Querying Spectrum Databases and Improved Sensing for Vehicular Cognitive Radio Networks

Abdulla Al-Ali,* Kaushik Chowdhury,* Marco Di Felice,† and Jarkko Paavola,‡
*Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, U.S.A.
E-mail: al-ali.a@husky.neu.edu, krc@ece.neu.edu
† Department of Computer Science, Bologna, Italy. E-mail: difelice@cs.unibo.it
‡ Turku University of Applied Sciences, Finland. E-mail: jarkko.paavola@turkuamk.fi

***Abstract*—Cognitive radio (CR) vehicular networks are poised to opportunistically use the licensed spectrum for high bandwidth inter-vehicular messaging, driver-assist functions, and passenger entertainment services. Recent rulings that mandate the use of spectrum databases introduce additional challenges in this highly mobile environment, where the CR enabled vehicles must update their spectrum data frequently and complete the data transfers with roadside base stations. As the rules allow local spectrum sensing only under the assurance of high accuracy, there is an associated tradeoff in obtaining assuredly correct spectrum updates from the database at a finite cost, compared to locally obtained sensing results that may have a finite error probability. This paper aims to answer the question of when to undertake local spectrum sensing and when to rely on database updates through a novel method of exploiting the correlation between 2G spectrum bands and TV whitespace. We describe experimental studies that validate our approach and quantify the cost savings made possible by intermittent database queries.**

***Index Terms*—cognitive radio, correlation, spectrum database, FCC database, query costs.**

I. INTRODUCTION

The landmark FCC ruling in November 2011 in the US mandated the use of spectrum databases, with rules of access for stationary and mobile cognitive radio (CR) nodes, as well as the consideration of specific capabilities such as geo-location [1]. These databases release information on the spectrum usage in the vicinity of the requesting node, which must be periodically refreshed to maintain updated information. However, the FCC also allows for local spectrum sensing, though such unassisted and unilateral sensing by a node must adhere to strict performance metrics. Identifying spatial regions or durations in which (i) the local sensing is likely to yield reliable and repeatable results (i.e., less random fluctuations on the signal imposed by the channel) and when (ii) accessing spectrum databases is a must, remains an open challenge.

The work described in this paper is focused on answering the following question: How and when must a node rely on spectrum databases as opposed to relying on self-generated measurements? To solve this problem, we propose a method for local sensing that considerably improves upon the accuracy of the sensing result by reducing the probability of mis-detection of the primary user (PU) signals that have priority of access. The main idea is to exploit the correlation between the received signal strength that exists at specific locations in entirely different and disconnected spectrum bands. Signals from transmitters that are located in close proximity are likely to experience a strong similarity in coarse channel behavior, owing to the common set of reflection, absorption causing objects, and the large-scale path loss. We explore an interesting concept through an experimental study, where we establish that at certain locations, the behavior of the received signal strength in cellular channels can be used to predict the accuracy of spectrum sensing for certain channels in the TV whitespace. In some cases, the TV transmitter and one or more cellular transmitters (operating on totally different bands) are located near each other, and when the signals from these two are received at a common location, we detect a strong correlation in the behavior of the Received Signal Strength Indicator (RSSI). Thus, any sudden change in the TV spectrum usage in such locations can be verified by comparing with the corresponding fluctuations in the cellular channels. This provides an additional layer of check and reduces mis-detection. Of course, such a close correlation occurs only for limited durations and at specific locations (based on the relative distances between these two different transmitter types), but it can potentially offset costly spectrum database queries at these times.

The main contributions of this paper are as follows:

- We propose a method that relies on correlation among the 2G signals and TV whitespace to improve the local spectrum sensing accuracy.
- We experimentally demonstrate this concept using actual vehicular measurements in the Boston area, using software defined radios scanning the TV channels and an Android based smartphone to obtain the 2G spectrum data.
- We devise a database query protocol that leverages this correlation, and minimizes the number of queries, thus also reducing cost of obtaining updated spectrum information.

