

# Modeling and Experimenting with Vehicular Congestion for Distributed Advanced Traveler Information Systems

Marco Rocchetti\* and Gustavo Marfia

Computer Science Department, University of Bologna  
Mura Anteo Zamboni 7, 40126 Bologna, Italy  
{roccetti,marfia}@cs.unibo.it  
<http://www.cs.unibo.it/>

**Abstract.** Advanced Traveler Information Systems, which have for long been regarded as one of the most promising future applications of wireless vehicular networks for use in the field of Intelligent Transportation Systems (ITS), are effectively becoming part of today's reality. Many drivers already access the information provided by such systems, for example checking for the state of the streets along a given route or reading traffic jam alerts on the displays of smart-phones or Personal Navigation Devices (PNDs). Based on such information, drivers, or their PNDs, select the best paths to reach their destinations. Clearly, in order to be effective, such systems are required to reliably estimate and forecast vehicular congestion states. Moreover, they should also be capable of efficiently utilizing the wireless channel resources, as the amount of information that may be exchanged by such systems in dense urban areas grows with the number of services supported by the onboard devices and the amount of vehicles that install them. To answer these challenges, we here discuss how a distributed ATIS can: a) implement an effective vehicular congestion detection and forecasting model, and, b) efficiently disseminate traffic information. The advantage of distributing an ATIS is that each vehicle can compute and redistribute accurate vehicular congestion information very rapidly, with little overhead and without resorting to a central entity. In order to validate our approach, we present the outcomes of a real world experimentation, as well as of multiple simulations.

## 1 Introduction

Advanced Traveler Information Systems (ATISs) are becoming established as an important part of all surface street transportation services that supplies urban areas. Just as those travelers who check the weather forecast before leaving for a journey, it is nowadays common for many drivers to consult vehicular traffic congestion information before and while moving through traffic. In fact, ATISs

---

\* This work was partially supported by the Italian FIRB DAMASCO project and by a grant of the Computer Science Department, University of Bologna.

information can be browsed online, as well as processed by the most advanced Personal Navigation Devices (PNDs), and used to decide when and how to reach a destination [1], [2].

However, a lack of sufficient traffic information has for long jeopardized the accuracy and reliability of ATISs. In fact, even today no more than the 1% of roads is monitored using induction loops or video cameras [3]. Although, what has changed compared to the past, is the availability of vehicles that mount a GPS, a digital map and a two-way communication interface. Such vehicles are able to play the role of real-time traffic probes. In Italy, for example, every 15 to 60 seconds 700,000 vehicles send to a centralized entity, using a cellular connection, a sample containing a timestamp and local traffic related information (i.e., GPS coordinates, speed and distance from the position where the previous probe was recorded), totalizing over 32,000,000 collected daily probes [4].

In Europe alone, in the near future, 300,000,000 vehicles could potentially be used to probe traffic conditions. Such huge amount of vehicles and the possibility that soon a number of different services (e.g., pollution management, entertainment, etc.) may rely on Vehicle-to-Infrastructure (V2I), as well as on, Vehicle-to-Vehicle (V2V) channels, demand for the design of jointly efficient congestion detection and communication schemes. In order to limit the consumption of wireless channels, we designed a system where vehicles advertise congestion events in a distributed fashion, sending probes only when congestion is effectively detected [5].

Clearly, the advantages that may derive from a distributed ATIS come at a cost. First, they require the design of effective distributed congestion detection and short-term forecasting algorithms, as false positive or false negative events could potentially harm the flow of traffic. Especially in metropolitan areas, tracking the trend of traffic flows is hard, since traffic jams can quickly build up. Second, they depend on the performance of distributed communication protocols that should preserve the consistency of the information and be efficient, minimizing the use of wireless resources, while expediting the spreading of traffic information.

At this point, the devise of a distributed ATIS walks through the design of an operative definition, which may be implemented at every vehicle, of what traffic congestion is. Although many of such definitions are available in literature, we were unable to find a one that was self-contained, which did not require any prior knowledge, nor an estimate of any road parameter [6], [7]. We therefore proposed a new definition, which derives from the observation that if two cars traverse the same street within some limited amount of time one from the other (e.g., say within a time span  $S$ ), they will probably experience the same traffic conditions. This follows from the observation that if a first car travels across a street while this is jammed, a second car will probably experience the same situation assuming that it entered the street not too far away in time from the first car. Such phenomenon is due to the inertia of vehicular queues, which causes a street to be seen as congested also by those vehicles that later enter it. This consideration allowed us to define congestion as a state that lasts for at least

$S$  units of time and during which travel times or delays exceed the time  $T^*$  normally incurred under light or free-flow travel conditions.

