Opportunistically-assisted parking search: a story of free riders, selfish liars and bona fide mules

Evangelia Kokolaki Georgios Kollias Maria Papadaki
Merkourios Karaliopoulos Ioannis Stavrakakis

Department of Informatics and Telecommunications
National & Kapodistrian University of Athens
Ilissia, 157 84 Athens, Greece
Email: {evako, grad1153, grad1140, mkaralio, ioannis}@di.uoa.gr

***Abstract*—In competitive autonomic networking environments, user nodes face a strategic dilemma: on the one hand, they need to cooperate to support the networking infrastructure and information flow; on the other hand they are tempted not to do so, *e.g.*, in order to conserve own system resources or create an advantage for themselves. In this paper we investigate a real-world scenario of parking assistance service that instantiates such environments. Under the nominal (altruistic) operation, the vehicles opportunistically collect and share information on the location and availability status of the parking spots they encounter. Yet the competition for parking spots may give rise to various facets of misbehaviors, such as deferring from sharing their information (*free riders*) and/or deliberately falsifying disseminated information so as to divert other drivers away from a particular area of own interest (*selfish liars*). Simulation results indicate a persistent *fate-sharing* effect, *i.e.*, misbehaving nodes fail to obtain any substantial performance advantage that would indeed motivate their misbehaviors. Furthermore, the overall performance of the system does not necessarily deteriorate as the intensity of misbehaviors increases. Misbehaviors rather tend to reduce the distance between the destination and the actual parking spot occupied for *all* vehicles at the expense of higher parking search times, which quickly become prohibitive when the vehicles' destinations overlap. Finally, the addition of mobile storage nodes (*bona fide mules*) compensates for the reduction of the information flow due to free riders but has almost no effect against selfish liars since the mobile storage nodes end up propagating the falsified information those nodes generate.**

Keywords*—vehicular networks, parking assistance systems, non-cooperative opportunistic dissemination

I. Introduction

In various mobile applications involving competition for scarce resources, networked entities (user nodes) have to autonomously decide whether to dispose private information about the resources. Information is essentially a kind of asset; sharing it, user nodes assist their potential competitors, in anticipation of their support in due course. Recent trends such as the smart city initiative [1] give rise to further settings, where truthful altruistic information sharing is required but not guaranteed. One of these settings, involving city-level parking assistance systems, is the subject of this paper.

In particular, advanced parking assistance systems have been proposed (*e.g.*, [2]), and in some cases realized (*e.g.*, [3] or [4], [5] via social networks), in an attempt to cope with the issue of parking space management in busy urban environments [6]. Fostered by recent advances in wireless networking, sensing and car navigation technologies have, these systems aim at helping drivers find vacant parking spots easier and faster by collecting and sharing information about the location and status (occupied/vacant) of parking spots. In *centralized* systems, a central server communicating with sensors at the parking spots coordinates the parking spot assignment process, by receiving the drivers' requests, reserving parking spots, and directing drivers thereto (*e.g.*, [7]). Whereas, in *opportunistic* systems, vehicles themselves serve as mobile sensing platforms that collect and store information about the location and status of parking spots and share it with other vehicles through vehicle-to-vehicle (V2V) communication technologies (*e.g.*, [8]). Opportunistic systems do not incur the upfront infrastructure cost of centralized systems, thus presenting a lighter and more scalable solution that leverages to-be-built-in vehicle equipment. On the other hand, opportunistic systems lack central coordination and rely on the drivers' willingness to let them share collected information, assuming that drivers have full or partial control over the information exchange process. This cannot be taken for granted since the sharing of information assists nodes by increasing their knowledge about parking space availability but, at the same time, *synchronizes* nodes' parking choices. This synchronization in turn increases the competition for the vacant parking spots, in particular when drivers' travel destinations overlap [9].

In this paper, we are, to the best of our knowledge, the first to question the robustness of opportunistic parking assistance systems to non-cooperative drivers' behaviors, which deviate from the purely *altruistic* paradigm of always *truthfully* sharing the cached information with encountered vehicles. Hence, we let nodes *misbehave* and study how this affects fundamental performance indices such as the parking search time and the distance of the acquired parking spots from the drivers' travel destinations. The dual question from a driver's viewpoint is whether nodes *do* have incentives to misbehave in that misbehaving lets them achieve better search times and/or parking spot-destination distances. Two intuitive instances of misbehaviors are considered. In the first one nodes defer from sharing parking information with other vehicles essentially acting as *free riders*. In the second one, they deliberately falsify information about the parking spots' status (*selfish liars*), *i.e.*, spots close to a misbehaving vehicle's destination are advertised as occupied whereas all others as vacant. The two misbehaviors essentially impair in different manner the *amount* and *accuracy* of information that is disseminated across the network.

The problem under consideration features strong spatiotemporal dynamics that are not always conducive to theoretical investigation. Hence, the study is carried out primarily through simulations, whereas modeling is the apparent next step for future work to make theoretical arguments about the simulation findings. The results do not lie always in line with intuition. Notably, in almost all cases misbehaving nodes fail to obtain distinctly better performance than cooperative nodes. Both types of misbehavior, through different mechanisms, tend to reduce the destination-spot distances and increase the parking search times for *all* vehicles, the latter increase becoming quickly prohibitive when drivers' destinations overlap. This *fate-sharing* effect essentially weakens vehicles' incentives to misbehave and increases the system resilience to selfishly-thinking drivers. On the other hand, neither of the two misbehaviors attenuates the *synchronization phenomena* emerging at the cache contents, and subsequently, the mobility patterns of vehicles when their destinations overlap. The introduction of mobile storage nodes in this case, which collect and share parking information with parking-seeking vehicles, has a sharply different impact on the two misbehavior instances. Whereas, in the presence of free riders, a few of them suffice to restore the information flow at the levels of a cooperative system, they have negligible impact in the presence of selfish-liars: even a few misbehaving vehicles suffice to overwrite the fresh information mobile storage nodes carry and convert them into relays of forged information (*bona fide mules*).

