

## A Network Tomography Approach for Traffic Monitoring in Smart Cities

Journal:	<i>Transactions on Intelligent Transportation Systems</i>
Manuscript ID	T-ITS-17-05-0511.R1
Manuscript Type:	Smart and Green Transportation for Smart Cities
Date Submitted by the Author:	13-Dec-2017
Complete List of Authors:	Zhang, Ruoxi; Missouri University of Science and Technology, Computer Science Newman, Sara; Missouri University of Science and Technology, Computer Science Ortolani, Marco; Universita' degli Studi di Palermo, Computer Science Silvestri, Simone; University of Kentucky, Computer Science
Keywords:	Traffic monitoring, Network Tomography, Smart Cities
Abstract:	<p>Traffic monitoring is a key enabler for several planning and management activities of a Smart City. However, traditional techniques are often not cost efficient, flexible, and scalable. This paper proposes an approach to traffic monitoring that does not rely on probe vehicles, nor requires vehicle localization through GPS. Conversely, it exploits just a limited number of cameras placed at road intersections to measure car end-to-end traveling times. We model the problem within the theoretical framework of network tomography, in order to infer the traveling times of all individual road segments in the road network.</p> <p>We specifically deal with the potential presence of noisy measurements, and the unpredictability of vehicles paths. Moreover, we address the issue of optimally placing the monitoring cameras in order to maximize coverage, while minimizing the inference error, and the overall cost. We provide extensive experimental assessment on the topology of downtown San Francisco, CA, using real measurements obtained through the Google Maps APIs, as well as on realistic synthetic networks. Our approach provides a very low error in estimating the traveling times of more than 95% of all roads even when as few as 20% of road intersections are equipped with cameras.</p>

# A Network Tomography Approach for Traffic Monitoring in Smart Cities

Ruoxi Zhang, Sara Newman, Marco Ortolani, and Simone Silvestri\*

**Abstract**—Traffic monitoring is a key enabler for several planning and management activities of a Smart City. However, traditional techniques are often not cost efficient, flexible, and scalable. This paper proposes an approach to traffic monitoring that *does not* rely on probe vehicles, nor requires vehicle localization through GPS. Conversely, it exploits just a limited number of cameras placed at road intersections to measure car end-to-end traveling times. We model the problem within the theoretical framework of network tomography, in order to infer the traveling times of all individual road segments in the road network. We specifically deal with the potential presence of noisy measurements, and the unpredictability of vehicles paths. Moreover, we address the issue of optimally placing the monitoring cameras in order to maximize coverage, while minimizing the inference error, and the overall cost. We provide extensive experimental assessment on the topology of downtown San Francisco, CA, using real measurements obtained through the Google Maps APIs, as well as on realistic synthetic networks. Our approach provides a very low error in estimating the traveling times of more than 95% of all roads even when as few as 20% of road intersections are equipped with cameras.

**Index Terms**—Traffic monitoring, network tomography, smart cities.

## I. INTRODUCTION

The Smart City paradigm is constantly gaining momentum, also thanks to innovative applications for the efficient management of a city’s assets leveraged by the availability of pervasive monitoring devices. According to the Smart City Council [1], efficient traffic monitoring is a key enabler to improve urban livability, and sustainability by optimizing traffic flow, road construction, and urban planning in general [2].

Collecting continuous and detailed traffic information is challenging. Previous work in this area has often relied on the deployment of traffic sensors, such as inductive loop detectors embedded in pavements [3], or cameras placed at strategic spots [4], but those approaches generally incur in high deployment and maintenance cost, and can only provide local information. Other strategies have been proposed that make use of probe vehicles, such as taxis, buses and Unmanned Aerial Vehicles [5]–[7]; however, probe vehicles generally do not represent the actual traffic flow and are not pervasively available in many cities. Crowd sensing technologies, such as smart phones, have also been proposed for the evaluation of urban dynamics [8]; although potentially effective, these approaches may pose a threat for the user’s privacy, and can be the target of cyber attacks [9].

An interesting alternative approach that has been discussed in the literature is based on *Network Tomography*, an efficient

theoretical tool to estimate the internal state of a network by relying only on end-to-end measurements [10]. This framework has been traditionally investigated in the context of computer networks, but it has also been applied to other contexts, including vehicular traffic estimation [11], [12]. A limited number of *monitors* are coupled to nodes at the edge of the network, so that they can exchange *probing packets* to collect end-to-end measurements of the network paths. The end-to-end measurements are related to the unknown individual link measurements by solving a linear system of equations. Clearly, monitors need to be strategically placed in order to maximize the number of *identifiable links*, i.e. the number of links for which a unique solution exists.

In this paper, we use network tomography to devise a novel approach for vehicular traffic monitoring in Smart Cities. We assume that cameras are statically placed at few selected road intersections, so that they will play the role of monitors in the network tomography terminology. Similarly, cars traversing the roads between pairs of cameras represent the probing packets. At each of the monitored intersections, images of car license plates are captured and sent to a centralized *Traffic Control Center* (TCC) for analysis, so that the corresponding end-to-end traveling times may be inferred. The TCC constructs and solves an optimization problem in order to infer the traffic conditions on all individual roads segments for which only end-to-end traveling times had originally been collected.

The application of Network Tomography to vehicular traffic is not straightforward and several challenges need to be addressed. Specifically, in communication networks probe packets are assumed to follow a predefined (possibly source-assigned) path [10]. This clearly does not hold for vehicles, which are completely independent, and may very well follow multiple paths between any pair of cameras. In addition, end-to-end traveling times are affected by non negligible noise, which might make the resulting linear system impossible.

In order to address those issues, we provide theoretical proofs that allow us to formulate an optimization problem for the camera placement problem to obtain maximum coverage and minimum error, while minimizing the cost of deploying cameras. Furthermore, we formulate a linear optimization problem to minimize the margin of error for measurements, so that it is actually possible to solve the linear system. Finally, we design a greedy approach inspired by the Kernighan-Lin graph partitioning algorithm [13] to assign end-to-end measurements to paths between cameras.

We test our approach on real network topologies of the downtown San Francisco area, and generate realistic traveling times for cars over such topologies by a script based on the Google Maps API [14]. Our results show that our approach is able to cover more than 95% of the road network when as few as 20% of road intersections are equipped with cameras.

R. Zhang and S. Newman are with the Dept. of Computer Science, Missouri University of Science and Technology. M. Ortolani is with the Dept. of Industrial and Digital Innovation, University of Palermo, Italy, and the Dept. of Computer Science, Missouri University of Science and Technology as a Visiting Scholar under a Fulbright Grant. S. Silvestri is with the Dept. of Computer Science of the University of Kentucky.  
\* Corresponding author: silvestri@cs.uky.edu.

## II. RELATED WORK

Traditional approaches to urban traffic monitoring rely on the city infrastructure to gather information about road conditions and usage; for instance, data related to traffic flow is often directly measured through inductive loop detectors embedded in pavements [3], or inferred by cameras placed at strategic spots, such as road intersections [4]. Even though extensive deployment may be devised, such approaches are clearly not very cost-effective, and can hardly obtain pervasive coverage.

Other strategies have been proposed that require some form of collaboration on part of the involved entities, either through devices installed on board of commercial or public transportation vehicles, such as taxis or buses; the authors of [5], for instance, describe a system for real-time estimate of the average speed of traffic through GPS sensors installed in the public transportation fleet of a Greek city. Direct control over the monitoring equipment clearly allows for more flexibility; however bus routes, for instance, are not necessarily representative of the general traffic conditions due to dedicated routes, frequent stops, and so on. Recently, the use of more advanced technology has been proposed; a system employing Vehicular Sensor Networks for traffic estimation and optimization has been proposed in [6], whereas the use of Unmanned Aerial Vehicles (UAVs) for real-time traffic monitoring and management is surveyed in [7]. The intrinsic drawback of such approaches is that an extensive investment needs to be made, which can be prohibitively expensive and might make testing impractical.

