WSN-108

Jointly Optimizing Data Acquisition and Delivery in Traffic Monitoring VANETs*

Antonios Skordylis
antonios.skordylis@comlab.ox.ac.uk

Niki Trigoni
niki.trigoni@comlab.ox.ac.uk

Oxford University Computing Laboratory
Wolfson Building, Parks Road
Oxford, OX1 3QD, United Kingdom

ABSTRACT

Vehicular ad hoc networks (VANETs) are envisaged to become a flexible platform for monitoring road traffic, which will gradually replace more cumbersome fixed sensor deployments. The efficacy of vehicle-assisted traffic monitoring systems depends on the freshness of traffic data that they can deliver to users, and the bandwidth used to do so. Clearly, high data freshness will allow users to estimate trip times accurately, and to select the fastest route to a destination. Low bandwidth utilization will allow the traffic monitoring application to coexist symbiotically with a wide variety of vehicle-based applications, ranging from road safety to advertising and entertainment.

In this paper, we investigate the problem of minimizing the bandwidth utilization of a vehicle-assisted traffic monitoring system, whilst adhering to user-defined requirements for data freshness. The novelty of our approach is that we jointly optimize two intertwined aspects of traffic monitoring: data acquisition and data forwarding. We investigate how their combined operation trades data freshness for bandwidth utilization, and we propose a novel mechanism that fine-tunes their parameters to optimize the overall system performance. Our mechanism is evaluated using realistic vehicular traces on a real city map.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Store and forward networks, wireless communication*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*routing protocols*

General Terms

Algorithms, Performance

Keywords

vehicular networks, data muling, multi-hop communication, delay-tolerant networks, sensor participation, traffic monitoring

1. INTRODUCTION

Recent trends in Intelligent Transportation Systems show that an increasing number of vehicles will be equipped with wireless transceivers that will enable them to communicate with each other and form a special class of wireless networks, known as vehicular ad hoc networks or VANETs. Researchers and automotive industries are envisioning the deployment of an *ambient traffic monitoring application*, wherein vehicles equipped with GPS detect local traffic and report it periodically to one of the stationary roadside units dispersed in the city. These units are referred to as access points (APs) and act as gateways to the city's traffic monitoring center and the outside world.

One of the most important attributes of traffic data is its *freshness*, i.e. the interval between the time that the data is generated by a vehicle on a particular road and the time that the data is made available to the user as a query response. Informally, data freshness indicates how stale the data is, and to what extent it can be used to estimate trip times or to select the fastest route to a destination in a reliable manner. Depending on the expected rate of change in traffic conditions, users may have different freshness requirements for different parts of the city, or for different times of the day. It is crucial that the ambient traffic monitoring application provides deterministic guarantees that the available traffic data satisfies the specified freshness requirements.

At the same time, the ambient traffic monitoring application will be sharing bandwidth resources with a wide variety of applications running on the same VANET, for example, applications that provide internet access to passengers, commercial applications that flood advertisements about nearby stores, safety applications that provide drivers with emergency braking services, and so on.

Thus, our high-level goal is to design an ambient traffic monitoring system that minimizes bandwidth utilization, whilst adhering to user-defined data freshness requirements. To achieve this goal, we investigate two intertwined aspects of traffic monitoring, *data acquisition* and *data delivery*, both of which significantly impact both data freshness and bandwidth utilization. Our contributions are the following:

1. We formulate a novel problem in the context of ambient traffic monitoring, that of minimizing the communication cost required to monitor traffic whilst providing deterministic guarantees of data freshness.

2. We propose the joint optimization of two closely coupled tasks: data acquisition and data delivery. We investigate how

∗. This work was supported by EPSRC Grant EP/C547640/1.

their combined operation trades data freshness for communication cost, and propose a mechanism to fine-tune their parameters in order to optimize the overall system performance.

3. We evaluate the benefits of our approach using realistic traffic traces on a real city map.

The remainder of this paper is organized as follows: Section 2 introduces the model we are considering and discusses our assumptions and objective. Section 3 reviews data acquisition and delivery algorithms and analyzes their incurred delay, and Section 4 presents our optimization techniques and results. We provide concluding remarks in Section 5.