The rest of this paper is organized as following: In Section II, we discuss the related works. In Section III, we showcase the experimental results and provide the motivation for our work. The proposed network architecture is discussed in Section IV, followed by our decision algorithm in Section V. Performance analysis are conducted in Section VI and finally, we conclude in Section VII.

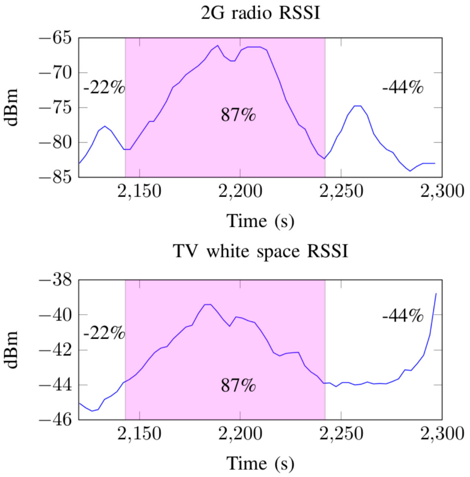

Fig. 1. Cross correlation between one TV channel band and one 2G radio tower RSSI values.

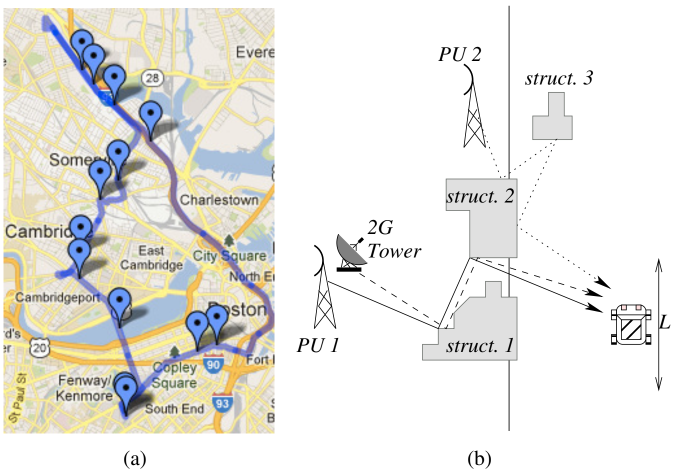

Fig. 2. The traversed path in our measurement experiment with markers indicating locations of high RSSI correlation between TV channels and 2G spectrum (a), and explanation for the correlation (b).

II. Related Work and Background

In [6], the authors devise a framework where vehicles in a CR network perform sensing, send the data to roadside units or a base station (BS), which in turn sends its data to a processing unit. The processing unit then infers which channels the vehicles are allowed to use based on the aggregate sensing information, and finally schedules this information to be broadcasted to passing vehicles alongside the same road segments periodically. Belief propagation techniques are used in [8], where vehicles combine different observations from surrounding vehicles, and spatial correlation is used to decide on channel availability. In [9], a framework for coordinated spectrum sensing method is proposed in the absence of any BS or a roadside static base station. Instead, some vehicles are temporarily assigned the role of a “master” vehicle that coordinates the sensing and schedules the transmission activity of surrounding vehicles. In [7], a framework called *Cog-V2V* is devised, which is a light-weight cooperative spectrum sensing scheme. Each node aggregates information it receives from surrounding vehicles to determine which channel to use in the current and future locations alongside the vehicles’ path. In a study of a different nature, [5] analyzes a major interstate freeway in the state of Massachusetts (I-90) for free spectrum availability along its length. The authors argue for a use of a centralized database that vehicles can use to access the free spectrum information. An algorithm is also proposed on how to reliably populate such a database based on the sensing measurements obtained along the I-90 interstate.

Our approach marks a radical departure from the above directions by making use of cross correlations between two different frequency bands to aid the decision on whether to sense or query the FCC database. To the best of our knowledge, this is the first effort that leverages this information in CR vehicular networks under the constraints of database querying.

