

# Parked Cars are Excellent Roadside Units

Andre B. Reis, Susana Sargento, and Ozan K. Tonguz

**Abstract**—A comprehensive implementation of the envisioned traffic safety and efficiency applications of the IEEE 802.11p and Wireless Access for Vehicular Environments (WAVE) standards assume the premise of the use of Dedicated Short-Range Communications (DSRC) technology both as on-board units and as roadside units (RSUs). The high cost associated with RSUs, however, has so far prevented massive deployment of RSUs. Finding alternative solutions to this longstanding problem is therefore very important. In this paper, we propose a self-organizing network approach to using parked cars in urban areas as RSUs. This self-organizing network approach enables parked cars to create coverage maps based on received signal strength and make important decisions, such as if and when a parked car should serve as an RSU. Our results show the feasibility and cost-effectiveness of the proposed approach, which is able to provide excellent coverage using only a small fraction of the cars parked in a city.

**Index Terms**—Intelligent vehicles, roadside units, self-organizing networks, vehicular ad hoc networks.

## I. INTRODUCTION

VEHICULAR networks require a minimum number of cars to be well connected and functional, which can fail to happen due to either low numbers of vehicles on the road or insufficient radio-equipped vehicles. Studies show that in areas of low vehicle density, important safety broadcasts can take more than 100 seconds to reach all nearby cars [1]. On the other hand, dense networks with too many vehicles can be overwhelmed with traffic and signaling [2], requiring careful coordination between the nodes to ensure proper operation.

One way to overcome both these problems is to supplement vehicle-to-vehicle (V2V) communications with vehicle-to-infrastructure (V2I) support systems, that is, by deploying infrastructure nodes known as Road-Side Units (RSUs) along the road, in addition to the Dedicated Short Range Communications (DSRC) / IEEE 802.11p units within the

vehicles [3]. These units can supplement a sparse network in a low-density scenario, and help coordinate and move data in dense scenarios.

The U.S. Department of Transportation (DoT) anticipated a nationwide deployment of supporting infrastructure of RSUs for vehicular networks to have happened by 2008 [4] – however, this forecast did not come to fruition due to difficulties in justifying the benefits of RSUs, lack of cooperation between the public and private sectors, but most importantly, a lack of funding for infrastructure whose widespread deployment was estimated to cost billions of dollars. A 2012 industry survey by Michigan’s DoT and the Center for Automotive Research reiterated that “one of the biggest challenges respondents see to the broad adoption of connected vehicle technology is funding for roadside infrastructure.” [5] and, in 2014, a nationwide study sponsored by the U.S. DoT reported average costs of \$17,680 per deployed Roadside Unit [6], for both hardware and installation. These prohibitive costs explain why one is unlikely to see substantial deployments of RSUs in the near future, despite their importance to vehicular networks.

One way to avoid the prohibitive expense of RSU deployment is to use the vehicles themselves as RSUs [7], and in urban areas the idea of leveraging parked cars to assist a vehicular network has been suggested [8], [9], although only for specific use cases (e.g., overcoming corner obstructions, or increasing node density in sparse areas). The objective of this paper is to propose a credible, low-cost alternative to a Roadside Unit deployment that can operate both independently and in conjunction with existing RSUs.

To this end, we introduce a *self-organizing network approach* that allows parked cars to function as RSUs forming a vehicular support network. We tackle the specific problems of selecting which parked cars should become part of the network, how to measure each car’s utility to be able to make informed decisions, what algorithmic steps should vehicles follow when they park, and how to deal with possible disruptions in connectivity when parked cars leave.

Simulation data from a unique platform with real-life data validate our proposed approach and its methods. The results show that an on-line, greedy decision process based on self-generated coverage maps and limited communication between vehicles can achieve an average coverage that is 93% to 97% as good as an optimum solution. The proposed algorithms can also optimize the number of parked cars that are designated as RSUs, selecting on average only 12% more cars than an optimum process. A second set of data shows, additionally, that parked cars bring tangible benefits in initial deployment stages, where insufficient numbers of DSRC-enabled vehicles cause the network to become sparse. In these scenarios, using small numbers of parked cars to act as RSUs can reduce the

Manuscript received May 30, 2016; revised November 18, 2016; accepted January 3, 2017. This work was supported in part by the Portuguese Foundation for Science and Technology through Pest-OE/EEI/LA008/13 under Grant UID/EEA/50008/2013, in part by the IT Internal Project SmartCityMules, in part by the CMU-Portugal Program through S2MovingCity: Sensing and Serving a Moving City under Grant CMUP-ERI/TIC/0010/2014, and in part by the Project MobiWise: from mobile sensing to mobility advising under Grant P2020 SAICTPAC/0011/2015. The Associate Editor for this paper was P. Kachroo.

A. B. Reis is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213-3890 USA, and also with the Instituto de Telecomunicacoes, Universidade de Aveiro, 3810-193 Aveiro, Portugal (e-mail: abreis@cmu.edu).

S. Sargento is with the Instituto de Telecomunicacoes, Universidade de Aveiro, 3810-193 Aveiro, Portugal (e-mail: susana@ua.pt).

O. K. Tonguz is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213-3890 USA (e-mail: tonguz@ece.cmu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TITS.2017.2655498

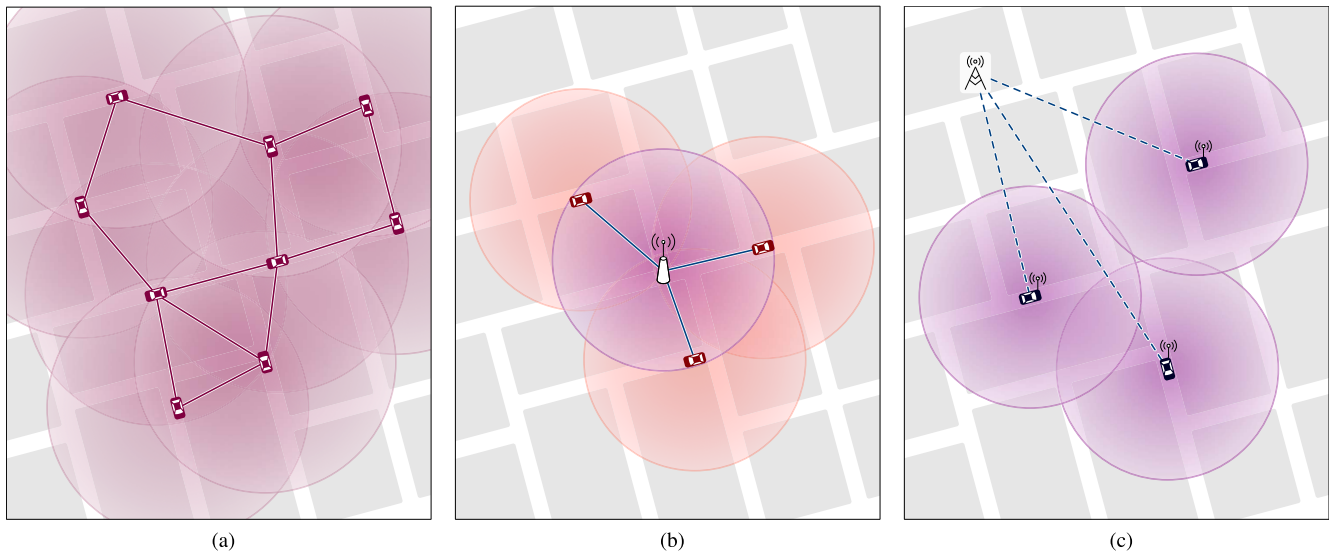


Fig. 1. Modes of operation for parked cars acting as RSUs. (a) Parked cars form a mesh network with point-to-point links to other parked cars in range, using 802.11p. (b) Parked cars extend the range of a fixed 802.11p RSU, acting as relays to it. (c) Parked cars with access to an uplink establish themselves as standalone RSUs (depicted: using a cellular network as an uplink).

time for emergency messages to be broadcast by 40-50%. Through these key results, we verify that parked cars can indeed serve both as RSUs and as an extension to vehicular infrastructure deployments, and can self-organize to do so with no significant overhead to the existing networks of moving cars. Combined, the two analyses in this work encompass both low-, medium- and high-density scenarios, and their corresponding challenges.

This work's main contributions can be summarized as follows.

- A low-cost solution to the problem of deploying RSUs for providing support to urban vehicular networks is proposed, by leveraging the parked cars in cities.
- A self-organizing network approach to selecting RSUs from a large pool of parked cars is formulated. We introduce a novel way for vehicles to assess their value to the network by listening to beacons transmitted by other cars on the road, and piecing together a map of their own coverage.
- A simple on-line algorithm is designed that maximises the coverage of the support network of parked cars, while minimizing the number of cars that are required to be enabled. Using the coverage maps, this algorithm requires only brief 1-hop exchanges between RSUs.
- A detailed study is provided for the benefits of the proposed approach at the initial stages of implementation, where the small market penetration rate of DSRC-equipped vehicles will imply that only a few parked cars in a given area will be able to become RSUs.