Finally, an alternative line of research has suggested the use of crowd sensing technologies, such as smart phones [8] for the evaluation of urban dynamics; such systems, however, impose extra costs to users, necessarily require user-provided information, which can be easily manipulated, and may also pose a threat with regards to users' privacy.

Our goals set our approach apart from those proposals. First of all we do not assume any explicit collaboration from the vehicles, that merely act as "passive" probes, and we do not require cars to be equipped with tracking devices, nor we assume any explicit collaborations from the users. We instead chose to rely on the theoretical grounds provided by network tomography [15]–[18]. Traditional approaches in this field have been proven successful in reconstructing the set of all end-to-end measurements by probing a basis of paths determined by rank decomposition techniques [19]. Nevertheless, these approaches are designed for communication networks, and probe packets are assumed to be source routed. This assumption clearly does not hold in our context, where drivers are free to chose their own path. Only few papers, e.g. [11], consider network tomography applied to vehicular traffic monitoring. However, these papers do not deal with the placement of cameras nor the presence of noisy measurements. In [18], a linear programming formulation is proposed in the context of network tomography in order to maximize the load carried on the network. This formulation differs from the objectives and constraints of the optimization problems studied in this paper, and therefore cannot be directly compared.

Artificial Intelligence approaches to infer individual road

segment delays given some, possibly noisy, traveling times obtained from cameras have been proposed in [12], [20], [21]. These approaches are based on kernel ridge regression [20] and inverse Markov chains [12], [21]. However, these solutions do not address the placement of cameras at intersections to obtain the end-to-end delays, and therefore cannot be directly compared to ours.

Overall, our framework advances previously proposed approaches by comprehensively solving the challenges that arise by applying network tomography to vehicular networks, in order to provide a cost efficient, flexible and accurate way for traffic monitoring in Smart Cities.

## III. BACKGROUND ON NETWORK TOMOGRAPHY

Network tomography was originally designed for communication networks, such as the Internet, to infer the internal network state through end-to-end measurements taken by monitors, typically strategically placed at the edge of the network [22]. A network is modeled as an undirected graph  $G = (V, E)$ , where  $V$  is the set of nodes, and  $E$  is the set of links. Some of the nodes are selected as monitors, and will exchange probing packets in order collect end-to-end measurements of the metric of interest (e.g. the overall delay, or packet loss, across the path between two monitors). With reference to the toy example shown in Figure 1, monitor nodes are indicated by shadowed circles, and thick lines highlight three of the possible paths between any pair of such monitors, namely  $p_1(m_1, m_3) = \{l_1, l_7\}$ ,  $p_2(m_1, m_9) = \{l_1, l_{15}\}$ , and  $p_3(m_9, m_3) = \{l_{15}, l_7\}$ .

The fundamental assumption is the *additive* nature of the metric to be inferred. As an example, the delay of an end-to-end path is the *sum* of the delays of its links. If we let  $b_i$  indicate the measurement obtained by probing path  $i$ , and by  $x_j$  the (yet unknown) delay on link  $j^{th}$ , measurements collected on the path between monitors  $m_1$  and  $m_9$  in Figure 1 would be expressed as:  $x_1 + x_{15} = b_1$ .

In general, let  $P = \{p_1, p_2, \dots, p_{|P|}\}$  represent the set of probing paths between monitors. The relation between paths and links is represented by a binary matrix  $R$ , of size  $|P| \times |E|$ , where each row refers to a specific probing path. More specifically, element  $(i, j)$  of matrix  $R$  is set to 1 if link  $l_j$  belongs to path  $p_i$ , and to 0 otherwise. The end-to-end measurements are stored in a  $1 \times |P|$  vector  $\mathbf{b}$ , whose elements  $b_i$  represent the end-to-end measurements over path  $p_i$ .

Hence, if we let  $x_j$  represent the delay along link  $l_j$ , then the overall delay on path  $p_i$  may be expressed as  $\sum_{j=1}^{|E|} r_{ij} x_j = b_i$ . We can easily extend this over all paths, and formulate the following linear system:

$$R\mathbf{x} = \mathbf{b} \quad (1)$$

where  $\mathbf{x}$  represents the individual link measurements.

In Figure 1, nodes 1 and 9 are selected as monitors. Adding node 3, and solving the corresponding linear system would completely determine the traveling times on all involved links which are thus said to be *identifiable*. Only a limited set of monitors was used. Node 5 in particular was not selected as monitor, so no camera would need to be deployed at the corresponding road intersection. It is




Fig. 1: A sample network with 10 nodes and 19 links; the linear system for the choice of nodes  $\{1, 4, 9\}$  as monitors, and multiple possible paths, and the corresponding matrix.

worth noting that a different choice of monitors, covering a larger part of the network, however, does not necessarily improve identifiability; for instance, with monitors  $\{1, 4, 9\}$  and paths  $p_1(m_1, m_9) = \{l_1, l_7, l_9\}$ ,  $p_2(m_1, m_9) = \{l_1, l_{15}\}$ , and  $p_3(m_9, m_4) = \{l_{15}, l_7, l_9\}$ , the system cannot produce a unique solution for links  $l_7$  and  $l_9$ . Considering additional paths would not guarantee an improvement in identifiability either. Several works focused on monitor placement and path selection to maximize identifiability in communication networks [17], [22]. However, applying network tomography to vehicular networks introduces novel and unique challenges.

#### IV. A FRAMEWORK FOR VEHICULAR TRAFFIC MONITORING BASED ON NETWORK TOMOGRAPHY

The goal of our framework is to exploit network tomography to provide an accurate estimation of the average traveling time for each road segment in a road network relying only on end-to-end measurements detected by cameras placed at road intersections. To this purpose, we begin by stating our model and assumptions, and then describe the three main modules of our framework. Table I summarizes the relevant notation used in the paper.

##### A. Model and Assumptions

Given the map of the Smart City's road network, we create a graph  $G = (V, E)$  where  $V$  is the set of nodes representing road intersections, and  $E$  the set of road segments connecting them. In the following we will use the term "road intersection" and "node" interchangeably; similarly we will use "road segment" and "link" as synonyms. With no loss of generality, we will also assume that all links are symmetrical. Each road segment  $l \in E$  is characterized by an average traveling time  $x_l$ , which represents the time for a car to traverse the link  $l$ , averaged over several measurements in order to account for variability in travel speed.

TABLE I: Summary of relevant notation.

Symbol	Description
$G = (V, E)$	Road network $G$ , $E$ road segments and $V$ intersections.
$V' \subseteq V$	Set of candidate intersections for camera placement.
$M \subseteq V'$	Set of selected intersections for camera placement.
$c_m$	Cost of placing a camera at intersection $m \in V$ .
$l \in E$	Road segment, or link.
$x_l, \mathbf{x}$	Traveling time for link $l$ , vector of link traveling times.
$b_i, \mathbf{b}$	Traveling time for path $p_i$ , vector of path traveling times.
$P$	Set of paths between the monitors in $V'$ .
$R$	Path matrix, $R[i, j] = 1$ if $l_j$ belongs to $p_i$ , 0 otherwise.
$P_{s,d}$	Set of paths between cameras $m_s, m_d \in M$ .
$T_{s,d}$	Set of measurements collected by cameras $m_s, m_d \in M$ .
$T_i$	Set of measurements assigned to path $p_i \in P_{s,d}$ .

Our only requirement in terms of hardware is that a limited number of cameras are deployed at road intersections in order to collect the images of license plates belonging to cars going through that intersection. We do not concern ourselves here specifically with the issue of automatic detection of license plates, as this is a well investigated problem in image analysis and can be performed with very high accuracy [23].

