Multimedia Traffic Information in Vehicular Networks[*]

Ouri Wolfson Bo Xu
Department of Computer Science
University of Illinois at Chicago, USA
{wolfson,boxu}@cs.uic.edu

Hyung Ju Cho
Daum Communications, Jeju, South Korea
hjcho@daumcorp.com

ABSTRACT

In this paper we consider a novel multimedia application, in which drivers may query multimedia clips captured by smartphones mounted on other vehicles. These multimedia clips visualize and voice-indicate the real-time traffic conditions on road segments ahead. We designed a systematic and exhaustive set of query processing strategies which differ from each other in terms of push versus pull, whether infrastructure communication is utilized, and whether metadata dissemination is separated from multimedia clip dissemination. We analyze these strategies theoretically and by simulations, and identify the one that is superior to the others.

Categories and Subject Descriptors
H.2.8 [**Database Applications**]: Spatial databases and GIS

General Terms
Algorithms, performance, design.

Keywords
Multimedia traffic information, vehicular networks, metadata, query processing, Wi-Fi, cellular.

1. INTRODUCTION

Have you ever been stuck in traffic and wondered whether the tie-up/accident is 200 meters or 5 miles ahead? Is it a stalled car that is being pushed to the shoulder, or is it an accident with several emergency vehicles on site? In other words, should you take the first possible exit, or stay on? The application of mobile wireless communication technology to provide traffic information can help answering these questions, and constitutes a challenging research area. Several recent papers present a variety of applications ranging from a reduction in the number of accidents by means of brake warning, intersection assistance, or collision avoidance systems, to offering guidance to available parking lots, discovering the traffic situation on a planned route, coordinating car flow and traffic lights, environmental monitoring and distributed surveillance, and dissemination of traffic information.

This paper addresses another challenging application of delivering traffic multimedia clips among moving vehicles in order to warn drivers about traffic jams and dangers. In this application, each vehicle periodically produces short multimedia clips (e.g. 2-4 seconds) consisting of audio and/or video of the traffic conditions surrounding it, e.g. a voice alert "Overturned car on I287 at Willow" produced by a passing driver, or a video clip of the incident. (A video clip can be captured without driver intervention, by mounting the smart phone equipped with a video camera on the car dashboard).

The multimedia clips are requested by trailing vehicles that issue traffic multimedia queries. Such a query issued at a vehicle requests clips produced on its route ahead. For example, a query may be: Get video clips produced by vehicles on my route 2 kilometers ahead of me.

In this paper we propose a system that disseminates and queries traffic multimedia clips by inter-vehicle communication. We refer to this system as *TrafficMedia*. Inter-vehicle communication can be either by Wi-Fi (short range), or by cellular infrastructure (long range), or by a combination of the two. The Wi-fi communication uses standard vehicular dissemination protocols and techniques.

TrafficMedia does not require a central database server, so the cost of establishing and maintaining it is avoided. This is a significant advantage over a cloud-computing solution such as ICEDB [6]. Furthermore, considering that a large city has millions of vehicles, a central-server solution will also have problems of bandwidth and performance. Also, remember that Wi-Fi communication is free. Thus, Trafficmedia is expected to be more scalable and cost-effective than a cloud architecture providing the same information.

Considering the potential volume of multimedia traffic information and the unpredictability of vehicular networks (e.g., partitions, high mobility), efficient querying is very challenging. We abstract the problem as follows. A database consisting of *multimedia reports* is geographically distributed among a set of *mobile nodes*. Each report consists of a *multimedia* sub-report, and a *metadata* sub-report which includes time and location information. Queries are disseminated as *query-reports*. The match (yes or no) between a query and a report is determined solely based on the *metadata* sub-report.