2. MODEL

2.1 Assumptions

We assume location-aware vehicles that obtain their geographical position from a GPS receiver or other location service and also have access to a digital map of the area. Using onboard sensors (GPS, laser, etc.) vehicles are able to estimate the average speed $\overline{u}$ and average vehicle density $\overline{d}$ on the road segment they are traversing.

We are considering an urban scenario where the network consists of mobile nodes (vehicles) and a few stationary access points (APs) that provide partial city coverage. Using short to mid range transceivers, vehicles can communicate with neighboring vehicles or APs within range. Vehicles are tasked with sensing traffic information and relaying it over multiple vehicles to the AP. We assume that once a message arrives at an AP, it immediately becomes available at the traffic monitoring center via a fixed high-bandwidth network.

The results of ad hoc network studies are heavily influenced by the mobility model utilized [2]. The random-waypoint mobility model is amongst the most commonly used, which however fails to capture the dynamics of the urban vehicular scenario. In this study we are basing our evaluation on realistic vehicular traces from the city of Zurich. The traces have been produced by a multi-agent traffic simulator that simulates public and private traffic over a real map, based on actual travel plans of individuals [7]. The size of the simulation area is $250km \times 260km$ with 260.000 vehicles involved. We have extracted a rectangular street area of size $20km \times 10km$, which covers the centre of the city and surrounding areas, and contains around 30000 distinct vehicle trajectories during a 60-minute interval in morning rush hour. We have uniformly distributed 150 stationary access points on road intersections in the area.

For our analysis we utilize a Java discrete event simulation environment developed with vehicular networks in mind. We have selected the simulation interval to coincide with morning rush hour in the traces and all simulations run during this 60 minute interval. We set the communication range to 250m and average results over 30 iterations.

2.2 Objective

We aim to minimize the bandwidth utilization of a traffic monitoring system, whilst adhering to user-defined data freshness requirements. In order to achieve this goal we investigate two system aspects that significantly impact both data freshness and bandwidth utilization: *data acquisition* and *data delivery*.

Data acquisition refers to the sampling of road traffic information by passing vehicles. High sampling rates can be achieved by having vehicles participate in the sampling process and generate

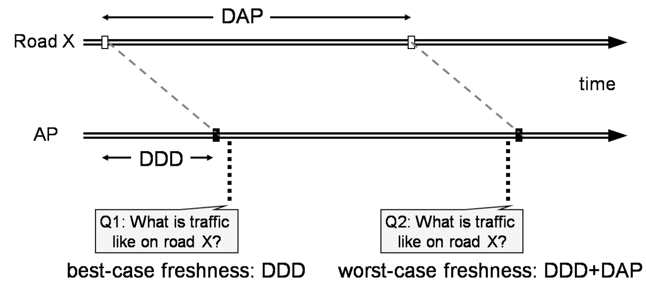

Figure 1: Depending on when a user asks an AP about traffic on a particular road, she will receive results of varying freshness. The worst-case freshness is DDD+DAP, where DDD denotes the data delivery delay, and DAP the data acquisition period.

traffic information messages with high frequency. The lower the data acquisition period, the fresher the traffic data that becomes available for each road, but the larger the number of traffic messages propagated through the network.

Data delivery refers to the propagation of traffic messages from the originating vehicle to one of the access points dispersed in the city. Traffic messages can be delivered either by wireless multihop forwarding, or by physically carrying messages at the vehicle's speed towards an AP. In recent work the authors propose hybrid algorithms that carefully combine multihop forwarding and data muling to achieve a desirable delivery delay [11]. Clearly the lower the data delivery delay, the fresher the traffic data available at the APs, but the higher the use of multihop forwarding and thus, the higher the communication cost.

