RISA: Distributed Road Information Sharing Architecture

Joon Ahn[†], Yi Wang[†], Bo Yu[♯], Fan Bai[‡], Bhaskar Krishnamachari[†]
[†] University of Southern California, Los Angeles, CA, USA
[♯] Wipro Technologies, Rochester, MI, USA
[‡] General Motors Global R&D, ECI Lab, Warren, MI, USA
{joonahn, wangyi}@alumni.usc.edu, {bo.yu, fan.bai}@gm.com, bkrishna@usc.edu

***Abstract*—With the advent of the new IEEE 802.11p DSRC/WAVE radios, Vehicle-to-Vehicle (V2V) communications is poised for a dramatic leap. A canonical application for these future vehicular networks is the detection and notification of anomalous road events (e.g., potholes, bumps, icy road patches, etc.). We present the Road Information Sharing Architecture (RISA), the first distributed approach to road condition detection and dissemination for vehicular networks. RISA provides for the in-network aggregation and dissemination of event information detected by multiple vehicles in a timely manner for improved information reliability and bandwidth efficiency. RISA uses a novel Time-Decay Sequential Hypothesis Testing (TD-SHT) approach in which event information from multiple sources is combined with time-varying beliefs. We describe our implementation of RISA which has been deployed and tested on a fleet of vehicles on-site at the GM Warren Technical Center in Michigan. We further provide a comprehensive evaluation of the aggregation mechanism using emulation of the RISA code on real vehicular mobility traces.**

I. Introduction

With the advent of the new IEEE 802.11p Dedicated Short Range Communications (DSRC) / Wireless Access for Vehicular Environments (WAVE) radios, Vehicle-to-Vehicle (V2V) communications is poised for a dramatic leap. Vehicular networks based on V2V communication are envisioned to be used for a broad range of safety and infotainment applications.

A canonical application for these future vehicular networks is the detection and notification of anomalous road events (e.g., potholes, bumps, icy road patches, etc.). State of the art approaches for this problem advocate the use of the centralized cellular infrastructure. However, increasing cellular utilization costs coupled with the required scale of dissemination make this an unattractive option.

We advocate in this work instead a distributed approach to road condition detection that we call the Road Information Sharing Architecture (RISA). With the distributed approaches, the system can be more flexible, being independent of more expensive and more regulated cellular networks. Moreover, the approach is well-aligned with sharing road information because the information has locality of interests so that vehicles are more interested in nearer information. In such a decentralized setting, with multiple vehicles potentially detecting each event, *it is important to aggregate the event detections within the network in a timely manner so as to improve accuracy of reports as well as network bandwidth efficiency*. We accomplish this in RISA through a novel mechanism called Time-Decay Sequential Hypothesis Testing (TD-SHT). The basic idea behind this mechanism is that each vehicle that hears a "rumor" about the event, maintains a time-decaying belief about it. Rumors from multiple vehicles are combined additively until they exceed a prescribed threshold, at which point they are converted to confirmed event reports. This threshold is determined to effect a desired tradeoff in information reliability, between the rate of false negatives and the rate of false positives. Both rumors and reports are distributed through the network in epidemic "gossip spread" fashion.

The key contributions of this study are the following:

1) We propose an in-network aggregation architecture (RISA) for road information gathering and dissemination that uses a novel time-decay sequential hypothesis testing (TD-SHT) approach in order to improve information reliability and bandwidth efficiency.
2) We implement and deploy the RISA software on a fleet of five GM vehicles equipped with DSRC/WAVE radios, and describe our experiments validating the functionality of the implementation.
3) We evaluate the performance of RISA with respect to various tunable parameters and key metrics in a comprehensive manner through code emulation based on real mobility traces collected from vehicular experiments under real driving conditions; the emulation results are compared against theoretical models.

II. Related Work

A majority of empirical studies on vehicular networks aim to provide "Drive-Through Internet" to moving vehicles via open APs deployed on the roads [1], [2], [3], [4], and are focused on optimizing the data transfer. Unlike these works, the focus of our empirical study is placed on developing an in-network data aggregation mechanism inside of vehicular networks.

There is a large body of research works on designing a participatory sensing system. One attractive example is traffic

Joon Ahn is currently with Ericsson; and Yi Wang is currently with Juniper.

monitoring application that uses a number of active probe vehicles to estimate traffic flow. The technical feasibility of this participatory sensing system was initially evaluated using embedded vehicle telematics systems ([5],[6]); realizing that the rapid growth of cellular phones could create alternative ways of implementing traffic probe system, the authors of [7], [8], and [9] proposed to use mobile phones (instead of embedded telematics systems) to detect traffic flow and predict travel time on a one-dimensional highway, while the authors of [10] suggested to do so for traffic estimation on two-dimensional surface streets. Unlike the cellular-centric data collection approach, the Cartel project [11] tries to leverage scattered deployments of hot-spots as a new "data pipe" for uploading traffic probe data. In all the above approaches, traffic probe data collected from the physical world is moved via ubiquitous wireless infrastructure to a cyber-space central server for further data processing (traffic flow estimation, traffic pattern identification, traffic anomaly detection), before the knowledge extracted from data is presented to end users [12]. In contrast, our work advocates a distributed in-network information processing protocol (including data aggregation, storage and replication) within the vehicular network, for such traffic/road monitoring applications.

Another line of work related to ours is how to further improve sensing capability in order to detect different physical phenomenon on the road accurately. The authors of [13] showed that a hybrid localization scheme using cellular and WiFi signal (in addition to GPS) could tolerate significant signal noise and/or GPS outages, achieving accurate travel time estimation. The authors of [14] and [15] used a rich set of inexpensive cellphone-equipped sensors (including accelerometers, microphone, GSM radios and GPS sensors) to detect potholes, bumps, braking and honking in less regulated traffic which is important to developing countries. The authors of [16] focused on detecting road pothole events using vibration and GPS sensors. Similarly, using GPS sensor and ultrasonic radar, the authors of [17] showed that a vehicular participatory system could discover empty road-side parking slots. (Needless to say, the accurate detection of an empty parking slot and its location is the cornerstone of such system.) Finally, via accessing fuel usage information of vehicle OBD-II systems, the authors of [18] demonstrated a participatory sensing system that improves fuel efficiency. The focus of our work being on network-level data aggregation of detected physical events among participants, our system could leverage these works that tackle the challenge of providing accurate detection results using a variety of on-board sensors.