In this paper we conduct an analysis of query-processing strategies in TrafficMedia. The reason for the multitude of strategies is that choices need to be made along the following design dimensions. First, the inter-vehicle communication may use purely Wi-fi or it may use both Wi-fi and cellular

[*] This research was supported by NSF DGE-0549489, IIS-0847680, 0513736, 0326284, and 0957394.

communication (i.e., hybrid). Pure cellular communication is not an option because it requires knowledge of network id of the receiver, whereas a TrafficMedia query originator initially only knows the time and location of the requested reports but not the network-id of the vehicles containing those reports. Second, the query processing may adopt push (data-to-query) or pull (query-to-data). In the push fashion, multimedia reports are proactively disseminated. In the pull fashion, queries are proactively disseminated; multimedia reports are disseminated as responses to received queries. Push and pull can also be combined. Finally, due to size-differences, the metadata and multimedia sub-reports of a given report may be disseminated independently, and in different ways.

In this paper **we define a set of query processing strategies** based on the above design dimensions. Then we analyze them and **identify four that dominate the others**. Finally, we **compare these four strategies by simulation**, and determine that one is superior by far. It is the one which disseminates queries and metadata sub-reports by Wi-fi; when a matching metadata and query meet at a vehicle, the corresponding multimedia report is transferred to the query originator by cellular communication.

2. THE MODEL

2.1 The Environment

The environment consists of a set of *vehicles* moving on a road network. Each vehicle is equipped with the following capabilities: (i) multimedia capturing; (ii) GPS positioning; (iii) short-range peer-to-peer communication capability Wi-fi; (iv) infrastructure based communication such as 3G cellular. All these capabilities can be provided by a smartphone that is mounted on the vehicle dashboard. Via the infrastructure, a smartphone is able to transmit messages to another smartphone by MMS (Multimedia Message Service) or TCP/IP communication. This is referred to as the *cellular-channel* or *cellular* communication. Each smartphone has a *network-id* that is used as its address for cellular communication, and this id is required in order to send a message to the vehicle via the cellular channel. The network-id can be a cell-phone number or an IP address. In addition, vehicles can communicate via the Wi-fi channel if they are within transmission range. The network id is not necessary for this purpose. With GPS, the time is automatically synchronized among all the vehicles without any extra communication.

2.2 Reports and Reports Databases

Each vehicle periodically captures short multimedia reports of the traffic conditions surrounding it[1]. Formally, a *multimedia report* R, is a couple *<Meta(R), Media(R)>*, where *Meta(R)* and *Media(R)* are the *metadata* and *multimedia* sub-reports, respectively. *Meta(R)* is a 16-byte identifier consisting of:
<producer, produce-time, produce-location>, where:

- *producer:* the network-id of the vehicle that produced R
- *produce-time*: the time when R starts to be captured
- *produce-location*: the (x,y) location where R starts to be captured

Media(R) is a multimedia blob. Its size in the simulations is taken to be 65K bytes. A vehicle is called the *producer* of the reports that it captures.

A vehicle may issue queries. A vehicle is the *producer* of all the queries it issues. Each query is represented by a *query-report* which is a 4-element tuple defined as follows:

<producer, time, location, target-region>

- *producer*: The network-id of the query producer
- *time*: The time at which the query is issued
- *location*: The location at which the query is issued
- *target-region*: A geographic region of interest. It indicates that the query producer is interested in receiving multimedia clips that started to be captured in this region.

Each multimedia report and each query report has a *lifetime* which defines the period of time during which it is considered valid. In our analysis it is assumed that all the reports have the same lifetime, e.g. 5 minutes. A report is dropped by the vehicles when its lifetime expires. A multimedia-report R, or its metadata-report, *satisfy* a query-report Q if (i) R is produced after the produce-time of Q; (ii) R.location falls within Q.target-region.

Each vehicle maintains a reports database that consists of three relations, namely a metadata-reports relation (MRR), a multimedia-reports relation (MuRR), and a query-reports relation (QRR). These relations store the metadata-reports, video-reports, and query-reports, produced by the vehicle or received from other vehicles, respectively.

