for AVCs

Although our ultimate goal is to produce a *unified* architectural framework for the AVC, the main goal of this section is to review several possible architectures, of increasing complexity that suit various particular manifestations of AVCs.

8.1. A static architecture

In some cases, an AVC may behave just as a conventional cloud. This is, no doubt, the case in static environments as the one we contemplate below. Indeed, consider a small business employing about 250 people and specializing in offering IT support and services. It is not hard to imagine that, even if we allow for car-pooling, there will be up to 150 vehicles parked in the company's parking lot. Day in and day out, the computational resources in those vehicles are sitting idle.

The company may proactively seek the formation of a static AVC by providing appropriate incentives to its employees who will rent the resources of their vehicles to the company on a per-day, per-week or per-month basis. The resulting (more or less) static AVC will harvest the corporate computational and storage resources of the participating vehicles sitting in the parking lot for the purpose of creating a computer cluster and a huge distributed data storage facility that, with proper security safeguards in place, will turn out to be an important asset that the company cannot afford to waste.

In the scenario above, the architecture of the AVC will be almost identical to the architecture of a conventional cloud, with the additional twist of, perhaps, limiting the interaction to week-days.

8.2. Interfacing with a static infrastructure

It is often the case that an AVC is created and evolves in an area instrumented by the deployment of some form of a static infrastructure supportive of the management of various activities. In an urban setting, such an infrastructure includes traffic lights, cameras as well as the utility or street lighting poles. On our roadways, the static infrastructure includes ILDs, the road-side units and other ITS hardware deployed in support of traffic monitoring and management.

We note here that in a not-so-distant future, a pre-deployed set of tiny sensors, even if not organized in a permanent sensor network, may play the role of the static infrastructure that the AVC may find it beneficial to inter-

act with. In fact, this view is consistent with our NSF-funded ANSWER project [6] where the place of the PSAR is taken by the AVC that is constantly interacting with the static infrastructure.

It is self-evident that the AVC benefits from the interaction with the existing static infrastructure. Consider for example, a city block where a minor traffic-related event has occurred and where, as a consequence, a number of vehicles are co-located. Once the traffic event has been cleared, relying on the existing scheduling of the traffic lights will not help dissipate the traffic backlog in an efficient way. We envision a solution to this problem where the vehicles themselves will pool their computational resources together creating the effect of a powerful super-computer that will recommend to a higher authority a way of rescheduling the traffic lights that will serve the purpose of de-congesting the afflicted area as fast as possible.

It is worth noting that in this particular instance, the scope of the traffic lights to reschedule is relatively modest and does not require the federation of several AVCs.

8.3. A simple dynamic architecture

Consider for example, a city block where a minor traffic-related event has occurred and where, as a consequence, a number of vehicles are co-located. Once the traffic event has been cleared, relying on the existing scheduling of the traffic lights will not help dissipate the traffic backlog in an efficient way. We envision a solution to this problem where the vehicles themselves will pool their computational resources together creating the effect of a powerful super-computer that will recommend to a higher authority a way of rescheduling the traffic lights that will serve the purpose of de-congesting the afflicted area as fast as possible.

It is worth noting that in this particular instance, the scope of the traffic lights to reschedule is relatively modest and does not require the federation of several AVCs. The architecture that will support the formation of this AVC will involve the following elements:

- A broker elected spontaneously among the vehicles that will attempt to spontaneously form an AVC.
- The broker will then secure a preliminary authorization from a higher (city) forum for the formation of an AVC. If several brokers attempt to secure such an authorization simultaneously, one will succeed and the others will possibly form a team that will coordinate the formation of the AVC. In the sequel we assume that there is a unique broker.
- The broker will inform the vehicles in the area of the received authorization and will invite participation in the AVC.
- The cars will/or will not respond to the invitation on a purely autonomous basis.