The remainder of this paper is organized as follows. Section II gives an overview of the roles parked cars can take in an urban network. Section III describes the self-organizing approach that is at the core of this work. Section IV begins with a description of our simulation platform, which is then followed by a study of the advantages this approach can bring to urban areas in initial stages of deployment, and afterwards with an analysis of the self-organization process

and how it fares against optimal decisions. Section V provides a discussion on the impact of our approach on a vehicle's battery life. Related work is presented in Section VII, and finally, concluding remarks are given in Section VIII.

## II. LEVERAGING PARKED CARS AS ROADSIDE UNITS

While the first suggestions for the use of cars as Roadside Units aimed at recruiting the moving vehicles for this purpose [7], in urban areas parked cars present themselves as equally compelling candidates. For one, parked cars share many of the benefits of fixed RSUs, of which one of the most significant is that *they do not move*.

As cars move, buildings and other obstructions enter and leave the line of sight between them, causing detrimental fading and shadowing effects on the communication channel. Eventually, the nodes move too far apart from each other, and the channel disappears. A parked car in an urban area has a fixed, known position for extended periods of time and consequently, a more stable communication channel with nearby cars and RSUs is possible. Their fixed location is also beneficial to applications that rely on *geocasting* (the broadcasting of messages to specific geographic areas), making it simpler to route such messages to their intended location. Furthermore, in areas of low vehicle density, the activation of parked cars can supplement the network by increasing the number of available nodes and, consequently, strengthening connectivity.

We introduce three methods of operation for parked cars to become a part of the vehicular network – they are depicted in Figure 1. When no fixed RSUs exist in the urban area, parked cars can form a **mesh network** with each other (see Figure 1a), each parked car connecting to its neighboring parked cars. This mesh network can support the moving vehicles by offering to relay and broadcast messages, with the benefit of having a stable node structure and stable links between nodes (moving cars do not enjoy this benefit).

When a limited deployment of fixed RSUs is already present, parked cars in the vicinity of an RSU can act as **relays** to it (see Figure 1b), extending the reach of the fixed unit. In this mode, the value of a deployed RSU can be increased by enlarging its area of coverage. Parked cars acting as relays advertise their ability to forward messages to an RSU, so moving vehicles use them as such. Around these relays, other parked cars can continue to form the previously-described mesh network, further amplifying the coverage of the original RSU via relays of relays.

Lastly, a parked car with the access to a backbone uplink can leverage that link to establish itself as a **standalone RSU** (see Figure 1c). In this mode, a parked car can communicate with other existing RSUs and uplink-enabled cars via the Internet, allowing it to send messages to distant locations in adequate time. Increasing numbers of vehicles now come equipped with Internet connectivity (using 3G or LTE/LTE-A cellular technologies), and there is the possibility for the car's electronics to be allowed the use of this link for selected purposes (e.g., safety messages).

In urban areas that have seen limited deployments of fixed infrastructure, parked cars can be used to fill in the gaps, providing coverage in areas not serviced by existing RSUs, directing messages to the RSUs as necessary. Parked cars can, therefore, serve both as the main solution to providing support to a mobile vehicular network and as a complimentary solution to the existing infrastructure.

The remainder of this work will focus on the crucial issue of deciding which parked cars should become a part of this vehicular support network.

### III. PROPOSED SOLUTION: A SELF-ORGANIZING NETWORK APPROACH TO CREATING RSUs

This section describes a self-organizing approach for constructing a vehicular support network from parked cars. At the core of this self-organizing approach lies a single decision: when a vehicle parks, should it become an RSU, or should it enter a sleep mode? The decision depends on what the support network aims to accomplish. When a parked car takes a Roadside Unit role, this has non-negligible costs both to the battery of that vehicle, as it must then power the DSRC electronics, and to the vehicular network, as active nodes will broadcast their presence and answer requests from other cars, causing overhead. So, for dense areas, the primary goal should be to maximize the reach of this support network, while minimizing the number of active RSUs.

With this objective in mind, we introduce an *on-line, greedy* algorithm that allows each car to decide whether to become an RSU or not. The process is necessarily on-line since the network is in constant evolution, as cars park and leave; our approach is also greedy, in the sense that we attempt to make a locally optimum decision whenever a vehicle is parked, aiming to approach a global optimum solution.

With full knowledge of each parked car at every location, one can evaluate all possible combinations of turning cars on and off, and reach a global optimum solution. Evidently this is not feasible, as not only it would quickly amount to millions

of computations, it would also need to repeat itself whenever a car parked or left. With this work, we aim for an approach that requires *minimal communication between nodes*.

#### A. Self-Observed Coverage Maps

To be able to optimize the number of parked cars that become active and take the role of RSUs, we require a new metric that can represent each vehicle's value to the network. The primary goal of this RSU network is to be able to reach as many locations in the city as possible – therefore, we are interested in knowing the *signal coverage* of each parked car, i.e., which areas it can reach, and how well it can do so.

Cars, however, can park in numerous distinct locations. Using propagation models, one could estimate the coverage map of each particular car, but doing so would also require the local roadmap, as well as knowing the shapes and structures of nearby buildings (to determine obstructions). Such a process would still fail to account for undocumented obstacles such as trees, trucks, billboards, etc.

With this work, we introduce a system whereby parked cars listen to beacons being broadcast by other nearby cars on the road and use those beacons to build maps of their coverage. By doing so, parked cars learn about which areas in their vicinity they can send and receive messages from. These beacons, which include position, speed, and bearing, are standardized and known as Cooperative Awareness Messages (CAMs), and are broadcast at rates no lower than 1 Hz [10]. When the signal strength of incoming beacons is made available from the lower protocol layers, parked cars can take advantage of this valuable information to track how strong coverage is at each location, building a distribution of the signal power of incoming messages and using it to estimate the quality of coverage.

In Figure 2, we illustrate the process of learning a coverage map, while also demonstrating coverage tracking in the shape of cells. In this work, we divide the urban area into a logical 2D cell map that is common to all cars. This approach stems from the need to have coverage maps that are easily representable, storable and shareable, and that can track coverage in irregular street maps with adequate precision. To ensure that all vehicles share the same grid division, we opt to align cell boundaries to GPS coordinates, a universal reference for all GPS-enabled cars.

Under this cellular division, a self-observed coverage map can be represented as a matrix containing signal strength values, with 5 meaning *excellent signal* and 0 meaning *no coverage*. In this work, we align cell boundaries to the nearest GPS second (a standard WGS84 coordinate point multiplied by 3600, rounded down),<sup>1</sup> which results in a cell that is approximately 30 meter wide. A coverage map for a vehicle with a maximum radio range of 150 meter will be  $11 \times 11$  cells large (with the vehicle positioned at the center cell); 3 bits per cell can represent the 6 degrees of coverage, which makes a map  $\approx 46$  bytes in size. To center the map, the latitude and longitude of the vehicle (2 bytes each) is sent along with

<sup>1</sup>E.g., the decimal WGS84 coordinate 43.91750000000000 can be expressed as 158 103 seconds, or N43°55'03.00" (hours (°), minutes ('), seconds (")).

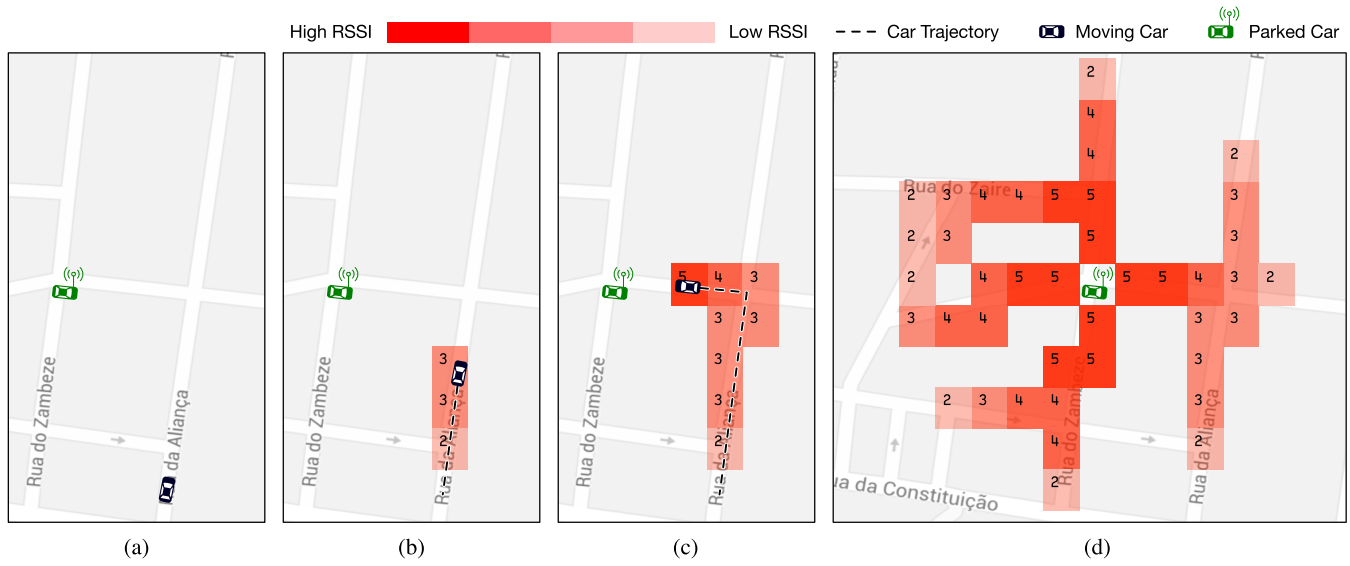


Fig. 2. In this example, Figures 2a, 2b, and 2c show how a parked car (the green car in the figure) observes parts of its coverage map by listening to beacons from a nearby moving vehicle. As a vehicle moves nearer, our parked car overhears its beacons and corresponding signal strength, tagging the cells in its local map. Figure 2d shows the complete, self-observed coverage map.

the data matrix. In this matrix form, coverage maps can be shared between vehicles through the WAVE Short Message Protocol (WSMP), a part of the WAVE standard [11].