III. Experimental study of Correlation between 2G and TV Channels

To study the correlation between the RSSI in the cellular 2G frequencies and the TV channels, we gathered a comprehensive set of measurements in the city of Boston over the span of one month. We used two devices to collect data simultaneously: a Universal Software Radio Platform (USRP) was used to sense the digital TV channels $21-51$ [3], and an Android Samsung Galaxy S3 smartphone was used to gather the following information: (i) RSSI values of nearby 2G cellular towers, (ii) current GPS coordinates, and (iii) the true TV spectrum availability queried at $60\,\mathrm{s}$ intervals. We used Spectrum Bridge [4] by writing an Android application that directly accessed their proprietary APIs to return the available/occupied channels and the signal strengths in the area of study. These devices were placed in a car, and the actual path traversed is shown in Figure 2(a). The path progresses along a counter-clockwise direction, starting and ending at Northeastern University campus (at the bottom of the figure).

The RSSI samples gathered by the moving car were stored and analyzed offline to detect whether a level of cross-correlation exists in any $40-220\ s$ moving window among the 2G spectrum and the TV channels. The points in Figure 2(a) show the locations where the cross-correlation were between $85-95\%$. Understandably, these correlations were detected at low building density areas: (a) bridges, (b) suburban low building density areas or (c) in broad street intersections, where the random multipath effects resulting from the neighboring structures was comparatively small.

Our motivation for conducting these experiments and exploring the correlation between dissimilar spectrum bands is shown in Figure 2(b). As the vehicle moves along a direction from the bottom to the top, there are specific regions for which the TV transmitters (i.e., PUs) in certain channels and the 2G spectrum towers are nearby, and their respective signals follow a somewhat similar path. This scenario exists for PU1 and the 2G tower. At this time, the signal from PU2, on a different TV channel, encounters an altogether different multipath environment. Thus, the signal fluctuations between 2G and PU1 towers are likely to be correlated in a small window of the traversal path, shown by L, and totally uncorrelated at other times when the structural neighborhood changes. Thus, on observing sudden changes in the subset of highly correlated TV channels and matching these changes with the channel behavior of the 2G spectrum, the CR vehicle can identify if these are caused by multipath effects or due to PU activity.

Figure 1 shows a randomly selected location point with two RSSI measurements plotted against time, one from a 2G cellular tower and the other from channel number 46 in the TV band, with intervals where strong correlation exists. The upper plot is obtained from measurements through the Android phone i.e., measuring signals in the 2G spectrum, and the bottom one is from the USRP i.e., in the TV channel. A moving average filter with a set span 11 was applied to both RSSI values for smoothening before calculating the correlation. The correlation is obtained using the $crosscorr$ function described in 1. Here, f^* denotes the complex conjugate of f and n is the lag of one set relative to the other. We use a default lag of $n = 20$ (i.e. from -10 to 10) and evaluate $max(f \star g)$ as our maximum cross-correlation value. We see in Figure 1 that the shaded time interval has 87% correlation, while intervals immediately next to this segment are highly uncorrelated, i.e., the cross-correlation result is between -22% and -44% for these time durations before and after the shaded region, respectively.

$$(f \star g)[n] = \sum_{m=-\infty}^{\infty} f^*[m]\; g[n+m] \tag{1}$$

In our proposed method, for certain locations and TV channels for which a strong correlation exists with the 2G spectrum, the CR vehicle may not require additional spectrum database updates to infer PU activity in these specific TV channels. Instead it may rely on local spectrum sensing alone, as for these channels it can track and attribute the cause of rapid signal fluctuations to multipath and other channel conditions, thereby lowering the chance of false alarm and missed detection. The FCC ruling recently authorized three types of devices for TV whitespace access: Fixed BSs (with direct connectivity to database), Mode II (with geolocation, direct connectivity to database), and Mode I (without geolocation, connectivity to fixed BS or Mode II device).