Our work represents the first distributed, in-network data aggregation system for vehicular networks which enables participating cars to collaborate on developing a consensus about detected physical events on the road. Such an in-network system does not necessarily require the engagement of ubiquitous wireless infrastructure, thus reducing the reliance of often congested cellular networks in the major metropolitan area.

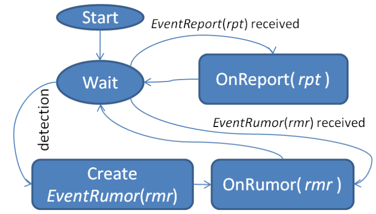

Fig. 1. RISA mechanism

III. RISA Architecture

In this section we describe our road information sharing architecture (RISA).

A. Overview

The key idea of RISA is to aggregate the event information announced by multiple sources within the delay tolerant network (DTN) to enhance both accuracy of the information and bandwidth efficiency of a vehicular ad-hoc network (VANET). The aggregation mechanism in RISA is inspired by the classical sequential hypothesis testing (SHT) technique for identifying events based on stochastic observations obtained sequentially over time [19].

In RISA, as each car encounters other cars that provide them with new event detection samples, it sequentially increases its belief about the road event until the detection becomes credible enough to become a confirmed observation of road events (similar to the traditional SHT). However, unlike SHT, the phenomenon to be detected in our setting is time-varying (e.g., a pothole may become filled, or an icy-patch may disappear due to warming), which discounts the credibility of older samples. In order to deal with this property, we allow for *belief value* of events to be decreasing over time; and sequentially add the time-decayed beliefs corresponding to encountered detection samples until they cross a predefined threshold. This time-decay SHT (TD-SHT) mechanism allows for event confirmations that are both reliable and time-sensitive.

RISA-capable vehicles monitor their environment for the interested road events. When a vehicle detects a road event, it creates and stores an *EventRumor* which contains the information about the detection sample for the event. It also broadcasts the rumor to its direct neighbor vehicles. When a vehicle meets another RISA vehicle, they transfer their knowledge of road events to each other, and the knowledge includes *EventRumor*s. When a vehicle gathers enough rumors of a certain road event, it produces an *EventReport* as a conclusion on the event by merging the rumors. The *EventReport* is propagated to a much larger set of vehicles, giving notification to drivers/passengers in the vehicles. While an individual *EventRumor* is treated as inconclusive, an *EventReport* is considered trustworthy information to report to people.

The high-level distributed algorithm for each vehicle of RISA is depicted in Figure 1. The details of each algorithm introduced in the figure are presented in the next subsections. The definitions of symbols used in the algorithms are presented in Table I.

Symbols	*Description*
$\mathrm{bv}(s,t)$	the belief value of s at time t
bv_{min}	the minimum belief value
bv_{tot}	the total belief value
th_{css}	the consensus threshold
$find\_event\_report(r)$	the function that returns the event report corresponding to rumor r

TABLE I
DEFINITIONS OF SYMBOLS USED IN ALGORITHMS

B. Road Event

In the RISA architecture, we are interested in discrete *road events* such as traffic congestion, potholes, ice patches, traffic light malfunction, obstacles on the road, etc. Road events are classified in our RISA system by the tuple: $<$**type**, **location**, **location precision**$>$. They are first classified based on their types – such as pothole or congestion – and the location points. The data structure of RISA does not explicitly distinguish two temporally disconnected road events on the same location, as long as they are of same type. However, since the focus of RISA is on providing time-sensitive event reports, information about sufficiently old events does not persist in the system due to the time-decaying beliefs adopted in our TD-SHT mechanism.

C. Core Mechanism

1) Aggregation Algorithm: We first introduce metrics to evaluate the credibility of information – *belief value* and *total belief value*.

The *belief value* is the amount of credibility of the road-event information in its associated *EventRumor* or *EventReport* at the given time. When a new *EventRumor* is created by a vehicle that detects a road event, a predetermined *initial* belief value b_0 is assigned to the rumor. As time goes, the belief value of the rumor is changed following a predetermined *decay function* in order to discount the aged information. Although the decay function may be any non-increasing function of the elapsed time from the creation of the rumor, we focus on the exponentially decreasing function in Section VII.

The *total belief value* is the aggregated belief value of a set S of *EventRumor*s at the given time. The total belief value represents how much one can believe that the associated event is true at the given time after obtaining information from multiple sources. We use the sum of individual belief values $\mathrm{bv}(s_i)$ of *EventRumor*s s_i for aggregation in this paper as given by Equation (1). But, our architecture can support an arbitrary function mapping a set of belief values into a scalar.

$$\mathrm{bv}_{tot}(S,t) = \sum_i \mathrm{bv}(s_i,t) \tag{1}$$

Now let us look at the aggregation algorithm MAKECONSENSUS which is the core algorithm of RISA architecture. The algorithm is the consensus mechanism for combining inconclusive event detections represented by *EventRumor*s from different nodes at different times into one trustworthy conclusion – *EventReport*. As can be seen in Algorithm 1, its basic operation is to evaluate the total belief value of the given road event and compare it with the predefined consensus threshold. If the total belief value of rumors is larger than the