Figure 1 shows that the freshness of traffic data is directly dependent on the data acquisition period (DAP) and the data delivery delay (DDD). Consider the example where users wish to query the speed of vehicles on a particular road. Let traffic messages concerning this road be generated every *DAP* time units, and let these messages take *DDD* time units to be delivered from the source vehicles to the AP. As shown in Figure 1, users that query traffic information immediately after the arrival of a traffic message get the freshest data, whereas users that pose their queries just before the arrival of a traffic message get the stalest data. The best-case freshness equals the data delivery delay (DDD), whereas the worst-case freshness equals the sum of the data delivery delay and the data acquisition period (DDD+DAP).

The question that arises is: given a freshness threshold that defines the maximum allowed gap between data capture and delivery to the user, how should we split it into a data acquisition period and a data delivery delay in order to minimize the total message transmissions in the network? Is it more bandwidth-efficient to generate traffic messages frequently, and allow extra delay during the delivery phase? Or is it preferable to generate traffic messages infrequently, and deliver them as fast as possible?

Our objective is to strike a good balance between the delay budgets allocated to data acquisition and data delivery, whilst keeping their sum below the freshness threshold.

3. ANALYSIS OF DATA ACQUISITION AND DELIVERY ALGORITHMS

In this section, we discuss the process of vehicle-assisted traffic monitoring step-by-step from data acquisition to data delivery. The goal of this section is not to propose novel algorithms, but to investigate how the parameters of existing algorithms used in the data acquisition and data delivery steps affect the actual delays incurred

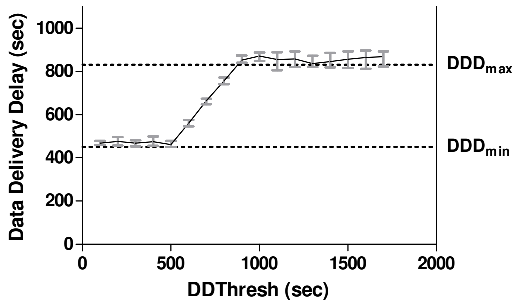

Figure 2: The effect of the D-Greedy parameter $DDThresh$ on the actual data delivery delay.

at each step. This information will be utilized in the next section to jointly optimize data acquisition and data delivery to minimize message transmissions in the network, whilst keeping their total delay below a freshness threshold.

3.1 Data Acquisition

For a traffic monitoring scheme to be successful in an urban environment, it must ensure complete coverage of the sensing field. In the scenario we are considering, this translates into providing regular traffic information updates for every road segment in the network. If a stationary sensor network were to be used, it would suffice to position one or more traffic sensors on each road, uniformly distributed across the road's length, and task them to generate traffic update messages with data acquisition period DAP. In our case, however, sensor nodes are mobile and we have no control over their mobility. We would like to task the mobile nodes in such a way, so that at least one traffic message per road is generated every DAP time units.

3.1.1 Background

A large part of the literature on sensor participation schemes for field coverage refers to stationary sensor networks [8, 3, 10, 6, 14]. Previous work that concentrates on mobile sensors operates on the assumption that sensor mobility can be controlled and, therefore, sensors can be moved on-demand to ensure coverage of the sensing field [4]. Other recent works discuss selection schemes, where the problem is to decide which sensor to move in order to compensate for node failures [9, 12]. In order to optimize node selection for a particular task, most of the above approaches present distributed algorithms which require message exchange between the mobile nodes.

3.1.2 Algorithm

As our main goal is to reduce the communication cost associated with traffic monitoring, we have opted to use a probabilistic sensor participation scheme, wherein each node independently and probabilistically decides whether to participate in the sensing task. Each node participates in sensing, i.e. generates a traffic information message, with probability P_g. The value of P_g is computed based solely on locally-available data.

We would like our mobile sensor network to provide an output similar to that of a stationary sensor network: one traffic message per road every DAP time units. Node mobility introduces two issues that need to be addressed: variable node position and variable traffic conditions. We address the first issue by only allowing vehicles to generate messages at a predefined fixed point on each road segment, e.g. the segment midpoint, effectively simulating a stationary sensor mounted on that point. In order to compensate for variable traffic conditions, we carefully tune the message generation probability P_g.