In our proposed network architecture, as vehicle already has access to 2G RSSI values, it is possible to assume that vehicles are effectively operating in Mode II, and are able to query the spectrum database directly without the need of a roadside base station or to perform local sensing. However, this is not always possible or optimal due to multiple factors: (a) phones might be able to detect RSSI values but that does not necessarily mean that they are accessing them to access the Internet. This is especially true given that a phone's carrier might represent only a subset of the total RSSI values obtained while the others are obtained from towers that the phone has no access to (e.g., while roaming), (b) having access to roadside base stations or local sensing as opposed to querying the database directly will save in the total overhead over the internet given that a BS needs to query the database once every 24 hour period. The BS can then mediate and share that information with passing vehicles, thereby saving on total traffic to the spectrum database over the internet, (c) the latency to retrieve such spectrum information is reduced because the BS is much closer to the vehicles than an online database and (d) given that the information is retrieved from the BS which queries the database once, the total cost per query incurred by each vehicle can be reduced because the BS queries the spectrum database once, and then each vehicle queries the BS for, possibly, a fraction of the total query cost.

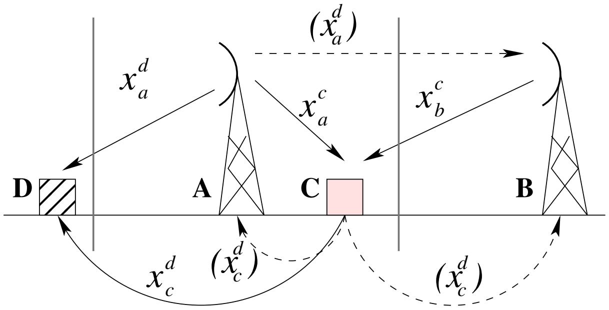

Fig. 3. Network architecture with two BSs A and B that have spectrum database access, and two vehicles C and D moving from left-right in a horizontal plane.

IV. Database Query Intervals and Method

In this section, we identify the intervals between the spectrum database updates for a vehicle moving along a one-dimensional path. The vehicles may choose to proceed with the query at these time intervals, if there is insufficient correlation among the 2G and TV spectrum at the location the vehicle is currently at, and hence, the spectrum sensing cannot be reliably undertaken.

The overall network architecture is shown in Figure 3. A and B are two BSs that have access to the spectrum database. Each BS can only provide reliable readings in a limited extent of space, shown by the boundary lines separating them. Each query to the database incurs a finite cost. CR enabled vehicles move along a straight line path from left-right, and two such nodes are indicated by C and D. When node D sends a request for new spectrum information, it sends its velocity, geo-coordinates and id. After sending this request, node D receives an update only from BS A, for an initial extent of its upcoming journey through the region served by A. This is shown by x_a^d. x_a^d includes the same node's ID, velocity and coordinates for use by tower B which can use this message via overhearing. As per the FCC ruling, a mobile node operating in Mode II must request a new set of spectrum availability readings after moving 100 m from its initial query point, or within a minute, whichever occurs earlier. Thus, a vehicle that has moved further into A's territory, beyond the first 100 m, must re-issue the spectrum query to A for the remainder of the journey. This spectrum information received by C is given by x_a^c. Similarly, BS B also updates C at this time with the spectrum occupancy in its' own control region, which C is expected to shortly enter. C evaluates, at this time, whether the local spectrum sensing performance is satisfactory by comparing the received spectrum update in the region, with the sensing results taken locally. Whether the RSSI is reliable or subject to indeterministic changes are also checked by comparing the TV channel signal fluctuations with the observed signal behavior in the cellular band. A high degree of correlation implies that the cellular channel can be a reliable metric to quantify the likelihood of sensing errors. If such an assurance is obtained, C transmits a message to D informing it that locally spectrum sensing may be adopted at this future location. This message (x_c^d) is overheard by both A and B (shown via the broken lines). If C suggests local sensing, then A and B cancel any future spectrum update for D, saving on the database query costs. In turn, when D reaches the location currently occupied by C, it re-assesses the channel using the cellular and TV signal strength correlation and issues an appropriate update to the future incoming vehicle.