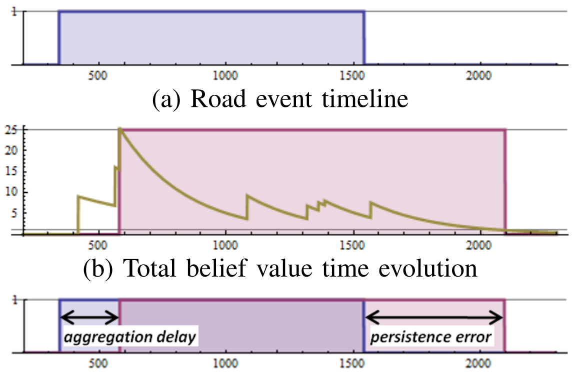

(a) Road event timeline

(b) Total belief value time evolution

(c) Timeline overlap of road event and RISA perception

Fig. 2. Belief value interpretation

threshold, it proclaims a conclusion on the considered road event and creates an *EventReport* for the event. Otherwise, the RISA system adds the rumor into its corresponding set and uses it in the future. The consensus conclusion is honored until its belief value falls below the minimum belief value. After that, the system abandons the report and regards the event as having disappeared. If a RISA car receives an *EventReport* from another RISA car when it has not reached a consensus, it acknowledges the road event as truthful and discards all the associated rumors in its storage. In this way, the distributed algorithm of RISA can converge quickly.

Figure 2 illustrates this aggregation mechanism in the time range of [300, 600] seconds. The data in the figure is from one of our experiments in Section VII. Figure 2.(a) shows the timeline of the road event, where value 1 means the occurrence of the event. Figure 2.(b) shows the time evolution of the belief values of the road event in a given car, and its perception of the event. This subfigure shows that the total belief value steps up with the detections or receptions of rumors while it decays otherwise. At $t = 579.2$, the value exceeds the consensus threshold (i.e. 25 in this case) leading to a consensus; the car acknowledges the road event from this moment.

Figure 2.(c) overlaps the timeline of the real occurrence of the road event and the perception of the car. We can see two lags – aggregation delay and persistence error. The aggregation delay is between the start time of the road event and the perception, and the persistence error is between the end of the road event and that of perception.

2) Phased Operation: The RISA mechanism can be divided into two phases for each road event, and the transition is triggered by the creation of an *EventReport*.

In Phase 1, when a RISA-capable vehicle detects a road event, it creates an *EventRumor* for the event and executes MAKECONSENSUS with the rumor and the set of previous rumors for the event in its storage, if the set exists. If the aggregation algorithm successfully makes consensus and produces an *EventReport*, RISA goes into Phase 2.

If the vehicle receives *EventRumor*s from another vehicle, it executes MAKECONSENSUS with them one-by-one. If it receives the *EventReport* of the road event, it keeps the report and refreshes its belief value with the largest belief value of its own rumors of the event if the latter is larger than the given

Algorithm 1 MAKECONSENSUS(rmr, S)

S := the set of *EventRumor*s of same road event as rmr in the storage
if S has already included rmr **then**
  return NullConsensus;
t := current time;
for each element s_i of S **do**
  if $\text{bv}(s_i, t) < \text{bv}_{min}$ **then**
    remove s_i from S;
Add rmr into S;
$\text{bv}_{tot} := \sum_i \text{bv}(s_i, t)$;
if $\text{bv}_{tot} > th_{css}$ **then**
  $rpt := event\_report(event\_info(rmr), \text{bv}_{tot}, t, node\_id)$;
  Remove S from the system;
  return rpt;
else
  return NullConsensus;

Algorithm 2 ONRUMOR(rmr)

rmr := new *EventRumor*;
$rep := find\_event\_report(rmr)$;
if rep is not empty **then**
  t := current time;
  $\text{bv}(rep) := max(\text{bv}(rep, t), \text{bv}(rmr, t))$;
  return
else
  S := local set of *EventRumor*s equivalent to rmr;
  $c := MakeConsensus(S, rmr)$;
  if c is a valid *EventReport* **then**
    save c in the node;
  return

Algorithm 3 ONREPORT(rep)

rep := new *EventReport*;
t := current time;
$rep_{local} := find\_event\_report(rep)$;
if rep_{local} is not empty **then**
  $\text{bv}(rep_{local}) := \max(\text{bv}(rep_{local}, t), \text{bv}(rep, t))$;
  return
else
  S := local set of *EventRumor*s equivalent to rmr;
  if S is not empty **then**
    Update S with decay function;
    Remove s_i from S if $\text{bv}(s_i, t) < \text{bv}_{min}$ for all s_i in S;
    $\text{bv}(rep) := \max(\max_i(\text{bv}(s_i, t)), \text{bv}(rep, t))$;
    Discard S;
  save rep in the storage;
  return

belief value of report. It does not add all individual belief values of the rumors to the report's value because it can make the report's belief value too high to make it stay too long in the system, potentially giving stale information to participants.

A pair of RISA vehicles exchange all rumors that exist in only one of them upon their encounter, which indicates that multiple rumors of the same event spread over the network without aggregation in this phase. However, the spread is expected to be confined near the origin of road event because all rumors are abandoned once they produce an *EventReport* or they meet the *EventReport* of the same event.

This phase may use more resources such as communication bandwidth and storage due to the lack of aggregation, but this is helpful to obtain more accurate aggregation at the end, avoiding double-counting of same rumors. In fact, we can improve the efficiency of Phase 1 by adopting an additional aggregation technique (*e.g.* FM-Sketch [20]) that is robust to double-counting even in this level. But, we focus on its basic functionality in this paper.

In Phase 2, the vehicle has only one *EventReport* for the road event, and propagates it to other vehicles until it becomes expired. When the vehicle receives another *EventReport* of the same event from another vehicle, it refreshes the belief value of its own report by updating it with that of the received report, if the latter is bigger than the former. It discards the received report after that as the way of aggregation. The procedure is similar when the vehicle receives an *EventRumor* – it refreshes the belief value of its own report and discards the rumor. These behaviors could be seen in Figure 2.(b) in [1000s, 1700s] interval. In this way of aggregation, the system saves the bandwidth and storage by keeping only one report and no rumors for a road event. When the report belief value falls below the minimum belief value, the system discards the report and regards the road event as being disappeared as in Figure 2.(b) at $t = 2097.18$.

