Light Pole Localization in a Smart City

Mike Holenderski, Richard Verhoeven, Tanir Ozcelebi, Johan J. Lukkien
Department of Mathematics and Computer Science, Technische Universiteit Eindhoven
Den Dolech 2, 5600 AZ Eindhoven, The Netherlands

***Abstract*—Many smart city lighting applications require information about the location of light poles, in particular about which light poles are neighbors along the street with respect to passing traffic. This paper addresses the problem of deriving the topology of light poles in a smart city, relying only on the data gathered from Passive Infrared Sensors attached to the light poles. A statistical algorithm is presented and evaluated based on a data set gathered from a real deployment in a residential area.**

I. Introduction

Modern streetlights are being equipped with LED luminaries, offering energy efficient lighting with fine grained control of color and intensity. Moreover, the light poles are being equipped with various sensors and a communication infrastructure allowing to control individual light poles. This provides an opportunity for many interesting urban applications aiming to improve energy efficiency, safety, productivity and quality of life in Smart Cities [1], [2].

A common example is the dynamic adjustment of street lighting following a car. To save energy and decrease light pollution in a city, the lights are by default dimmed to a minimum setting. However, when a moving car is observed, the lights in front of and behind the car are turned up to maximum brightness. As the car moves along the street, the light follows the car, increasing the light intensity only where it is needed. When the car approaches an intersection, the lights on all possible exists are turned up, and after the car has left the intersection they are dimmed.

Implementing such urban lighting applications requires the knowledge of the street light topology in order to send commands to the right light pole addresses. In particular, given the light pole identities, one needs to know which are the neighboring light poles along the street with respect to passing traffic, illustrated in Figure 1.

Manual methods for determining the light pole topology during commissioning of a street lighting system, such as sending out a technician who maps a location to a light pole address, are expensive and error prone. Due to human error a significant percentage of the light poles are miss labeled, making it very difficult for a technician to accurately assign a location to the correct light pole address. The ability of a system to automatically find its topology, by fusing data from various sensors, can ease the commissioning and make it less error prone.

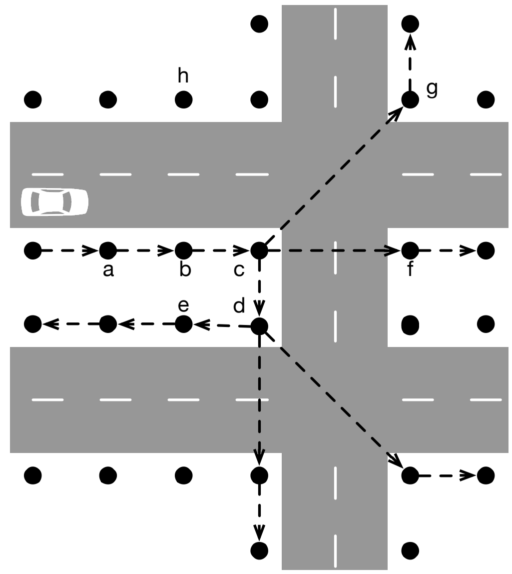

Fig. 1. Example of neighboring light poles along a street.

In this paper we propose a method for determining the light pole topology based on presence data gathered from sensors attached to the light poles. Different sensors can be used to measure presence, e.g. camera, microphone, or passive infrared sensor (PIR). We chose PIRs, since they are cheap, easily deployed and controlled and guarantee anonymity of the observed objects.

Section II discusses related work. Section III specifies the problem in more detail. The proposed method for deriving the light pole topology is described in Section IV and evaluated in Section V. Section VI concludes the paper and presents opportunities future work.

II. Related work

Much of the work on node localization is done in the area of wireless sensor networks where the problem is to find the absolute or relative geographic locations of nodes in a 2D plane. The algorithms presented in [3], [4] rely on Radio Signal Strength Indicator (RSSI) data for estimating distances between the nodes. This information is then used for localizing a set of nodes using triangulation, provided the location of a (small) subset of nodes is given.

The authors in [5] take a different approach to the localization problem, and replace the location computation by the location assignment. Given a set of known locations, they provide an iterative algorithm based on RSSI for assigning the known locations to the nodes. They assume that the mapping of a subset of nodes to their locations is given. Since RSSI is a very crude measure of distance, their algorithm provides only an approximate solution. The accuracy of the solution depends on the choice of nodes for which the location is known initially.

Finding the topology of a set of nodes, i.e. a graph representing which nodes are neighbors, is complementary to finding their geographic location in a 2D space. To the best of our knowledge, this paper is the first to consider the problem of finding the topology of a set of nodes, where the neighbor relationship is defined with respect to passing traffic derived from presence data.