- The broker decides if a sufficient number of vehicles have volunteered and will then announce the formation of the AVC.
- The AVC will pool their computational resources to form a powerful super-computer that, using a digital map of the area, will produce a proposal schedule to the higher (city) forum for approval and implementation.
- Once the proposal has been accepted and implemented, the AVC is dissolved.

While the scenario above and the resulting architecture are slightly more complex than those of a conventional cloud, we note that, in general, the solution cannot involve only a handful of traffic lights but will require rescheduling the traffic lights in a large area. This motivates the collaboration of several AVCs.

9. Concluding remarks

The main goal of this paper was to put forth a novel concept whose time has come: namely that of AVCs. AVCs are emerging from the convergence of advances in mobility, powerful embedded in-vehicle resources, ubiquitous sensing and cloud computing. When fully realized and deployed, AVCs would yield significant enhancements in safety, security and economic vitality of our modern society. They would enable non-conventional applications that go far beyond what people are expecting from today's VANET and ITS. Not surprisingly, the practical realization of our vision and the production of AVC standards will require tackling numerous novel technical challenges, whose resolution will certainly involve adopting a clean-slate approach. New research and development programs are needed to build AVC reference models, architectures and protocols, to address emergent trust and trust assurance issues, to provide AVC-driven cyber-physical resource coupling and coordination, to realize broader benefits through AVC federations, among many others. In addition to terrestrial vehicles the AVC concept also applies to aerial and aquatic vehicles or any hybrid combination thereof.

Acknowledgements. This work was supported, in part, by the following NSF Grants CNS 0721523, CNS 0721563, CNS 0721586 and CNS 0721644. A preliminary version of this paper has appeared in the Proceedings of Second International ICST Conference on Ad Hoc Networks, 18–20 August 2010, Victoria, British Columbia, Canada.

References

- [1] NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION (2006) *Traffic Safety Facts*, <http://www-nrd.nhtsa.dot.gov>.