The basic operation of the opportunistic parking assistance system and the two obvious ways selfish nodes may try to manipulate it are reviewed in Section II. The simulation environment and our methodology are described in Section III. We present and discuss the simulation results in Section IV, outline the related research in Section V and conclude our work in Section VI.

II. Opportunistically-assisted parking search and imperfect cooperation

According to the current common practice in search for parking space, drivers wander around their travel destination and sequentially check the availability of encountered parking spots. Typically, the search is initially carried out within an area around the drivers' travel destination (*initial parking search area*), whose size depends on the drivers' attitude and sense of traffic load and parking demand thereby. The radius of the search area then grows progressively as parking search time increases until drivers find a vacant parking spot and occupy it. This, essentially *blind*, search practice gives often rise to congestion problems and results in fuel/time wastage, especially around popular travel destinations such as the centers and business districts in big cities.

Recent progress in wireless communication, sensing and navigation technologies promise to make the parking search process smarter and more efficient. One way to do this is by equipping vehicles with sensors and standard wireless interfaces (*e.g.*, 802.11x) in ad-hoc mode that let them collect and share information about parking spots' location and status as they drive around. Such information can be further filtered across time (*aging*) and space through the use of timestamps and the geographic addresses (*e.g.*, via GPS) of individual parking spots. With such information at hand, vehicles can make more informed decisions. Rather than wandering randomly in the parking search area, a vehicle can now direct its search towards selected parking spots that are listed in its cache as the closest vacant ones to its travel destination. If the spot is actually vacant when it arrives at it, it occupies it; otherwise, it repeats the spot selection process, being also prompt to occupy any vacant spot it may find on its way to the candidate spot.

Critical for the efficiency of this *opportunistically-assisted* parking search are the *amount* and *accuracy* of the information that is stored in the vehicles' caches and shared among them. Both are subject to strong spatiotemporal effects: vehicles generally possess partial rather than global information about parking space availability and as the status of parking spots changes over time, stored data are potentially outdated after some time interval. Moreover, vehicular nodes have good reasons to hide information from other, potentially competitor, vehicles. Overall, the processes of information dissemination (benefiting discovery of parking spots and their availability) and competition growth (reducing the chances to acquire a spot) are coupled and counter-acting. Indeed, the faster information circulates across the wireless opportunistic networking environment, the more similar (accurate or not) data are stored in the caches of vehicles. Thus, depending on the travel destinations of users, the movement patterns of individual vehicles get synchronized and sharpen the effective competition for given parking spots[1] [9]. This additional level of competition, this time for information at the "service discovery" level, motivates various deviations from the perfectly cooperative (altruistic) behavior.

In this paper, we consider in detail two variants of imperfect cooperation, hereafter called misbehaviors for the sake of brevity. In the first variant, misbehaving nodes defer from sharing their own information with other vehicles, while readily accepting such information from other vehicles that make it available. These free riders reduce the amount of disseminated information but also its accuracy since vehicles' caches are less frequently updated with fresh information about the spots' occupancy status. On the contrary, the second misbehavior instance involves the dissemination of falsified information about the status of parking spots. Nodes do so in order to create zones free of competition around their travel destinations by diverting encountered vehicles away from them. Compared to the first misbehavior instance, this one affects only the accuracy of the disseminated information.

Inferring *a priori* the impact of these rather common misbehaviors is not straightforward for two main reasons. The first one is the aforementioned spatiotemporal effect. For example, misbehaving nodes that forge information may inadvertently correct outdated information (*i.e.*, turn the availability status

[1]Similar synchronization effects emerge from traffic congestion information systems, as well (*e.g.*, Google Maps with Traffic Layer); broadcasting information about traffic congestion within particular city areas or blocks discourages drivers from travelling there, yet increases traffic pressure in other road sections.

of the advertised parking spots to their real up-to-date value) and, thus, end up assisting the process. The second reason relates to the cache synchronization effects that emerge as the frequency of information updates rises. It may be argued that the two types of misbehaviors can serve as regulators for the synchronization phenomena and the resulting competition. We explore these aspects in detail in Section IV.

III. PERFORMANCE EVALUATION METHODOLOGY

A. Simulation Environment

Our study is carried out in the simulation environment developed for [9]. In what follows, we outline its features that are critical to our study.

Road grid and parking spots: The simulator implements a grid of two-lane roads (one lane in each direction) with roundabouts connecting up to four roads. Parking spots are uniformly distributed across roads' lanes of the grid.

Vehicle movement: The vehicle mobility model comes under the broad category of behavioral mobility models. Two levels of behavior can be identified: the *global*, determining how destinations are selected and the way the vehicles choose the route towards them; and the *local*, addressing how the vehicle moves within the roads comprising the route.