III. Model

Figure 1 illustrates the notion of neighboring light poles. The black dots represent light poles. A dashed arrow from a to b means that light pole b is an immediate successor of light pole a with respect to a car traveling on the street. Two light poles are considered neighbors if there is an arrow between them. The figure shows the order in which a car will pass by the light poles when it travels along different possible paths. Light poles along a straight road have a single successor, e.g. b is the only successor of a. Some light poles, however, will have several successors, e.g. d, f, and g are all successors of c. The figure also illustrates the difference with the notion of neighbors in a 2D space. While light poles b and e are geographically close to each other, they are not neighbors along the street.

We assume that each light pole which we want to localize is equipped with a PIR presence sensor. This presence sensor is facing the street and logs the times when it observes an object passing in front of it. We assume that the time is synchronized between the PIR sensors of different light poles.

PIR sensors can only measure presence and cannot differentiate between different kinds of traffic, e.g. cars and pedestrians. We therefore assume mixed traffic, comprised of different kinds of objects.

The problem of finding the light pole topology based on the presence data can be formalized as follows. Given a set of light poles equipped with presence sensors, where each sensor generates a sequence of times when it has observed an object in front of it, create a directed graph (V, E), where node $a \in V$ represents a light pole a, and edge $(a, b) \in E$ means that light pole b is an immediate successor of light pole a with respect to an object traveling on the street.

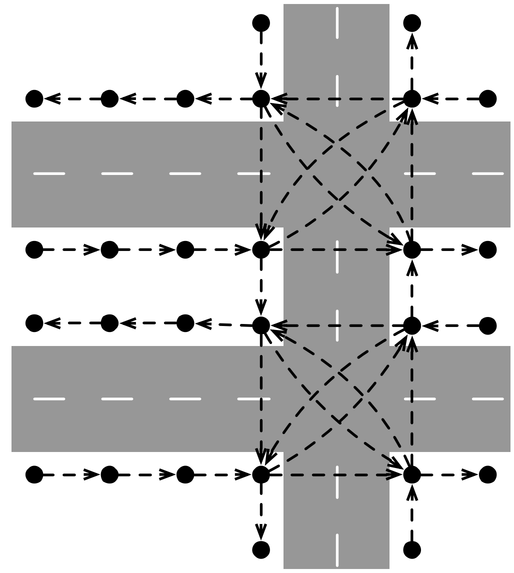

Fig. 2. Example of a light pole topology graph.

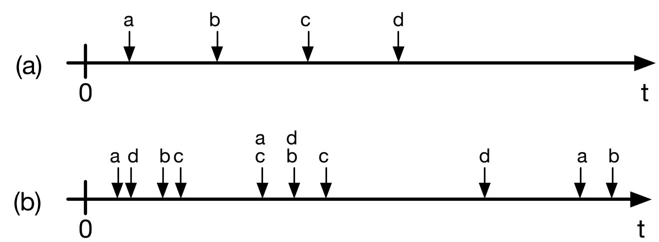

Fig. 3. Example of presence data for (a) single object (b) multiple objects.

Figure 2 shows the graph that we would like to derive for the example in Figure 1. Depending on the range of the presence sensors, light pole h in Figure 1 may or may not be a successor of light pole a. We assume that the presence sensors are calibrated for a single lane.

IV. Method

We consider two different kinds of traffic: a single object moving through the streets and multiple objects.

A. Single object

Each sensor attached to a light pole will generate a sequence of time stamps when it has observed an object in front of it. Since the time is synchronized between the sensors, we can plot all the observations on a single timeline, indicating where and when objects were observed, as shown in Figure 3.a.

Identifying the neighboring light poles in the case of a single object is trivial: two consecutive marks on the timeline indicate that the corresponding light poles are neighbors. From Figure 3.a we can identify the following edges in our topology graph: $\{(a,b),(b,c),(c,d)\} \subseteq E$. To identify all edges in the graph one would need to cordon off entire city neighborhoods and have a single object move through all possible streets.

B. Multiple objects

Having a single object move through the streets is unrealistic. Therefore, we need to assume that multiple objects are moving at any given time. Moreover, the number of objects may vary and is outside of our control.

The example in Figure 3.b illustrates a timeline produced by the presence sensors with multiple objects. Inferring the topology of the light poles from the timeline is no longer trivial, because it is no longer clear which observations belong to which object.

We propose a statistical method for extracting the topology from the timeline. The idea is that if we collect a large enough sample set of the presence data, then the light poles which are close to each other on the street should *on average* be close to each other on the timeline.