These double-phased operations are controlled by ONRUMOR and ONREPORT. Their pseudo codes are presented in Algorithms 2 and 3.

D. Tunable Parameters

The RISA architecture has several tunable parameters. In this section, we introduce them and describe their first-order contributions.

Consensus Threshold (th_{css}): This is the minimum total belief value such that its associated set of *EventRumor*s is to be promoted to an *EventReport* by being recognized as trustworthy and conclusive information. This value contributes to timing of phase transition in RISA.

Initial Belief Value (b_0): This is the initial and maximum possible value for an individual *EventRumor*. This value affects the lifetime of the *EventRumor* and the consensus timing.

Minimum belief value (b_m): This value is the smallest value of belief of an event rumor/report for them to survive and be propagated over the network. If an event rumor/report has the belief value decayed to less than this value, it is considered expired and will be discarded.

Decay function ($d(\tau)$): This function decides how the belief value decays over time, affecting the consensus timing and the lifetime of event rumors and reports. The decay function can be different between *EventRumor* and *EventReport* even for the same type of road event.

The initial and minimum belief values and decay function jointly decides the lifetime of *EventRumor*s, given by Equation (2). The lifetime of *EventReport*s with no refresh is decided with one more parameter which is the consensus threshold, given by Equation (3). However, its overall lifetime

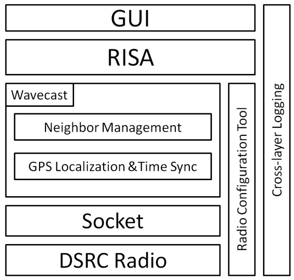

Fig. 3. RISA Software System

also depends on the refresh process, which is hard to derive because it depends on the random process of detections and the communication delay from the detector to the report holder.

$$\tau_{rumor} = \arg\min_{\tau} (\max\{b_0 d(\tau) - b_m, 0\}) \quad (2)$$

$$\tau_{report} < \arg\min_{\tau} (\max\{(th_{css} + b_0) d(\tau) - b_m, 0\}) \quad (3)$$

We note that the parameters are not independent, but the same effects can be drawn by fixing some parameters and adjusting others properly. For example, the same lifetime can be achieved for a rumor by doubling the initial belief value, when the minimum belief value is also doubled. The effects of these parameters are further investigated with the results of performance evaluation in Section VII.

IV. Implementation

We developed a prototype system composed of hardware and software in order to experiment and evaluate the RISA architecture. This prototype system was installed on a fleet of five GM testing vehicles for empirical experiments in real driving environments.

A. Software and Hardware System

Hardware System: Each of the GM research vehicles that we used for the experiment is equipped with a Linux laptop running our prototype RISA software, a DSRC-compatible Denso WRM radio operating the IEEE 802.11p WAVE BSS mode on the 5.9 GHz frequency band, transmitting at 20 dBm, 6 Mbps data rate, with an omnidirectional 0 dB gain antenna mounted on the roof, and a GPS unit with 1m accuracy, used for location and time synchronization. Note that our testing vehicles are not equipped with real event detection sensors, as we do not focus on the local event detection problem in this paper. Instead we use virtually generated event detections.

Software System: The software system is built on top of a customized software system called *WaveCast* [21], which is used to support vehicular opportunistic communication research. It consists of the entire DSRC protocol stack (up to network-layer), based on GrooveNet [22]. As a modularized test platform, *WaveCast* is flexible enough to support the development of a variety of vehicular communication applications, using standard network-layer APIs. The architecture of our software system is depicted in Figure 3.

Our software system also provides a set of Graphic User Interfaces, which are able to display a map, vehicle dynamic information, neighboring vehicles, road events (such as RISA events), as well as radio and GPS status information. Our software system is composed of about 50,000 lines, of which about 12500 lines correspond to RISA module.

Our developed software system is able to run in two modes: real experiment mode and emulation mode.

Real Experiment mode: On top of *WaveCast* system, we implemented the RISA mechanism introduced in the previous section. Each testing vehicle is equipped with its own *WaveCast* system and RISA protocol (as a single node). Using 5 such testing vehicles as a small-scale testbed, we are able to run experiments with the full functionalities of RISA protocol in the real driving environments.

Emulation mode: We also modified our *WaveCast* software system to support the emulation mode that support multiple cars to run RISA protocol on a single machine, as other traditional simulators/emulators. In addition we added a vehicle mobility module into our software system that takes GPS traces as input and moves cars accordingly in the simulation mode. We use this feature for RISA emulation as discussed in Section VI. Note that the exact same RISA code is used for both real experiments and the emulation mode, so that we believe that our emulation results using empirical vehicle traces are trustworthy.

V. Real Experiment

We first run a representative set of real experiments with our prototype system on vehicles to demonstrate the correctness of the architecture and its software implementation. Figure 4 shows the results of two samples of the experiments. In the experiments we had three cars, whose paths are shown as colored curves, and one pothole road event (shown as a star).

The first experiment *Exp1* is the case in which the cars that detect the road event make consensus, and let other cars know. As shown in Figure 4.(a), Car 1 (red) detects the pothole at 03:40.522 and broadcasts the *EventRumor*. Car 2 (yellow) that follows Car 1 within the radio range receives the rumor at 03:41.501 and keeps it. Soon it also detects the pothole and creates an *EventReport* for it because the total belief value exceeds the consensus threshold. It broadcasts the report and Car 1 receives it at 04:16.261, which is transferred to Car 3 (blue) at 04:30.863. Figure 4.(b) shows how the perception of each car for the pothole changes over time in *Exp1*. Although the perception is either on (1) or off (0), we put the state of 0.5 in the plot to identify the time duration that the car has rumors only. In the figure we can see that Car 2 receives the rumor Car 1 created almost immediately because they are within the radio range at that time. But, Car 1 is out of range when Car 2 creates the report so that we can see a time lag between the time Car 2 confirms the pothole and the time of Car 1. Because Car 3 receives the report without any rumors, its perception goes from 0 to 1 directly.