At the global level, every time a vehicle frees a parking spot, it chooses a new destination (geographical coordinates within the bounds of a city road grid) and drives towards it. Once it reaches adequately close to the destination (*initial parking search area*), the parking search process is initiated. The initial parking search area is circular; it is centered at the travel destination with radius equal to half the distance between two adjacent road intersections. Where the vehicle drives next depends on the information stored in its memory. The stored records (parking spot, status, timestamp) are filtered both temporally, to exclude information that is outdated (*i.e.*, coupled with a timestamp that is beyond a threshold value), and spatially, to retain as candidates only spots in the current search area. Out of the remaining spots, the user picks up the nearest-to-her-destination available one (*Full use of Memory, FM*). If no record survives the spatiotemporal filtering step, the driver chooses randomly one spot within the parking search area and moves towards it (*Random use of Memory, RM*). In the absence of *any* information about parking spots within the current area of interest, the vehicle circulates blindly/randomly within the area (*No Memory, NoM*). In all cases, vehicles move along shortest routes to their destinations and occupy the first available parking spot on their way to them rather than pursing closer-to-destination, yet non-guaranteed, parking options. If the driver finds a spot vacant, either a memory-selected or a randomly met one, it occupies it for a time interval that may follow different probability distributions. By the end of this interval, she vacates the spot and selects another destination. Otherwise, upon a failured attempt, the user will check anew her memory and repeat the attempt, as aforedescribed. After a particular number of failured attempts in the current parking search area, the driver increases its range.

At local level, the position of each vehicle by the next simulation time step depends on its current position and velocity.

TABLE I. SIMULATION PARAMETERS

Parameters	Values
Simulation grid	$1200 \times 1200m^2$
Simulation time duration	10^5 sec
Number of uniformly distributed spots, P	25
Number of vehicles, V	$5-70$
User maximum speed	$14m/s \sim 50km/h$
Vehicle - spot sensor commun. range	$15m$
Vehicle - Vehicle commun. range	$70m$
Exponential parking time with mean	$1800sec$
Distance between adjacent roundabouts	$300m$
Linear increase step of parking search area	$150m$
Radius of Interest, RoI	$150, 350, 500m$
Ratio of misbehaving nodes, p	$0, 0.3, 0.5, 0.8, 1$

More specifically, the vehicles adapt their speed according to their distance from: (a) the front vehicles (they are not allowed to overtake one another); (b) the next intersection; and (c) the nearest parking spot, assuming that they decelerate when encountering parking spots to check their status. Their speed is zeroed when they get stuck in traffic jam, enter a round about intersection, or park. Finally, the vehicles are not allowed to stop or move in the reverse direction of the traffic flow.

Cooperative *vs.* misbehaving vehicles: All vehicles inform their memory cache every time they hit a parking spot sensor. Well-behaving (cooperative) vehicles *share truthfully* stored information about the location and status of parking spots each time they encounter other vehicles. On the other hand, misbehaving vehicles realize the two misbehavior instances described in Section II:

Information Denial: Upon encounters with other nodes, they suppress information they store about the location and availability of parking space, whereas they update their cached information with all the new knowledge offered. During their search, they use the cached information the same way as cooperative nodes.

Information Forgery: They advertise all parking spots within a specific distance from their destinations (*Radius of Interest, RoI*) as occupied, and all others as vacant, while setting the relevant timestamps to fresh values. Being more suspicious about falsified information, they persist more when searching around their destinations; namely, they run additional random trips (in the RM or NoM mode) over the initial parking search area before they decide to increase the range of their search.

B. Simulation set-up and performance metrics

Unless otherwise stated, the simulations are run with the parameter values (value ranges) shown in Table I. Indeed, similar research efforts provide clues on the vehicle densities that reflect a realistic simulation environment. For instance, in [8] the authors study a networking environment, where the vehicular nodes' (not stable) density is drawn near $33veh/km^2$. Similarly, [10] explores the performance of a V2V communication platform, assuming $50veh/km^2$ moving according to the Manhattan Model. Motivated by these values, we end up with vehicle densities ranging from 3.5 to $45veh/km^2$. Furthermore, the real effective curbside is far from $20km$ (if margins around street corners are accounted for) and much less if we consider typical center areas of big European cities, where parking is completely forbidden in whole areas of road blocks. This is why we primarily focus our research on ratios $V/P > 1$.

Performance Metrics: The two main performance metrics throughout our study are the average time spent for searching available parking place (*Parking search time,* T_{ps}) and the average geographical distance between the vehicles' travel destinations and the selected parking spots (*Destination - Parking spot distance,* D_p). In addition, at a more microscopic level, we extract results for the amount and the profile of the information that is stored in vehicles' caches as well as the way vehicles use it and benefit from it, by plotting statistics about the percentage of time (total efforts) the vehicles search in *FM* and *RM* mode.

IV. Simulation results - Experimentation

In all plots, we compare the metric values under perfectly cooperative operation with those under different misbehavior intensities for various levels of parking demand. Each point in the plot results from averaging parking events over either the full set of nodes, or, separately, cooperative (denoted by 'C') and misbehaving (non-cooperative) ones (denoted by 'NC'). Drivers are assumed to be persistent in their search. Alternatively, they could abandon their effort to park, *e.g.*, stop looking for on-street parking and head for a more expensive parking lot, once the parking search time exceeds an upper bound. A red line in the plots indicates a timeout for the parking search process at 1800 seconds.