2.3 Wi-fi Communication

Wifi communication at a vehicle is implemented by a *Wifi Communication Method (WCM)* that is invoked by the query processing strategy executing at the vehicle. Each invocation provides the Wifi communication method with a set of reports to be transmitted. Consequently, at any point in time the WCM stores a *transmission set*, i.e. reports to be transmitted to neighbors. Addition of a report to the transmission set is an idempotent operation, i.e. adding the report again after its first addition does not change anything. The WCM may service multiple applications, some different than TrafficMedia, and thus the WCM transmits the reports in the transmission set with a frequency that will optimize communication for all the applications. Furthermore, due to bandwidth limitations a transmission may not be able to send all the reports in the transmission set, thus the reports are prioritized by the WCM. The transmission frequency and priority are outside the control of the query processing strategy, and are executed by the WCM such that the Wifi communication efficiency is maximized. When the lifetime of a report expires, then it is dropped from the transmission set.

Due to Wifi communication errors and limited bandwidth, it is more likely that long (multimedia) reports get lost or delayed than short ones (queries and metadata sub-reports). Thus a short

[1] For example, a 2-second video clip is captured every 30 seconds at the vehicle. Our empirical study has shown that a 2-second video clip with the size of 65K bytes is able to provide a satisfactory perception of traffic condition to a viewer [5].

report may propagate differently from a long one, even if initially they are both broadcast simultaneously from the same vehicle.

3. Query Processing Strategies

Since queries do not have the network id of the target vehicle, all query processing strategies must start with some Wifi dissemination to neighboring vehicles. The dissemination may be of the query, the multimedia reports, the metadata sub-reports, or some combination. When a match is found, it may be followed by additional cellular or Wifi communication (e.g. the match is between a query and a metadata report, the multimedia sub-report may be located at another vehicle and has to be transferred to the query originator). This is the (Wifi-communication, Match, Communication) paradigm, called WiMaC, and all query processing strategies discussed in this paper are special cases of WiMaC.

This section discusses the WiMac strategies, and is organized as follows. In §3.1 we present the structure of the design space, resulting in 13 WiMaC query processing strategies. In §3.2 we prove that four strategies dominate the others.

3.1 Structure of the Design Space

The structure of the design space is depicted in Table 3.1, and explained as follows. As aforementioned, the first stage of the WiMaC paradigm is a Wi-Fi dissemination which finds a match. The Wi-Fi dissemination is possibly followed by a second stage to complete query processing. There are seven design choices for the first stage, depending on whether the query, the multimedia reports, the metadata sub-reports, or some combination thereof are disseminated (see left column of Table 3.1 which indicates the disseminated combination).

In the (media, meta) and (media, meta, Q) choices, multimedia reports and metadata sub-reports are disseminated independently of each other. Observe that a multimedia report contains its metadata sub-report, and the reason for disseminating them independently is that, as discussed in §2.3 they propagate differently in the vehicular system, with the metadata report propagating faster thus possibly meeting more queries within a given time period. In any case, the two sub-reports of a disseminated multimedia report always correspond to each other.

The design choices for the second stage depend on whether the second stage follows at all, and if so whether Wi-Fi or cellular communication is used for this stage (see right column of Table 3.1). Observe that when a query and matching metadata report meet at a vehicle V, then the network id's of both the query originator and the report producer are known to V since they are in the query and metadata report, respectively. Thus the second stage can be conducted via the cellular infrastructure.

Observe that if only multimedia reports are disseminated at the first stage (i.e., the (media) choice), then there is no need for the second stage. This is because in this case the match must be found at the query originator, and the multimedia report constitutes the answer.

Even if the strategy has a second stage, it does not mean that the second stage always executes after a match at a vehicle. For example, for the (media, meta) case both multimedia and metadata reports are disseminated in the first stage of WiMaC. If a match involving a metadata report occurs, then a second stage is necessary to get the actual multimedia report. But if the match involves a multimedia report, then the second stage is not necessary because the multimedia report is already at the query originator. Similarly, for the (media, Q), and (media, meta, Q) cases, a second stage may not occur after a match at a vehicle.