We assume that cameras can be deployed at a subset  $V' \subseteq V$  of all possible intersections. For each intersection  $m \in V'$  the cost of placing a camera is  $c_m$ . The actual set  $M \subseteq V'$  of intersections where cameras are placed is determined by the approach discussed in Section IV-B. If we consider a pair of cameras deployed at intersections  $m_1, m_2 \in M$ , as a vehicle passes through  $m_1$  and  $m_2$  the corresponding time stamps are sent to the TCC which can thus compute the traveling time for that specific vehicle.

##### B. Camera Placement Problem

The first problem we address is the selection of the set  $M$  of intersections where cameras need to be deployed. Ideally, our choice of nodes should incur minimum cost, while providing maximum coverage and minimum estimation error.

Let us consider the linear system  $R\mathbf{x} = \mathbf{b}$  resulting by placing cameras at all possible intersections in  $V'$ . Obviously, this solution would provide maximum coverage and minimum error, however it would incur a very high cost. By using network tomography, we are able to use a much smaller set  $M \subseteq V'$  and provide both same coverage and error, for a significantly lower cost. To this purpose, we recall the concept of *basis* of the matrix  $R$ .

**Definition 1** (Basis [24]). *Given a binary matrix  $R$  of size  $|P| \times |E|$ , a basis  $B$  is a maximal subset of linearly independent rows (paths).*

In the following theorems, we show that any basis of the matrix  $R$  has two relevant properties, namely it provides both maximum coverage and minimum error. In the following, we use path matrices such as  $R$  as sets of vector, to ease the notation.

**Theorem 2.** *Given a binary matrix  $R$  of size  $|P| \times |E|$ , for any basis  $B$  of  $R$ , if there exists a path in  $P$  that covers a link  $l \in E$ , then there exists a path in  $B$  that covers  $l$ .*

*Proof.* We prove the statement by contradiction. Let us consider a basis  $B = \{v_1, \dots, v_n\}$ , and assume that there exists a link  $l$  that is not covered by any vector in  $B$ . However, since  $l$  is covered in  $R$ , then there must exist at least a vector  $\hat{v} \in R$  that covers  $l$  and, since  $B$  is a basis, it should be possible to express  $\hat{v}$  as a linear combination of  $\{v_1, \dots, v_n\}$ . However,



we assumed that  $l$  is not covered in  $B$ , i.e.,  $v_i[l] = 0$  for each  $i = 1, \dots, n$ . Therefore,  $\hat{v}$  cannot be expressed as a linear combination of the paths in  $B$ , thus  $B$  is not a basis of  $R$ , which leads to a contradiction.  $\square$

We now focus on the properties of a basis relative to the inference error. By solving the linear system  $R\mathbf{x} = \mathbf{b}$ , some links are identifiable. Therefore the error for those links is zero, provided that the measurements in  $\mathbf{b}$  are accurate. However, for links that are not identifiable, multiple values exist that would satisfy the system. Our algorithm for inferring values for unidentifiable links is described in detail in Section IV-C, but intuitively the larger the space of possible choices, the larger the potential inference error for the unidentifiable links.

More formally, let  $Q \subseteq R$  be any set of paths whose end points we are planning to monitor through cameras. Furthermore, let  $Vol(Q)$  be the volume of the polyhedron of feasible solutions for the linear system considering only the equations in  $R\mathbf{x} = \mathbf{b}$  corresponding to the paths in  $Q$ . We denote such smaller system as  $Q\mathbf{x}_Q = \mathbf{b}_Q$ . The larger  $Vol(Q)$ , the higher the error we may cause by picking a point in the polyhedron, conversely if  $Vol(Q)$  is smaller the inference error is also smaller<sup>1</sup>.

Considering all paths in  $R$  obviously allows us to obtain the minimum number of unidentifiable links, as well as the minimum volume of  $Vol(R)$ . It is well known in linear algebra that if we restrict ourselves to consider a basis  $B$  of  $R$ , and solve the corresponding linear system  $B\mathbf{x}_B = \mathbf{b}_B$ , we obtain the same set of unidentifiable links [25]. In the following theorem, we further show that such system implies the same volume for the polyhedron of possible solutions for unidentifiable links (i.e.,  $Vol(B) = Vol(R)$ ), and therefore it provides the same minimum inference error.

**Theorem 3.** Consider a binary matrix  $R$  of size  $|P| \times |E|$  and an end-to-end measurement vector  $\mathbf{b}$ . For any basis  $B$  of  $R$ , the polyhedron of possible solutions for the system  $R\mathbf{x} = \mathbf{b}$  is the same of the polyhedron of the solutions of the system  $B\mathbf{x} = \mathbf{b}_B$ .

*Proof.* A solution  $\mathbf{x}$  for the system  $R\mathbf{x} = \mathbf{b}$  exists if the vector  $\mathbf{b}$  lies in the span of the vectors in  $R$  [24]. Since  $B$  is a basis of  $R$ , they both span the same vector space. As a result,  $\mathbf{x}$  is also a solution for  $B\mathbf{x} = \mathbf{b}_B$ .  $\square$

By combining Theorems 2 and 3 we can conclude that in order to obtain maximum coverage and minimum error, while minimizing the cost of deploying cameras, we need to look for the *basis of the matrix  $R$  with minimum cost*. This may be formally stated as an optimization problem, as follows:

$$\begin{aligned} & \underset{Q \subseteq R}{\text{minimize}} && C(Q) \\ & \text{subject to} && Q \text{ is a basis of } R \end{aligned} \quad (2)$$

where  $C(Q)$  is the cost of deploying monitors at the intersections identified by the set of paths  $Q$ . The solution  $Q$  of the above problem can be easily translated into a set of intersections  $M$  by placing cameras at the end points of each path in  $Q$ .

<sup>1</sup>Note that, upper and lower bounds for the link delays can be easily added to the system to ensure  $Vol(Q) < \infty$ .

---

#### Algorithm 1: Greedy camera placement algorithm.

---

**Input:** Matrix  $R$

```

1  $Q = \emptyset$ ;
2 while  $rank(Q) \leq rank(R)$  do
3    $p^* = \text{path in } R \setminus Q \text{ with minimum cost and linearly independent}$ 
    $\text{from the paths in } Q$ ;
4    $Q = Q \cup p^*$ ;
5   foreach Path  $p \in Q \setminus R$  do
6     Update cost  $c_p$  considering the cameras already required by  $Q$ 
7 return  $Q$ 

```

---

We note that  $C(\cdot)$  is a *submodular* function [26]. Additionally, the linearly independent subsets of  $R$  form a *matroid*. As a result, our problem is the *minimization of a submodular function over a matroid constraint*. The unconstrained minimization of a submodular function can be solved in polynomial time using the Lovász extension [26]. However, even simple cardinality constraints can make the problem hard [26]. For this reason, in this paper we propose a greedy algorithm to solve the optimization problem in Eq. (2).

The algorithm takes as input the matrix  $R$  obtained by placing cameras at all intersections in  $V'$ . It then starts with an empty solution  $Q$  (line 1), which is iteratively expanded by the while loop (lines 2-6). The loop iterates as long as  $Q$  is not a basis (i.e., its rank is less than that of  $R$ ). At each iteration, the path with least cost, which is linearly independent from the paths in  $Q$ , is selected and added to  $Q$  (lines 3-4). The inner for loop (lines 5-6) updates the costs of the paths not in  $Q$  by taking into account the cameras required by the paths in  $Q$ . This is a necessary step to take into account the submodularity of the objective function.

**Complexity Analysis.** The while loop runs at most  $rank(R)$  times. Calculating the rank can be performed in  $O(|E|^3)$  using Gaussian Elimination. Finding the path with minimum cost can be done in  $O(|P|)$ . Finally, the for loop to update the costs performs  $O(|P|)$  iterations for each iteration of the while loop. As a result, the overall complexity is  $O(rank(R) \times (|E|^3 + |P|))$ .