Figures 4.(c) and (d) pertain to the second experiment *Exp2*, which shows a case when the consensus is made in the car (Car 3) that does not detect the road event at all. This happens when the cars that detect the event (Car 1 and Car 2) keep out of the radio range of each other, but encounter the same other car (Car 3) at later times.

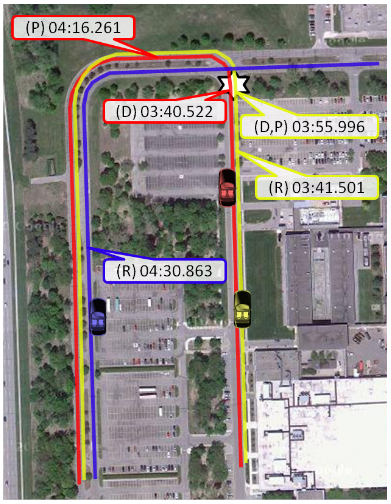

(a) Exp1-Map

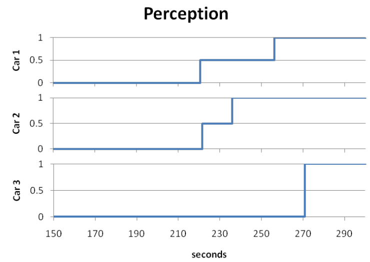

(b) Exp1-Perception

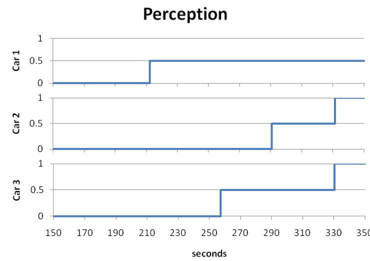

(d) Exp2-Perception

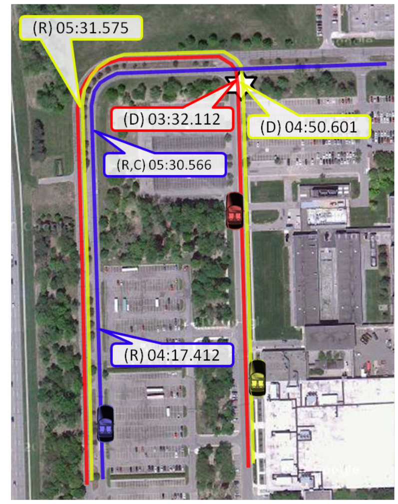

(c) Exp2-Map

Fig. 4. Real Experiments with RISA

VI. Trace Collection Experiments and Evaluation Methodology

Having already validated the functionality of the RISA protocol on real vehicles as described in section 5, we now wish to comprehensively evaluate the RISA system over many different parameter configurations. For this purpose, we did not run the RISA protocol directly on the cars; instead, we have emulated the RISA protocol offline based on the above-mentioned mobility traces from vehicles driving through real roads. Because we are more interested in the early phase of smart vehicle deployment, we focus on the low density of RISA cars, which make them operate as a delay tolerant network.

We used a fleet of five GM research vehicles for the experiment to collect movement traces and interactions between cars in real traffic. We drove the fleet of cars circling on Mound Rd, Warren, MI, USA from 12 Mile Rd to 14 Mile Rd, which is one of the major streets in the region carrying a large traffic volume.

We chose to circle within the same street to include the interaction between cars in opposite directions, with a small number of research cars. We conducted the experiments in high traffic time (*i.e.* 7am-9am, Oct. 17th, 2010) and low traffic time (*i.e.* 1pm-3pm, Oct, 18th, 2010), making for 20 driving hours in total (across all vehicles). The vehicles are inserted one-by-one into the real road traffic with an initial inter-vehicle time interval of 90 seconds. We did not control manually the distance between cars, but let each car follow the real traffic accordingly. So, the experiment experienced the inversion of the order of cars sometimes. During the experiment, we collected the GPS traces of the vehicles; each car sampled the GPS location every second.

In our evaluations described in Section VII, we assume that a pair of cars can communicate with each other if their distance is less than a predefined radio range. We use 200 meters for the radio range, which is more conservative than suggested by [23], [24]. In the emulation we have injected virtual road events on the road. The position of the event is uniform at random over the experimented street, and the arrival time of the road events follows Poisson distribution with the mean of 180 seconds, unless stated otherwise.

VII. Evaluation

We evaluate empirically the RISA system in this section, by examining the effect of various RISA system parameters on two major performance metrics – aggregation delay and consensus accuracy.

A. Aggregation Delay

In this section we evaluate the consensus delay of RISA mechanism. We will see how the delay is related with tunable parameters such as consensus threshold, initial and minimum belief values, and decay function; and untunable parameters such as the arrival rate of RISA cars to the event region and the miss probability of the local detection.

1) Evolution of Total Belief Value: The aggregation delay is dependent on the instantaneous total belief value of rumors. It is the duration until the total belief value crosses the consensus threshold in the first time from the time the road event have occurred.

We have looked into the evolution of the total belief value in this empirical study. Figure 5 visualizes how the value evolves as time goes under different rumor lifetime regimes. Although the plots are from a set of realizations of the random process, it seems that the total belief value is upper-bounded by a finite value no matter how long the process continues, except the no decay case.