A. Uniformly distributed travel destinations

1) Information Denial: The first remark out of Fig. 1 is that the system exhibits remarkable robustness to this type of misbehavior. Neither the average parking search time (Fig. 1(a)) nor destination-spot distance (Fig. 1(b)) are penalized even when half the vehicular nodes defer from sharing information. An increase in parking search time becomes visible when 80% of the nodes misbehave and evolves to a striking tradeoff when all nodes misbehave; namely, if all vehicles defer from information sharing, they end up acquiring spots closer to their destinations at the expense of higher search time. The reason for this can be traced in the combination of Fig. 1(c) and Fig. 2(a). Without information sharing, the caches of nodes are primarily populated with records of spots around their destination (initial parking search area), encountered during their first attempts. As these spots are occupied (for medium- to high demand), and although vehicles gradually increase the range of their search, they still end up randomly selecting one of these spots (RM mode) with high probability. Contrary to when even a few nodes share information, their caches are not refreshed with records of more distant spots communicated by other vehicles[2] (Fig. 1(c)). Instead they are only occasionally enriched with some randomly encountered spot in the destination proximity, where their search ends up being restricted. Reading the system robustness the other way round, equally remarkable is the failure of selfishly misbehaving nodes to attain better performance, when compared to what cooperative nodes achieve (ref. to ingraphs in Fig. 1(a), 1(b)).

On the other hand, Fig. 2 gives clear insights into a fundamental inefficiency of opportunistically-assisted search, the coupling of information sharing (*about* parking spots) with the generated competition (*for* parking spots). The ratio of searching attempts in FM mode (Fig. 2(b)) starts from low levels at small demand, where anyway it is easier for a vehicle to find a spot, and decreases as the number of competing vehicles grows, where more spots are occupied, more vehicles are parked, and the flow of information is yet too slow to fill the vehicles' caches with adequately fresh information about vacant spots. When the demand grows even more and more vehicles end up cruising around, the information flow (at least for low or moderate intensity of misbehavior) is strengthened. Vehicles find fresh records about vacant spots in their caches, yet these are only a few and the competition for them so sharp that this information rarely results in a successful attempt (Fig. 2(c), 2(d)). For higher intensity of misbehaviors, both the frequency and success rate of search in FM mode decrease.

2) Information Forgery: Under Information Forgery, the vehicular nodes try to spontaneously generate competition-free zones around their travel destinations. For small RoI values, these zones are narrow and disjoint. Since misbehaving nodes advertise parking spots outside these zones as vacant and the drivers' destinations are uniformly distributed, the (cooperative) nodes end up (incorrectly) listing spots around their own travel destinations as vacant for most of the time. These spots emerge as top choices out of the spatiotemporal filtering step (FM mode) and attract repeated parking attempts (Fig. 3(c)). As a result, the vehicles park closer to the destination at the expense of higher search times. As misbehaving nodes become more aggressive and the zones they try to induce start to overlap ($RoI = \{350, 500\}$), most spots in the vehicles' caches are reported as occupied, the vehicles exercise more the RM mode, and a tradeoff emerges between destination-spot distances and parking search times, as shown in Fig. 3(a) and Fig. 3(b).

Contrary to the Information Denial misbehavior, under Information Forgery the misbehavior intensity and its impact do not only depend on the number of misbehaving nodes but also on the population of cooperative nodes. The latter inadvertently propagate forged information across the network once they get infected with it upon encounter with a misbehaving node. This has two direct consequences. First, the destination-spot distance *vs.* parking search time tradeoff is now milder, as shown in Fig. 4(a) and Fig. 4(b); for given RoI even a small ratio of misbehaving nodes suffices to populate the vehicles' caches with supposedly vacant spots and steer their attempts to spots around their travel destinations (Fig. 4(c)). Secondly, with a small exception for low parking demand levels ($V < P$), misbehaving nodes cannot gain any substantial performance advantage over cooperative nodes (ref. to ingraphs in Fig. 3(a), 3(b), 4(a), 4(b)) since the manipulated information they generate, bounces back to them after one or more hops over cooperative nodes.

[2] As the demand increases and larger amounts of information are circulated, the vehicles' caches store information about more spots. Eventually, for high inter-vehicle communication rates, caches store information about all spots and the average distance between drivers' destinations and stored parking spots approximates the expected distance of two randomly selected points within the square area (*square line picking problem*), known to equal $0.52 \times l$, where l denotes the square side length.

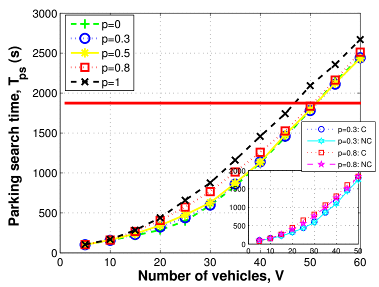

(a) Average parking search time

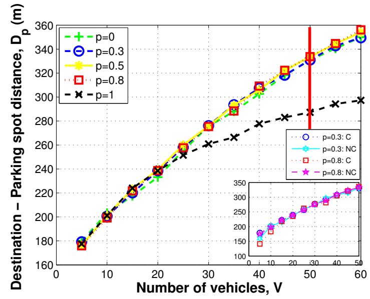

(b) Average destination-spot distance

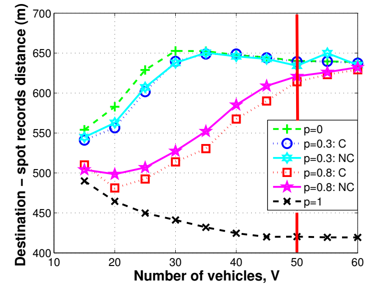

(c) Average distance from destination of spot records in vehicles' caches

Fig. 1. Robustness of the opportunistically-assisted parking search to *Information Denial*: uniformly distributed destinations.