### B. Procedures for Newly Parked Cars

Each newly parked car must follow a logical set of actions to allow it to decide whether it should become an RSU. A newly parked car begins by listening for beacons being broadcast by other vehicles in its vicinity, as seen in Figure 3a. During this process, the vehicle builds its coverage map, as described in the last section. We provide reference values for the duration of this step later on, in Section IV-C.

Once the step of building a map of coverage is complete, the vehicle requests the coverage maps of neighboring active RSUs, using WAVE Short Messages (WSMs). The decision process only requires maps from fixed RSUs and active parked cars. The outcome of this algorithm ultimately decides whether the parked car should become an RSU or switch to a power-saving (sleep) mode. The decision made by each car depends on its observed coverage and the maps it receives from other RSUs in its 1-hop neighborhood. We opt to keep the decision down to each vehicle, while relying only on information from the car's 1-hop neighbor nodes, to minimize the associated network overhead. We describe this decision process in more detail in the next section.

The process that newly parked cars must follow is straightforward. However, it is possible for a car's coverage map to change after the decision step, which can occur if no vehicles passed through a nearby road during the listening step. Coverage can also worsen if an obstruction appears near a parked car, e.g. if a large truck decides to park next to it. To account for this possibility, when a car detects that its coverage map has changed significantly, it reevaluates its decision, by following the steps in Figure 3b. A delta value,  $\Delta_{cov}$ ,

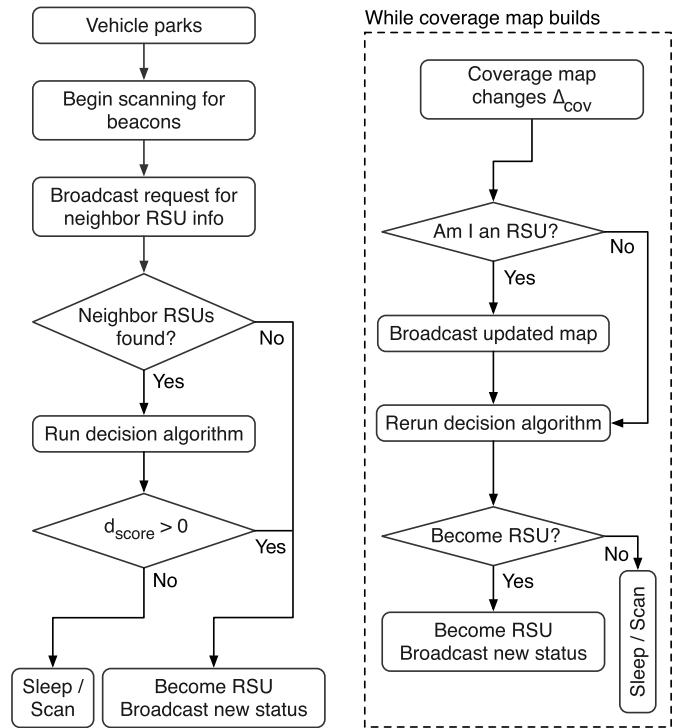


Fig. 3. (a) Flowchart of actions taken by a newly parked car in order to decide whether to become an RSU. (b) A parked car may see a substantial update to its coverage map, triggering a new decision process.

tracks how many cells have changed since the last decision, and crossing a pre-specified threshold triggers a new decision step.

### C. Making a Decision

At the core of the decision process lies the *decision score* ( $d_{score}$ ), a measure of each car's added value to

TABLE I  
NOTATION REFERENCE

Definition	Notation
Matrix indices (row, col.)	$i, j$
Geographic indices (lat., lon.)	$x, y$
Self-observed coverage map	$[SCM] = (scm_{ij})$
Center coordinates of an SCM	$x_{center}, y_{center}$
Deciding vehicle's own SCM	$SCM_0$
Collection of neighbor SCMs	$\mathcal{N} = \{SCM_1, SCM_2, \dots, SCM_n\}$
Local map of coverage	$[LMC] = (lmc_{ij})$
Local map of saturation	$[LMS] = (lms_{ij})$

the network. As mentioned in the previous section, each car must make a decision when it parks, once it has inferred its coverage map and received the maps from nearby RSUs (which can be other active parked cars, as well). With this data in hand, the now-parked vehicle estimates what will happen to the network should it decide to become an RSU at the location where it parked. In essence, it measures its net benefit against its perceived detriment.

The coverage matrixes from neighboring RSUs are combined to obtain two maps of the local area: a first map with the best signal coverage available at every cell, and a second map with the number of RSUs that serve each cell. When the decision process is triggered, the new parked car overlays its coverage map with these local area maps. For each cell in this map, the parked car will see one of the following occur:

- **Establishing new coverage.** If the new parked car is covering a cell that no other RSU in its vicinity can reach, then becoming an RSU will provide new coverage to the network. This action is a major benefit.
- **Improving existing coverage.** If the newly parked car can cover an existing cell at a signal level better than any other nearby RSU, then becoming an RSU will improve the strength of the existing coverage provided by the RSUs. This action also benefits the network.
- **Adding redundant coverage.** A parked car that adds coverage to an already-covered cell, but does so at a signal strength level equal to or lower than what is already provided by other RSUs, is adding unnecessary weight to the network. We refer to this as *saturating* the network.

Each of these three effects is then quantified, and a weighted sum of each results in the decision score. In this work,  $d_{score}$  is given by

$$d_{score} = \kappa \cdot d_{new} + \lambda \cdot d_{boost} - \mu \cdot d_{sat}, \quad (1)$$

where  $d_{new}$ ,  $d_{boost}$  and  $d_{sat}$  are metrics for *new coverage*, *improved coverage*, and *excess coverage*, respectively. The coefficients  $\kappa$ ,  $\lambda$  and  $\mu$  weight each component of the decision score, adjusting the balance between improvement and degradation to the network. Later on, in Section IV-C, we will provide reference values for these coefficients.

We now formalize each of these metrics and how they are calculated. All notation is summarized in Table I. A newly parked vehicle on its decision step has its own observed coverage map,  $[SCM_0]$ , whose elements  $scm_{ij}$  represent geographic

#### Algorithm 1 BuildLocalMaps

---

**Data:**  $\mathcal{N} = \{SCM_1, SCM_2, \dots, SCM_n\}$   
**Result:** Local maps LMC, LMS

```

1  $\triangleright$   $lmc_{xy}$  and  $lms_{xy}$  are initialized to 0
2 foreach  $SCM_n \in \mathcal{N}$  do
3   foreach  $scm_{n[xy]} \in SCM_n$  do
4     if  $scm_{n[xy]} > lmc_{xy}$  then  $lmc_{xy} \leftarrow scm_{n[xy]}$  if
        $scm_{n[xy]} > 0$  then  $lms_{xy} \leftarrow lms_{xy} + 1$ 
5   end
6 end
```

---

#### Algorithm 2 ScoreMetrics

---

**Data:**  $SCM_0$ , LMC, LMS  
**Result:** A tuple of score metrics ( $d_{new}$ ,  $d_{boost}$ ,  $d_{sat}$ )

```

1  $d_{new} \leftarrow d_{boost} \leftarrow d_{sat} \leftarrow 0$ 
2 foreach  $scm_{xy} \in SCM_0$  do
3   if  $scm_{xy} > 0$  then
4     if  $lmc_{xy} = 0$  then
5        $d_{new} \leftarrow d_{new} + scm_{xy}$ 
6     else if  $lmc_{xy} < scm_{xy}$  then
7        $d_{boost} \leftarrow d_{boost} + (scm_{xy} - lmc_{xy})$ 
8     end
9    $d_{sat} \leftarrow d_{sat} + lms_{xy}$ 
10 end
11 end
```

---

cells. At this point, the vehicle also holds a collection of its neighbors' coverage maps,  $\mathcal{N} = \{SCM_1, SCM_2, \dots, SCM_n\}$ . Coverage matrixes are square matrixes of order  $n$  (but not symmetric, i.e.,  $scm_{n1} \neq scm_{1n}$ ) where  $n$  is always odd, so that a single cell exists at the center of the matrix that corresponds to the vehicle's location.