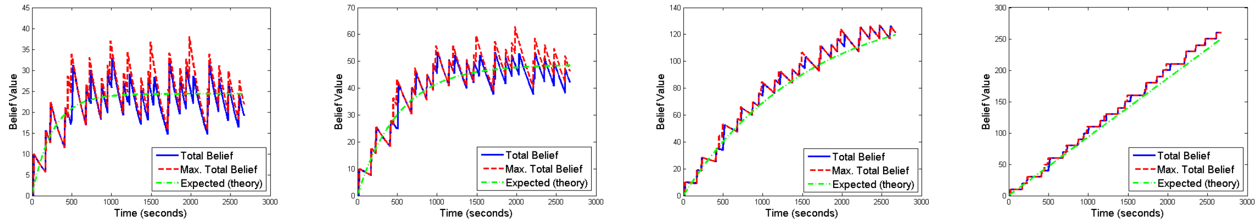

(a) $\tau_{rmr} = 10$ mins (b) $\tau_{rmr} = 20$ mins (c) $\tau_{rmr} = 60$ mins (d) $\tau_{rmr} = \infty$

Fig. 5. Time evolution of total belief value of *EventRumor*

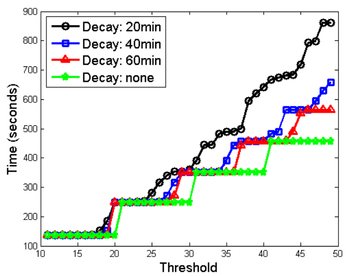

Fig. 6. Aggregation Delay

In fact, the expected total belief value turns out to have the closed-form expression as in Equation (4).

$$\mathcal{B}_{tot}(\tau) = \begin{cases} \lambda_D b_0 \gamma^{-1}(1 - e^{-\gamma\tau}), & \text{if } \gamma > 0 \quad \text{(4a)} \\ \lambda_D b_0 \tau, & \text{if } \gamma = 0 \quad \text{(4b)} \end{cases}$$

where λ_D is the arrival rate of RISA cars, b_0 is the initial belief value and γ^{-1} is the decay function. It can be easily derived with the theorem in Appendix and the time evolution function of the belief value of an individual rumor $g(t) = b_0 e^{-\gamma t}$.

The equation indicates that the expected total belief value is always upper-bounded by, and converges to, the finite value $\lambda_D b_0 \gamma^{-1}$ in the practical cases where $\gamma > 0$. Hence, if the consensus threshold is set much above this value, the system would suffer from excessively large consensus delays, and consensus failures would start to occur even if there are no detection misses from individual vehicles.

Note that the evolution process is determined by the detection arrival process given system parameters, because there is no randomness in the system except the arrival process.

2) Average Delay: Figure 6 shows the delay in terms of consensus threshold in several cases of rumor lifetimes. The initial belief value is set to 10 and the minimum belief value is set to 1 while the decay rate is adjusted to make the lifetime as in the figure legend. We only vary the threshold values because it gives scaling effects on the result to vary the values of other contributing tunable parameters, as long as the rumor lifetime is the same.

As can be seen in the figure, the delay plot looks like a step function when the belief value of rumors never decays. If it does not decay at all, the relationship between the consensus threshold and the initial belief value can be interpreted so that the system requires a certain number of rumors, which is $\lceil th_{css}/b_0 \rceil$, to conclude the consensus regardless of their freshness. So, the expected delay of consensus is the expected time duration between detections multiplied by the required number of rumors, given by Equation (5). This is why the plot has step width of $10 = b_0$ and the average step height of $106.95 \approx 1/\lambda_D$.

$$E[\tau_{css} | \text{no decay}] = \lceil th_{css}/b_0 \rceil / \lambda_D \quad (5)$$

As the lifetime of a rumor decreases, the step shape gets smoother and smoother to get more like linear function. This is because the probability that you need more rumors is increasing as the threshold approaches each multiple of b_0. And we can notice that the delay is upper bounded by a linear function whose slope depends on the rumor lifetime.

The shorter rumor lifetime also incurs significant increase in the delay in general. It is because the rumors are more likely to die out too early, which prevents useful information from spreading over the network. Although we omit the details due to shortage of space, the aggregation delay profile is similar when we also take into account the possibility of missed detections.

B. Consensus Accuracy

In this section we look into the accuracy of consensus. We examine the rate of false *EventReport* generation for the false positive in accuracy, and the probability of consensus failures for the false negative. While we use probability as the metric for the false negative, we use the occurrence rate for the false positive because *EventReport*s are generated in the continuous time domain; hence, probability is not defined well in this domain.

1) False Positives of Consensus: In order to examine the false positive of the system, we emulate false detections by injecting false road events into each car independently. The inter-injection duration follows the Poisson distribution with varying rates. Higher rate represents higher detection errors made by low-quality sensors and local detection mechanism.

Figure 7 shows the consensus false positive rate, that is the network-wise number of false *EventReport* generation in a minute (Hz/60), caused by accumulated false detections of multiple cars in the similar region. The red line labeled as *no_aggr* shows the performance when there is no aggregation mechanism in the system; each local detection is regarded as perfect and shared in the network. Other curves show the performance of RISA with different parameters; various rumor lifetimes are examined in Figure 7(a) with a fixed consensus threshold, $th_{css} = 15$, while Figure 7(b) considers various consensus threshold values with fixed rumor lifetime $\tau_{rmr} =$

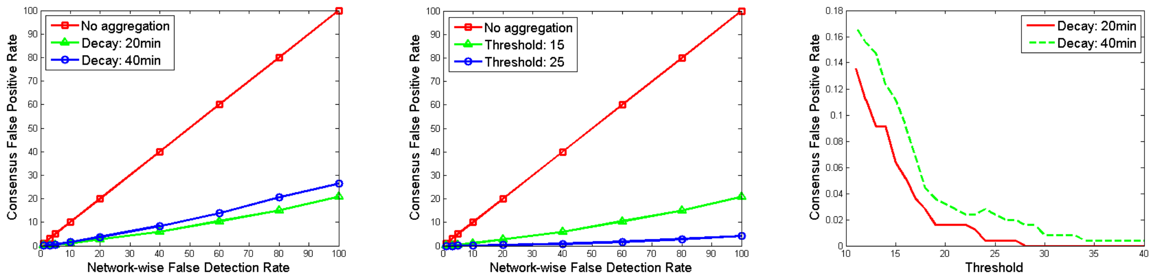

(a) $th_{css} = 15$ (b) $\tau_{rmr} = 20$ mins (c) Net-wise false detection rate $= 5$ Hz$/60$

Fig. 7. Consensus false positive rate (Hz/60)

20 mins. Figure 7(c) shows in more detail how the consensus false positive rate changes in terms of the consensus threshold value under different rumor lifetime regimes.