#### C. Solving a Linear System of Noisy Measurements

In this section we address two specific challenges arising when applying Network Tomography to vehicular traffic monitoring. First, traffic conditions, drivers' behavior and other external factors will likely influence the traveling times between any pair of monitored intersections, so we must assume that the corresponding delay measurement will be affected by unpredictable noise. Such noise may result in an inconsistent linear system, with no solution. Second, given a camera placement, there may still be unidentifiable links, and therefore multiple solutions for the linear system. Therefore, in this section we discuss our approach to deal with noisy measurements and to identify a solution that is consistent with the available information and provides a small inference error of the traveling times for unidentifiable links.

1) *Dealing with Noisy Measurements:* Let us consider a given set of cameras  $M \subseteq V'$  obtained as discussed in Section IV-B. As previously mentioned, each pair of cameras  $m_s, m_d \in M$  identifies a set  $P_{s,d}$  of paths; for each path

$p \in P_{s,d}$  the Traffic Control Center (TCC) will compute the average traveling time  $b_p$  which will be used in the linear system. However, not only might  $b_p$  be intrinsically affected by a non-negligible amount of noise due to different cars spending different times to traverse the same path, but cars might as well follow different paths altogether between the same two intersections. As a result, the average  $b_p$  for a path  $p$  may be affected by a significant noise. If the standard network tomography approach were to be used, the linear system in Eq. (1) might have no solution.

In order to deal with this problem, we modify the original formulation to account for the possible presence of noise in the measurements. Specifically, we define a parameter  $\Delta$  representing the allowed margin of noise, and substitute each of the equations in the linear system by two inequalities:

$$\sum_{j=1}^{|E|} r_{ij} x_j \leq b_i + \Delta; \quad \sum_{j=1}^{|E|} r_{ij} x_j \geq b_i - \Delta \quad (3)$$

A sufficiently large value for  $\Delta$  allows to make sense of the available end-to-end measurements, so that a solution for the system may exist. However, an excessively large value for this parameter would result in an inaccurate solution, differing significantly from the actual values of the traveling times of the road segments.

Our goal is thus to find the minimum value of  $\Delta$  that enables a solution of the linear system while also minimizing the error in inferring the road segments traveling times. To this purpose, we define the following linear optimization problem:

$$\begin{aligned} & \text{minimize} && \Delta \\ & \text{subject to} && Rx \leq \mathbf{b} + \Delta, \quad Rx \geq \mathbf{b} - \Delta \\ & && \mathbf{x} \geq 0 \end{aligned} \quad (4)$$

where  $R$  is the binary matrix of size  $|P| \times |E|$ . Note that, here the set  $P$  accounts for the possibility of multiple paths between any two intersections where cameras are installed.

The problem can be solved efficiently by well-known linear optimization algorithms such as the Simplex method [27], and the solution would provide a value  $x_j$  for each road segments  $l_j \in E$ , that we use as an estimate of the average traveling time for that link if  $l_j$  is identifiable.

Note that we may extend the above formulation to include vehicle classes. It may be possible to derive a linear system for each class to obtain more accurate information. However, this may come at the expense of additional noise due to potentially inaccurate vehicle classification, or inherent noise in each class. We will consider vehicle classes as a future extension of our work.

2) *Solution for Unidentifiable Links*: It may be the case that only some road segments are identifiable. For such segments the optimizer would return a unique solution; conversely, the values for non identifiable links would not be not unique. In general, the optimizer provides a solution in the form of a point on the surface of the convex polyhedron  $H$  defined by the constraints of the optimization problem. We already discussed in Section IV-B how to place cameras so that such polyhedron is as small as possible. However, unidentifiable links cannot be avoided in general, so we need a method to pick a “good”

solution within the polyhedron that provides a small inference error.

The intuition behind our approach is that even though the vertices of the polyhedron may be admissible solutions, they are likely not convenient for us. Such solutions are characterized by very high traveling times assigned to few links, and very low ones to others. Such disproportionate polarization is not realistic, therefore in our approach the goal is to return a solution which corresponds to the *centroid* of the polyhedron.

**Definition 4** (Centroid). *Given a polyhedron  $H \subseteq \mathbb{R}^n$  with vertices  $= \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ , the centroid  $\mathcal{B}(H) \in \mathbb{R}^n$  is calculated as  $\mathcal{B}(H) = \frac{1}{n} \sum_{i=1}^n \mathbf{v}_i$ .*

Note that since  $H$  is convex, as it results from the intersection of convex polyhedrons defined by the linear inequalities, the centroid  $\mathcal{B}(H)$  always lies within  $H$ , and therefore it is a feasible solution for the linear system. Calculating the centroid, however, requires the knowledge of all the vertices of the polyhedron. Identifying such vertices is feasible using a similar approach as the one adopted by the Simplex algorithm [27]. However, the number of such vertices can be exponential with the number of constraints. Therefore we propose an alternative approach to provide an approximated centroid.

Our approach consists in generating a set  $W$  of random points that lie within the polyhedron  $H$  identified by the linear system in Eq. (4). We iteratively generate random points, and accumulate in  $W$  those that verify the constraints of the linear system. Subsequently, we calculate the centroid of  $\mathcal{B}(W)$  as an approximation of the centroid  $\mathcal{B}(H)$ . We show in Section V that this approach provides a very accurate estimation of the traveling times of the links in the road network.

#### D. Pairing End-to-End Measurements to Paths

We now focus on the last challenge, and deal with the possible different paths that drivers follow between two monitors. Note that this problem does not occur in the original formulation of network tomography for communication networks, since routing is either source-based or known a priori [28].

As an example of this problem, with reference to the linear system in Figure 1, traveling times  $b_1$  and  $b_2$  are both computed from measurements collected by the same pair of monitors ( $m_1$  and  $m_4$ ); however,  $b_1$  should be associated to path  $p_1 = \{l_1, l_7, l_9\}$ , whereas  $b_2$  to path  $p_2 = \{l_1, l_2, l_8, l_9\}$ .

In general, let  $P_{s,d}$  indicate the set of all paths that may be potentially traveled by cars having two monitored road intersections  $m_s, m_d \in M$  as source and destination, respectively. Paths in  $P_{s,d}$  may differ in length and traffic conditions, and consequently be characterized by significantly different traveling times. Moreover, we do not expect car to be an active part of our framework; in particular, they are not assumed to be equipped with any tracking device, such as GPS, that might enable us to infer which of the paths in  $P_{s,d}$  was actually traversed by a specific car. As a consequence, we are faced with the problem of correctly assigning a delay measurement to a specific path.

Formally, the TCC may compute the traveling time  $t_j$  for each vehicle  $j$  traveling from  $m_s$  to  $m_d$ . Additionally,

---

**Algorithm 2:** Assignment of end-to-end measures to paths between a pair of monitors.

---

```

1 Generate a random assignment  $T_1, T_2, \dots, T_k$ ;
2 Calculate the empirical distributions  $\hat{\mathbb{P}}_i$  for each  $X_i$ ;
3 do
4    $\delta^* = 0$ ;
5   foreach pair of samples  $t_1, t_2 \in T_{s,d}$  do
6     Calculate  $\delta_{1,2}$  as the increase in  $\mathbb{P}(T_1, \dots, T_k)$  by switching
7      $t_1$  and  $t_2$ ;
8     if  $\delta_{1,2} > 0$  and  $\delta^* < \delta_{1,2}$  then
9        $\delta^* = \delta_{1,2}$ ;
10    Switch  $t_1$  and  $t_2$  that provided  $\delta^*$ 
11 while  $\delta^* > \delta_{min}$ ;
12 return the current assignment  $T_1, T_2, \dots, T_k$ 

```

---

in order to build the linear system described in Eq. (4), it needs to compute the average traveling time  $b_i$  for each path  $p_i \in P_{s,d}$ . Let  $T_{s,d} = \{t_1, t_2, \dots, t_r\}$  be the set of end-to-end measurements collected for any of the paths in  $P_{s,d}$  between monitors  $m_s$  and  $m_d$ , and let  $|T_{s,d}| = r$  and  $|P_{s,d}| = k$ .