But this addresses only part of the problem of designing a distributed ATIS. We also need a communication protocol that efficiently distributes traffic information throughout a platoon of vehicles in a realistic wireless propagation setting. An algorithm that fits such requirements exists and works as follows [8]. Traffic congestion alerts are broadcast throughout a vehicular network asking, at each hop, to the farthest spanning relay, the relay that can reach the farthest vehicle in a given direction, to retransmit it. This strategy has been proven to be optimal, as it minimizes the number of hops required to disseminate a message in a vehicular network, even when wireless channels are asymmetric and transmission ranges vary in time.

The scope of this paper is to present a novel distributed ATISs architecture, which is able to efficiently and pervasively disseminate traffic information throughout a vehicular network and effectively detect vehicular congestion situations, as well as their duration, within every vehicle. The novelty of our architecture is that it works in a distributed fashion, with no central point of failure, for any type of road and without any prior knowledge, while it is able to perform short-term congestion forecasting by simply analyzing the information gathered by probing vehicles.

Between 2008 and 2009 we carried out over 450 miles of road experiments, in Los Angeles (CA) and Pisa (Italy), to validate our approach. In Los Angeles we were also able to compare our congestion detection algorithm to the results provided by Google Traffic.

The rest of our paper is organized as follows. In Section II we provide a succinct review of the state of the art on ATISs. In Section III we sketch the architecture of our system. We provide results on our congestion detection algorithm Section IV. We finally conclude with Section VI.

## 2 Related Work

Although a wealth of work has been performed in the area of ATISs, we here focus on describing the approaches that fall closest to the aim of this paper.

The idea of implementing a distributed ATIS has been studied for long. Most of the approaches we are aware of, however, study how to plainly disseminate road section travel time information within a vehicular network using known broadcast and aggregation schemes [9]. While such strategies can be beneficial to enable the routing module of each vehicle compute the fastest path from an origin to a destination, on a large scale such approaches can be bandwidth consuming and cost inefficient, as travel time information is sent for each road section. Only relevant traffic information, concerning traffic jams or car accidents for example, should be broadcast, thus requiring each vehicle to distinguish a congested from a non-congested road section.

Seeking now for a congestion detection scheme that may be utilized in the implementation of a distributed ATIS, we found that the authors of [10] proposed

a scheme that permits a centralized entity to identify congestion on traffic light controlled streets. In brief, by processing the GPS traces sent by the vehicles that traverse a road section, a centralized entity classifies the given section as congested if one of the two following situations occurs: vehicles waste their time by repeatedly moving in a stop and go pattern, or by stepping in queue for at least one full red light cycle. Although this approach provides a road section traffic state classification method, it is limited to traffic light controlled streets.

The state of the art of ATISs is represented by systems such as Google TrafficView, which can be accessed from both a web interface and from a cell phone [2]. Such type of systems displays the map of a given area with streets that are colored depending on their congestion situation. For example, a street that is drawn with an intermittently black and red color means that it is highly congested, red that it is congested, yellow lightly congested, while green non-congested. The main drawback of such systems is that if accessed while driving, they come at a cost, since their information is received using a cellular network infrastructure. Moreover, how they classify the state of a street should be clarified, since with our experiments on two different road sections in Los Angeles we found many discrepancies between the traffic states they advertised and the travel time we experienced.


Summarizing, we believe that our proposed distributed ATIS architecture, which can sense when a given road section is congested or not and spread a congestion alert when congestion is found, can overcome all the main issues we here highlighted for previously proposed systems that aim at performing similar functionalities.

### 3 Issues of a Distributed ATIS: Congestion Detection and Information Spreading

We begin by recalling how a centralized ATIS works. Simply said, on traversing a given road  $R$  a vehicle sends its traversal time sample to a centralized entity, which in turn saves such information into an internal data structure, until a sufficient number of observations have been collected. Once the entity stops harvesting data, it builds a picture of the congestion states characterizing that given street (typically based on the computation of a congestion threshold  $T^*$ ). This activity may usually last a day or for the time necessary to collect a few hundreds of observations. If any time later a vehicle traverses  $R$  and sends its traversal time to the centralized entity, this can recognize whether  $R$  is congested by comparing this value to the threshold  $T^*$ . Further, it can exploit this information to, for example, alert all the vehicles that are moving in that area.