Adhering to a constant data acquisition period DAP requires messages to be generated with frequency $f_g = \frac{1}{DAP}$. The vehicle can locally derive its average speed $\overline{u}$ as well as the average vehicle density $\overline{d}$ for the road it is traversing using onboard sensor information. Assuming uninterrupted flow conditions, we can derive the average flow q of vehicles on each road as follows: $\overline{q} = \overline{u} \cdot \overline{d}$. The following gives the desired probability $P_g = \frac{f_g}{q} \Rightarrow P_g = \frac{1}{DAP \cdot \overline{u} \cdot \overline{d}}$. Intuitively, the higher the flow of vehicles over the road midpoint where sensing is performed, the lower the value of P_g necessary to maintain a constant sensing period DAP.

3.2 Data Delivery

Once the traffic information message has been generated, the underlying routing protocol will forward it to an AP. The routing protocol is responsible not only for the message delivery delay, but also for the number of transmissions until successful delivery occurs.

3.2.1 Background

Several routing protocols have been proposed in the literature with vehicular networks in mind. Briesemeister et al. [1] proposed an epidemic-style protocol to multicast messages about an accident to cars with a specific role (e.g. geographic location, speed and direction), limiting message propagation to a certain number of hops. MOVE [5] considers the scenario where location-aware mobile nodes attempt to deliver information to a stationary destination whose position is globally known, not unlike our model's access points. It relies on the relative velocity of a node and its neighbors to make forwarding decisions and assumes that a node will maintain its heading until it reaches the destination. MDDV [13] aims to route information to receivers that have expressed an interest for it. Zhao and Cao [15] design vehicle assisted data delivery (VADD) protocols taking into account traffic patterns over a predefined road layout and aim to identify lowest-delay delivery paths. In very recent work, Skordylis et al. [11] study the tradeoff between routing delay and bandwidth utilization in urban environments; they devise routing algorithms that attempt to exhaust a delay requirement in order to minimize message transmissions.

3.2.2 Algorithm

We have chosen to utilize the D-Greedy algorithm presented in recent work [11] for two reasons: it allows us to adjust the desired delivery delay for our messages and it always attempts to minimize message transmissions by exhausting that delay. This is achieved by carefully alternating between two strategies, Multihop Forwarding and Data Muling. Multihop Forwarding refers to relaying messages wirelessly from one vehicle to another towards the direction of an AP, while Data Muling refers to buffering messages in local memory and carrying them at the vehicle's speed. D-Greedy allows vehicles to oscillate between the two strategies taking into account the delay budget for message delivery. As long as the delay budget is high, data muling is preferred since it incurs lower communication cost; as soon as the delay budget tightens, the algorithm reverts to multi-hop communication to ensure timely data delivery at a higher communication cost.

The delay budget that is initially available to a message is an algorithm parameter that the user can vary, called *Delay Threshold* ($DDThresh$). D-Greedy attempts to deliver the message to the closest AP within the user-defined delay threshold $DDThresh$. In fact, it endeavors to deliver as close to $DDThresh$ as possible, by aggressively utilizing multihop forwarding for low values

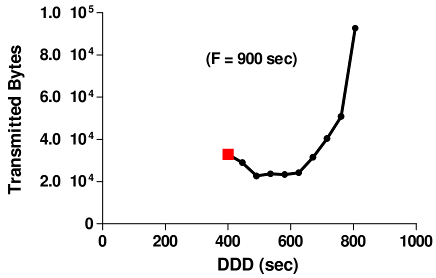

Figure 3: Communication cost across the valid DDD range for a single road.

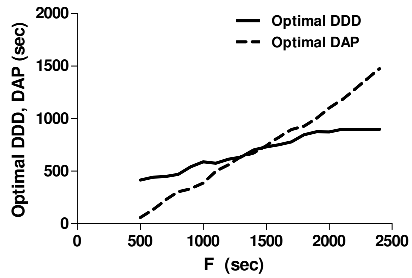

Figure 4: DDD, DAP pairs for optimal communication savings, single road.