- [2] US DEPARTMENT OF TRANSPORTATION (2008) *National Transportation Statistics*.
- [3] NATIONAL HIGHWAY TRAFFIC SAFETY ADMINISTRATION (2010, March) *Traffic Safety Facts—Preliminary 2009 Report*, <http://www-nrd.nhtsa.dot.gov/Pubs/811255.pdf>.
- [4] ELBATT, T., GOEL, S., HOLLAND, G., KRISHNAN, H. and PARIKHAN, J. (2006) Cooperative collision warning using dedicated short range wireless communications. In *Proceedings of ACM VANET* (ACM Press).
- [5] SENGUPTA, R., REZAEI, S., SHLAVODER, S.-E., CODY, D., DICKEY, S. and KRISHNAN, H. (2006, May) Cooperative collision warning systems: concept definition and experimental implementation. *California PATH Technical Report UCB-ITS-PRR-2006-6*.
- [6] VIRGINIA DEPARTMENT OF TRANSPORTATION (2006) *Commonwealth of Virginia's Strategic Highway Safety Plan, 2006–2010*, http://virginiadot.org/info/resources/Strat_Hwy_Safety_Plan_FREPT.pdf.
- [7] FONTAINE, M. (2009) Traffic monitoring. In OLARIU, S. and WEIGLE, M.C. [eds.] *Vehicular Networks: From Theory to Practice* (Boca Raton, FL: Taylor and Francis), 1.1–1.28.
- [8] SIGHTLINE (2009), http://www.sightline.org/research/energy/res_pubs/analysis-ghg-roads.
- [9] SREEDEVI, I. and BLACK, J. (2001, February) *Loop Detectors*, California Center for Innovative Transportation, http://www.calccit.org/itsdecision/serv_and_tech/Traffic_Surveillance/road-based/in-road/loop_report.html.
- [10] ROESS, R.-P., PRASSAS, E.-S. and MCSHANE, W.-R. (2004) *Traffic Engineering* (Erewhon, NC: Pearson Prentice Hall), 3rd ed.
- [11] US DEPARTMENT OF TRANSPORTATION (2008, August) *Federal-Aid Highway Program Guidance on High Occupancy Vehicle (HOV) Lanes*, <http://ops.fhwa.dot.gov/freewaygmt/hovguidance/index.htm>.
- [12] VARAIYA, P., LU, X.-Y. and HOROWITZ, R. (2006, October) *Deliver a Set of Tools for Resolving Bad Inductive Loops and Correcting Bad Data*, http://path.berkeley.edu/~xylu/TO6327/TO6327_SEMP.pdf.
- [13] UNIVERSITY OF VIRGINIA CENTER FOR TRANSPORTATION STUDIES (2005, November) *Virginia Transportation Research Council, Probe-Based Traffic Monitoring State-of-the-Practice Report*.
- [14] LEE, U., CHEUNG, R. and GERLA, M. (2009) Emerging vehicular applications. In OLARIU, S. and WEIGLE, M.C. [eds.] *Vehicular Networks: From Theory to Practice* (Boca Raton, FL: Taylor and Francis), 6.1–6.30.
- [15] LOCHERT, C., SCHEUERMANN, B., WEWETZER, C., LUEBKE, A. and MAUVE, M. (2008, September) Data aggregation and roadside unit placement for a VANET traffic information system. In *Proceedings of ACM VANET* (ACM Press).
- [16] YANG, Y. and BAGRODIA, R. (2009, September) Evaluation of VANET-based advanced intelligent transportation systems. In *Proceedings of ACM VANET*, Beijing, China (ACM).
- [17] LE, L., FESTAG, A., BALDESARI, R. and ZHANG, W. (2009) CAR-2-X communications in Europe. In OLARIU, S. and WEIGLE, M.C. [eds.] *Vehicular Networks: From Theory to Practice* (Boca Raton, FL: Taylor and Francis), 4.1–8.32.
- [18] MISENER, J.A., DICKEY, S., VANDERWERF, J. and SENGUPTA, R. (2009) Vehicle-infrastructure cooperation. In OLARIU, S. and WEIGLE, M.C. [eds.] *Vehicular Networks: From Theory to Practice* (Boca Raton, FL: Taylor and Francis), 3.1–8.35.
- [19] US FEDERAL COMMUNICATIONS COMMISSION (FCC) (2003, September) *Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems—5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications* (Washington, DC).
- [20] ANDA, J., LEBRUN, J., GHOSAL, D., CHUAH, C.-N. and ZHANG, M. (2005, May) VGrid: vehicular ad hoc networking and computing grid for intelligent traffic control. In *Proceedings of IEEE Vehicular Technology Conference—Spring* (IEEE Press), 2905–2909.
- [21] CZAJKOWSKI, K., FITZGERALD, S., FOSTER, I. and KESSELMAN, C. (2001) Grid information services for distributed resource sharing. In *Proceedings of 10th IEEE International Symposium on High Performance Distributed Computing* (New York: IEEE), 181–184.
- [22] ERIKSSON, J., BALAKRISHNAN, H. and MADDEN, S. (2008, September) Cabernet: vehicular content delivery using WiFi. In *Proceedings of 14th ACM International Conference on Mobile Computing and Networking* (San Francisco, CA: MobiCom'2008).
- [23] LOCHERT, C., SCHEUERMANN, B., CALISKAN, M. and MAUVE, M. (2007, January) The feasibility of information dissemination in vehicular ad-hoc networks. In *Proceedings of 4th Annual Conference on Wireless On-demand Network Systems and Services (WONS '07)* (ACM Press), 92–99.
- [24] OTT, J. and KUTSCHER, D. (2004) Drive-thru Internet: IEEE 802.11b for automobile users. In *Proceedings of IEEE INFOCOM* (IEEE Press).
- [25] OTT, J. and KUTSCHER, D. (2005) A disconnection-tolerant transport for drive-thru Internet environments. In *Proceedings of IEEE INFOCOM* (IEEE Press).
- [26] TAN, W.-L. LAU, W.C. and YUE, O.-C. (2009, September) Modeling resource sharing for a road-side access point supporting drive-thru Internet. In *Proceedings of ACM VANET*, Beijing, China (ACM).
- [27] ABUELELA, M. and OLARIU, S. (2009) Content delivery in zero-infrastructure VANET. In OLARIU, S. and WEIGLE, M.C. [eds.] *Vehicular Networks: From Theory to Practice* (Boca Raton, FL: Taylor and Francis), 8.1–8.15.
- [28] RYBICKI, J., SCHEUERMANN, B., KIESS, W., LOCHERT, C., FALLAHI, P. and MAUVE, M. (2007, September) Challenge: peers on wheels—a road to new traffic information systems. In *Proceedings of 13th Annual ACM International Conference on Mobile Computing and Networking* (Montreal: ACM).
- [29] RYBICKI, J., SCHEUERMANN, B., KOEGEL, M. and MAUVE, M. (2009, September) PeerTIS—a peer-to-peer traffic information system. In *Proceedings of ACM VANET* (Beijing, China: ACM).
- [30] FOLEY, J. (2008, August) Private clouds take shape. In *InformationWeek*.
- [31] HODSON, S. (2008, May) *What is Cloud Computing?* <http://www.winextra.com/2008/05/02/what-is-cloud-computing>.