Table 3.1. Design space of the WiMaC paradigm

Notation:
media = multimedia-report, meta = metadata sub-report, Q = query, cell = cellular.
Strategy Names: 1 is (media), 2a is (meta)-WiFi, 3b is (Q)-cell, etc.

No.	Type of reports disseminated in the first stage (always via Wi-Fi)	Communication medium in the second stage
1	(media)	No second stage
2	(meta)	WiFi (2a), cell (2b)
3	(Q)	WiFi (3a), cell (3b)
4	(media, meta)	WiFi (4a), cell (4b)
5	(media, Q)	WiFi (5a), cell (5b)
6	(meta, Q)	WiFi (6a), cell (6b)
7	(media, meta, Q)	WiFi (7a), cell (7b)

In the following subsection we present the 13 WiMaC strategies. Each strategy is denoted as follows. The denotation consists of the strategy number as defined in Table 3.1 and the strategy name. The strategy name is formed as follows. If there is not the second stage, then the strategy is named by the first stage, i.e., (media). If there is a second stage, then the strategy is named by the two stages connected by a "-". For example, 2b (meta)-cell denotes the 2b strategy which disseminates metadata sub-reports in the first stage and uses cellular communication in the second.

3.2 Strategy Dominance Analysis

Let a vehicle receive an answer multimedia report at time t. The *response-time* of the answer is the length of the time period since the answer is produced until t.

We say that a strategy A is *dominated* by another strategy B if the following 4 conditions are satisfied for every multimedia report M:

(1) For every query that M answers, if the answer is received in B, it is also received in A;

(2) For every query that M answers, the response-time of the answer M in B is no higher than that in A.

(3) The Wi-Fi communication cost of M in B is not higher than that in A.

(4) The cellular communication cost of M in B is not higher than that in A.

Intuitively, if A is dominated by B, then the performance and the efficiency of A are no better than those of B and therefore A is not worth further studying. In this subsection we identify the strategies that are dominated.

In the dominance analysis, the communication cost (but not the delay) of query-reports and metadata sub-reports is ignored for Wi-Fi communication. Similarly, the communication cost of these reports is ignored for cellular communication. This is because query-reports and metadata sub-reports are very short. However, the simulations take into account the communication overhead of the query-reports and metadata sub-reports (see §4.1).

We say that strategy A is *weakly dominated* by strategy B if the above dominance relationship only satisfies conditions 1-3, i.e. the cellular communication cost of B may be higher. Weak dominance is appropriate for unlimited data plans offered by some cellular service providers.

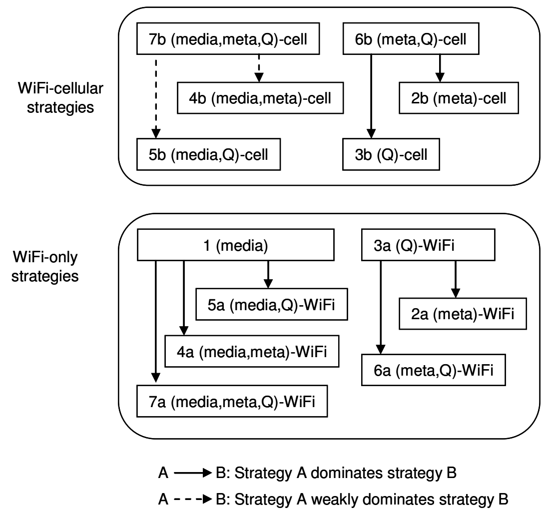

Figure 3.11. Dominance relationship among strategies.

We have proven the dominance relationship depicted in Figure 3.8. Each dominated strategy is dominated by a strategy from the same group and thus is not worth further studying.

