

Efficient Vehicle-To-Pedestrian Exchange of Medical Data: An Empirical Model with Preliminary Results

Gustavo Marfia
University of Bologna
Mura Anteo Zamboni 7
40127 Bologna (Italy)
+39-051-2094503
marfia@cs.unibo.it

Marco Rocchetti*
University of Bologna
Mura Anteo Zamboni 7
40127 Bologna (Italy)
+39-051-2094503
rocchetti@cs.unibo.it

Claudio E. Palazzi
University of Padua
Via Trieste 63
35131 Padua (Italy)
+39-049-8271426
cpalazzi@math.unipd.it

Alessandro Amoroso
University of Bologna
Mura Anteo Zamboni 7
40127 Bologna (Italy)
+39-051-2094503
amoroso@cs.unibo.it

ABSTRACT

Ambulances and emergency vehicles (buses and taxis as well), if equipped with wireless devices, can be exploited to harvest medical data during unexpected events and also on a daily basis, from all those patients that require a constant monitoring of health conditions. Ambulances can be utilized as trusted intermediaries to transport medical information, at little cost, to hospital central servers. Patients equipped with physiological sensors connected to wireless devices could dump, during each contact, all the medical information collected so far, thus utilizing emergency vehicles as data mules. Inevitably, contact times may be short and not sufficient to transfer all of the information collected from a patient's medical sensors. In such cases, computing in advance, or during the very initial phase of a data transfer, an estimate of how long a contact time will last is key to maximize the utility of any successfully transmitted chunks, in general of different sizes and priorities, of medical data. In this paper we address the problem of predicting patient-vehicle contact times, through an empirical model based on real-world experiments focused on the key sections of a road, which most influence the average speed of an emergency vehicle that traverses it. Our preliminary results are encouraging, as they indicate that it is possible to predict the time an emergency vehicle will spend traversing a given road segment within one third of its traversal.

Categories and Subject Descriptors

J.3 [Computer Applications]: Life and Medical Sciences—health, medical information systems.

General Terms

Algorithms, Measurement, Performance, Design, Experimentation.

Keywords

Biomedical networks, vehicular networks, opportunistic networking.

Corresponding author: Marco Rocchetti. Authors are indebted towards the International FIRB Project DAMASCO.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission from the authors.

MobileHealth'11, May 16, 2011, Paris, France.

Copyright © 2011 ACM 978-1-4503-0780-2/11/05... \$10.00.

1. INTRODUCTION

New wearable medical sensors are expected to revolutionize the healthcare industry in the next few years, leading to a scenario where thousands of people will be constantly equipped with devices that are capable of collecting relevant physiologic data which is in turn sent to central entities located at hospitals or clinics. In such situation, it is key to guarantee a reliable and secure delivery of information, while this traverses different network technologies and hosts.

Without any doubt, cellular networks represent the ideal candidates for all those situations where a fast response is required (e.g., sending an SMS with main vital sign values when an emergency is detected). However, such type of networks hardly suits circumstances where large amounts of data are continuously uploaded (e.g., signal waveforms, images, videos, etc.) by several devices within a cell or, also, when telecommunication infrastructures fail (e.g., terrorist attacks, earthquakes, etc.).

In such cases, a viable solution is provided by the opportunistic exploitation of ambulances and public safety vehicles in general, for the purpose of collecting information via reliable and secure hosts. When an ambulance traverses a given road, for example, it can receive and store, as it moves, medical information from a patient that is walking along that same road. Later, while in proximity of a hospital (e.g., when delivering a patient or at the end of its shift), the same ambulance can connect, through a secure hospital connection to a central server, dumping all the information it received during its service hours.

A successful implementation of such solution heavily relies on an efficient transfer of chunks of medical data during the contact events that occur between a patient and any given ambulance. When a patient and an ambulance meet, two scenarios might occur. The first one results as an outcome of those situations where an ambulance moves so fast that no data can be transferred. The second, contrarily, occurs when an ambulance moves regularly or even slowly, letting a patient's device transfer most of the chunks of data that it recorded so far. In either scenario, an algorithm running

within a patient's medical device can take the right decision if: (a) it knows the size and the priority of each chunk of medical data that is locally stored, and, (b) it knows for how long a contact with an ambulance will last. While we will assume that (a) can in general be known, as it is possible to estimate the importance of the information that is stored in several different ways (e.g., based for example on the disease that is being monitored), we will here focus on (b), showing how it could be possible to predict the contact time between a patient and an ambulance.