In this work, we use the subscripts  $(i, j)$  to refer to matrix rows and columns, and the subscripts  $(x, y)$  to refer to geographic coordinates. Coverage matrixes are always accompanied by the coordinates  $(x_{center}, y_{center})$  of the center cell  $scm_{\lceil n/2 \rceil \lceil n/2 \rceil}$ , which allow any element to be mapped to a latitude/longitude pair – when  $(x, y)$  is used, the mapping is implied, since it is straightforward.

We begin by constructing two cell maps of the local neighborhood: a local map of signal coverage  $LMC = (lmc_{ij}) \in \{0, 1, \dots, 5\}$  and a local map of saturation (RSU redundancy)  $LMS = (lms_{ij}) \in \mathbb{N}_0$ . These maps give a picture of the existing RSU support network: LMC tracks the best coverage being provided at each cell, and LMS counts the number of RSUs covering that cell at the same time. Algorithm 1 shows how these local maps are built.

Local maps span from the lowest to the highest latitudes/longitudes seen in all SCMs, i.e., they are as large as the total geographic area covered by the underlying self-observed coverage maps.

With the LMC and the LMS, the metrics in Equation (1) can now be calculated as described in Algorithm 2 below.  $d_{new}$  sums the strength of new coverage: e.g., a new cell covered with signal '4' will add '4' to  $d_{new}$ . This way,

the decision process can distinguish between new coverage that can be stronger or weaker.  $d_{\text{boost}}$  measures the improvements to already-covered cells, summing the delta between the existing and new signal levels.  $d_{\text{sat}}$  tracks redundant coverage (network saturation): for every cell that the new vehicle covers, sum the number of neighbors that already service it. For example, if a given cell that the car can reach is already serviced by 5 RSUs, add '5' to  $d_{\text{sat}}$ . This additive increase process helps keep a balanced number of RSUs that serve each cell.

Figure 4 shows two contrasting examples of the self-organizing decision process, in a simplified environment with no obstructions. The first, in Figures 4a and 4b, shows a new car parking in an advantageous location and evaluating its decision score. The algorithms reveal that, of the 93 cells in the car's coverage map, 86 will add new coverage to the network while 7 will bring no improvement to the existing coverage. This car will, therefore, have a positive  $d_{\text{score}}$  and become an RSU. With two RSUs now active, Figures 4c and 4d give a second example where another vehicle parks in between the RSUs. This new vehicle sees, from its decision algorithms, that it can add coverage to 10 new cells, and improve existing coverage to 25 cells – however, it will also be adding to network saturation to 58 cells, and may therefore have a negative  $d_{\text{score}}$ , and enter a sleep state.

The decision process builds from coverage matrixes that reflect the real-life status of the RSU support network and returns a Yes/No decision on each newly parked car. With the procedure for gathering coverage maps, the algorithms for newly parked cars, and the decision score equation, the self-organizing approach is fully functional.

#### D. Substituting Displaced Cars

We now outline an optional mechanism to preserve the structure of the support network whenever an RSU is displaced. When a parked car that is taking a Roadside Unit role becomes mobile again, we refer to this situation as the RSU having been *displaced*. Because decisions only occur when cars park, the service that was being provided by that RSU is not reestablished until a new car parks in the area.

One optimization to the self-organizing approach is to allow inactive parked cars to wake up periodically and listen to the 802.11p Control Channel (CCH), so they can react to changes in the network or respond to requests to come back on-line. We will first show how to achieve this periodic wake-up, and then propose a limited communication system for inactive vehicles to elect a replacement to a displaced RSU.

1) *Periodic Wake-Up*: A periodic wake-up allows inactive parked cars to be recalled in case of need. For a replacement protocol to be able to act upon all viable candidates, the wake-up events must be synchronized, which can be achieved with ease through 802.11p's mandatory time synchronization.

Through Time Advertisement messages and GPS timecode data, all On-Board Units (OBUs) of a vehicular network are synchronized to UTC (Universal Time Coordinated). We can then implement modular arithmetics based on the local time  $t_{\text{OBU}}$  to ensure all OBUs wake up in parallel. A modulo operation triggers a wake-up when the remainder of  $(t_{\text{OBU}} \bmod N)$

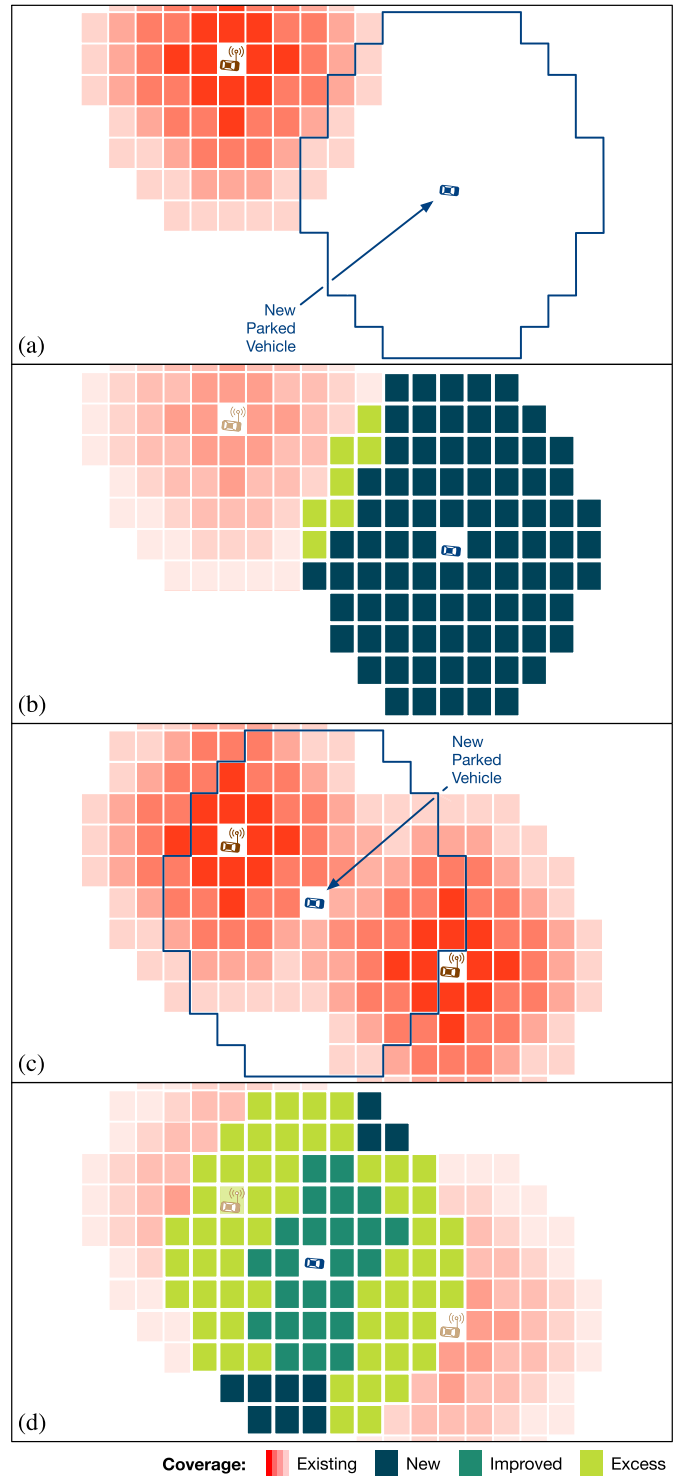


Fig. 4. A new vehicle parks in an advantageous location (Figure 4a), in the vicinity of a single existing RSU. The decision algorithms (Figure 4b) show that it can cover a substantial new area, and so the car establishes itself as a second RSU. Later on, another vehicle parks in between the two existing RSUs (Figure 4c). The algorithms now show (Figure 4d) that this car would be a net loss to the network, adding more redundant cells than new and improved cells. This vehicle does not become an RSU. Both images depict an obstruction-less scenario for illustrative purposes.

is zero: e.g., with  $t_{\text{OBU}}$  in seconds and  $N = 15$ , the OBUs will wake 4 times a minute. On each wake-up, the OBU must listen for a single CCH interval (50 ms), so with a






Fig. 5. The proposed election process, occurring on every inactive car that detects the loss of an RSU. This process occurs exclusively in IEEE 1609.3's Continuous Channel Access, where the radios remain on the Control Channel (CCH), and do not switch to the Service Channels (SCH). Slots correspond to CCH Intervals, standardized as 50 ms.

15-second interval inactive cars to have their OBUs active 0.3% of the time, with a corresponding energy savings. There is therefore a tradeoff between the reaction time of inactive parked cars and the energy required to allow inactive cars to be contacted.

2) *Electing a Replacement*: An inactive parked car detects a displaced RSU when it fails to hear a beacon coming from said RSU during its periodic wake-up interval. Beacons in 802.11p do not benefit from guaranteed delivery, so a car should only react once it fails to hear multiple beacons in a row. Figure 5 illustrates the election process, which works by quickly eliminating as many candidates as possible with no communication between them. This silent elimination ensures that a network in an area with hundreds of inactive cars (e.g. in a parking lot) is not flooded with messages every time an RSU is displaced.