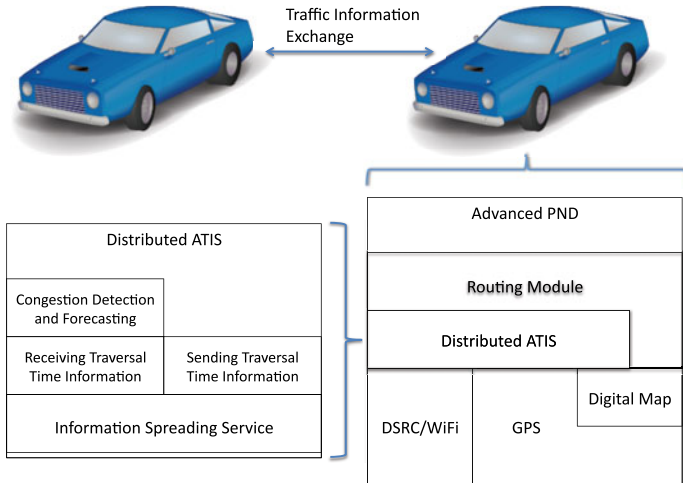
To describe how a distributed ATIS architecture could implement similar functionalities, we resort to Figure 1 which depicts a platoon of vehicles traversing a given road and a few cars approaching it. In particular, vehicle A ends traversing the given street and broadcasts a message containing its traversal time to all the cars that are in range. On receiving such message, vehicle B performs the same actions a centralized entity would (i.e., data harvesting and, eventually,

detecting congestion states), and retransmits it, based on a given information spreading algorithm. Hop by hop, from one vehicle to another, this message propagates until all the vehicles connected within a certain distance from the originating point (vehicle A) receive it. At a cost of this message exchange, each vehicle is now in the condition of detecting possible congestion states regarding roads belonging to its potential route. Hence, each vehicle can take informed decisions on the route to follow based on the information other vehicles have provided it with. Contrary to the centralized approach, this happens without the need of maintaining a central entity, as each vehicle is made able to compute the traffic congestion situation pertaining it on the basis of the received information.



**Fig. 1.** Distributed ATIS application scenario

The cited scenario can be implemented with the architecture shown in Figure 2. In particular, the rightmost software architecture displays an advanced PND where the routing module discovers the fastest driving path to destination by relying on a digital map and on traffic information processed by a distributed ATIS. The distributed ATIS, instead, relies on a congestion detection and forecasting algorithm to analyze all the traffic data received from other vehicles and on an information spreading service to disseminate traversal time samples. All this requires, at the lowest stack layer, that each vehicle mounts a GPS and a V2V communication interface (e.g., DSRC, WiFi), as shown in the Figure.



**Fig. 2.** An advanced PND-distributed ATIS stack

Summarizing, it is possible to implement a distributed ATIS relying on two main building blocks: a vehicular congestion detection and forecasting algorithm and an efficient information spreading scheme. We proceed describing the details of the former in Section 4 and then sketching the idea of one possible candidate for the latter in Section 5.

## 4 Congestion Detection and Forecasting: A Model with Its Algorithm

Let us begin with the congestion definition we have devised. That is we define congestion as a state that lasts for at least  $S$  units of time and during which travel times or delays exceed the time  $T^*$  normally incurred under light or free-flow travel conditions. Based on this definition, we present a model with which we can now compute the congestion threshold  $T^*$  and the minimum time span  $S$  for which a given state of congestion or non-congestion lasts on a road section. The model can be summarized as follows. A road  $R$  is congested if a value of  $T^*$  exists for which, when a vehicle traverses  $R$  requiring more than  $T^*$  units of time, the majority of cars (e.g., 80%) that follow, later entering  $R$  (say within a time span  $S$ ), still require more than  $T^*$  units of time to traverse it. If, alternatively, only a small percentage of subsequent cars (e.g., much below 80%) took more than  $T^*$  units of time to traverse  $R$ ,  $R$  would clearly be leaving a state of congestion. In a similar manner,  $R$  is non-congested if a vehicle traverses it in less than  $T^*$  units of time and the majority of the vehicles that follow (e.g., 80%) still require less than that time. In the case, instead, that a low percentage (e.g., much below 80%) of the vehicles that followed experienced a traversal time below  $T^*$ , this would indicate  $R$  was transitioning into

a congested state. At the basis of this model lies a very intuitive consideration which is as follows. Consider now two cars traversing the same street within some limited amount of time one from the other (e.g., say within a time span  $S$ ), they will probably experience the same traffic conditions. This follows from the observation that if a first car travels across a street while this is jammed, a second car will probably experience the same situation assuming that it entered the street not too far away in time from the first car. Such phenomenon is due to the inertia of vehicular queues, which causes a street to be seen as congested also by those vehicles that later enter it. Extending to many does not modify this intuition.