From this set of experiments, we can see that the RISA system can filter out most of individual false detections due to the inevitable sensor imperfection, thus improving the accuracy greatly. The resultant decrease of false *EventReport* can save the network bandwidth and the storage in each vehicle that would be consumed to disseminate and store the report, otherwise. We can also see that the relatively small value of consensus threshold is enough to make the consensus false positive very low although a larger threshold value gives better results.

2) False Negatives of Consensus: In order to examine the probability of consensus misses, we let each of the vehicles miss the detection with the varying miss probabilities p_{miss} in the emulation.

Figure 8 summarizes the experiment results; its plots show the probability that the system eventually fails to conclude a consensus over a given road phenomenon, in terms of the local detection miss probability of an individual vehicle. We have fixed the consensus threshold to be $th_{css} = 25$ in Figure 8(a) while fixing the rumor lifetime as $\tau_{rmr} = 20$ mins in Figure 8(b).

The red lines in the figure correspond to non-RISA mechanisms that have no aggregation of information independently gathered by individual vehicles. Other curves visualize the performance of the RISA architecture with different parameter values. As can be seen, RISA radically eliminates the complete detection failure in the reasonable detection misses of individual cars: $p_{\text{miss}} < 0.2$. Its performance is better with smaller consensus threshold value and larger rumor lifetime, as opposed to the case of false positive of consensus. The latter benefits from larger threshold values and smaller rumor lifetime are discussed in VII-B1.

VIII. Conclusion

In this paper, we have proposed a distributed architecture, called *Road Information Sharing Architecture* (RISA), in order to exploit the wisdom of crowd in the setting of vehicular delay tolerant networks for monitoring the road environment and sharing the information among participants. Among its core mechanisms are the aggregation mechanism, which is inspired by the traditional sequential hypothesis testing (SHT) in statistics, to combine potentially inconsistent sets of detection information from multiple cars. The main difference of our aggregation mechanism from SHT is that ours has a notion of time decay when dealing with detection samples by discounting old information.

We have implemented the RISA architecture and tested it on the streets around the GM R&D campus in Warren, Michigan, on a fleet of research vehicles. In order to comprehensively evaluate the performance of RISA, we collected GPS traces from five research vehicles, and emulated the RISA mechanism using the same code, with the traces for a wide range of parameter configurations.

We have examined the aggregation delay, the delay to capture the disappearance of road event, and the perception accuracy of RISA system in terms of false negatives and false positives. From the evaluations, we have found that (a) the low consensus threshold (which means a small number of detection samples) is sufficient to achieve high perception accuracy of detected road events in terms of both false negatives and positives; (b) the RISA makes the perception accuracy much higher than detections made by individual cars, which can be interpreted that lower-quality lower-cost sensors can be used in individual cars for the given accuracy requirements; and (c) the aggregation delay is reasonably small – an order of several minutes (with the aforementioned low consensus threshold). It is clear that the idea of aggregating on-the-road observations made by a number of individual vehicles in a distributed manner is technically feasible.

From both information reliability and delay perspectives, intuitively the RISA mechanism performance should improve with a higher density of cars. From this perspective, our study showing decent performance with only 5 cars suggests that the RISA architecture is a promising approach. The quantification of bandwidth usage and scalability of RISA mechanism are topics of future work.

References

[1] J. Ott and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for Automobile Users," in *IEEE INFOCOM 2004*.

[2] J. Eriksson, H. Balakrishnan, and S. Madden, "Cabernet: Vehicular Content Delivery Using WiFi," in *ACM MobiCom 2008*.

[3] T. Zahn, G. OShea, and A. Rowstron, "Feasibility of Content Dissemination Between Devices in Moving Vehicles," in *ACM CoNEXT 2009*.

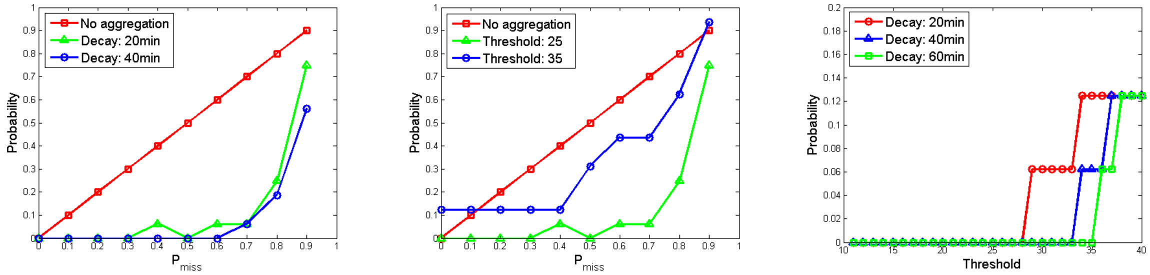

(a) $th_{css} = 25$ (b) $\tau_{rmr} = 20$ mins (c) $p_{\text{miss}} = 0.1$

Fig. 8. Probability of consensus failure

[4] D. Hadaller, S. Keshav, T. Brecht, and S. Agarwal, "Vehicular Opportunistic Communication: Under the Microscope," in *ACM MobiSys 2007*.

[5] X. Dai, M. Ferman, and R. Roesser, "A simulation evaluation of a real-time traffic information system using probe vehicles," in *In Proceedings of the IEEE Intelligent Transportation Systems*, 2003.

[6] M. Ferman, D. Blumenfeld, and X. Dai, "A simple analytical model of a probe-based traffic information system," in *In Proceedings of the IEEE Intelligent Transportation Systems*, 2003.

[7] C. Claudel and A. Bayen, "Guaranteed bounds for traffic flow parameters estimation using mixed lagrangian-eulerian sensing," in *In Allerton Conference on Communication, Control, and Computing*, 2008.

[8] C. Claudel, A. Hofleitner, N. Mignerey, and A. Bayen, "Guaranteed bounds on highway travel times using probe and fixed data," in *In 88th TRB Annual Meeting Compendium of Papers*, 2009.