We anticipate now how we pursued such task. In particular, performing (b) entails predicting the minimum time a generic vehicle and a pedestrian spend, together, on the same given road segment (i.e., a stretch of a road that lies between two intersections), as, after an intersection, a vehicle may rapidly slip out of sight (abruptly interrupting a communication link). Hence, since a walking pedestrian, in most cases, may be approximated as still during the contact time with a vehicle, our problem translates in predicting the time an emergency vehicle will spend, driving through traffic, over a given road segment. We studied how this could be done performing a wide set of experiments where we repeatedly drove a vehicle equipped with a GPS receiver through a road segment that ended with a traffic light. We then correlated the time we spent on initial subsections of the road with the time taken to exit it to build an empirical model whose results reveal that it is possible to predict how long a vehicle will need to drive through a road while remaining in line of sight with a pedestrian.

This paper is organized as follows. In Section 2 we review the most relevant approaches available in literature that fall closest to ours. In Section 3 we describe our methodology, while in Section 4 our experiments and results. We conclude with Section 5.

2. RELATED WORK

The contemporary introduction of advanced sensor and communication devices has paved the way to the devise of triage and first responder systems thought to rapidly collect information from locations where emergencies have occurred [1]-[3]. In the following we will review the approaches, which, from different perspectives, fall closest to ours.

Authors of [4] are among the first to identify the opportunity of utilizing VANETs for the collection of information from emergency situations. Using their system, when something unexpected occurs and no cellular connections are available, ambulances exploit their farthest relays (i.e., the farthest vehicles equipped with communication devices) to send video streams of the surrounding scene, through multiple vehicles, to the central entities that coordinate first responders. The particular contribution of this work is the devise of a mechanism, which supports each vehicle in identifying its farthest relay,

by dynamically estimating its transmission range. However, little attention is given in this paper to the problem of harvesting information from any physiologic sensors worn by patients.

The work presented in [5], instead, assesses the practical feasibility of collecting any medical information collected by physiologic sensors from a generic vehicle. To do this, the authors of this work performed experiments where they observed the amount of chunks of medical data that could be transferred from a patient to a vehicle moving at a fixed speed, utilizing an 802.11g communication link. This study, however, does not take into account the fact that vehicles rarely traverse roads at a fixed speed, due to the presence of vehicular queues and traffic light phases. Hence, this work did not consider that the amount of information that it may be possible to transfer during a contact might greatly vary, depending on traffic conditions.

The estimation of the length of contact times for a medical monitoring system has been first addressed in [6]. In this work, authors aim at disseminating medical data towards a central entity exploiting the contacts that occur between patients, thus taking advantage of the wealth of research that has so far investigated human mobility models. However, the authors limit their investigation to pedestrian mobility models, not considering the opportunities that might emerge from the use of emergency vehicles.

In summary, to the best of our knowledge, there is no trace in literature of approaches that predict the contact time between a moving emergency vehicle and a patient in order to optimize their exchange of data and, thus, lower the time taken to send the information gathered by medical sensors to a central entity of a hospital or clinic.

3. EXPLOITING CONTACT TIMES

Assuming an ambulance is able to predict the time it will spend over a given road segment, since the very first moment it enters it, a communication protocol with the medical devices worn by a patient can operate as follows.

As soon as the ambulance enters the road, it sends a broadcast message advertising its presence on that road and its position. Any wireless medical sensor, worn by a patient that is walking along that road, answers, utilizing a transmission mechanism which privileges a patient, based on a given priority. Whether using a contention or a priority mechanism (as discussed in the literature), this does not have any effect on the validity of our approach and can be left open to future investigation. On receiving an answer from the selected patient, the other medical devices abort transmitting their answers, whereas, the ambulance in reply sends a message that allows only the selected wireless sensor to dump its chunks of data. As no information is known, yet, regarding the expected contact time with the ambulance, the patient's medical device begins sending chunks assuming the contact will last the time taken to

traverse the road while driving at the speed limit (the length of this time window can be provided by the ambulance within its reply, for example). In the meanwhile, the ambulance moves and predicts how long it will spend over that given road segment. As soon as this information is known, it is provided to the medical device, which can, in turn, maximize, based on a function that maps the priority and size of each chunk into a *utility value* of the transferred data within the given time window (i.e., remaining contact time).

The problem, at this point, results in predicting the time an ambulance will spend upon a given road segment as explained in the following Section where an empirical model has been derived to this aim.

4. AN EMPIRICAL MODEL

Differently from highways, urban roads present traffic patterns that can radically change over a short time range, thus complicating the task of predicting the time that could be spent to traverse them. For this reason it is not possible to assess these values statically, by, for example, simply considering the length and the speed limit of a given road. Moreover, it is usually not even possible to directly compute their traversal times from the time taken, for example, to traverse their initial 10%, as in urban scenarios most of the time is usually spent in the final section, waiting for a red traffic light cycle.