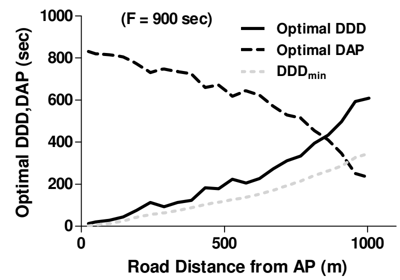

Figure 5: Optimal DDD, DAP pairs as a function of road distance from the AP.

of $DDThresh$ and using data muling when $DDThresh$ is high.

3.2.3 Data Delivery Delay Analysis

Whether D-Greedy can achieve the $DDThresh$ delay target inevitably depends upon the underlying network topology; it may be impossible for messages generated far from an AP to be delivered within certain low $DDThresh$ thresholds, while messages originating near an AP might be delivered much sooner than $DDThresh$ even if data muling is used for the duration of the routing phase. Recall that in our scenario (Figure 1) we would like to know the actual delay DDD incurred by the routing algorithm. Knowing the value of DDD will allow us to allocate the remaining data freshness budget to DAP.

We have examined the effect of the D-Greedy $DDThresh$ parameter on the actual delivery delay incurred for different roads. Figure 2 shows the effect of $DDThresh$ on the maximum delivery delay incurred for a road. 95% confidence intervals are shown as a result of 30 iterations with a different set of participating vehicles chosen from our traces. We observed that for every road there is a lower bound DDD_{min} on how fast the data can be propagated and an upper bound DDD_{max} above which the routing algorithm cannot further delay messages to save extra bandwidth. We also observed that for $DDD_{min} \leq DDD \leq DDD_{max}$, D-Greedy always achieves the $DDThresh$ target, resulting in a linear relationship between $DDThresh$, the algorithm parameter, and DDD, the resulting delay. By storing DDD_{min}, DDD_{max} as well as the slope a and intercept b of the least squares fit between the two points, we now can predict not only the range of allowable DDD values per road, but also the corresponding $DDThresh$ parameter of the D-Greedy algorithm that results in the desired DDD. To aid our optimizations in the next section, we preload the street map with the values DDD_{min}, DDD_{max} as well as the slope a and intercept b for each road.

4. JOINT OPTIMIZATION

A user query with a data freshness requirement of F provides an upper bound for the worst-case freshness allowed by the system. Recall from Figure 1 that the following needs to be satisfied:

$$DDD + DAP \leq F \quad (1)$$

In Sections 3.1.2 and 3.2.3 we discussed how increasing the value of either DDD or DAP will result in less message transmissions in the network. Thus, in order to keep the number of message transmission to a minimum, we need to maintain the sum $DDD+DAP$ as close to F as possible, in an attempt to exhaust the available freshness budget. The naive approach to splitting the budget between DDD and DAP would be to select DDD_{min} for the data delivery delay, i.e. route data as fast as possible, and utilize the full remaining budget ($F - DDD_{min}$) to slow down data acquisition. We refer to this basic approach as *Rapid Delivery*. In other words, *Rapid Delivery* aims to reduce the rate of traffic information generation as much as possible.

This basic approach does not necessarily yield optimal communication savings. We investigate whether we can outperform *Rapid Delivery* by jointly optimizing the data acquisition and data delivery tasks as follows: In Section 4.1 we examine how to divide the freshness budget into DDD and DAP in search for the optimal balance that minimizes communication; we measure how this balance is affected by different freshness budgets and by road proximity to the AP. In Section 4.2 we compare the communication savings of *Rapid Delivery* against those of our joint optimization approach.

4.1 Algorithm Tuning

In Section 3.2.3 we noticed that the actual data delivery delay DDD incurred by the routing algorithm lies within a certain interval for each road $[DDD_{min}, DDD_{max}]$. We measure the communication cost incurred, in the form of transmitted bytes, for DDD values within this interval and their corresponding DAP values, where $DAP = F - DDD$. Each DDD value corresponds to a $DDThresh$ value used to set up the routing algorithm (Section 3.2.3), while DAP values control the data acquisition rate for each road (Section 3.1.2).