To this end, we propose an initial backoff period inversely related to each car's decision score, as described earlier in Section III-C. Each car detecting a displaced RSU computes its  $d_{\text{score}}$  setting aside the coverage map of the displaced RSU. Its backoff time  $t_{\text{backoff}}$  is given by

$$t_{\text{backoff}} = \left\lfloor \left( 1 - \frac{d_{\text{score}}}{d_{\text{score}_{\text{max}}}} \right) \times N_{\text{backoff}_{\text{slots}}} \right\rfloor \times t_{\text{CCH}}, \quad (2)$$

where  $d_{\text{score}_{\text{max}}}$  is the largest value a decision score can take (which is found empirically),  $N_{\text{backoff}_{\text{slots}}}$  specifies how many intervals the backoff process is allowed to last for, and  $t_{\text{CCH}}$  is the duration of a CCH Interval, 50 ms. In this calculation, the car's decision score is used as a ratio to the number of backoff intervals. Larger (better) decision scores will then lead to shorter wait times. As the backoff timer on the best candidate expires, it begins to broadcast new RSU beacons, advertising itself as the winner of the process and instructing other candidates to return to sleep.

The number of backoff slots determines the balance between how quickly a replacement RSU is found and how many candidates are excluded by the process. When multiple listeners are assigned the same expiration timer, contention at the Medium Access layer ensures that only one will broadcast a beacon first, but at this point, the choice is random.  $N_{\text{backoff}_{\text{slots}}}$  should therefore be sufficiently large to eliminate most candidates based on their decision scores (e.g., 40 backoff slots can exclude 97.5% of all candidates in 2 seconds).

#### IV. PERFORMANCE ANALYSIS

Here, we present a study of the benefits that parked cars can bring to sparse urban areas and an analysis of the performance of the self-organizing approach shown in Section III. To do so, we run simulations on a platform that integrates real road topologies, realistic vehicle mobility, real maps of obstructions, and empirical signal measurements with 802.11p hardware, all of which are gathered from and applied to the city of Porto, Portugal. The platform, which was custom-developed for this work, is described next.

##### A. An Urban Simulation Platform Using Real Data

The urban environment introduces a number of challenges to a vehicular network. The movement of cars is dependent on factors such as the time of day, the type of road the car is on, and the activity of intersection traffic lights; the propagation of radio waves is affected by obstructing buildings of varied sizes and composition, signal reflection and diffraction on multiple surfaces, and even by other cars themselves blocking the communication path; moreover, constant node mobility causes these conditions to change at a rapid pace.

To better characterize the complexities of urban areas, we set out to design simulation scenarios that are highly realistic, so that the resulting data is reliable and indicative of what can happen in a real-life scenario. With this goal in mind, we developed a platform with the following features:

- Realistic vehicle mobility and traffic light patterns, obtained with the well-known open-source vehicle simulator SUMO [12].
- Real urban street layouts generated from publicly-available city maps that include lane numbers, speed limits and traffic lights [13].
- An accurate, vectorial model of urban obstructions, made available by the Porto City Council. With this obstruction data, in the form of a *shapefile*, we created a Geographic Information System (GIS) [14] to determine Line-Of-Sight status between any two vehicles.
- Realistic modeling of core wireless network metrics such as Bandwidth and Packet Loss, as a function of node distance and Line-Of-Sight status, derived from empirical measurements gathered with actual 802.11p-equipped vehicles circulating in Porto as well [15].

In our simulations, the signal strength between two cars is modeled with the results presented in [15]. The data from this study are used to assign levels of coverage quality while a parked vehicle is building its coverage map, through the process described earlier in Section III-A. This classification criteria is shown in Table II. In a real-life scenario, coverage strength can be determined with received signal strength measurements from the DSRC radios.

##### B. Improving Broadcast Delay in Sparse Networks

In the initial stage of a vehicular network deployment, insufficient numbers of DSRC-equipped vehicles will cause the network to become sparse. A sparse network is one where some of its nodes are too far apart from their neighbors for

TABLE II  
QUALITY MEASURE CRITERIA FOR COVERAGE MAPS

Quality	Distance (LOS)	Distance (NLOS)
5	70 m	58 m
4	115 m	65 m
3	135 m	105 m
2	155 m	130 m

communication to occur, leading to network disconnection. In an urban area, this causes virtual clusters of cars to form, where cars in a cluster can talk to one another, but are unable to send messages to other neighboring clusters. This sparse network problem has been well studied on highway vehicular networks [16] and, in these environments, connected Roadside Unit deployments bring substantial benefits [17], [18]. While this issue is prevalent in scenarios of low market penetration, it can also occur in periods of low traffic, even with full market penetration of DSRC hardware.

Here, we study a scenario where the urban area sees a low density of DSRC-enabled vehicles, which in turn also results in small numbers of parked cars that can join the network as RSUs. Our goal is to determine whether the approaches introduced in this paper bring tangible benefits in these sparse scenarios. For the problem at hand, we simulate an accident that occurs at a random point in the city, automatically triggering the dispatch of an emergency message. In real life, these messages have a dual purpose: to reach any nearby emergency vehicles, hospitals or police stations, and to quickly inform other drivers of the accident, so they can anticipate danger, expect congestion, and take alternative routes. The latter may also aid in the former, by reducing traffic in the vicinity of a crash, which helps emergency services en route.

This scenario serves as a tool to study the overall performance of a message broadcast in an urban vehicular network, and is not exclusive to safety applications. Our metric of interest is then the *message reachability*: the rate at which the message spreads to nearby vehicles, and how quickly everyone in the immediate neighborhood is informed. To broadcast the message, we select UV-CAST, a well-known urban broadcasting protocol designed specifically for vehicular networks [1]. When the need to disseminate a message occurs, UV-CAST comes into play by locating the edge nodes of the cluster the message is coming from, with a gift-wrapping algorithm. These edge nodes are then designated as the ideal message broadcasters and continue to rebroadcast the message as they meet new vehicles on the road. Figure 6a shows an example of the algorithm at work.

In this analysis, we adapt UV-CAST to support multiple message origins, by making use of the support network of parked cars. An example of this multi-origin system is seen in Figure 6b – here, active parked cars in the area assist the rebroadcasting of messages to other parts of the network. At locations where other active parked cars exist, UV-CAST's algorithms are executed again, and new sets of broadcasters are selected, effectively increasing the numbers of vehicles that redistribute our emergency message, as well as their geographical dispersion.

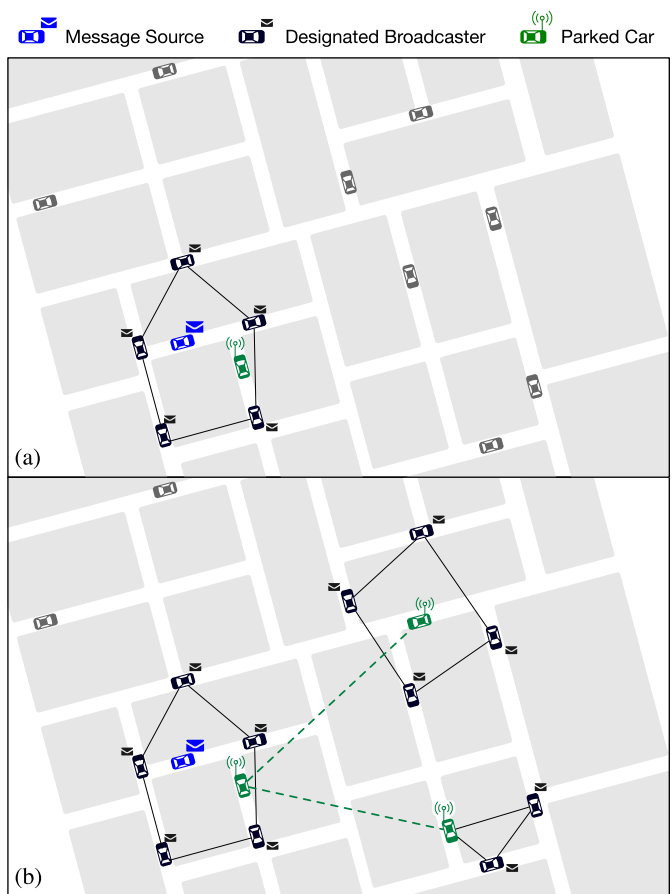


Fig. 6. UV-CAST selecting forwarders for a message broadcast with (a) no RSUs, and (b) 3 parked cars active in standalone mode.

For simulation purposes, we set up a 1 square kilometer area in the city of Porto, as described earlier in Section IV-A, featuring a diverse mix of road structures, speed limits, and functional traffic lights. We predefine a ratio of 1 parked car acting as an RSU for every 10 cars on the road (a 1:10 ratio) and analyse three densities of vehicles: 20, 40, and 80 cars per square kilometer. The first two densities match low-density and early-morning scenarios, while the third is a medium-density scenario where network sparsity is expected to be less problematic. Parked cars are placed randomly in the area, and the 1:10 ratio between parked cars and moving cars is chosen to be conservative – often, close to 80% of all cars in a given area are parked. The reduced number of parked cars that are able to operate as RSUs in standalone mode is intentionally chosen to emulate environments where uplink availability is scarce.