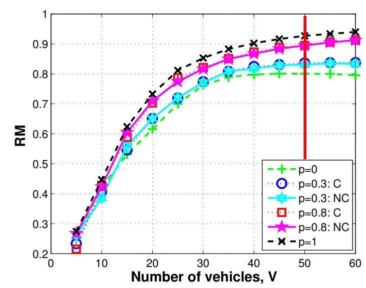

(a) Ratio of parking attempts in RM mode

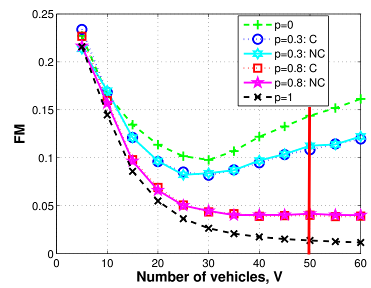

(b) Ratio of parking attempts in FM mode

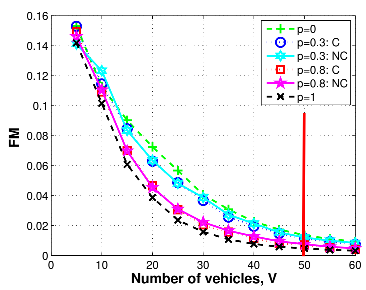

(c) Ratio of successful parking attempts in FM mode

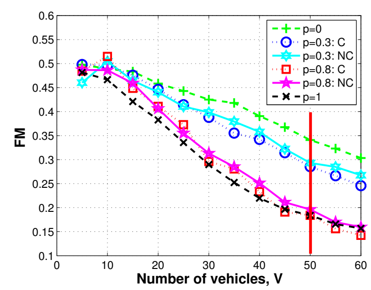

(d) Ratio of successful parking attempts in FM mode over all successful attempts

Fig. 2. Search mode and parking attempt success rates under *Information Denial*: uniformly distributed destinations.

B. *Hotspot scenario*

Under a fully cooperative setting, the spatial concentration of vehicles' travel destinations has two direct consequences on the information stored in their caches. First, as all vehicles cruise along the hotspot area and encounter each other more frequently, they tend to synchronize their caches with records about the same set of spots. Secondly, and most importantly, they rank these spots identically. Hence, at least as long as drivers let the system direct their attempts, their trips get synchronized, competition sharpens and parking search times increase substantially [9].

1) Information Denial: In the hotspot setting, the Information Denial has a double-edged effect. On the positive side, the system is shown to be resilient to the free rider behavior; even when half the nodes defer from sharing information, the average parking search times and spot-destination distances are almost intact, as shown in Fig. 5(a) and Fig. 5(b), respectively. Furthermore, misbehaving nodes do not gain in both performance indices by hiding information (ref. to ingraphs in Fig. 5(a) and Fig. 5(b)). On the other hand, this misbehavior does not manage to break the inherent synchronization effects and drive the system to a better-than-nominal performance level. When eventually, with most nodes in the network misbehaving, differentiation is achieved at the vehicles' caches, it is outweighed by a substantial decrease of disseminated information. Vehicles do not get informed about and do not take advantage of vacant parking spots further away from their common destinations (Fig. 5(c)). They rather end up parking closer to them, yet at the expense of unacceptable cruising times, even under moderate parking demand levels.

2) Information Forgery: In the hotspot scenario, the zones that misbehaving vehicles try to clear from competition overlap and all vacant spots beyond a distance equal to RoI are advertised as vacant by misbehaving nodes. For small RoI, vehicles persistently direct their attempts towards the few spots lying close to their common destinations so that their caches are not enriched with information about vacant spots further away, as shown in Fig. 6(c). The synchronization/competition effect is stronger and vehicles waste even more time in myopically searching for a parking spot around the hotspot road (Fig. 6(a)). However, as a result of this search mode, the vehicles

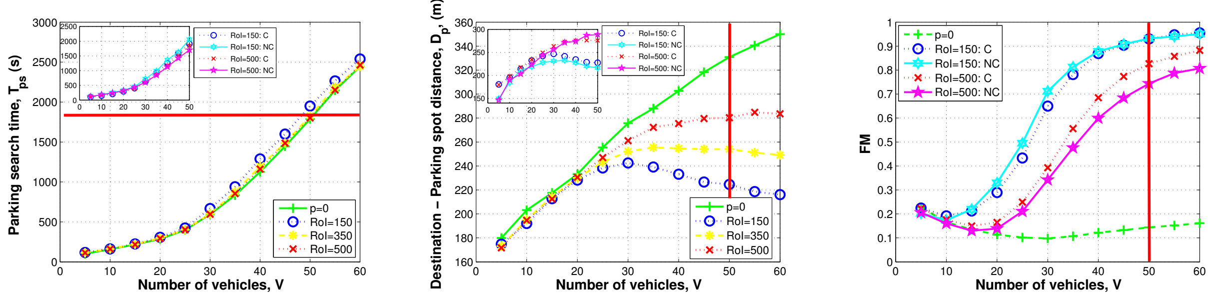

(a) Average parking search time (b) Average destination-spot distance (c) Ratio of parking attempts in FM mode

Fig. 3. Robustness of the opportunistically-assisted parking search to *Information Forgery*: uniformly distributed destinations, $p = 0.3$.