Based on the above, we are able to informally define the following sets of pairs of vehicles which enter a given road  $R$  within a time span  $S$ .  $HC_{T_1^*}$  includes all those pairs whose traversal times are both above  $T_1^*$  (high congestion), whereas  $N1_{T_1^*}$  identifies those pairs of vehicles where only the first that enters  $R$  experiences a traversal time above  $T_1^*$  (leaving congestion). Similarly,  $NC_{T_2^*}$  contains all those pairs whose traversal time are both below  $T_2^*$  (low congestion), and  $N1_{T_2^*}$  those pairs where only the first vehicle traverses  $R$  in less than  $T_2^*$  units of time. Accordingly, it is possible to test the membership of a given pair of vehicles  $(i, j)$  defining the indicator functions  $I_{HC_{T_1^*}}(i, j)$ ,  $I_{N1_{T_1^*}}(i, j)$ ,  $I_{NC_{T_2^*}}(i, j)$  and  $I_{N2_{T_2^*}}(i, j)$  on each of the defined sets, respectively.

Given the aforementioned sets and their respective indicator functions, we can now provide the propositions that verify if, within a platoon  $P$ , the percentage of cars that experience congestion or not is above or below the 80% threshold. Please remind 80% is a representative value that can be modified based on the degree of congestion one wants to detect.

**Proposition 1.** (*Congestion*). *A given road segment R is congested during a period S if the following holds:*

$$\frac{\sum_{(i,j) \in P \times P} I_{HC_{T_1^*}}(i, j)}{\sum_{(i,j) \in P \times P} I_{HC_{T_1^*}}(i, j) + \sum_{(i,j) \in P \times P} I_{N1_{T_1^*}}(i, j)} \times 100\% \geq 80\%. \quad (1)$$

The same can be drawn for a non-congested state, as follows:

**Proposition 2.** (*No Congestion*). *A given road segment R is not congested during a period S if the following holds:*

$$\frac{\sum_{(i,j) \in P \times P} I_{NC_{T_2^*}}(i, j)}{\sum_{(i,j) \in P \times P} I_{NC_{T_2^*}}(i, j) + \sum_{(i,j) \in P \times P} I_{N2_{T_2^*}}(i, j)} \times 100\% \geq 80\%. \quad (2)$$

Now, an efficient way to determine the values of  $T_1^*$  and  $T_2^*$  is to search for the pair  $(T_1^*, T_2^*)$  which maximizes the size of the  $HC(T_1^*)$  and  $NC(T_2^*)$  sets and, at the same time, minimizes the size of the  $N1(T_1^*)$  and  $N2(T_2^*)$  sets.

**Proposition 3.** *A congestion threshold  $T_1^*$  and a non-congestion threshold  $T_2^*$ , for a given road section  $R$  observed during both congested and non-congested states, can be obtained as:*

$$\begin{aligned}
 (T_1^*, T_2^*) &= (T_1, T_2) \text{ s.t.} \\
 &\{ \max_{T_1, T_2} \sum_{(i,j) \in P \times P} I_{HC_{T_1}}(i, j) + \\
 &+ I_{NC_{T_2}}(i, j) + \\
 &- I_{N1_{T_1}}(i, j) + \\
 &- I_{N2_{T_2}}(i, j) \}. \tag{3}
 \end{aligned}$$

Obtained  $T_1^*$  and  $T_2^*$ , their values should be checked to verify that the inequalities expressed in Propositions 1 and 2 are satisfied.

#### 4.1 Implementation

Our distributed ATIS embodies a congestion detection and forecasting algorithm, which processes the traversal times observed on a given street by probing vehicles (Table 1) and finds a congestion threshold  $T^*$  and a minimum time span  $S$  for which a congestion event is expected to last (Table 2).

Table 1 shows the pseudo-code of the `ATIS()` function that processes an incoming message. In particular, a message could contain more than one traversal time for more than one road (e.g., all the information cache of a vehicle), therefore two nested loops (lines 1 and 2) iterate through all the roads and their corresponding traversal times. However, before adding a new traversal time sample to an internal data structure, a check is carried out to verify whether more data for the given road is required (line 3). Once a sufficient number of samples are collected, the ATIS can stop gathering more information for this road (lines 5 and 6) and  $S$ ,  $T_1^*$  and  $T_2^*$  can be computed resorting to the `CDF()` function (line 7).