[9] B. Hoh, M. Gruteser, R. Herring, J. B. A, D. Work, J. carlos Herrera, R. Bayen, M. Annavaram, and Q. Jacobson, "Virtual trip lines for distributed privacy-preserving traffic monitoring," in *ACM ACM MobiSys 2008*.

[10] J. Yoon, B. Noble, and M. Liu, "Surface street traffic estimation," in *ACM MobiSys 2007*.

[11] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden, "Cartel: a distributed mobile sensor computing system," in *ACM SenSys 2006*.

[12] J. Goldman, K. Shilton, J. Burke, D. Estrin, M. Hansen, N. Ramanathan, S. Reddy, V. Samanta, M. Srivastava, and R. West, "Participatory sensing: A citizen-powered approach to illuminating the patterns that shape our world," in *White Paper by Center for Embedded Networked Sensing*.

[13] A. Thiagarajan, L. Ravindranath, K. Lacurts, S. Toledo, J. Eriksson, S. Madden, and H. Balakrishnan, "Vtrack: Accurate, energy-aware road traffic delay estimation using mobile phones," in *ACM SenSys 2009*.

[14] R. Sen, B. Raman, and P. Sharma, "Horn-ok-please," in *ACM MobiSys 2010*.

[15] U. M. Smartphones, P. Mohan, V. N, and R. Ramjee, "Nericell: Rich monitoring of road and traffic conditions," in *ACM MobiSys 2008*.

[16] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan, "The pothole patrol: Using a mobile sensor network for road surface monitoring," in *ACM MobiSys 2008*.

[17] S. Mathur, T. Jin, N. Kasturirangan, J. Ch, W. Xue, and M. Gruteser, "Parknet: Drive-by sensing of road-side parking statistics," in *ACM MobiSys 2010*.

[18] R. K. Ganti, N. Pham, H. Ahmadi, S. Nangia, and T. F. Abdelzaher, "Greengps: a participatory sensing fuel-efficient maps application," in *ACM MobiSys 2010*.

[19] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," in *Ann. Math. Stat.*, vol. 19, ch. 3, pp. 326–339.

[20] P. Flajolet and G. Nigel Martin, "Probabilistic counting algorithms for data base applications," *Journal of Computer and System Sciences*, vol. 31, no. 2, pp. 182–209, 1985.

[21] B. Yu and F. Bai, "ETP: Encounter Transfer Protocol for Opportunistic Vehicle Communication," in *IEEE INFOCOM 2011*.

[22] R. Mangharam, D. Weller, R. Rajkumar, P. Mudalige, and F. Bai, "GrooveNet: A Hybrid Simulator for Vehicle-to-Vehicle Networks," in *In Third Annual International Conference on Mobile and Ubiquitous Systems: Networking Services*, 2006.

[23] F. Bai and H. Krishnan, "Reliability analysis of dsrc wireless communication for vehicle safety applications," in *IEEE ITSC 2006*.

[24] K. Hong, D. Xing, V. Rai, and J. Kenney, "Characterization of dsrc performance as a function of transmission power," in *ACM VANET 2009*.

APPENDIX

Theorem A.1. *Suppose the arrivals of detections follow the Poisson process with rate λ_D and $g(t)$ is the time evolution function of an individual rumor such that it gives the belief value of the rumor at time t given that the rumor is created at $t = 0$, i.e. $g(0) = b_0$. Then, the expected maximum total belief value at time t is given by,*

$$\mathcal{B}_{tot}(t) = \lambda_D \int_0^t g(\tau) d\tau. \tag{6}$$

Proof: The maximum total belief value at time t, $B_{\text{tot}}(t)$, can be expressed as follows:

$$B_{\text{tot}}(t) = \sum_{k=1}^{\infty} B_k(t) \tag{7}$$

where $B_k(t)$ is the belief value of the rumor created by k-th detection, at time t. Note that it is maximum because the expression considers all rumors generated by the time. This is different from the total belief value of individual vehicles because the latter may not have received some rumors in time. If k-th detection has not occurred by the time t, $B_k(t') = 0, \forall t' \geq t$. Because of the linearity of expectation, the expected value $\mathcal{B}_{\text{tot}}(t)$ becomes as follows:

$$\mathcal{B}_{\text{tot}}(t) \doteq E[B_{\text{tot}}(t)] = \sum_{k=1}^{\infty} E[B_k(t)] \tag{8}$$

The expected k-th belief value $\mathcal{B}_k(t)$ is given by

$$\mathcal{B}_k(t) \doteq E[B_k(t)] = \int_0^t f_{T_k}(\tau)\, g(t-\tau) d\tau \tag{9}$$

where T_k is the random variable that represents the time of k-th detection, and $f_{T_k}(\cdot)$ is its probability density function (PDF). Because the detection process is Poisson with rate λ_D, T_k follows Erlang distribution with rate λ_D and shape k. Hence, its PDF is as follows:

$$f_{T_k}(\tau) = \frac{{\lambda_D}^k \tau^{k-1} e^{-\lambda_D \tau}}{(k-1)!} \tag{10}$$

Substituting Equation (10) into Equation (9), again substituting into Equation (8), we have the following:

$$\begin{aligned} \mathcal{B}_{\text{tot}}(t) &= \sum_{k=1}^{\infty} \mathcal{B}_k(t) = \int_0^t g(t-\tau) \sum_{k=1}^{\infty} f_{T_k}(\tau) d\tau && (11) \\ &= \lambda_D \int_0^t g(t-\tau) d\tau = \lambda_D \int_0^t g(\tau) d\tau && (12) \end{aligned}$$

The last equation in Equation (11) involves the interchange of limit and integral, which holds because $f_{T_k}(\tau)$ is uniformly convergent on $[0, t]$. The proof of the uniform convergence is omitted in this paper because it is not very difficult to derive. The first equation in Equation (12) is because $\sum f_{T_k}(\tau) = \lambda_D$. ■