Our problem consist in determining a partition  $T_1, T_2, \dots, T_k$  of  $T_{s,d}$  such that each  $T_i$  contains only the end-to-end measurements that belong to path  $p_i \in P_{s,d}$ . Once the sets  $T_1, T_2, \dots, T_k$  are identified, the average end-to-end traveling time for a path  $p_i$  will be computed as  $b_i = \frac{1}{|T_i|} \sum_{t \in T_i} t$ . This value can be used as the right-hand side for the inequalities related to path  $p_i$  in the linear system described in Section IV-C.

In order to find the partitioning of  $T_{s,d}$ , we assume that the probability  $\mathbb{P}_{s,d}(p_i)$  with which a car follows a path  $p_i$  between  $m_s$  and  $m_d$  is known, so that we know that the expected number of measurements belonging to set  $T_i$  is given by  $r \times \mathbb{P}_{s,d}(p_i)$ , where we recall  $r = |T_{s,d}|$ . Note that, in Section V-E we relax this assumption and experimentally show that our approach is effective even when there is an inaccurate knowledge of the probability distribution.

We model the traveling times along a path  $p_i$  by means of a random variable  $X_i$ , whose probability distribution is not restricted to have any specific shape.  $T_i$  is then the set of realizations of such random variable. If  $\mathbb{P}(X_i = t)$  indicates the probability that a sample  $t \in T_{s,d}$  has been generated by  $X_i$ , the *likelihood* of an assignment  $T_1, \dots, T_k$  is given by:

$$\mathbb{P}(T_1, \dots, T_k) = \prod_{i=1}^k \prod_{t \in T_i} \mathbb{P}(X_i = t) \quad (5)$$

and the best assignment  $T_1, \dots, T_k$  is one that maximizes such likelihood. In the ideal case when the distributions are known for each variable  $X_i$ , the best partitioning can be determined by assigning each measurement  $t$  to the variable  $X_i$  for which  $\mathbb{P}(X_i = t)$  is the highest.

In our case, however, the distributions are not known, so we resort to an iterative greedy approach inspired by the Kernighan-Lin graph partitioning algorithm [13]. The idea is to start from a random assignment respecting the cardinality constraints for each set  $T_i$ , and iteratively make use of the samples in  $T_i$  to estimate the *empirical distribution*  $\hat{\mathbb{P}}_i$  of each variable  $X_i$ . Then, for each pair of samples  $t_1 \in T_i$  and  $t_2 \in T_j$  we calculate  $\delta_{1,2}$  which is defined as the increase in the likelihood of the assignment in Eq. (5) by swapping  $t_1$  and




Fig. 2: Example of quadtree based synthetic topology [29].

$t_2$ , i.e. by assigning  $t_1$  to  $T_j$  and  $t_2$  to  $T_i$ . The pair of samples with highest  $\delta$  are then switched at the current iteration. The algorithm terminates as soon as the increase drops below a threshold  $\delta_{min}$ .

The pseudo-code for the algorithm is shown in Algorithm 2. Each iteration of the algorithm has  $O(r^3)$  complexity, since each pair of samples  $t_1, t_2 \in T_{s,d}$  must be analyzed, and, for each of these pairs,  $\delta_{1,2}$  must be computed, an operation whose complexity is upperbounded by  $O(r)$ .

In order to ensure convergence of the algorithm,  $\delta_{1,2}$  is calculated as follows. For  $t_1 \in T_i$  and  $t_2 \in T_j$ , we simulate the swapping of  $t_1$  and  $t_2$ , obtaining the corresponding sets  $T'_i = T_i \setminus \{t_1\} \cup \{t_2\}$  and  $T'_j = T_j \setminus \{t_2\} \cup \{t_1\}$  and the new empirical distributions  $\hat{\mathbb{P}}'_i$  and  $\hat{\mathbb{P}}'_j$ , for  $X_i$  and  $X_j$ .  $\delta_{1,2}$  is defined as:

$$\delta_{1,2} = \max\{(\prod_{t \in T'_i} \hat{\mathbb{P}}'_i(X_i = t) - \prod_{t \in T_i} \hat{\mathbb{P}}_i(X_i = t)), 0\} \times \max\{(\prod_{t \in T'_j} \hat{\mathbb{P}}'_j(X_j = t) - \prod_{t \in T_j} \hat{\mathbb{P}}_j(X_j = t)), 0\} \quad (6)$$

The rationale behind the above formula is that whenever  $t_1$  and  $t_2$  are exchanged, a strictly positive increase in the probability of the new assignment occurs even if the empirical distributions of  $X_i$  and  $X_j$  may change.

## V. EXPERIMENTAL ASSESSMENT

### A. Experimental Setup

In this section we experimentally evaluate our framework through simulations. We consider both synthetic and real road networks. Synthetic road networks have been generated using the realistic model proposed in [29]. The model is based on quad trees and has two parameters: the number of intersections  $f$ , and  $q$  which controls the amount of sprawl. In our experiments  $f$  is varied to consider different scenarios and  $q$  is set to 1. Traveling times are randomly generated in the interval  $[3, 10]$ , which represents the ground truth for our inference. An example of the topology generated with this model is shown in Figure 2.

The real road network topology is instead taken from downtown San Francisco, CA, as shown in Figure 3. When required, we generate networks of varying size by considering only the part of the road network comprised within a specified radial distance from geographical center of the map. We developed a script based on the Google Maps API [14] to obtain real traveling times. Specifically, our script uses the APIs to obtain the current traveling time in every road segment in the downtown San Francisco map. Figure 4 shows the



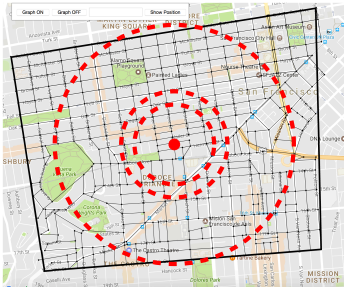


Fig. 3: Road topology used for the experiments.




Fig. 4: Cumulative Density Function (CDF) of delays observed in downtown San Francisco, CA.

Cumulative Density Function (CDF) for the observed delays at 2PM on 8/8/2017<sup>2</sup>. We used these values as ground truth for our experiments with the real road network topology.

For both synthetic and real networks we generated the cost  $c_m$  of deploying a camera at intersection  $m \in V'$  as a random number in the interval  $[1, 10]$ .

In order to account for the fact that in a real-life scenario, drivers may decide to follow alternate, but not entirely dissimilar, routes for each pair of intersections  $m_s, m_d$ , we consider all paths in the set  $P_{s,d}$  whose number of hops is less than or equal to a factor  $\theta$  of the length of the shortest path between  $m_s$  and  $m_d$ . The set  $P_{s,d}$  was constructed through a modified version of the transitive closure of the adjacency matrix of the road map graph  $G$  [30]. In our experiment, we show the impact of the parameter  $\theta$  on the performance of our approach. We also assumed that collected end-to-end traveling times may be randomly corrupted by some amount of noise, and we performed experiments to show the impact of such noise.