The details of how the `CDF()` function works are shown in Table 2. In brief, it basically implements the mechanism proposed in Proposition 3, searching for the values of  $T_1$  and  $T_2$  that maximize the size of the congested and non-congested sets ( $HC_{T_1}$  and  $NC_{T_2}$ ), while minimizing the number of pairs contained in the others ( $N1_{T_1}$  and  $N2_{T_2}$ ) (lines 2 and 5). Once the values of  $T_1^*$  and  $T_2^*$  have been computed, a check is performed to verify if the conditions expressed by Propositions 1 and 2 are satisfied (lines 3 and 7). If the positive case, the function ends returning the values of  $S$  and  $T_1^*$ . As our experiments show that  $T_1^*$  and  $T_2^*$  are always confined within a 3% difference, we chose to return a unique congestion threshold value  $T^* = T_1^*$ . Nevertheless, the checks could fail. A motivation for this to happen could be the selection of a too large duration  $S$  for the state of congestion of interest, as the congested (or non congested) state a first vehicle incurs in does not last in time, as a second vehicle does not find the same state any longer. However, this could be a problem simply concerned with the duration of the  $S$  we have chosen, while a smaller value for  $S$  could exist for which all subsequent cars incur in the same state of congestion. The idea is then to



**Table 1.** Distributed ATIS Algorithm: Processing Traversal Time Information

```

input: Message containing road traversal times, including  $R$  and  $T$  information.
output: Road congestion characterization  $(S, T_1^*, T_2^*)$ .
0: void ATIS()
1: foreach road R do
2:   foreach traversal time T do
3:     if collectingData == true then
4:       traversalTimes.Add(T);
5:       if R.observationTime > 12 h and traversalTimes.Length > 100 then
6:         collectingData = false;
7:          $(S, T_1^*, T_2^*) \leftarrow \text{CDF}(\text{traversalTimes})$ ;
8:       end
9:     end
10:   end
11: end

```

**Table 2.** Congestion Threshold Detection and Forecasting Function

```

input: A list of traversal times.
output:  $S$  and  $T^*$ .
0: void CDF()
1:  $S \leftarrow \Delta_{max}$ ;
2:  $(T_1^*, T_2^*) \leftarrow (T_1, T_2)$  s.t.  $\text{Max}(T_1, T_2)$ ;
3: while  $\neg\text{Check1}(T_1^*) \wedge \neg\text{Check2}(T_2^*) \wedge S > \Delta_{min}$  do
4:    $S \leftarrow S - \delta$ ;
5:    $(T_1^*, T_2^*) \leftarrow (T_1, T_2)$  s.t.  $\text{Max}(T_1^*, T_2^*)$ ;
6: end
7: if  $\neg\text{Check1}(T_1^*) \wedge \neg\text{Check2}(T_2^*)$  then
8:   return null;
9: else
10:   return  $(S, T_1^*)$ ;
11: end

```

seek for such value, by reducing  $S$  until a situation is captured where both the subsequent vehicles of the pair experience a similar state of congestion (or no congestion). This motivates the iterative structure of the  $\text{CDF}()$  function (lines 3 through 6).

## 4.2 Evaluation and Comparison

We carried out a number of experiments with a fleet of probing vehicles in 2008 and 2009 in Los Angeles, CA, and in Pisa, Italy. For obvious reasons, we were not able to implement the full distributed ATIS architecture producing a scenario as the one shown in Figure 1, however we were still able to test the validity of our congestion detection and forecasting algorithm. In our experiments, each participating vehicle carried an onboard system consisting of a laptop, a GPS receiver and an EVDO interface and shared its traversal time samples with the

**Table 3.** Experiment information: location, road section, road length, free flow traversal time (FFTT), traffic light cycle time (CT) and green time (GT)

	Road	Section	Length [m]	$T_{FFTT}$ [s]	CT [s]	GT [s]
1	Via B. Croce	Guerrazzi-Queirolo, left	380	34	85	55
2	S. Monica Blvd	Veteran-Sepulveda, left	380	61	120	15
3	S. Monica Blvd	Wilshire-Roxbury, straight	280	17	90	54
4	S. Monica Blvd	Wilshire-Bedford, right	390	30	90	54
5	Lincoln Blvd	Fiji-Venice, back	2300	205	120	60
6	Wilshire Blvd	Midvale-Westwood, right	130	7	150	80
7	S. Monica Blvd	Roxbury-Bedford, right	100	7	90	54
8	Wilshire Blvd	Veteran-Westwood, right	340	33	150	80
9	S. Monica Blvd	Westwood-Sepulveda, right	680	75	120	50