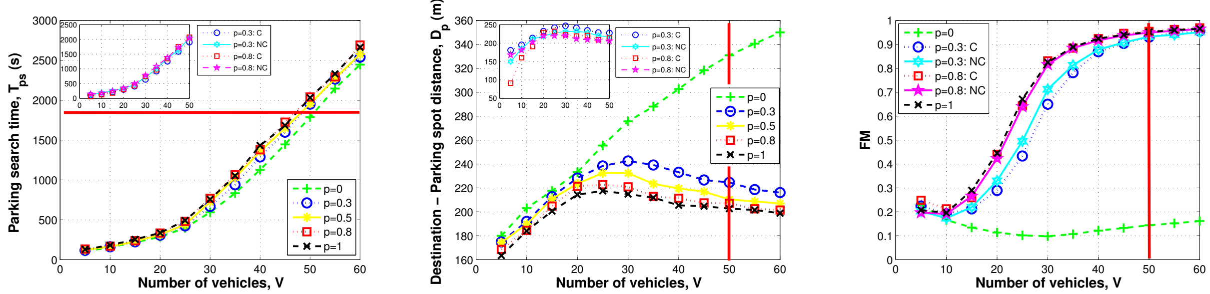

(a) Average parking search time (b) Average destination-spot distance (c) Ratio of parking attempts in FM mode

Fig. 4. Robustness of the opportunistically-assisted parking search to *Information Forgery*: uniformly distributed destinations, $RoI = 150m$.

park closer to their destination (Fig. 6(b)). Interestingly and rather counter to intuition, as misbehaving nodes become more aggressive and try to clear from competition larger areas (*i.e.*, $RoI = \{350, 500\}$), the parking search times improve for all vehicles. The reason is that vehicles are steered by the content of their caches to expand their search further away from the hotspot area and have the chance to encounter and, potentially occupy, spots they were not aware of. Essentially, the movement of vehicles in a broader area helps alleviate, though not resolve, the synchronization effect. Again, as with uniformly distributed travel destinations, misbehaving nodes cannot attain some performance advantage since the falsified information returns back to them, this time even faster due to more frequent encounters between vehicles (ref. to ingraphs in Fig. 6(a), 6(b)).

C. Mobile Storage Nodes for the hotspot scenario

The Mobile Storage Nodes (MSNs) can be either dedicated or normal vehicles, *e.g.*, city cabs, equipped with wireless interfaces that allow them to collect parking information from the entire area and share it with other vehicles and MSNs. By relaying information, MSNs indirectly increase the effective contact opportunities between vehicles and thus, the speed of information spread. The efficiency of MSNs as a countermeasure for the two types of misbehaviors is very different.

1) Information Denial: In this case, even a very small number of MSNs restore the information flows at the levels (and even better) of the fully-cooperative system. They render both the average parking time and the spot-destination distance independent of the number of free rider vehicles, as can be clearly seen in Fig. 7. Even when vehicles do not exchange at all information with each other, the communication with MSNs suffices to achieve better parking search times than those under the fully cooperative system. The addition of more MSNs (we experimented with 15 MSNs) does not bear visible changes to the performance metrics; on the other hand, similar results are obtained with even one MSN. In fact, a single encounter with MSN informs nodes about almost *all* parking spots in the area, helping them expand their search in a broader area around the hotspot road and partly randomize their driving patterns. Yet, the synchronization phenomena due to the vehicles' overlapping travel destinations are not fully eliminated and retain the parking search times at significantly higher levels than under uniformly distributed destinations.

2) Information Forgery: When nodes misbehave this way, the MSNs are a far less efficient solution. Although they collect and store up-to-date information about the actual status of all parking spots as they move randomly within the grid, this information is *rewritten* upon encounters with misbehaving nodes (or even otherwise cooperative nodes that have been polluted with falsified information). Thus, the MSNs end up further fostering the diffusion of falsified information that synchronizes the vehicles' caches making the synchronization effects even stronger and the decrease of the search times thanks to additional fresh information, marginal (Fig. 8).

V. Related Work

Misbehaviors and challenges in securing systems have been explored in the broader context of VANETs with respect to a wide range of safety, traffic management, and infotainment applications [11]. Primitives for secure applications and properties that can support secure systems are discussed in [12], [13]; while particular paradigms for authentication mechanisms and security protocols are presented in [14], [15].

Parking assistance applications lie at the intersection of traffic management and infotainment applications. *Opportunis-*

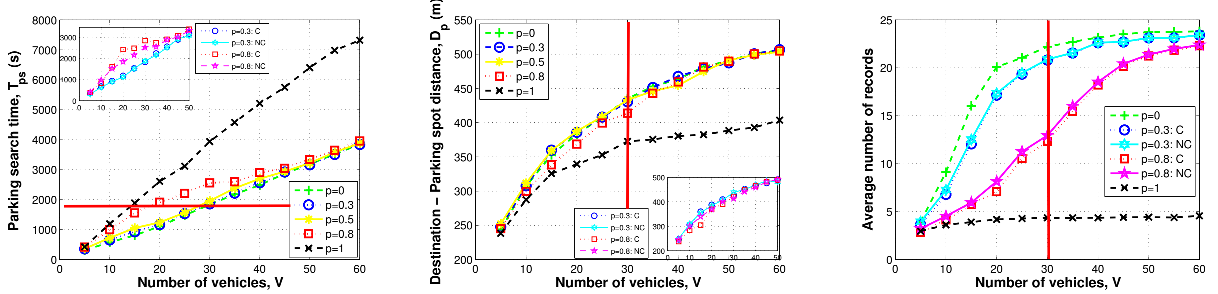

(a) Average parking search time (b) Average destination-spot distance (c) Average number of records in memory

Fig. 5. Robustness of the opportunistically-assisted parking search to *Information Denial*: hotspot road.