Figure 3 shows the bytes transmitted for different values of DDD for a single road when the freshness requirement F is set to 900 seconds. We observe that for $DDD \cong 500$ the bandwidth utilization is minimized for this road. This essentially means that for a specific freshness budget, it is worth allocating part of the budget to slow down data delivery, rather than using it all to slow down data acquisition. Observe the square point on the graph that corresponds to DDD_{min} and thus to the *Rapid Delivery* algorithm: by jointly optimizing we achieved a 30% reduction in communication cost compared to *Rapid Delivery* for this road.

For the same road, Figure 4 shows the optimal DDD value as we vary the freshness budget. The corresponding optimal DAP value that results from the choice of DDD is also shown. A comparison of DDD and DAP slopes reveals that as the freshness budget increases, we should allocate proportionally more delay to data acquisition than data delivery for optimal behavior. Observe that after the $2000sec$ mark, DDD ceases to increase, as DDD_{max} has been reached; from that point onwards, the extra freshness budget is exclusively absorbed by DAP.

An important variable affecting the behavior of the routing algorithm is road distance from the closest AP. For roads that are further away, messages need to travel longer distances and over more hops to reach the AP. Figure 5 shows the optimal (DDD, DAP)

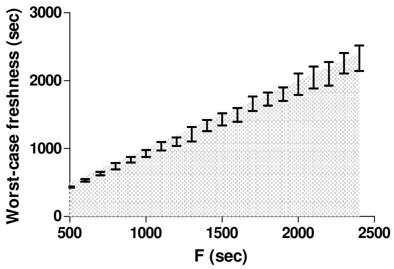

Figure 6: Worst-case freshness of traffic data for a road 840 meters away from the AP.

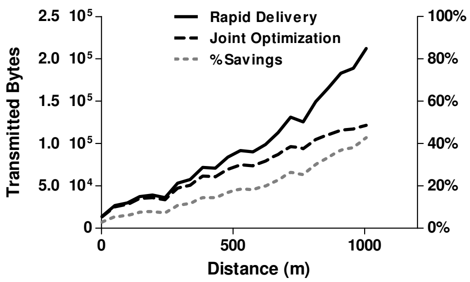

Figure 8: Transmitted bytes for roads at various distances from the AP.

pairs for roads at different distances from the AP. DDD_{min} is also shown here, which corresponds to the delivery delay that *Rapid Delivery* incurs. There are several conclusions that can be derived from Figure 5: The optimal (DDD, DAP) pairs are almost linearly dependent on distance, which provides us with a mechanism to assign (DDD, DAP) pairs to any road based solely on its distance from the gateway. An interesting observation is that, for roads closer to the AP, the freshness budget should be mostly allocated to data acquisition. For roads further from the AP, the freshness budget should be increasingly allocated to data delivery. Notice that for roads further away from the AP, the optimal delay for data delivery (DDD) is significantly larger than the minimum possible delay for data delivery (DDD_{min}).

Figure 6 shows the worst-case freshness achieved for a specific road when using our optimization scheme. For different values of F, we measured the worst case freshness i.e. the one a user would receive if she issued a query just before the arrival of a new message at the AP. The shaded area represents the freshness budget F. Our scheme performs as desired, since it comes very close to exhausting the available freshness budget.

4.2 Benefits

Figures 7 and 8 depict the benefits of the joint optimization. In Figure 7 we observe increasing benefits of our approach over the one that uses *Rapid Delivery* for routing as we relax the freshness requirement, that reach up to 38%. Figure 8 outlines the benefits as a function of road distance from the AP. As anticipated following our observations on Figure 5, joint optimization saves more communication cost compared to *Rapid Delivery* for roads that are further away from the AP, reaching up to 42% for the furthest roads. Our approach is not very effective for roads close to the AP. We could thus omit optimization for roads surrounding the AP without significantly impacting the number of bytes transmitted.