**Table 4.** Road data: number of loops,  $T^*$ ,  $S$ ,  $N$  and  $H$ 

	Road	Section	# of loops	$T^*$ [s]	$S$ [s]	$N$	$H$	$\hat{T}$ [s]
1	Via B. Croce	Guerrazzi-Queirolo, left	111	93	362	92%	84%	64
2	S. Monica Blvd	Veteran-Sepulveda, left	134	175	608	80%	87%	166
3	S. Monica Blvd	Wilshire-Roxbury, straight	77	62	987	94%	99%	53
4	S. Monica Blvd	Wilshire-Bedford, right	77	82	987	92%	99%	63
5	Lincoln Blvd	Fiji-Venice, back	30	354	900	100%	97%	265
6	Wilshire Blvd	Midvale-Westwood, right	71	36	454	39%	98%	77
7	S. Monica Blvd	Roxbury-Bedford, right	77	42	987	46%	83%	43
8	Wilshire Blvd	Veteran-Westwood, right	71	74	454	37%	100%	103
9	S. Monica Blvd	Westwood-Sepulveda, right	67	121	493	90%	54%	145


rest of the fleet. The details regarding the roads observed during the experiments are listed in Table 3 (name, section, length, free flow traversal time, full and green traffic light cycles).

The results of our congestion detection and forecasting algorithm are provided in Table 4 which gives for each street: the number of times it was traversed, its congestion threshold  $T^*$ , its congestion time span  $S$  and the values of  $N$  and  $H$  (specifically, the ratio of vehicles that experienced stably congested and non-congested conditions, in the mentioned order). In such table we compare  $T^*$  to the time a vehicle would wait when traversing a street and being in queue for a full red light,  $\hat{T}$ , supposing this an acceptable traversal time with no congestion. The results shown in Table 4 let us state that streets 1 through 5 all experienced successive situations of congestion and non-congestion. The  $N$  and  $H$  values both confirm this, both exceeding the 80% threshold. Moreover, we also observe that in each of these experiments the value of  $T^*$  surpasses  $\hat{T}$ , which means that the vehicles that traverse these streets exceeding  $T^*$  really experience congestion and queuing. Streets 6 through 8, instead, are known to be only seldom congested. Our results confirm such knowledge, since  $T^* < \hat{T}$  (cars almost always enjoy a smooth drive) and  $N$  is small (no stable congestion was visible over those streets). Finally, street # 9 behaves in a different manner. A high value of  $N$


and a small value of  $H$  seem to reveal a stable high congestion state, however, surprisingly,  $\hat{T}$  is greater than  $T^*$ . This phenomenon can be explained observing that the traffic light on Sepulveda Blvd. allows right turns on red, thus cars not very often waited for a full red light time. A more extensive set of experiments and results may be found in [11].

As a further term of comparison, we also contrast our results to those produced by Google TrafficView on street # 3 (Figures 3 and 4). Each circle in Figures 3 and 4 represents the traversal time experienced by one of our vehicles at the time indicated on the x-axis, while the background colors reproduce those reported on the <http://maps.google.com/> website at the time of our experiments. A red background represents congestion, yellow corresponds to mild congestion and green to no congestion. Reminding we found  $T^* = 62$  and  $S = 987$ ) for street # 3, we discover there are many points for which the two methods do not agree on when the road is congested or not. In addition, interestingly, we also find that Google TrafficView can contradict itself, as in Figure 3 green and yellow areas correspond to traversal time samples that exceed 100 seconds, while in Figure 4 a red area corresponds to traversal time samples well below 70 seconds.

In conclusion, we believe our congestion detection and forecasting algorithm is able to meet our expectations, it provides clear results and it is very easy to implement. For this reason, we believe it is the ideal candidate to be integrated in an advanced distributed ATISs.



**Fig. 3.** Comparison between Google TrafficView and traffic samples on S. Monica Boulevard between Wilshire Boulevard and Roxbury Avenue on the 3rd of April 2009



**Fig. 4.** Comparison between Google TrafficView and traffic samples on S. Monica Boulevard between Wilshire Boulevard and Roxbury Avenue on the 26th of April 2009.