V. DECISION BETWEEN LOCAL VS DATABASE SENSING

In the previous section, we discussed *when* the database query must happen (if at all). Using the architecture described in Figure 3, we describe our approach of deciding whether spectrum database access *should occur* instead of local spectrum sensing in this section at these query instants.

When a CR vehicle, say D in Figure 3 comes near the coverage area of a BS, it transmits a query including its own ID and velocity. Subsequently, this vehicle (as well as the BS) begins a timer after receiving (or transmitting) the first spectrum database information from the BS, or a notification of local spectrum sensing by another upstream vehicle. Depending on which event occurs earlier, i.e., whether a distance of greater than $100\,\mathrm{m}$ was traversed (calculated using the velocity of the vehicle), or the elapsed time exceeded $60\,\mathrm{s}$, the vehicle must seek a subsequent spectrum database update. A vehicle is such a position, say C, may receive this new update. At this point, C is tasked with an additional function: it may inform the downstream vehicle, here D, that the latter may undertake local spectrum sensing, instead of relying on the scheduled database update which comes at a finite overhead of both the query cost and the resulting additional channel usage for control information.

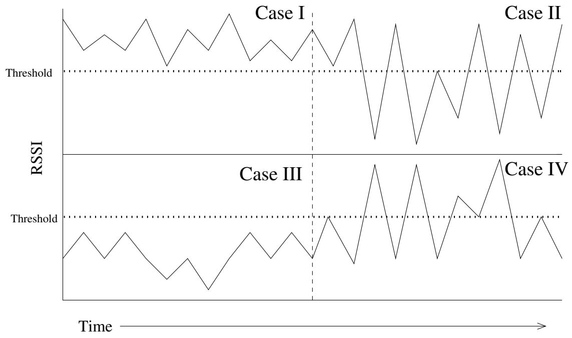

Fig. 4. Four different types of RSSI readings performed by sensing.

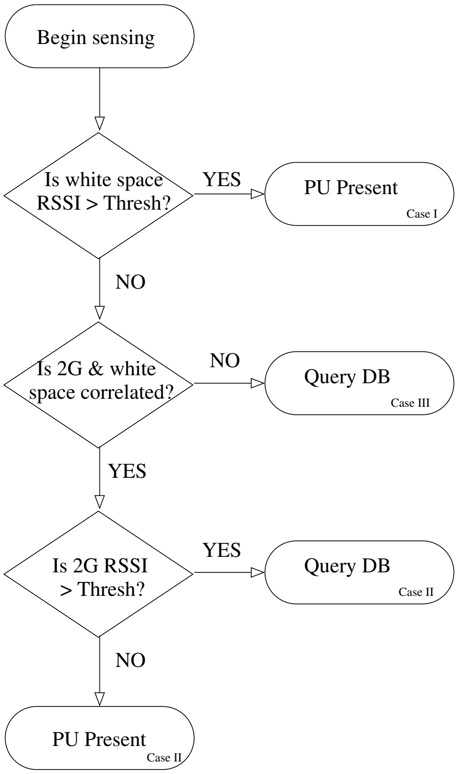

Fig. 5. Sensing vs. Querying decision flow chart

After collecting the RSSI values of the 2G and whitespace towers spectrum, the data was analyzed offline to determine the points on the map where a high correlation exists between any pair of channels in the two data sets (Figure 2(a)). Nodes that traverse that path in the future (Node C) will then leverage this information to determine with higher accuracy whether a PU does or does not exist, or if a spectrum database query is still needed due to indeterminate results and informs the downstream vehicle D of that decision. To explain this concept further, we list four cases that demonstrate the types of readings that a sensing vehicle may obtain at any given point. These cases are shown in Figure 4. We will use the flowchart in Figure 5 to step through the decision process and explain the reasoning behind it.