We decided, hence, to study how the traversal time of a given road can be predicted from the time that is spent to drive through a significant initial portion of the same road, performing a wide set of experiments, driving several times through the same street while recording the GPS traces of our traversals.

The chosen experiment site is a 430 meters long street in Pisa, Italy (via Benedetto Croce). The street ends with two traffic lights (approximately distant 30 meters one from the other) and is composed of three lanes, of which vehicles can traverse only the central one from east to west. The timing of both traffic lights depends on pedestrian crossing requests. If no pedestrian requests to cross the street, the traffic light is constantly green. When a pedestrian requests to cross the street, the two traffic lights behave in a slightly different way. Both of them have a complete cycle of 85 seconds (i.e., red plus green light duration), but a red light lasts for 20 seconds at the first one, and for 30 at the second (in order of appearance). When pedestrians request to cross the street, both traffic lights become red at the same time.

4.1 Predicting Traversal Times

The histogram shown in Figure 1 shows the distribution of the traversal times experienced during the experimental campaign we performed over via Benedetto Croce. The wide range of values taken by this variable confirm that the number of chunks of data that can be transferred during a contact can greatly vary, depending on traffic conditions.

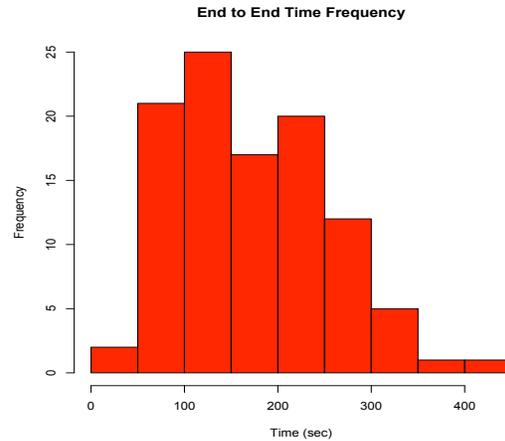


Figure 1. Traversal Time Distribution.

Now, in order to identify significant subsections of the road under investigation, we defined the two following functions:

1. $avg(speed_{x,y})$, the mean value of the average speeds experienced by our vehicle during each traversal, between road positions x and y ;
2. $std(speed_{x,y})$, the standard deviation of the average speeds.

At the basis of our model lies the idea of identifying that section of road whose traversal time highly correlates with the total amount of time needed to drive through the entire road. With an estimate of the time spent on that given portion of road, we can suggest to a running ambulance how long it will remain on that road, and hence the contact time with the patient. Obviously, the sooner this estimate is available, the more useful. Practically, this translates into the problem of finding that portion of road that is closest to the entrance and has the maximal correlation, in terms of traversal time, with the entire road. To this aim, we generated four different objective functions obtained as a combination of the average speed and its standard deviation. As shown in Table 1, it is interesting to maximize or minimize them with aim of identifying the estimate mentioned before. In fact, the meaning of searching for locations where the average speed took high or low values is motivated by the fact that low values are in general experienced near the traffic light, while high values are located typically between the beginning and the midpoint of the road.

For what concerns the standard deviation, we can expect that low values could be seen where, independently from traffic conditions, speeds did not vary (e.g., at the end of a road, close to its traffic light). We can, instead, predict that high values be observed over those portions where the vehicle's speed was most unpredictable, hence over those sections that might have been occupied by queues when, during rush hours, high volumes of traffic entered the road.

In particular, Table 1 shows, from left to right, the optimization made to our objective function, the

Table 1. Optimizations.

Optimization	Section	Avg. speed, std. dev. [m/s]	Corr.
$\max_{x,y}\{avg(speed_{x,y})\}$	[$x = 0.00$, $y = 0.14$]	6.23, 2.93	0.45
$\max_{x,y}\{std(speed)\}$	[$x = 0.20$, $y = 0.57$]	5.03, 4.01	0.89
$\min_{x,y}\{avg(speed)\}$	[$x = 0.91$, $y = 0.92$]	2.11, 0.90	-0.03
$\min_{x,y}\{std(speed)\}$	[$x = 0.98$, $y = 0.99$]	2.20, 0.67	0.08
$\max_{x,y}\{avg(speed) + std(speed)\}$	[$x = 0.15$, $y = 0.36$]	5.63, 3.92	0.75
$\max_{x,y}\{avg(speed) - std(speed)\}$	[$x = 0.08$, $y = 0.14$]	5.63, 2.10	0.20
$\min_{x,y}\{avg(speed) + std(speed)\}$	[$x = 0.98$, $y = 0.99$]	2.19, 0.67	0.08
$\min_{x,y}\{avg(speed) - std(speed)\}$	[$x = 0.45$, $y = 0.84$]	3.32, 2.77	0.77

normalized portion of road (i.e., 0 corresponds to the beginning, 1 to the end of the road) where the functions take their optimal values, the average and standard deviation of the speeds over those portions and, the linear correlation between the time spent over those portions and the time spent over the whole road, respectively.