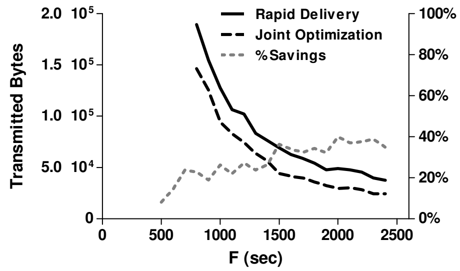

Figure 7: Transmitted bytes (single road) as a function of F .

5. CONCLUSIONS AND FUTURE WORK

We have defined the problem of minimizing the communication incurred by traffic monitoring systems whilst providing deterministic guarantees of information freshness. We have proposed a framework that jointly optimizes the two key processes associated with monitoring traffic, data acquisition and data delivery. Our results have shown that the optimal allocation of freshness budget to these processes depends on the freshness budget itself and the distance of the monitored road from the closest AP. Roads further away from the AP are those that benefit the most from our optimization. By striking an optimal balance between data acquisition and data delivery delays we obtain communication savings of up to 42% compared to the basic approach. Due to scarcity of VANET datasets, this study is based on traces produced by a traffic simulator. We plan to extend our evaluation as more datasets become available in the future, in order to identify whether our optimizations yield similar results for different topological and mobility parameters.

6. REFERENCES

[1] L. Briesemeister and G. Hommel. Role-based multicast in highly mobile but sparsely connected ad hoc networks. In *ACM Mobihoc*, 2000.

[2] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *WCMC: Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.

[3] M. Cardei and D. zhu Du. Improving wireless sensor network lifetime through power aware organization. *ACM Wireless Networks*, 11:333–340, 2005.

[4] K. S. Kwok, B. J. Driessen, C. A. Phillips, and C. A. Tovey. Analyzing the multiple-target-multiple-agent scenario using optimal assignment algorithms. *J. Intell. Robotics Syst.*, 35(1):111–122, 2002.

[5] J. LeBrun, C.-N. Chuah, D. Ghosal, and M. Zhang. Knowledge-based opportunistic forwarding in vehicular wireless ad hoc networks. In *VTC 2005-Spring. 2005 IEEE 61st*, volume 4, pages 2289–2293, May/June 2005.

[6] J. Lu, L. Bao, and T. Suda. Coverage-aware sensor engagement in dense sensor networks. In *EUC 2005*, 2005.

[7] V. Naumov, R. Baumann, and T. Gross. An evaluation of inter-vehicle ad hoc networks based on realistic vehicular traces. In *MobiHoc*, 2006.

[8] M. A. Perillo and W. B. Heinzelman. Optimal sensor management under energy and reliability constraints. In *IEEE WCNC*, pages 1621–1626, 2003.

[9] A. Sekhar, B. S. Manoj, and C. S. R. Murthy. Dynamic coverage maintenance algorithms for sensor networks with limited mobility. In *PERCOM '05*, pages 51–60, Washington, DC, USA, 2005. IEEE Computer Society.

[10] K.-P. Shih, Y.-D. Chen, C.-W. Chiang, and B.-J. Liu. A distributed active sensor selection scheme for wireless sensor networks. *IEEE ISCC '06.*, pages 923–928, June 2006.

[11] A. Skordylis and N. Trigoni. Delay-bounded routing in vehicular ad-hoc networks. In *MobiHoc '08*, pages 341–350, New York, NY, USA, 2008. ACM.

[12] G. Wang, G. Cao, and T. L. Porta. A bidding protocol for deploying mobile sensors. In *ICNP '03*, page 315. IEEE Computer Society, 2003.

[13] H. Wu, R. Fujimoto, R. Guensler, and M. Hunter. Mddv: a mobility-centric data dissemination algorithm for vehicular networks. In *VANET '04*, pages 47–56, New York, NY, USA, 2004. ACM Press.

[14] T. Yan, T. He, and J. A. Stankovic. Differentiated surveillance for sensor networks. In *SenSys '03*, pages 51–62, New York, NY, USA, 2003. ACM.

[15] J. Zhao and G. Cao. Vadd: Vehicle-assisted data delivery in vehicular ad hoc networks. In *Conf. on Computer Communications (INFOCOM)*, 2006.