We focus on three main metrics to evaluate our approach. *Cost* refers to the total cost of deploying cameras at the selected intersections. *Coverage* is defined as the fraction of links for which an estimation is provided with respect to the total number of links in  $E$ . The inference error is measured using the *Mean Square Error* (MSE). Specifically, let  $x_1, \dots, x_{|E|}$  be the ground truth values of delays in the network, and let  $\hat{x}_1, \dots, \hat{x}_{|E|}$  be the inferred values. The MSE is defined as  $MSE = \frac{1}{|E|} \sum_{i=1}^{|E|} (\hat{x}_i - x_i)^2$ . Note that, for uncovered links we assume that the inferred delay is zero, so that the corresponding error is maximized. The availability of additional information is available, e.g. road lengths and speed limits, would result in better inference.

<sup>2</sup>We also considered different times and days, and obtained similar results.

Results are averaged over several runs in order to obtain a reliable confidence interval. The plots show the average and standard deviation of the considered metrics.

### B. A Comparison Approach based on Vertex Cover

Previous works on vehicular traffic estimation based on cameras [11], [12], [21] do not address the camera placement problem. Similarly, monitor placement in communication networks is based on routing assumptions that do not hold for vehicular networks [16], [22], [31], [32]. For these reasons we chose instead to compare our method to an approach inspired by the Weighted Vertex Cover (WVC) problem [33], which has been adopted in several coverage problems in communication and vehicular networks [34]. Specifically, given the road network  $G = (V, E)$  and the costs  $c_m$  for each  $m \in V$ , we assume that if a camera is placed at an intersection  $m \in V$ , the traffic along all the links in  $E$  incident to  $m$  can be accurately monitored. In terms of the WVC problem, we say that  $m$  covers the links adjacent to it. Under this setting, we are looking for a set  $M_{WVC} \subseteq V$  to place cameras such that it incurs minimum costs and each edge in  $E$  is covered by at least one camera in  $M_{WVC}$ . A similar formulation of the problem is straightforward for the case when cameras are allowed to be placed only at a subset  $V' \subseteq V$  of intersections.

The WVC problem is NP-Complete, however there are several heuristics that provide a provable approximation bound. In this paper we use the heuristic proposed by Clarkson et Al. that provides a 2 approximation bound with respect to the optimal solution [33]. This approach is referred to as WVC.

### C. Synthetic Networks

**Experiment I.** In the first set of experiments we allow cameras to be potentially deployed at any intersection in the network, i.e.,  $V' = V$ . In addition, we consider three values for  $\theta$ , namely 1, 1.2, and 1.5, and we consider no errors in the measurements nor in the assignment of measurements to paths.

Figure 5(a) shows the total cost of deploying cameras at intersections returned by our approach, and at those selected by WVC, under different network sizes. Our approach clearly outperforms WVC. The strength of our solution is the ability to exploit the correlation between end-to-end measurements through the solution of the linear system. Conversely, each camera in WVC is only able to provide a local estimation, although unbiased, of its adjacent links. As a result, WVC requires more cameras and therefore it incurs a higher cost. In addition, our solution requires a lower cost as the value of  $\theta$  increases, since fewer cameras are necessary to cover the network. Nevertheless, even in the extreme case of a single path between each pair of cameras ( $\theta = 1$ ), we still significantly outperform WVC.

Note that, by allowing cameras to be potentially deployed at every intersection, our approach is able to uniquely identify all links in the network. Similarly, WVC returns a set of nodes for which every link is adjacent to at least a node in that set. As a result, both approaches achieve full coverage and incur a null inference error.



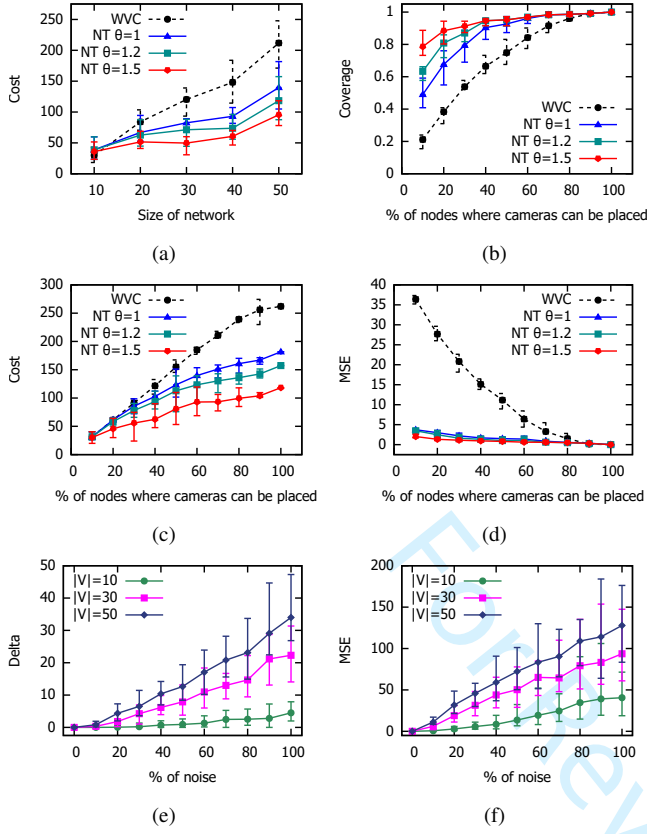


Fig. 5: Synthetic Networks: Cost vs. size of the network (a), Coverage (b), Cost (c) and MSE (d) vs. % of nodes where cameras can be placed,  $\Delta$  (e), MSE (f) vs. % of noise.

**Experiment II.** In this set of experiments, we allow cameras to be placed only at a subset of intersection  $V' \subseteq V$  and we study the performance upon increasing the size of  $V'$ . The intersections in  $V'$  are chosen randomly, and again we assume no errors are present.

Figure 5(b) and (c) show the coverage and the cost of camera deployment versus the size of  $V'$ , expressed as the percentage of nodes in  $V$  where cameras can be placed. Remarkably, our approach is able to provide higher coverage at a smaller cost compared to WTC by exploiting end-to-end measurements. As a numerical example, our approach provides a coverage higher than 90% coverage of the network, with  $\theta = 1.2$ , when having only 30% of intersections in  $V'$  and incurring a cost less than 75. On the contrary, WTC requires at least 70% of intersection available to achieve a similar coverage, and it incurs a cost of about 150.

The lower coverage also has an impact on the quality of the inference, as shown in Figure 5(d). Our approach incurs a very low error for all settings of  $\theta$ , again highlighting the inference capabilities of network tomography when applied to vehicular networks. Conversely, WTC results in an error which is significantly higher than our approach.

**Experiment III.** In the third and final set of experiments, we focus on the effect of noisy measurements. For this purpose, we consider the end-to-end measurements to be affected by a random noise  $e \in [0, 1]$ , and we study the performance by

increasing the value of  $e$ . Specifically, depending on  $e$  we alter the ground truth  $b_i$  value for the delay of path  $p_i$  in Eq. 4, with a random value in the interval  $[(1 - e) \times b_i, (1 + e) \times b_i]$ . This makes the linear system likely to have no solution. Therefore a value of  $\Delta > 0$  is necessary to satisfy all the constraints given the noisy measurements. In these experiments we allow cameras to be potentially deployed at all of intersections.

Figure 5(e) shows the values of  $\Delta$  corresponding to increasing values for the error parameter for network topologies of size 10, 30, and 50 nodes. As expected, a higher value of  $\Delta$  is required as the noise increases, and also a higher value is necessary to satisfy the constraints for bigger topologies. Intuitively, the bigger the topology the higher the number of paths, the higher the impact of noise on the feasibility of the linear system. Figure 5 (f) shows the inference error under different magnitude of noise. The error increases with the amount of noise and the size of the network. This increase is due to the higher values of  $\Delta$ , that result in larger polyhedrons, and consequently to larger inference errors. Nevertheless, it is worth noting that although the error increases, our approach is effectively able to achieve a small inference error in all the considered settings.