Referring to Figure 4: *Case I* is when a vehicle reads RSSI values that all surpass a certain threshold λ [2] on average. *Case II* is when the average drops below that threshold due to significant fading. *Case III* represents very low RSSI readings, which is the case for vacant channels. These RSSI readings are well below the threshold and finally, *Case IV* demonstrates the case where the average RSSI value is above the threshold when the channel is really vacant. Case IV's outcome, however, can be minimized because contrary to current detection schemes that have a probability of false alarm rate [2], [10], [11] due to uncertain noise power knowledge, in our proposed architecture, nodes can have perfect noise power knowledge by studying the RSSI values of channels that are known to be vacant by the database and setting the threshold to a value above the highest observed RSSI value. After minimizing the probability of false alarm in *Case IV*, vehicle C will have to make a decision based on which of the remaining three cases of RSSI readings it is seeing.

- *Case I:* Based on the flowchart in Figure 5, when a node C enters a new location, it first checks whether the white space RSSI values are above the sensing threshold λ. If they are, then the node can safely determine that the PU is present and active in this region.
- *Case II:* If the average RSSI value is below the threshold, the vehicle first checks whether a high correlation existed previously with another 2G channel. If there existed a high correlation, then the vehicle will first need to check whether the 2G spectrum is below the 2G threshold as well. This threshold λ for the 2G spectrum may be different than the threshold we use for the TV whitespace spectra. The threshold λ can be can be determined by the minimum probability of error decision rule as $f(\lambda|H_1)P_{on} = f(\lambda|H_0)P_{off}$ where $f(y|H_1)$ and $f(y|H_0)$ are probability density functions of the received RSSI through the occupied and the idle spectrum, respectively. P_{off} and P_{on} are the probability of the period unused and used by the primary user, respectively [2]. If the correlated 2G RSSI values are also below their set threshold, then the node infers that the values are fluctuating in both channels due to significant fading and the PU is therefore actively occupying this channel. If the average 2G RSSI values are above the threshold however, then the node cannot infer the channel availability because the correlated channels do not exhibit the same relative strength. In this case, the vehicle will need to query the spectrum database, if it wishes to use this channel.
- *Case III:* If there is no high-correlation detected at this point, then we infer this is Case III. Here, the vehicle cannot determine the outcome with certainty, and will need to poll the spectrum database for the exact spectrum occupancy information.

We note that $Case\ II$ is the critical point, where our network architecture leverages the spectrum information and the correlation obtained with the 2G network, and effectively minimizes the probability of miss-detection. Thus, the non-intuitive conclusion can be derived for this case that a PU exists even though the RSSI readings show otherwise. This decision process is made possible through by checking the 2G network signals at those specific locations of high correlation.

While the BSs supply information to the CRs, there is also an associated inter-BS data exchange to schedule a future spectrum update as the vehicle moves along the road. BS A lets B know when it should transmit the latter's own spectrum information next by suggesting a future transmission slot.

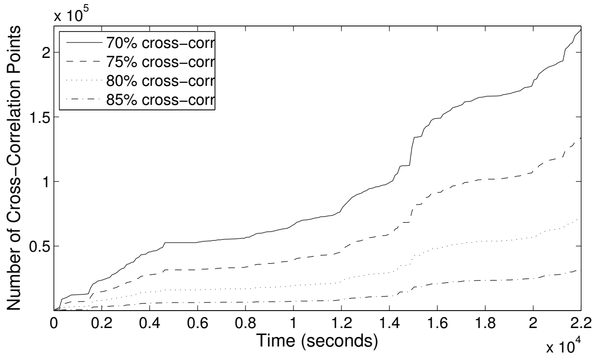

Fig. 6. Number of cross-correlation points as the time of the experiment increases