Our model reveals that the portion of road where the standard deviation of speeds takes a high value corresponds to the section where highest is the recorded correlation with the traversal time (second line of Table 1). Unfortunately, this estimate can be taken only after more than half of the road has been traversed, thus reducing its practical utility.

This fact is also witnessed by Figure 2, where we can appreciate the linear correlation between the times spent over the section where the highest standard deviation is experienced, divided by its length, on the x-axis, and the time spent over the whole road, on the y-axis. This section stretches well beyond the midpoint of the road under consideration. A better estimate can be obtained based on the optimization reported on the fifth line of Table 1. It shows that the section where the sum of the average speed and the standard deviation is highest also reveals a strong correlation with the traversal time (0.75). This estimate may be obtained within approximately one third of the road, thus enabling a quicker prediction of the traversal time of the whole segment.

5. CONCLUSION

Computing during the very initial phase of a data transfer an estimate of how long a contact time will last is key to successfully optimize patient to ambulance

communications. We provided an intuition of how this can be accomplished, supported by preliminary results.

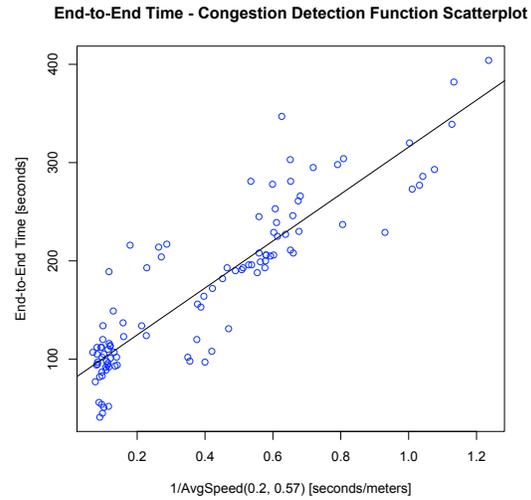


Figure 2. Correlation.

6. REFERENCES

- [1] Voskarides, S., Pattichis, C.S., Istepanian, R., Kyriacou, E., Pattichis, M.S., Schizas, C.N. 2002. Mobile health systems: a brief overview. In Proceedings of the SPIE AeroSense 2002.
- [2] Gao, T., Massey, T., Selavo, L., Crawford, D., Chen, B., Lorincz, K., Shnyder, V., Hauenstein, L., Dabiri, F., Jeng, J., Chanmugam, A., White, D., Sarrafzadeh, M., Welsh, M. The Advanced Health and Disaster Aid Network: A Light-weight Wireless Medical System for Triage. IEEE Trans. on Biom. Circ. and Syst. (IEEE TBCAS), Vol. 1, No. 3, pp. 203-216. September 2007.
- [3] Khan, A.M., Lee, Y-K., Lee, S., Kim, T-S. Accelerometer's position independent physical activity recognition system for long-term activity monitoring in the elderly. Med. Biol. Engineering and Computing, Vol. 48, No. 12, pp. 1271-1279. 2010.
- [4] Rocchetti, M., Gerla, M., Palazzi, C.E., Ferretti, S. and Pau, G. 2007. First Responder's Crystal Ball: How to Scry the Emergency from a Remote Vehicle. In Proceedings of the IEEE International Conference on Performance, Computing and Communications (New Orleans, AL, April 11 - 13, 2007). IPCCC '07.
- [5] Noshadi, H., Giordano, E., Hagopian, H., Pau, G., Gerla, M. and Sarrafzadeh, M. B. 2008. Remote Medical Monitoring Through Vehicular Ad Hoc Network. In Proceedings of the IEEE Vehicular Technology Conference (Calgary, BC, September 21 - 24, 2008). VTC '08.
- [6] Massey, T., Marfia, G., Stoelting, A., Tomasi, R., Spirito, M., Sarrafzadeh, M., Pau, G. Leveraging Social System Networks in Ubiquitous High Data Rate Health Systems. IEEE Trans. Inform. Tech. Biom., 2011, in press.