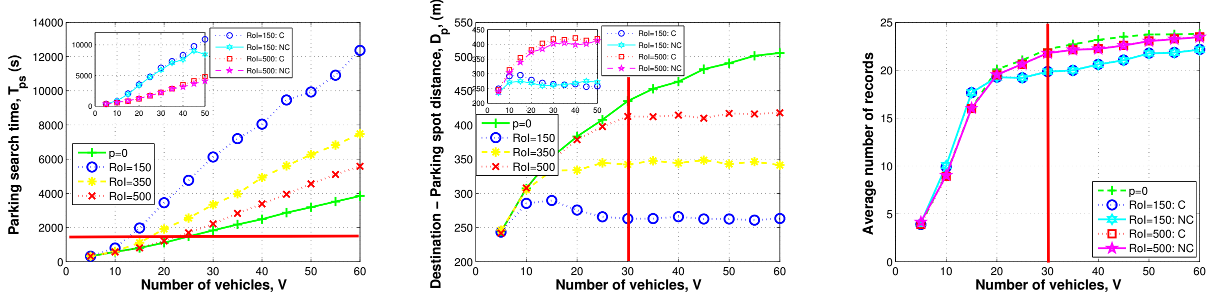

(a) Average parking search time (b) Average destination-spot distance (c) Average number of records in memory

Fig. 6. Robustness of the opportunistically-assisted parking search to *Information Forgery*: hotspot road, $p = 0.3$.

tic parking assistance systems, in particular, are proposed in [8], [16] and [17]. In [8], a scalable information dissemination algorithm is presented where the vehicles are allowed to exchange aggregate parking information of variable accuracy. In a similar work, the vehicles exchange information and solve a variant of the Time-Varying Travelling Salesman problem while dynamically planning the best feasible trip along all (reported-to-be) vacant parking spots [16]. In a different approach, Delot *et al.* propose in [17] a distributed virtual parking space reservation mechanism, whereby vehicles vacating a parking spot selectively distribute this information to their proximity. Hence, they mitigate the competition for the scarce parking spots by opportunistically controlling the diffusion of the parking information among drivers. Interestingly, the systems in [4] and [18] realize almost the same idea for parking management in the cities of Athens (Greece) and New York, respectively. In particular, in the absence of any realized on-board system for intervehicle communication, both applications leverage the social network element: users can offer their parking spot to the rest of the users or find a parking spot for themselves by claiming a spot another user is offering. A rating mechanism on drivers' sharing and reserving habits, shapes parking spot seekers' likelihood to be chosen by a parking spot sharer (defender) in [4] or get informed about a vacancy prior to other seekers in [18].

Common to all these studies is that the parking assistance systems are proposed under the assumption of full cooperation of vehicles. To the best of our knowledge, our study is the first one that considers the impact of imperfect cooperation on the operation of opportunistic parking assistance systems. We have particularly focused on the ways different nodes may try to impede or manipulate the flow of information in order to better serve their own interests and whether the introduction of storage nodes may compensate for these misbehaviors.

VI. Conclusions

The paper has looked into the vulnerability of opportunistic parking assistance systems to drivers' selfish behaviors. In our study drivers are let behave as free riders that benefit from information other vehicles collect and share but do not share theirs; and selfish liars that falsify information in their caches in order to increase their chances to find a spot close to their destinations.

Interestingly and counter to intuition, our results reveal a persistent fate-sharing effect; namely the misbehaving nodes fail to obtain any substantial performance advantage over what the cooperative nodes achieve, irrespective of the distribution of travel destinations. On the contrary, misbehaviors tend to increase parking search times, sometimes (overlapping travel destinations) to unacceptable levels, and reduce the distance between parking spot and travel destination. Both misbehaviors deteriorate the synchronization phenomena that emerge with respect to the information stored by vehicles and their movement patterns when travel destinations overlap. Mobile storage nodes can compensate the impact of free riders and improve the system performance beyond that of the fully cooperative scheme. On the contrary, they have almost no effect when confronting selfish liars since they end up propagating the falsified information those nodes generate. A natural direction for future work is to validate these experimental results against analytical models.