The four non-dominated strategies are: 1 (media), 3a (Q)-WiFi, 7b (media,meta,Q)-cell, and 6b (meta,Q)-cell (see Table 3.1). We compared them by simulation. The results of the simulation are given in Figure 4.2, in terms of throughput (average number of answered queries) as a function of the fraction of vehicles that participate in TrafficMedia query processing. They indicate that strategy 6b (meta,Q)-cell performs better than the other three by far. We have also shown that its total communication cost is lower.

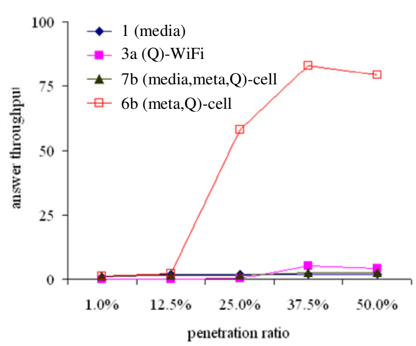

Figure 4.2. Answer throughput in the light-congestion scenario

4. Relevant Work

Video streaming in VANET's. Guo et al. [1] propose a vehicle-to-vehicle live video streaming architecture called V3. V3 adopts a query-to-data paradigm for query processing and it uses only Wi-Fi communication. Thus the query processing method in V3 is similar to the (Q)-WiFi strategy. However, [1] does not compare the query-to-data paradigm with other WiFi-only strategies and with WiFi-cellular strategies, whereas we do so here. PAVAN [2] is a vehicular video streaming system. It uses central servers for the collection and dissemination of metadata information, and uses WiFi inter-vehicle communication for video streaming. Thus PAVAN relies on central servers whereas our system is purely peer-to-peer.

Data dissemination in vehicular ad-hoc networks (VANET's). Data dissemination in VANET's has been extensively studied (see e.g., [3]). Our present work is orthogonal to these papers in the sense that the methods proposed in these papers can be used for Wi-Fi communication in the WiMaC strategies.

Zhu et al. [4] compare WiFi-only dissemination and cellular augmented Wi-Fi dissemination by theoretical analysis using epidemic models. Their application environment is sensor networks. Their work relies on the assumption that each peer knows the network-id's of each other peer, and thus is able to transmit a message to all the other peers via the cellular channel. This assumption is not reasonable in our environment.

5. Conclusion

In this paper we developed and compared strategies for querying traffic multimedia clips in vehicular ad hoc networks, where cellular communication may be available. We designed a systematic approach that considers all the possible design choices in terms of push versus pull, whether infrastructure communication is utilized, and whether metadata dissemination is separated from multimedia clip dissemination. This approach generates a list of 13 possible query processing strategies. By analysis, we identified 4 strategies that dominate the others. Via extensive simulations we found that (i) The strategy with the highest answer throughput and the lowest communication overhead is the one that combines push and pull, uses cellular, and separates metadata dissemination from its multimedia report; (ii) When only WiFi communication is used, push is better than pull when the WiFi network is sparse and pull is better than push otherwise.

6. REFERENCES

[1] M. Guo, M. Ammar, E. Zegura. V3: A vehicle-to-vehicle live video streaming architecture. *PerCom* 2005.

[2] S. Ghandeharizadeh, S. Kapadia, and B. Krishnamachari. PAVAN: A policy framework for content availability in vehicular ad-hoc networks. *VANET*, 2004.

[3] U. Lee, et al. Dissemination and Harvesting of Urban Data Using Vehicular Sensing Platforms. IEEE Transactions on Vehicular Technology, (58)2, 2009, pp. 882-901.

[4] X. Zhu, B. Xu, O. Wolfson. Spatial Queries in Disconnected Mobile Networks. *ACMGIS* 2008.

[5] http://www.cs.uic.edu/~boxu/video_clips/video_13126071744_11-21-24.asf

[6] Y. Zhang, et al. ICEDB: Intermittently-Connected Continuous Query Processing. ICDE 2007.