VI. PERFORMANCE EVALUATION

In this section, we describe the results from the experimental studies to validate the benefits of our approach. We focus on the number of cross-correlated points in the entire trial, and an evaluation of the efficiency and accuracy gains through our approach. The setup is similar to that described in Section III, where a USRP device within a moving vehicle emulates the CR vehicle. The Android phone accesses the 2G spectrum, maintains a history of the path traveled via the in-built GPS capability, as well as queries the Spectrum Bridge Inc. database through software APIs every 100m traversal, or when 60s elapse, whichever is earlier. As a subsequent check, all measured sensing results obtained via energy detection through the USRP are validated with offline verification through the database, at the end of the experiment. A subsequent trailing vehicle may schedule its own query, unless the vehicle ahead transmits a beacon that informs it to cancel the database query and rely on local sensing, using the procedure adapted from Section V. The observed signal fluctuations in the 2G spectrum as well as the TV whitespace are saved in a laptop computer and correlations are calculated via a continuously running MATLAB program.

A. Number of Cross Correlation points

Figure 6 shows the number of cross-correlated points as the vehicle moves along the map shown in Fig 2(a). In this figure, we can see that the number of such cross-correlated points is higher when the correlation level is set at 70%, as compared to 85%, by about 4 times. This demonstrates the trade-off that exists between higher accuracy (imposed by higher correlation demands) between the 2G spectrum and TV whitespace, and total number of cross correlation points where this approach can actually be used.

B. Number of queries

We give the total number of database queries that a vehicle engages in the path that we traversed to collect the spectrum information in Figure 2(a). As the distance traversed increases, the number of queries that have to be performed also increase. In Figure 7, we plot the number of queries that are undertaken by the CR vehicle if no cross-correlation information is leveraged, as opposed to when it utilizes the cross-correlation information to reduce the number of direct queries.

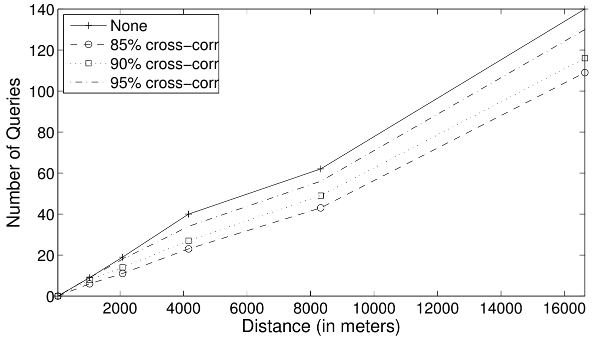

Fig. 7. Number of queries for various cross correlation percentages

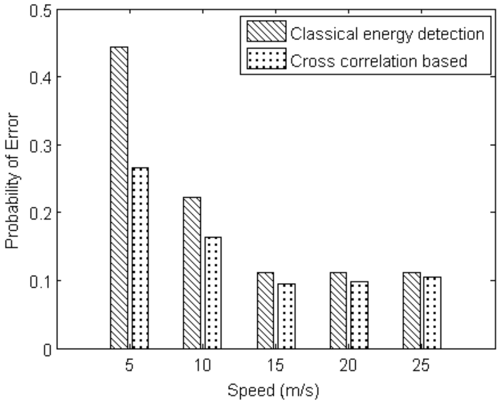

Fig. 8. Accuracy of sensing

One can observe that the lower the allowed cross-correlation percentage, the less queries that vehicle will undertake. In the case of cross correlation = 70%, the vehicle only needs to query the database 108 times instead of 140 without our proposed scheme. That's a saving of up to 23%. This number can be easily scaled with appropriate multipliers, such as the monetary cost of spectrum usage time imposed by carriers, the channel utilization, resulting data congestion, among others to obtain our tangible impacts of reducing the number of database queries.

C. Accuracy

In Figure 8, we plot the probability of error with (P_{err}^x) and without (P_{err}^c) our proposed scheme. We define P_{err} as the ratio of the number of times the query resulting in mis-detection to the total number of queries.