VII. Acknowledgements

This work has been supported in part by the European Commission IST-FET project RECOGNITION (FP7-IST-257756), the European Commission under the Network of Excellence in Internet Science project EINS (FP7-IST-

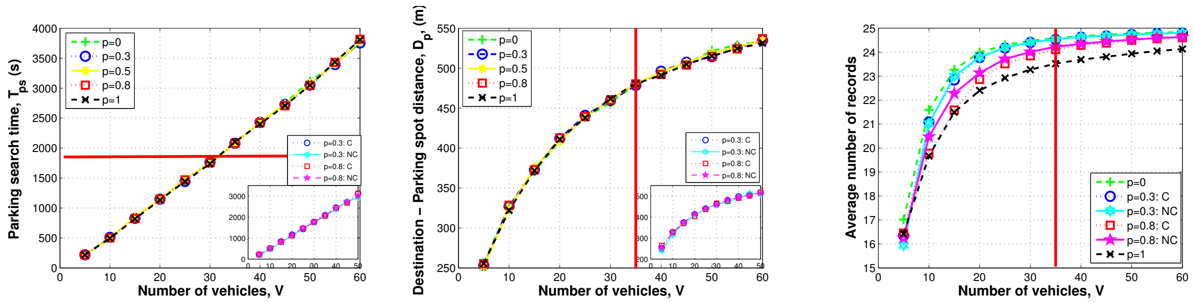

(a) Average parking search time (b) Average dest-spot distance (c) Average number of records in memory

Fig. 7. Mobile Storage Nodes and *Information Denial*: hotspot road.

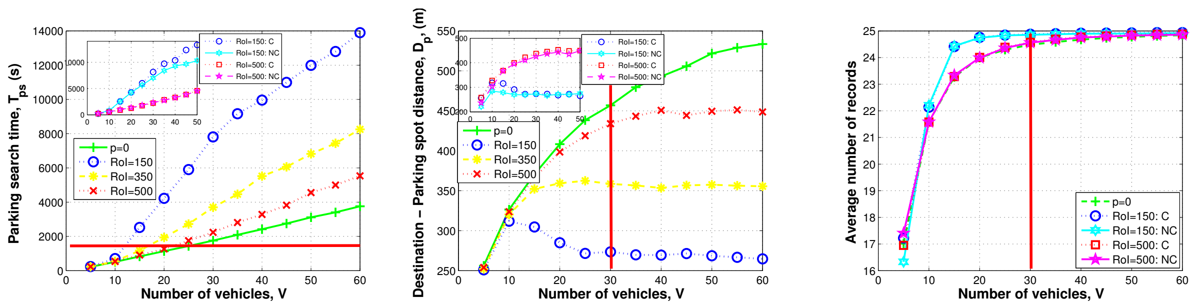

(a) Average parking search time (b) Average dest-spot distance (c) Average number of records in memory

Fig. 8. Mobile Storage Nodes and *Information Forgery*: hotspot road, $p = 0.3$.

288021) and the University of Athens Special Account of Research Grants no 10812.

REFERENCES

[1] A. Caragliu, C. Del Bo, and P. Nijkamp, "Smart cities in Europe," *Series Research Memoranda 0048*, 2009.

[2] S. Mathur, T. Jin, N. Kasturirangan, J. Chandrasekaran, W. Xue, M. Gruteser, and W. Trappe, "Parknet: Drive-by sensing of road-side parking statistics," in *Proc. 8th ACM MobiSys*, San Fran., USA, 2010.

[3] SFpark: parking system for San Francisco, available online in http://sfpark.org/.

[4] Parking Defenders mobile application, available online in http://www.parkingd.com.

[5] Parkomotivo: parking system for the city of Lugano, available online in http://www.bmob-park.com.

[6] http://www-03.ibm.com/press/us/en/pressrelease/35515.wss.

[7] H. Wang and W. He, "A reservation-based smart parking system," in *Proc. 5th Inter. Workshop on Cyber-Physical Networking Systems (colocated with IEEE INFOCOM)*, April 2011.

[8] M. Caliskan, D. Graupner, and M. Mauve, "Decentralized discovery of free parking places," in *Proc. 3rd VANET (in conjunction with ACM MobiCom)*, USA, 2006.

[9] E. Kokolaki, M. Karaliopoulos, and I. Stavrakakis, "Opportunistically assisted parking service discovery: Now it helps, now it does not," *Pervasive and Mobile Computing*, vol. 8, no. 2, 2012.

[10] I. Carreras, I. Chlamtac, H. Woesner, and H. Zhang, "Nomadic sensor networks," in *Proc. 2nd European Workshop on Wireless Sensor Networks (EWSN)*, Istanbul, Turkey, February 2005.

[11] M. Gerla and L. Kleinrock, "Vehicular networks and the future of the mobile internet," *Computer Networks*, vol. 55, no. 2, 2011.

[12] M. Raya, P. Papadimitratos, and J.-P. Hubaux, "Securing vehicular communications," *IEEE Wireless Communications*, 2006.

[13] B. Parno and A. Perrig, "Challenges in securing vehicular networks," in *Workshop on the Hot Topics in Networks (HOTNETS-IV)*, 2005.

[14] M. Raya and J.-P. Hubaux, "Securing vehicular ad hoc networks," *J. Comput. Secur.*, vol. 15, no. 1, 2007.

[15] B. Aslam and C. Zou, "Distributed certificate and application architecture for Vanets," in *Proc. 28th IEEE MILCOM*, Piscataway, USA, 2009.

[16] V. Verroios, V. Efstathiou, and A. Delis, "Reaching available public parking spaces in urban environments using ad hoc networking," in *Proc. 12th IEEE MDM*, Beijing, China, 2011.

[17] T. Delot, N. Cenerario, S. Ilarri, and S. Lecomte, "A cooperative reservation protocol for parking spaces in vehicular ad hoc networks," in *Proc. 6th Inter. Conference on Mobile Technology, Application and Systems*, ser. MOBILITY, 2009.

[18] ParkShark: parking system for New York City, available online in http://www.parkshark.mobi/www/.