## 5 Efficient Information Spreading: A Broadcast Approach

Many different strategies can be adopted to implement a distributed ATIS (e.g., multicast, peer-to-peer, etc.), however, a broadcasting scheme appears as the ideal solution to rapidly convey traffic information to all the vehicles that are moving in a given area.

Table 5 shows the functionalities a broadcasting scheme should implement to rapidly spread traffic information as far as possible. This is a complex issue which can be summarized as follows [12]-[20]. Every time a given vehicle receives

**Table 5.** Distributed ATIS Algorithm: Receiving a message

<p><b>input:</b> Traversal time message with <math>R</math> and <math>T</math>.</p> <p><b>output:</b> Deliver information to the ATIS application and retransmit the message if a relay.</p> <p><b>0:</b> void BroadcastService()</p> <p><b>1:</b> ATIS(traversalTime);</p> <p><b>2:</b> if this.vehicle == message.currentRelay then</p> <p><b>3:</b>     nextRelay = informationService.NextRelay;</p> <p><b>4:</b>     SendMessage(message, nextRelay);</p> <p><b>5:</b> end</p>
---

a message, this is delivered to the ATIS application which will permit to that vehicle to construct a picture of the traffic situation concerning the roads surrounding it (line 1). Once this step is performed, the vehicle should in some way discover whether it should relay the message or not. For the sake of brevity, we assume that such information is contained in the message (line 2), which means that each current relay elects the next relay at each transmission round. In case the given vehicle has been chosen as a relay, it finds the next relay (line 3) and retransmits the message (line 4), otherwise nothing more should be done.

The problem at this point is how the next relay can be identified at each transmission round which is able to span as far as possible. The good news is that this problem has been solved, as the best relay has been recognized to be the farthest spanning relay (i.e., the possible relay that spans farthest in the desired direction among all) and a distributed mechanism has been devised to unveil it

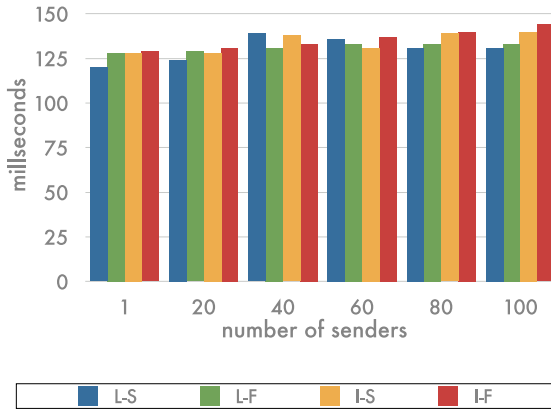


Fig. 5. End-to-end delay

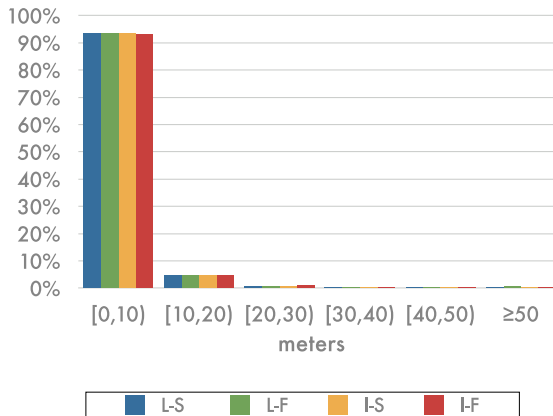



Fig. 6. Percentage of vehicles whose position is identified with an error of  $[x, y]$  meters



**Fig. 7.** Percentage of vehicles whose transmission range is estimated with an error of  $[x, y]$  meters

at each hop. Again, to be concise we do not provide here additional details on this subject. The interested reader can refer to the following paper [8]. Rather, it is more interesting to report here that such kind of broadcast mechanism may be very effective in spreading traffic information at a high speed. This is confirmed by the simulation results reported in Figure 5 which show that, selecting the farthest spanning relay at each hop, a message can reach all the vehicles in an eight kilometers long platoon within 150 milliseconds at most. In particular, in Figure 5 the chosen performance parameter (i.e., the end-to-end delay) has been measured with the combination of two broadcasting policies. Specifically, *L* stands for Lazy, which means each traffic message is sent at the rate of 1 Hz. *I*, instead, means Intensive, that is each message is sent with a frequency of 2 Hz. *S* and *F* stand, respectively, instead, for Slim and Fat which, in turn, mean that each message is long either 1 KB or 2 KB.