Figure 7 shows the evolution of an emergency message's *reachability* for the three densities under consideration. Here, reachability denotes the number of nodes that have received the message – e.g., for the 40 vehicle scenario, the maximum reachability is 40. We compare six scenarios, three with and three without parked car support, and average the results over 50 repetitions each. The data show improvements of 47%, 45% and 41% in the time required for full reachability, for the 20-, 40- and 80-vehicle densities respectively. These are substantial






Fig. 7. Multi-origin UV-CAST message reachability over time, with 10% of the cars on the road parked and active as standalone RSUs.

gains that effectively cut an emergency message's broadcast delay in half, in these sparse scenarios.

The improvements diminish, proportionately, as the number of cars on the road increases, which is an expected result given that higher vehicle density translates to improved network connectivity (less sparseness). For example, in the medium-density scenario (80 cars), because the network is less sparse, the emergency message immediately reaches  $\approx 60\%$  of all cars with the first broadcast – a consequence of the improved connectivity is to have fewer and larger clusters, and so a higher proportion of vehicles will be in the same cluster as the broadcast source. Seeing smaller gains in well-populated networks does not, however, mean that RSUs are not needed when the network is not sparse – in higher-density scenarios, the benefits of RSUs will come in the form of controlled broadcasting (e.g., preventing *broadcast storms*) and improved network efficiency, among others.

The simulation results shown in Figure 7 indicate that even a small number of parked cars randomly-distributed throughout an urban area can bring substantial improvements to the broadcast of important emergency messages, which is one of the most important applications of a vehicular network.

### C. Performance of Self-Organization With Minimal Data

We now turn our attention to the performance of our self-organizing approach. To do so, we run a series of simulations in a 1 sq. km. region in the city of Porto using the simulation platform detailed earlier, and analyze the behavior of the algorithms and mechanisms that were shown in Section III.

1) *Time to Build Coverage Maps*: The coverage maps that the vehicles build are a fundamental part of the proposed approach, but due to the randomness inherent in a vehicular network, a parked car may not have a reliable way to know if it has overheard enough beacons to form a reliable coverage matrix. A simulator is an ideal tool to analyze this observation step, as it can determine the complete map of a specific car beforehand, and track its evolution. We run a series of simulations where cars are parked at random locations in the city, introduce traffic at three different densities, and track the evolution of each car's coverage map. Figure 8 plots the maps' percentage of completeness as a function of time.




Fig. 8. Probability that a car's self coverage map has been completely observed, as a function of elapsed time, in low- and medium-density scenarios.

The data confirm the intuition that with more moving cars a reliable map will take less time to build. For 80% of the map to be statistically complete, in these low-density scenarios, the newly parked car must receive beacons for 498, 392, and 330 seconds, for the 20, 40 and 80 vehicle densities, respectively. A conservative approach for a real-life scenario is then to instruct cars to build coverage maps for 500 to 600 seconds (i.e., for about 10 minutes) after they park, to ensure a sufficiently complete map. A decision may be executed earlier, particularly if large numbers of beacons are received from multiple directions (suggesting a dense network) – in the worst case, the car will temporarily take an unnecessary RSU role, and once more coverage is learned, the decision process is retaken. In higher-density scenarios, our tests show that coverage maps are typically built in 30 seconds to a minute.

2) *Decision Algorithm vs. an Optimal Solution*: Given the cell map division proposed in Section III-A, our core metrics for a support network of parked cars are: the mean signal level available to each cell (network coverage); the average number of RSUs covering each cell (network saturation); and a count of how many parked cars take on the RSU role. In dense networks where considerable numbers of parked cars exist, the selection of which cars should take on Roadside Unit roles becomes the primary concern. The optimal solution to a given set of parked cars can be determined by evaluating each possible combination of active and inactive vehicles,

TABLE III  
REFERENCE OPTIMAL METRICS

Scenario	Signal Coverage	RSU Saturation	RSU Count
All RSUs active	4.13	8.28	24
Optimal selection	4.06	1.63	5
Set1 ( $\mu = 0.1$ )	3.93	2.42	7.2
Set2 ( $\lambda = 8.0$ )	3.76	2.02	5.6

with complete knowledge of the candidates and their coverage maps. Here,  $2^{\#cars}$  scenarios are possible, easily exceeding millions of computations, which makes optimal decisions infeasible in real-life.

We design a simulation scenario of high vehicle density where 24 cars are instructed to park randomly in a small  $0.18 \text{ km}^2$  section of the map. This constraint forces the cars to park on nearby streets, so that their coverage overlaps and a decision is made on which cars to keep active. For comparison purposes, we first determine the optimal solution by bruteforcing the  $2^{24}$  possible combinations (approx. 16,7 million runs), and we also calculate the effects of simply activating all available parked cars. The results are shown in Table III. Predictably, the activation of all available vehicles leads to the best coverage, but also causes each cell to see an average of 8 RSUs, which is inefficient and may be problematic. The optimal solution provides roughly the same level of signal coverage, now with only 1.6 RSUs seen at each cell, and just 5 out of the 24 cars left active. These results show that, given perfect decisions, 19 of these 24 parked cars are not needed and can enter an energy-efficient sleep state.

Next, we analyze our self-organizing approach, with sharing of coverage maps and decisions taken at each node, to determine whether it can approach the outcome of an optimal solution. To do so, we vary the weights of each component in the decision criteria  $d_{score}$ , rerun the simulation scenario, and plot the resulting coverage, saturation and RSU count metrics. In Figure 9, each of the coefficients  $\kappa$ ,  $\lambda$  and  $\mu$  is adjusted individually from 0.1 to 10.0, causing a corresponding decrease/increase in the weight that the  $d_{new}$ ,  $d_{boost}$  and  $d_{sat}$  metrics have in the final decision score of each parked car. The resulting city-wide signal coverage, RSU saturation, and total number of RSUs is shown. Figure 10 then plots these same city-wide metrics for two preselected parameter sets: *Set1* is geared towards better coverage, and *Set2* aims for fewer active RSUs.

The data in Figures 9 and 10 show that a decision process optimized towards improved coverage can reach a global signal strength of 3.93, which is only 3% worse than the Optimal selection. This particular scenario occurs by setting  $\mu = 0.1$  (the  $d_{sat}$  coefficient). However, the result is also that 7.2 parked cars remain active, which is 2.2 cars more than the optimal solution. A second set of coefficients can be selected to reduce the number of active cars to a minimum: with  $\lambda = 8$ , the resulting signal strength is of 3.76 (7% worse than Optimal), and only 5.6 parked cars remain active.

The data presented in this section show that our decision process can obtain near-optimal results without the need for




Fig. 9. Effect of varying the coefficients  $\kappa$ ,  $\lambda$  and  $\mu$  (and the corresponding components of the decision score,  $d_{new}$ ,  $d_{boost}$  and  $d_{sat}$ ) on signal strength, RSU saturation, and number of active RSUs. A single coefficient is varied at a time, while the others remain fixed at 1.0.




Fig. 10. (a) Average strength of coverage, and (b) number of parked cars activated, for two coefficient sets with different optimization goals, against an optimal selection. *Set1* attempts to maximize the quality of service coverage, while *Set2* is aimed at reducing the number of active parked cars.

an extensive knowledge of a network's RSUs, and that it can also be directed towards specific goals, such as better signal strength or fewer active parked cars. Extrapolating the results to a  $1 \text{ sq. km.}$  area, we observe that about 28-30 parked cars per sq. km. may be sufficient to form an extensive vehicular support network. This number is an order of magnitude smaller than the typical density of parked cars in a city.

## V. CONSIDERATIONS ON VEHICLE BATTERY LIFE

The main requirement for a vehicle to work as a Roadside Unit is to be able to maintain power to the DSRC electronics.

TABLE IV  
AVERAGE POWER DRAW OF MAJOR OBU SYSTEMS

Base system	1.64 W	Powering DSRC radio	2.42 W
CPU under 100% load	0.45 W	Powering WiFi radio	0.49 W
DSRC transmission at 100% duty cycle	1.75 W	DSRC reception at 100% duty cycle	0.21 W
DSRC transmission at 50% duty cycle	0.75 W	DSRC reception at 50% duty cycle	0.10 W

Ordinary passenger cars have two sources of energy: when the engine is running, an alternator generates electricity, powering the car's electronics and recharging its battery; when the engine is off, power is sourced from that same battery. DSRC electronics, when the car is parked and its engine is off, must be powered from the latter. The amount of energy that is drained should be controlled, as other applications (such as security systems and keyless entry systems) will also be draining the battery at the same time.

To understand the impact that DSRC hardware will have on the battery of a parked vehicle, we have performed a comprehensive study of the power consumption of the *NetRider*, an OBU/RSU prototype hardware unit developed by the company *Veniam*. The *NetRider* devices consist of a Single-Board Computer equipped with GPS and separate DSRC and WiFi radios. The units that were generously provided to us for this study are used at the core of the company's testbed deployments.