#### D. Real Networks

We now describe the results obtained using the road network of downtown San Francisco, CA. We performed similar experiments with respect to the synthetic networks.

**Experiment I.** Figure 6(a) shows the cost of deploying cameras at selected intersections as a function of the network size. The results confirm the outcome obtained for synthetic networks, and also in this case our approach outperforms WTC incurring a lower cost. The advantage becomes more evident as the network size increases, since the effect of exploiting end-to-end measurements is more prominent.

**Experiment II.** We now consider the performance when increasing the size of the set  $V'$  where cameras can be placed. As figures 6(b), (c) and (d) show, WVC reveals its weaknesses in real settings too, achieving lower coverage, at a higher cost, which results in a higher inference error for the link delays. Note that the higher values of the MSE compared to the synthetic networks are due to the higher delays of the links in this dataset, in the interval  $[10, 300]$  as obtained by the script based on the Google APIs. Conversely, the delays for the synthetic networks are in the interval  $[3, 10]$ .

**Experiment III.** In the final set of experiments we added noise to the measurements as discussed for the case of synthetic networks. Figure 6(e) and (f) show the value of  $\Delta$  and the inference error as a function of the magnitude of added noise. Similarly to the synthetic network case, a higher  $\Delta$  and a higher error occur as we increase the noise and the size of the network. Nevertheless, even in this case our approach shows great tolerance to noisy measurements, as the MSE, which is a quadratic error, increases only linearly with the noise.

#### E. Assigning Measurements to Paths

In this set of experiments, we aim to evaluate the performance of the algorithm presented in Section IV-D whose goal is to assign the measurements  $T_{s,d}$  collected between two intersections  $m_s$  and  $m_d$ , to the paths  $P_{s,d}$  between them. We

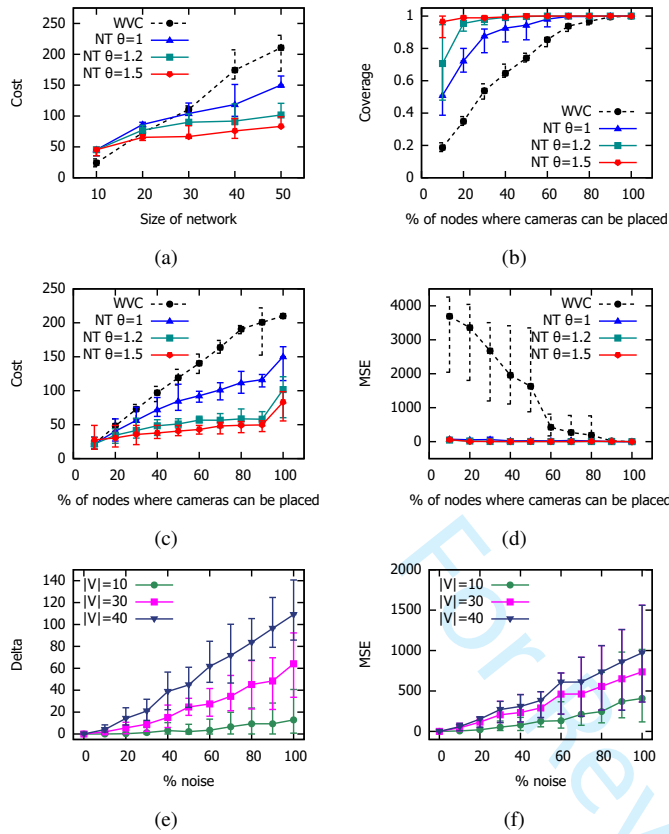


Fig. 6: Real Networks: Cost vs. size of the network (a), Coverage (b), MSE (c) and Cost (d) vs. % of nodes where cameras can be placed,  $\Delta$  (e), MSE (f) vs. % of noise.

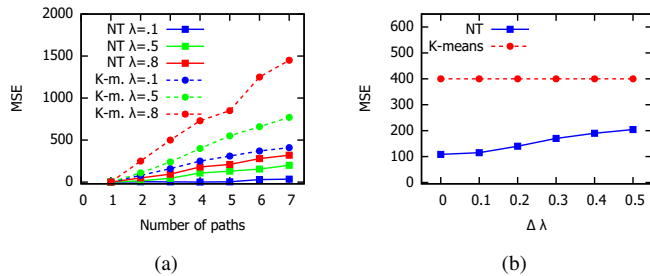


Fig. 7: Assigning measurements to paths, with (a) perfect and (b) inaccurate knowledge of the driver preference distribution.

assume that each path  $p_i \in P_{s,d}$  has a *ground truth* average delay  $\mu_i$  and that the traveling times over  $p_i$  follow a Gaussian distribution  $\langle \mu_i, \sigma \rangle$ . The value of  $\mu_i$  is picked randomly in the interval  $[1, 100]$ , while  $\sigma$  is set to 3. We additionally assume that the driver preference for the available paths follows a geometric distribution with parameter  $\lambda \in [0, 1]$ , meaning that path  $p_i \in P_{s,d}$  is picked with probability  $(1 - \lambda)^i$ . This distribution models the real life scenario in which most drivers tend to concentrate over few paths.

Using these assumptions, we simulate 800 cars traveling from  $m_s$  to  $m_d$ . Each car picks a path  $p_i$  in  $P_{s,d}$  according to the geometric distribution. Then, we add a measurement in the set  $T_{s,d}$  as realization of the Gaussian distributed random

variable with parameters  $\langle \mu_i, \sigma \rangle$ .

We run our algorithm to partition the set  $T_{s,d}$  into  $T_1, T_2, \dots$ , where  $T_i$  contains the measurements assigned to the path  $p_i$ . The estimated average traveling time for  $p_i$  is  $\hat{\mu}_i$ , obtained averaging the measurements in  $T_i$ . We use the Mean Square Error (MSE) between the actual averages and the estimated averages to evaluate the accuracy of our approach, specifically  $MSE = \frac{1}{|P_{s,d}|} \sum_{p_i \in P_{s,d}} (\hat{\mu}_i - \mu_i)^2$ .

We compare our approach with the clustering algorithm K-means [35]. This algorithm starts from  $K$  randomly assigned centroids, where for us  $K = |P_{s,d}|$ , and it assigns the measurement in  $T_{s,d}$  to the closest centroid. At each subsequent iteration, centroids are updated based on the previous measurement assignment, and measurements are assigned again to the closest updated centroid. The process is repeated until a convergence criteria is met. We refer the reader to [35] for more details about K-means.

We initially assume that the parameter  $\lambda$  of the geometric distribution of the driver preference is known. Figure 7(a) shows the results increasing the number of paths in  $T_{s,d}$ . As expected, the availability of a larger number of paths makes the problem harder. Nevertheless, our approach incurs in a small error, and it outperforms K-Means by providing a more accurate estimation of the traveling times in all settings.

Next, we relax the assumption on perfect knowledge of the driver preference distribution. Specifically, we set  $\lambda = 0.5$  but run the algorithm assuming a distribution with parameter  $\lambda + \Delta\lambda$ , where  $\Delta\lambda$  represents the amount of knowledge inaccuracy. Figure 7(b) shows the MSE for 4 paths increasing the value of  $\Delta\lambda$ . K-Means is not affected by  $\Delta\lambda$ , since it does not make use of the underlying driver preference distribution. Our approach shows very high robustness as it is able to provide better accuracy than K-means, even under severe inaccurate knowledge.

## VI. CONCLUSIONS

This paper addressed the issue of vehicular traffic monitoring in a smart city. We described a method for inferring the traveling times on each road segment of the city, while requiring just a minimum number of monitoring devices to be deployed at selected intersections. The theoretical grounds for our method are provided by the Network Tomography approach. We formulated an optimization problem for the optimal placement of monitoring cameras, and exploited linear algebra to propose an efficient greedy solution. We additionally adapted the classical network tomography approach by allowing for noisy measurements and unpredictability of vehicles paths. Experimental results on real and synthetic networks show that our approach provides a cost efficient deployment of cameras that allows to achieve full coverage of the road network and a very low inference error.