We first calculate P_{err}^c by assuming Rayleigh fading over a period of time for the gathered RSSI readings (e.g. Figure 9), and see whether the average RSSI values obtained over the sensing time t_s are above or below the RSSI threshold λ. Here, we chose $t_s = 0.022s$ and Y was empirically observed over the total simulation time [2] to be -0.5 (amplitude dB). We then calculate the fraction of times that the average RSSI value was below the threshold over the total number of sensing intervals to give us the classical energy detection probability of error P_{err}^c. This is in turn plotted against P_{err}^x:

$$P_{err}^x = (1 - \frac{x_j}{X}) \times P_{err}^c \tag{2}$$

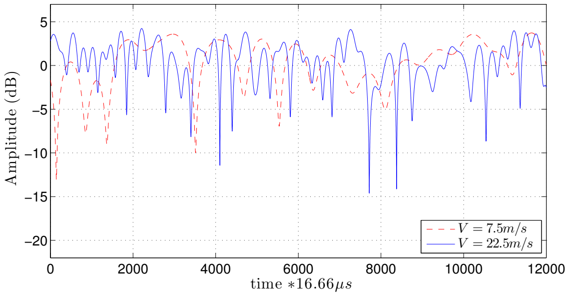

Fig. 9. Signal fluctuations at speeds of 7.5 m/s and 22.5 m/s

where x_j is the number of cross correlations that occurred at speed j, and X is the total number of cross correlations obtained over the entire experiment.

VII. CONCLUSION

In this paper, we have described a new paradigm for spectrum database access, which allows querying the database only when needed. The resulting method exploits the correlation that exists among two entirely different spectrum bands at specific locations, thereby improving the performance of local sensing and reducing the costs associated with repeated database queries. Results reveal about 23% reduction of queries, making it attractive for practical spectrum database deployments. This non-trivial approach can potentially open up a hitherto unexplored direction in spectrum sensing, and future work will be focused on building the protocol suite that enables quick and efficient exchange of spectrum data between the vehicles and the BSs.

ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation under research grant CNS-1265166.

REFERENCES

[1] FCC, Second Memorandum Opinion and Order. *ET Docket No. 04-186, DA 11-131*, Jan. 2011.

[2] W.-Y. Lee and I. F. Akyildiz, Optimal spectrum sensing framework for cognitive radio networks. *IEEE Transactions on Wireless Communications*, vol. 7, no. 10, pp. 3845-3857, Oct. 2008.

[3] Ettus Research. [Online] http://www.ettus.com/

[4] Spectrum Bridge. [Online] http://www.spectrumbridge.com/

[5] S. Pagadarai, A. M. Wyglinski and R. Vuyyuru, Characterization of vacant UHF TV channels for vehicular dynamic spectrum access. *Proc. IEEE Vehicul. Networking Conf.*, Oct. 2009

[6] K. Fawaz, A. Ghandour, M. Olleik and H. Artail, Improving Reliability of Safety Applications in Vehicle Ad hoc Networks through the Implementation of a Cognitive Network. *Proc. IEEE ICT*, Apr. 2010

[7] M. Di Felice, K. R. Chowdhury and L. Bononi, Analyzing the potential of Cooperative Cognitive Radio Technology on inter-vehicle communication. *Proc. of IFIP Wireless Days*, 2010.

[8] H. Li and D. Irick, Collaborative Spectrum Sensing in Cognitive Radio Vehicular Ad hoc Networks: Belief Propagation on Highway. *Proc. of IEEE VTC Spring*, pp. 1-5, Taipei, 2010.

[9] X. Y. Wang and P. H. Ho, A Novel Sensing Coordination Framework for CR-VANETs. *IEEE Transaction on Vehicular Technology*, vol. 59 no. 4, pp. 1936-1948, 2010.

[10] H. Urkowitz, Energy detection of unknown deterministic signals. *in Proc. of IEEE*, vol. 55, no. 4, pp. 523-531, 1967.

[11] A. Sonnenschein and P. M. Fisherman, Radiometric detection of spread spectrum signals in noise of uncertainty power. *in IEEE Trans. on Aerospace and Electronic Systems*, vol. 28, no. 3, pp. 654-660, 1992.