In a laboratory setting, we attached digital multimeters to a pair of NetRiders and measured the power draw of these equipment in the following activity states: powering the base hardware; powering the unit's DSRC radio; powering the unit's WiFi radio; simulating a 100% computational load on the processor; transmitting data on the DSRC interface at 100% and 50% duty cycles; and receiving data on the DSRC interface at 100% and 50% duty cycles. Power consumption measurements for each activity state were then taken for a duration of 30 seconds, at a rate of 10 measurements per second, and averaged. Table IV reports the power that was required by each activity.

From the data in Table IV, we estimate the power consumption of a hypothetical RSU scenario where the DSRC channel is averaging 50% utilization, the unit is transmitting and receiving similar amounts of data, and the mean CPU load nears 30% (a number commonly seen in our own testbeds), to be of approximately 4.62 W.

We can now make the estimation that, for the average 6.64 hours that a car spends parked every day [19], having an active DSRC unit will take a 4.2% toll on a car's battery. This is calculated as follows: a standard 12-volt automotive battery on a passenger vehicle has a capacity of 60 Ah (amperes/hour), or 720 Wh (watts/hour). The energy consumed by an OBU over the  $\approx 6.64$  hours the vehicle is parked is equal to  $E = P_{OBU} \times t_{parked} = 30.68$  Wh, and  $30.68/720 \approx 4.2\%$ . Battery capacity degrades with age and usage, and car batteries reach End-Of-Life (EOL) when less than 50% of the original capacity remains. At this EOL point, the battery toll would be of  $\approx 8.4\%$ , still less than one tenth of the total capacity.

We underline the fact that the hardware units we were able to obtain are prototypes meant for testbed deployments, and that production hardware is expected to draw under 3 watts in operation, further decreasing battery use. Despite these reassuring figures, parked cars acting as RSUs should nevertheless limit their activity, to safeguard excessive drain on the car's battery.

## VI. DISCUSSION

Leveraging parked vehicles as roadside units is a relatively new and unexplored concept, and the mechanisms and algorithms we present in this work should be considered as initial explorations of this idea. To the best of our knowledge, this work is the first to approach the problem of assigning RSU roles to a large pool of candidates, while optimizing the network in a self-organized way, thereby showing the viability of the concept in an urban environment. Our results show that a self-organizing approach with limited information can indeed approach a globally optimal coverage solution, indicating that parked cars are able to form a vehicular support network without the need for centralized control and decisions. The following are some of the open problems and possible areas of improvement.

In this work, we restricted exchanges of coverage maps to a conservative 1-hop neighborhood. Relaxing this limitation may yield more efficient algorithms (e.g., issuing a map request to the 2-hop neighborhood, or geocasting that same request), at the expense of incurring a higher overhead on the network.

The second part of our study analyzed how the network establishes itself in areas where no network existed before. This initial setup represents a transient state in the support network. In urban areas where parked vehicles exist at every hour of the day, once the initial RSUs are established a steady state is reached, and the replacement mechanism provided in Section III-D takes over. The proposed mechanism aims to be overhead-efficient, and preserve the initially-established structure by picking the replacement that is most similar to the departing vehicle. Here, steady state coverage may be improved by picking the replacement car that optimizes the network at the local neighborhood instead. This, however, may require large numbers of inactive parked vehicles to communicate so an informed decision can be performed.

The use of the car's battery is the ultimate concern in this work. We have shown, with empirical data, that an OBU operating for the average parking duration will draw under 5% of the available capacity, but other factors such as ageing batteries, prolonged parking events or the presence of other power-hungry electronics may require that this power use be further reduced. With the mechanism shown in Section III-D, that allows inactive cars to react to displaced RSUs and automatically select a replacement, an active parked car can intentionally disable itself after a pre-specified amount of time (e.g., after 1 hour, the power draw will be under 1%) and be automatically compensated for by neighboring cars. More advanced mechanisms can elect replacements beforehand and execute a soft handoff, however, a detailed study of such mechanisms is beyond the scope of this paper.

## VII. RELATED WORK

The concept of using cars as Roadside Units started to gain traction as it became apparent that the cost of these units would pose a significant barrier to their deployment. Leveraging moving cars as temporary RSUs in urban areas, by requesting that they make brief stops to aid in emergency message broadcasts, was shown to bring a measurable improvement to such goals [7]. Similar concepts, when applied to highway vehicular networks, were able to reduce broadcast delay when opposite-lane traffic was insufficient to relay packets [20].

The specific idea of *parked cars* as members of the vehicular network is first introduced in [8]. This work suggested activating parked cars to increase the number of nodes in a sparse network, improving its connectivity. It reported a 3.3x improvement in node density when 10% of available parked cars became active. These results follow from tenuous assumptions of stable 250 meter radio ranges in the urban area, which our empirical data has shown to be improbable.

The interesting work in [21] suggests using parked cars as a means to overcome the signal degradation that occurs when buildings block the line of sight between two vehicles. By activating parked cars at key intersection points, other cars can use the parked car at the intersection as an unobstructed message relay. With this approach, cars were shown to be able to receive nearby emergency messages up to 17 seconds faster. A follow-up to this work [22] also demonstrated how parked cars can be used to aid existing RSUs in content downloading. By caching content from fixed RSUs, bandwidth demands on the main RSU for distributing content could be alleviated. These works presented data obtained under the assumption of 200 meter radio ranges, irrespective of obstructions.

The existing body of work on parked cars has revealed the interesting possibilities that can be brought by leveraging parked cars as RSUs, albeit in specific, limited scopes. While these works describe some interesting use cases for parked cars in urban vehicular networks, they do not address the fundamental question of how these entities should be selected, managed, and replaced using a self-organizing network approach, nor whether they can assume more flexible support roles in the network, instead of being relegated to these specific uses. Our work aims to show that using a self-organizing network approach, parked cars can serve as a low-cost and very efficient alternative to deploying RSUs in urban areas.

Interestingly, the general idea of using a self-organizing network approach for solving several important problems in traffic safety and traffic efficiency has been demonstrated before. For example, the well-known “Virtual Traffic Lights (VTL)” idea is based on a similar self-organizing network approach whereby vehicles approaching an intersection can communicate among themselves using vehicle-to-vehicle communications for electing a “Leader” for managing the right-of-way and the traffic flow at that intersection without any need for infrastructure-based traffic lights [23]–[25]. It was shown before that VTL can reduce the commute time by more than 30% during rush hours, which is pretty significant. In another recent study, it was shown that such a self-organizing network approach can also handle priority of emergency vehicles such as ambulances and fire trucks at

intersections if one uses a different set of *local rules* for each vehicle [25], [26]. Clearly, the underlying local rules that are needed for performing different functions (such as serving as RSUs, serving as a Leader in the VTL scheme that obviates the need for infrastructure-based traffic lights, providing connectivity, priority management, sensing, etc.) in vehicular networks would have to be carefully designed for each specific task at hand or the specific function to be performed. However, the overall benefits of using this self-organizing network approach are compelling and the same approach has proven to be profitable in several instances for different tasks. As such, it holds great promise for solving several other challenging engineering problems in a cost-effective and efficient manner.


## VIII. CONCLUSION

It has been shown that, using a self-organizing network approach, the large numbers of parked vehicles in urban areas can be leveraged to serve as RSUs and provide support to the networks of moving vehicles. These networks of parked cars can serve as low-cost, efficient alternatives to expensive deployments of fixed Roadside Units. To function as RSUs, they require only the ability to keep DSRC electronics powered while vehicles are parked. Our proposed self-organizing network approach introduced novel ways for cars to determine their ability to act as RSUs, and to decide whether to become RSUs from that knowledge. Our analysis showed how a support network that provides excellent coverage is possible using only a small fraction of the cars that park in a city.

## REFERENCES

- [1] W. Viriyasitavat, O. K. Tonguz, and F. Bai, “UV-CAST: An urban vehicular broadcast protocol,” *IEEE Commun. Mag.*, vol. 49, no. 11, pp. 116–124, Nov. 2011.
- [2] S. Eichler, “Performance evaluation of the IEEE 802.11p WAVE communication standard,” in *Proc. IEEE 66th Veh. Technol. Conf. (VTC)*, Sep. 2007, pp. 2199–2203.
- [3] *IEEE Standard for Information Technology—Local and Metropolitan Area Networks—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments*, IEEE Standard 802.11p-2010, Jul. 2010.
- [4] M. Freitas, “Vehicle infrastructure integration,” U.S. Dept. Transp., Washington, DC, USA, unpublished presentation, Oct. 2005.
- [5] “Connected vehicle technology industry delphi study,” Center Autom. Res., Ann Arbor, MI, USA, Tech. Rep., Sep. 2012. [Online]. Available: [https://www.michigan.gov/documents/mdot/09-27-2012\\_Connected\\_Vehicle\\_Technology\\_-\\_Industry\\_Delphi\\_Study\\_401329\\_7.pdf](https://www.michigan.gov/documents/mdot/09-27-2012_Connected_Vehicle_Technology_-_Industry_Delphi_Study_401329_7.pdf)
- [6] J. Wright, C. J. Hill, J. K. Garrett, and R. Rajbhandari, “National connected vehicle field infrastructure footprint analysis: Deployment scenarios,” Amer. Assoc. State Highway Transp. Off., Washington, DC, USA, Tech. Rep. FHWA-JPO-14-125, Jun. 2014. [Online]. Available: [http://stsmo.transportation.org/Documents/Task%206a%20AASHTO\\_CV\\_Footprint\\_Deployment\\_Scenarios\\_v2.pdf](http://stsmo.transportation.org/Documents/Task%206a%20AASHTO_CV_Footprint_Deployment_Scenarios_v2.pdf)
- [7] O. K. Tonguz and W. Viriyasitavat, “Cars as roadside units: A self-organizing network solution,” *IEEE Commun. Mag.*, vol. 51, no. 12, pp. 112–120, Dec. 2013.
- [8] N. Liu, M. Liu, W. Lou, G. Chen, and J. Cao, “PVA in VANETs: Stopped cars are not silent,” in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 431–435.
- [9] A. B. Reis and S. Sargento, “Leveraging parked cars as urban self-organizing road-side units,” in *Proc. IEEE 82nd Veh. Technol. Conf. (VTC)*, Sep. 2015, pp. 1–5.
- [10] *Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service*, document ETSI TS 102 637-2, European Telecommunications Standards Institute, Apr. 2010.