Let S be a sequence of presence observations, i.e.

$$S = ((t_1, n_1), (t_2, n_2), (t_3, n_3), \ldots)$$

where (t, n) means that at time t an object was observed at the light pole n. For each pair of light poles (x, y) we first compute the set of time distances between x and y,

$$D_{x,y} = \{t_j - t_i : (t_i, x) \in S \wedge (t_j, y) \in S \wedge t_j - t_i \leq T\}$$

where T is a conservative bound on the time distance between two neighboring light poles, considering the geographical distances between light poles and speed limits. We then compute the mean $\mu_{x,y}$ and standard error $SE_{x,y} = \frac{s_{x,y}}{\sqrt{|D_{x,y}|}}$ for every $D_{x,y}$, where $s_{x,y}$ is the sample standard deviation of $D_{x,y}$. The smaller the $\mu_{x,y}$ the smaller the distance between x and y. The smaller the sample standard deviation $s_{x,y}$ and the larger the set $D_{x,y}$, the more confident we are that $\mu_{x,y}$ is an accurate estimate of distance between x and y. Therefore, the smaller the $\mu_{x,y}$ and $SE_{x,y}$ the more confident we are that light pole y is a successor of x.

We propose a simple metric for deciding whether an edge belongs to the graph:

$$(x, y) \in E \iff \mu_{x,y} \leq \theta_\mu \ \wedge \ SE_{x,y} \leq \theta_{SE}$$

where θ_μ and θ_{SE} are the thresholds for the mean and standard error, respectively. The choice of the threshold values has a big impact on which edges are included in the graph.

V. Results

We evaluate our proposed method on a data set gathered from a deployment of presence sensor enabled light poles in a residential neighborhood in Eindhoven, Netherlands. The data set contains a sequence of 480,000 presence observations for the month of December 2013. Figure 4 shows a map of the residential area with a superimposed topology graph derived from the presence data, for $T = 60\text{s}, \theta_\mu = 15\text{s}, \theta_{SE} = 0.26\text{s}$.

In general, the edges in the topology graph follow the streets in the neighborhood. There are, however, some false positives (edges between light poles that are not direct neighbors on the street) and false negatives (missing edges between actual neighbors). The false positives $(26, 30)$ and $(35, 33)$ are between relatively close light poles. Excluding them from the graph would require lowering θ_μ below $\mu_{26,30}$ and $\mu_{35,33}$. However, for this dataset, this would result in more false negatives.

The false negatives on the 19, 18, 16, 15 street can be explained by the low number of measurements in that street (about 10 times fewer samples than for most other light poles), resulting in a high standard error value. The missing edge $(13, 14)$ can be explained by (bike) traffic on the perpendicular street to the right of 14, which would result in additional presence measurements not related to the neighboring light pole 13.

The parameters $T, \theta_\mu, \theta_{SE}$ were selected by generating the topology graphs for different parameter combinations and selecting those resulting in a "good graph". Without a formal definition of a "good graph" we decided to chose the parameter combination manually, by visually inspecting the generated graphs. An automated method for selecting the parameters is part of future work.

VI. Conclusion and future work

We presented a statistical algorithm for deriving the topology of light poles in a city relying only on presence data gathered from PIR sensors attached to light poles. We applied the algorithm to a data set gathered from a real deployment in a residential area. The initial results are promising and provide interesting possibilities for future work. The presented algorithm is generic and can be applied also in other domains for identifying object flows through a network of presence sensors.

While the results presented in this paper validate the proposed method, they also contain false negatives and false positives. We would like to investigate how these can be eliminated by applying filtering to the raw data (e.g. consider only presence data gathered during low traffic and construct the graph incrementally), other statistical measures than simple average time distances, or machine