- [32] HOOVER, J.N. and MARTIN, R. (2008, June) Demystifying the cloud. In *InformationWeek Research & Reports*, 30–37.
- [33] KIM, W. (2009, January–February) Cloud computing: today and tomorrow. *J. Object Technol.* **8**(1): 65–72, http://www.jot.fm/issues/issue_2009_01/column4/.
- [34] OLARIU, S., KHALIL, I. and ABUELELA, M. (2011) Taking VANET to the Clouds. *Pervasive Comput. Commun.* **7**(1): 7–21.
- [35] FELDSTEIN, D. and STILES, M. (2005, September). Too many people and no way out. In *The Houston Chronicle*.
- [36] US DEPARTMENT OF TRANSPORTATION (2006, June) *Catastrophic Hurricane Evacuation Plan Evaluation: A Report to Congress*, <http://www.fhwa.dot.gov/reports/hurricane-evacuation/>.
- [37] ELTOWEISSY, M., OLARIU, S. and YOUNIS, M. (2007). ANSWER: autonomous networked sensor system. *J. Parallel Distrib. Comput.* **67**(1): 111–124.
- [38] XU, Q., MAK, T., KO, J. and SENGUPTA, R. (2004, October) Vehicle-to-vehicle safety messaging in DSRC. In *Proceedings of 1st ACM International Workshop on Vehicular Ad Hoc Networks* (ACM Press).
- [39] YAN, G. OLARIU, S. and WEIGLE, M.C. (2008) Providing VANET security through active position detection. *Comput. Commun.* **31**(12): 2883–2897.
- [40] AIJAZ, A., BOCHOW, B., DÖTZER, F., FESTAG, A., GERLACH, M., KROH, R. and LEINMÜLLER, T. (2006, March) Attacks on inter-vehicle communication systems: an analysis. In *Proceedings of International Workshop on Intelligent Transportation* (Hamburg, Germany: WIT'2006).
- [41] TROPOS NETWORKS (2010) <http://www.tropos.com/pdf/solutions/Parking-Final.pdf>.
- [42] YAN, G. OLARIU, S. and WEIGLE, M.C. (2009) Providing location security in vehicular ad-hoc networks. *IEEE Wireless Commun.* **16**(6): 48–55.
- [43] Automated Parking Management System at New Hyderabad International Airport, http://www.inrnews.com/realestateproperty/india/hyderabad/automated_parking_management_s.html.