- [11] *IEEE Standard for Wireless Access in Vehicular Environments (WAVE)—Networking Services*, IEEE Standard 1609.3-2016, Apr. 2016.
- [12] D. Krajewicz *et al.*, “Recent development and applications of SUMO—Simulation of urban mobility,” *Int. J. Adv. Syst. Meas.*, vol. 5, no. 3–4, pp. 128–138, 2012.
- [13] OpenStreetMap. *OpenStreetMap Foundation and OpenStreetMap Contributors*, accessed Feb. 2014. [Online]. Available: <http://www.openstreetmap.org/>
- [14] *PostGIS: Spatial and Geographic Objects for PostgreSQL*, accessed Feb. 2014. [Online]. Available: <http://postgis.net/>
- [15] A. Cardote, F. Neves, S. Sargento, and P. Steenkiste, “A statistical channel model for realistic simulation in VANET,” in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Nov. 2012, pp. 48–55.
- [16] N. Wisitpongphan, F. Bai, P. Mudalige, V. Sadekar, and O. Tonguz, “Routing in sparse vehicular ad hoc wireless networks,” *IEEE J. Sel. Areas Commun.*, vol. 25, no. 8, pp. 1538–1556, Oct. 2007.
- [17] A. B. Reis, S. Sargento, F. Neves, and O. K. Tonguz, “Deploying roadside units in sparse vehicular networks: What really works and what does not,” *IEEE Trans. Veh. Technol.*, vol. 63, no. 6, pp. 2794–2806, Jul. 2014.
- [18] S.-I. Sou and O. K. Tonguz, “Enhancing VANET connectivity through roadside units on highways,” *IEEE Trans. Veh. Technol.*, vol. 60, no. 8, pp. 3586–3602, Oct. 2011.
- [19] C. Morency and M. Trepanier, “Characterizing parking spaces using travel survey data,” CIRRELT, Montreal, QC, Canada, Tech. Rep. 2008-15, May 2008.
- [20] A. B. Reis, S. Sargento, and O. K. Tonguz, “Parameters that affect safety message delay in sparse infrastructure-less vehicular networks,” in *Proc. IEEE ICC*, Jun. 2014, pp. 2568–2574.
- [21] C. Sommer, D. Eckhoff, and F. Dressler, “IVC in cities: Signal attenuation by buildings and how parked cars can improve the situation,” *IEEE Trans. Mobile Comput.*, vol. 13, no. 8, pp. 1733–1745, Aug. 2014.
- [22] F. Malandrino, C. Casetti, C. F. Chiasserini, C. Sommer, and F. Dressler, “Content downloading in vehicular networks: Bringing parked cars into the picture,” in *Proc. IEEE PIMRC*, Sep. 2012, pp. 1534–1539.
- [23] M. Ferreira *et al.*, “Self-organized traffic control,” in *Proc. 7th ACM Int. Workshop Veh. InterNETw.*, Sep. 2010, pp. 85–89.
- [24] O. K. Tonguz, W. Viriyasitavat, and J. M. Roldan, “Implementing virtual traffic lights with partial penetration: A game-theoretic approach,” *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 173–182, Dec. 2014.
- [25] O. K. Tonguz and W. Viriyasitavat, “A self-organizing network approach to priority management at intersections,” *IEEE Commun. Mag.*, vol. 54, no. 6, pp. 119–127, Jun. 2016.
- [26] O. K. Tonguz and W. Viriyasitavat, “Methods and software for managing vehicle priority in a self-organizing traffic control system,” U.S. Patent 9536427, Jan. 3, 2017.



**Andre B. Reis** received the B.S. and M.Sc. degrees in electronics and telecommunications engineering from the University of Aveiro, Portugal, in 2009, in collaboration with the Eindhoven University of Technology, The Netherlands. He is currently pursuing the Ph.D. degree in electrical and computer engineering with Carnegie Mellon University (CMU), Pittsburgh, PA, USA, and the University of Aveiro under the CMU-Portugal Program. His current research is focused on infrastructure support systems for vehicular networks in challenging scenarios.



**Susana Sargento** received the Ph.D. degree in electrical engineering from the University of Aveiro, in 2003. She was a Guest Faculty with the Department of Electrical and Computer Engineering, Carnegie Mellon University, USA, in 2008, where she performed Faculty Exchange in 2010/2011. In 2012, she co-founded a vehicular networking company, Veniam'Works, which builds a seamless low-cost vehicle-based internet infrastructure. She joined the Department of Computer Science, University of Porto, in 2002, and is an Associate Professor with Habilitation, Universidade de Aveiro, and with the Institute of Telecommunications, where she is leading the Network Architectures and Protocols Group. She has been involved in several national and international projects, taking leadership of several activities in the projects, such as the network of vehicles on buses and trucks, as well as several fixed stations, counting with over 600 vehicles in the city of Porto, Portugal. Her main research interests are in the areas of next generation and future networks, more specifically QoS, routing, mobility, self- and cognitive networks. She was a recipient of the 2016 EU Prize for Women Innovators. She has been the TPC-Chair, has organized several international conferences and journals, such as *IEEE Wireless Communications*, *IEEE Network*, and *IEEE Communication*. She regularly acts as an expert for European research programs.



**Ozan K. Tonguz** is currently serving as a tenured Full Professor in the Department of Electrical and Computer Engineering of Carnegie Mellon University. Before joining academia, he was with Bell Communications Research (Bellcore), doing research in optical communications systems and optical networks. His current research interests include vehicular networks, ad hoc wireless and sensor networks, computer networks, Internet of Things, Smart Grid, security, self-organizing networks, artificial intelligence, and machine learning.

He has authored about 300 technical papers in IEEE journals and conference proceedings. In addition, he is an inventor or co-inventor on 15 patents (issued or pending). He is well known for his contributions to vehicular networks, wireless networks and sensor networks, and optical communications systems and networks. He co-authored (with G. Ferrari) the book “Ad Hoc Wireless Networks: A Communication-Theoretic Perspective” (Wiley, 2006). His work and discoveries on Virtual Traffic Lights (VTL) has been featured in over 100 countries in different continents by CNN, New Scientist, Discovery Channel, GizMag, CBS, AUDI Technology Magazine, The Atlantic, Metro World News, Eyewitness News in Johannesburg, South Africa, many radio and TV channels, and Internet media. His industrial experience includes periods with Bell Communications Research, CTI Inc., Harris RF Communications, Aria Wireless Systems, Clearwire Technologies, Nokia Networks, Nokia Research Center, Neuro Kinetics, Asea Brown Boveri, General Motors, Texas Instruments, and Intel. He currently serves as a consultant or expert for several companies, such as Aria Wireless Systems, Harris RF Communications, Clearwire Technologies, Nokia Networks, Alcatel, Lucent Technologies, major law firms such as Jones Day, Williams and Connolly, WilmerHale, Heller Ehrman, Baker Botts, Soroker-Agmon, Dinsmore&Shohl, and Carlson and Caspers, and government agencies such as NSF, DARPA, and DoT, in the USA, Europe, and Asia in the broad area of telecommunications and networking. He is the Founder and President of VTL, LLC, a CMU start-up that was launched in December 2010 for providing solutions to several key transportation problems related to safety and traffic information systems, using vehicle-to-vehicle and vehicle-to-infrastructure communications paradigms. He also served as the Co-Director (Thrust Leader) of the Center for Wireless and Broadband Networking Research, Carnegie Mellon University.

Prof. Tonguz has served as a member of the technical program committees of several IEEE conferences and symposia, such as INFOCOM, SECON, GLOBECOM, ICC, VTC, and WCNC. He served as an Associate Editor or Guest Editor of several IEEE Journals and Transactions, such as the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, the IEEE JOURNAL OF LIGHTWAVE TECHNOLOGY, the IEEE/OSA JOURNAL OF LIGHTWAVE TECHNOLOGY, and the IEEE Communications Magazine.