To conclude this section, we provide information concerning the degree of precision with which this mechanism is able to identify the best relay. In particular, Figures 5 and 6 demonstrate that in the 90% and 70% of cases respectively, the positions and the transmission ranges of all the vehicles in the platoon are accurately identified.

## 6 Conclusion

We here discussed how a distributed ATIS can implement an effective vehicular congestion detection and forecasting model, and, efficiently disseminate traffic information. The advantage of distributing an ATIS is that each vehicle can compute and redistribute accurate vehicular congestion information very rapidly, with little overhead and without resorting to a central entity. In order to validate our approach, we present the outcomes of a real world experimentation, as well as of multiple simulations.

## References

1. Garmin, <http://www.garmin.com/>
2. Google TrafficView, <http://www.google.com/mobile/navigation>
3. Leduc, G.: Road traffic data: Collection methods and applications. Technical Report, Institute for Prospective Technological Studies, EU (2008)
4. Octotelematics, <http://www.octotelematics.com/>
5. Rybicki, J., Scheuermann, B., Kiess, W., Lochert, C., Fallahi, P., Mauve, M.: Challenge: Peers on Wheels - A Road to New Traffic Information Systems. In: 13th ACM International Conference on Mobile Computing and Networking, pp. 215–221. ACM, New York (2007)
6. Bertini, R.: You are the traffic jam: an examination of congestion measures, Technical Report, Department of Civil and Environmental Engineering, Portland State University (2005)
7. Dowling, R.G.: Arterial Speed Study. Technical Report, Southern California Association of Governments (2005)
8. Rocchetti, M., Marfia, G., Amoroso, A.: An Optimal 1D Vehicular Accident Warning Algorithm for Realistic Scenarios. In: IEEE Symposium on Computers and Communications. IEEE Press, New York (2010)
9. Yang, Y., Bagrodia, R.: Evaluation of VANET-based advanced intelligent transportation systems. In: 6th ACM International Workshop on Vehicular InterNetworking, pp. 3–12. ACM, New York (2009)
10. Yoon, J., Noble, B., Liu, M.: Surface Street Traffic Estimation. In: 5th ACM International Conference on Mobile Systems, Applications and Services, pp. 220–232. ACM, New York (2007)
11. Marfia, G., Rocchetti, M.: Vehicular Congestion Detection and Short-Term Forecasting: A New Model with Results. Technical Report, University of Bologna (2010)
12. Palazzi, C.E., Ferretti, S., Rocchetti, M.: An Inter-Vehicular Communication Architecture for Safety and Entertainment. IEEE T. Int. Tran. Syst. 11, 90–99 (2009)
13. Amoroso, A., Rocchetti, M., Nanni, M., Prati, L.: VANETS without limitations: an optimal distributed algorithm for multi-hop communications. In: 6th IEEE Conference on Consumer Communications and Networking Conference, pp. 1307–1311. IEEE Press, New York (2009)
14. Amoroso, A., Ciaschini, C., Rocchetti, M.: The farther relay and oracle for VANET. preliminary results. In: 4th Annual International Conference on Wireless Internet, pp. 1307–1311. IEEE Press, New York (2008)
15. Wisitpongphan, N., Tonguz, O.K., Parikh, J.S., Mudalige, P., Bai, F., Sadekar, V.: Broadcast Storm Mitigation Techniques in Vehicular Ad Hoc Networks. IEEE W. Com. 14, 84–94 (2007)
16. Zanella, A., Pierobon, G., Merlin, S.: On the limiting performance of broadcast algorithms over unidimensional ad-hoc radio networks. In: Wireless Personal Multimedia Communications, pp. 165–169 (2004)
17. Fasolo, E., Furiato, R., Zanella, A.: Smart broadcast algorithm for inter-vehicular communication. In: 2005 Wireless Personal Multimedia Communication (2005)

18. Blum, J.J., Eskandarian, A.: A reliable link-layer protocol for robust and scalable intervehicle communications. *IEEE T. Int. Tran. Syst.* 8, 4–13 (2007)
19. Alshaer, H., Horlait, E.: An optimized adaptive broadcast scheme for inter-vehicle communication. In: *61st Conference on Vehicular Technology*, pp. 2840–2844. IEEE Press, New York (2005)
20. Korkmaz, G., Ekici, E., Ozguner, F., Ozguner, U.: Urban multi-hop broadcast protocol for inter-vehicle communication systems. In: *1st ACM International Workshop on Vehicular Ad Hoc Networks*, pp. 76–85. ACM, New York (2004)