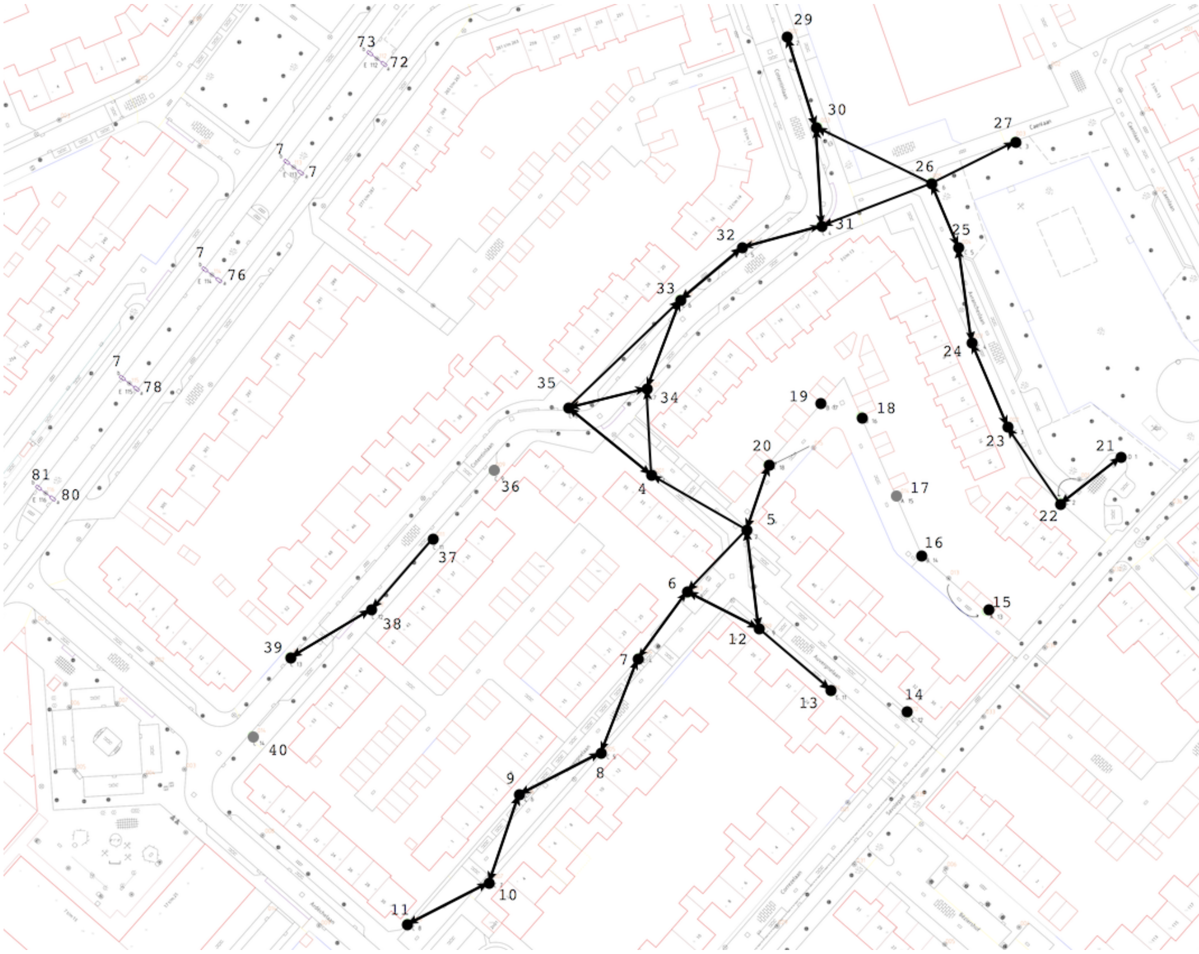

Fig. 4. Light pole topology derived from presence data gathered in a residential neighborhood in Eindhoven superimposed on the street map. Black dots represent light poles equipped with presence sensors. Gray dots represent light poles without presence sensors. Black arrows represent the edges in the derived topology graph.

learning techniques. Also, we would like to gain better insight into how the quality of the collected data can be improved, e.g. by identifying and addressing common problems in hardware deployment or configuration.

The results show that the proposed method works in spite of the mixed traffic assumption. We would like to investigate whether fusing the presence data from PIR sensors with data from other sensors, e.g. cameras, can improve the accuracy of the derived topology, e.g. by differentiating between different kinds of traffic.

The localization accuracy of the algorithm presented in [5] depends on the choice of light poles for which the location is known. The intuition is that the accuracy of their algorithm can be improved by choosing the leaf nodes on the perimeter of the space where the light poles are located. A topology graph derived using our algorithm can identify such leaf nodes. We would like to see whether combining these two algorithms can improve the accuracy of the geographic localization of light poles.

ACKNOWLEDGEMENTS

We would like to thank Pieter Cuipers for the insightful discussions, and EIT ICT Labs for partly funding this work.

REFERENCES

[1] *Lighting the Cities – Accelerating the Deployment of Innovative Lighting in European Cities*. Office des publications de l'Union européenne, 2012.

[2] C. Atc, T. Ozcelebi, and J. Lukkien, "Exploring user-centered intelligent road lighting design: a road map and future research directions," *IEEE Transactions on Consumer Electronics*, vol. 57, no. 2, pp. 788–793, May 2011.

[3] N. Patwari and A. O. Hero, III, "Using proximity and quantized rss for sensor localization in wireless networks," in *ACM International Conference on Wireless Sensor Networks and Applications*, 2003, pp. 20–29.

[4] J.-H. Choi, J.-K. Choi, and S.-J. Yoo, "Iterative path-loss exponent estimation-based positioning scheme in wsns," in *International Conference on Ubiquitous and Future Networks (ICUFN)*, July 2012, pp. 23–26.

[5] N. Jovanovic, J. Lukkien, R. Verhoeven, and Z. Aleksovski, "Iterative node localization for intelligent street lighting," in *International Conference on Ubiquitous and Future Networks (ICUFN)*, July 2013, pp. 679–684.