## REFERENCES

- [1] Smart Cities Council, "Smart cities readiness guide: The planning manual for building tomorrow's cities today," *Smart Cities Council, Redmond*, 2014.
- [2] M. Batty, K. W. Axhausen, F. Giannotti, A. Pozdnoukhov, A. Bazzani, M. Wachowicz, G. Ouzounis, and Y. Portugali, "Smart cities of the future," *The European Physical Journal Special Topics*, vol. 214, no. 1, pp. 481–518, 2012.

- [3] S. Oh, S. Ritchie, and C. Oh, "Real-time traffic measurement from single loop inductive signatures," *Transportation Research Record: Journal of the Transportation Research Board*, no. 1804, pp. 98–106, 2002.
- [4] S. R. E. Datondji, Y. Dupuis, P. Subirats, and P. Vasseur, "A survey of vision-based traffic monitoring of road intersections," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 10, pp. 2681–2698, 2016.
- [5] E. Denaxas, S. Mpollas, D. Vitsios, C. Zolotas, D. G. Bleris, G. M. Spanos, and N. P. Pitsianis, "Real-time urban traffic information extraction from GPS tracking of a bus fleet," in *IEEE CIVTS*, 2013.
- [6] R. Du, C. Chen, B. Yang, N. Lu, X. Guan, and X. Shen, "Effective urban traffic monitoring by vehicular sensor networks," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 1, pp. 273–286, 2015.
- [7] K. Kanistras, G. Martins, M. J. Rutherford, and K. P. Valavanis, "Survey of unmanned aerial vehicles (uavs) for traffic monitoring," in *Handbook of Unmanned Aerial Vehicles*, Springer, 2015, pp. 2643–2666.
- [8] K. Sohn and K. Hwang, "Space-based passing time estimation on a freeway using cell phones as traffic probes," *IEEE Transactions on Intelligent Transportation Systems*, vol. 9, no. 3, pp. 559–568, 2008.
- [9] K. Atherton, *Israeli students spoof waze with fake traffic jam*, <http://tinyurl.com/p4gcgkv>, 2014.
- [10] T. Bu, N. Duffield, F. Lo Presti, and D. Towsley, "Network tomography on general topologies," in *ACM SIGMETRICS Performance Evaluation Review*, ACM, vol. 30, 2002, pp. 21–30.
- [11] S. Santini, "Analysis of traffic flow in urban areas using web cameras," in *IEEE WACV*, 2000, pp. 140–145.
- [12] T. Morimura, T. Osogami, and T. Idé, "Solving inverse problem of markov chain with partial observations," in *Advances in Neural Information Processing Systems*, 2013, pp. 1655–1663.
- [13] B. W. Kernighan and S. Lin, "An efficient heuristic procedure for partitioning graphs," *The Bell system technical journal*, vol. 49, no. 2, pp. 291–307, 1970.
- [14] Google, *Google maps apis*, <https://developers.google.com/maps/>, 2017.
- [15] A. Gopalan and S. Ramasubramanian, "On identifying additive link metrics using linearly independent cycles and paths," *IEEE/ACM Transactions on Networking (TON)*, vol. 20, no. 3, pp. 906–916, 2012.
- [16] R. Kumar and J. Kaur, "Practical beacon placement for link monitoring using network tomography," *IEEE Selected Areas in Communications*, vol. 24, no. 12, pp. 2196–2209, 2006.
- [17] S. Tati, S. Silvestri, T. He, and T. La Porta, "Robust network tomography in the presence of failures," in *IEEE ICDCS*, 2014.
- [18] A. Medina, N. Taft, K. Salamati, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: Existing techniques and new directions," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 161–174, 2002.
- [19] Q. Zheng and G. Cao, "Minimizing probing cost and achieving identifiability in probe-based network link monitoring," *IEEE Transactions on Computers*, vol. 62, no. 3, pp. 510–523, 2011.
- [20] T. Idé and M. Sugiyama, "Trajectory regression on road networks," in *AAAI*, 2011.
- [21] T. Idé, T. Katsuki, T. Morimura, and R. Morris, "City-wide traffic flow estimation from a limited number of low-quality cameras," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 4, pp. 950–959, 2017.
- [22] L. Ma, T. He, A. Swami, D. Towsley, and K. K. Leung, "On optimal monitor placement for localizing node failures via network tomography," *Performance Evaluation*, vol. 91, pp. 16–37, 2015.
- [23] S. Du, M. Ibrahim, M. Shehata, and W. Badawy, "Automatic license plate recognition (alpr): A state-of-the-art review," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 23, no. 2, pp. 311–325, 2013.
- [24] D. Lay, *Linear Algebra and Its Applications*. Pearson Education, 2002.
- [25] Y. Chen, D. Bindel, H. Song, and R. H. Katz, "An algebraic approach to practical and scalable overlay network monitoring," *ACM SIGCOMM Comp. Com. Rev.*, vol. 34, no. 4, pp. 55–66, 2004.
- [26] A. Krause and D. Golovin, *Submodular function maximization*. 2014.
- [27] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.
- [28] N. Bartolini, T. He, and H. Khamfroush, "Fundamental limits of failure identifiability by boolean network tomography," in *IEEE INFOCOM*, 2017.
- [29] D. Eisenstat, "Random road networks: The quadtree model," in *SIAM ANALCO*, 2011.
- [30] M. Newman, *Networks: an introduction*. Oxford university press, 2010.
- [31] L. Ma, T. He, K. K. Leung, A. Swami, and D. Towsley, "Inferring link metrics from end-to-end path measurements: Identifiability and monitor placement," *IEEE/ACM Transactions on Networking (TON)*, vol. 22, no. 4, pp. 1351–1368, 2014.
- [32] D. Zad, T. He, and T. La Porta, "Parsimonious tomography: Optimizing cost-identifiability trade-off for probing-based network monitoring," *IFIP Performance*, 2017.
- [33] K. L. Clarkson, "A modification of the greedy algorithm for vertex cover," *Information Processing Letters*, vol. 16, no. 1, pp. 23–25, 1983.
- [34] A. Y. Lam, Y.-W. Leung, and X. Chu, "Electric vehicle charging station placement: Formulation, complexity, and solutions," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2846–2856, 2014.
- [35] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge university press, 2003.



**Ruoxi Zhang** is a graduate student and research assistant at the University of Missouri Science and Technology in the Department of Computer Science. She received her Bachelor degrees in Computer Science and Mathematics at the Missouri Southwest Baptist University. Present research interest include network tomography applied to vehicular and communication networks.



**Sara Newman** is an undergraduate student and research assistant at Missouri University of Science and Technology in the Department of Computer Science. She is an active member and officer for her university's student chapter of the Association for Computing Machinery (ACM). Present research interests include network tomography applied to vehicular networks.



**Marco Ortolani** received his Ph.D. degree in Computer Engineering in 2004 from the University of Palermo, Italy. Since 2008 he is an Assistant Professor at the same university. Since January to September 2017 he was a Visiting Research Scholar, at the Comp. Sci. Dept. of the Missouri Univ. of Science and Technology under a Fulbright grant. His current research interests include distributed systems, and intelligent data analysis techniques applied to Ambient Intelligence and the Internet of Things.



**Simone Silvestri** graduated with honors and received his PhD in computer science at Sapienza University of Rome, Italy. He is an Assistant Professor in the Department of Computer Science of the University of Kentucky. His recent research interests include cyber-physical-human systems, internet of things, hybrid sensor networks, and network management. His research is funded by several agencies such as NIFA, NSF and NATO.