An Event-triggering Approach for Bus Tracking based on Multimodel Mobility Prediction

M. Michalopoulou, P. Kolios, C. Panayiotou, and G. Ellinas

***Abstract*— Tracking a fleet of buses that serve routes across a city is an essential functionality for the successful utilization of modern transportation networks. However, to date, little attention has been paid to the effectiveness of the implemented tracking algorithms, which typically rely on periodic signaling messages sent by the buses in motion to indicate the vehicle location. However, as we will demonstrate in this work the latter approach entails a large number of unnecessary signaling messages to be communicated in order to achieve a high level of tracking accuracy.**

The alternative approach we present in this work is based on a novel event-triggering strategy that substantially improves tracking and reduces significantly the number of messages that need to be sent out compared to periodic signaling. Time-series traces are used to extract and update mobility models which are then used to estimate future vehicle locations. The estimate, and actual arrival times at predetermined locations are assessed by onboard units and an event is triggered whenever the deviation exceeds the desired tracking accuracy.

I. Introduction

The uncertainty related to the information provided by public transport services (especially public buses) is the primary factor for the hesitation of the public to use the public transportation system. As a result, the general public still tends to prefer private vehicles to public transport when given the choice. As already shown in [1], the current state-of-the-art public transport systems implement rather simplistic tracking strategies based on periodic triggering. Evidently, increasing the frequency of communicating tracking messages will not improve the estimated arrival times but merely provide a more updated instantaneous view of where the vehicles are and do not account for the congestion that vehicles may experience ahead. In addition to not-addressing the inherent uncertainty of traversing the road network, such strategies have higher communication costs and entail high computation overhead that unnecessarily raises complexity [2]. Further, periodic communication makes the system significantly more vulnerable to communication delays.

In this work, we investigate how modeling the vehicle behavior and using the derived models to track the vehicle progress in time can result to better travel time estimates at desired locations (i.e., bus stops) and can minimize the signaling overhead associated with tracking those vehicles. More specifically, the remote host that seeks to track the fleet of buses will use the mobility models to estimate arrival times while a local host onboard the bus will trigger resynchronization (and model-switching) events whenever the actual bus position (and behavior, respectively) deviates from that anticipated by the model.

This event-triggering (ET) approach has the potential to improve the tracking accuracy and significantly reduce the signaling overhead [3]. From a wireless communications perspective, periodic triggering will generate a large amount of messages for a fleet of buses. Thus, even though the packets to be sent are typically small, the number of packets is very large creating challenging conditions in terms of channel contention. With ET, computation/communication actions take place only when a particular *event*, or after a certain *series of events* have taken place. Events represent changes, abnormalities or faults of the process/system that alter its state. As opposed to continued and periodic triggering, ET ensures that normal operation is interrupted only after a particular event has taken place. In doing so, ET ensures that the available resources (including, battery capacity and processing) are thriftily used, while communicated information reduces from raw data to simple event descriptions [4].

The primary contribution of this work is on multimodel data-driven ET in which the mobility models are derived explicitly from the collected traces. Model-switching occurs whenever a triggered event indicates a substantial deviation from the assumed model and for which a switching to a more representative model would result to fewer event interrupts in the future. Within this context, a novel mathematical framework is derived to capture the three phases of ET, namely behavior modeling, event detection, and model switching. The applicability of this framework is then examined under realistic scenarios of bus tracking.

The rest of the paper is organized as follows. Section II includes related work, placing extra attention on data-driven approaches. The proposed multimodel approach is elaborated in Section III and Section IV provides extensive numerical evaluations. Finally Section V provides some concluding remarks.

II. Related Work

While several challenges exist in implementing vehicle tracking systems as indicated in [5]-[7], in this work we assume that such a system is in place and focus on the challenges faced when trying to devise accurate and efficient tracking algorithms. The classical approach followed by most city operators involves surveys using samples, and calibration of the outcomes using sensory data. However, the approach is neither accurate (due to infrequent surveying) nor efficient (due to resource-intensive responses) [8].

The authors are with the KIOS Research and Innovation Center of Excellence (KIOS CoE) and the Department of Electrical and Computer Engineering, University of Cyprus {michalopoulou.maria, pkolios, christosp, gellinas}@ucy.ac.cy.

Hence, numerous automated solutions have been proposed to compensate for these shortfalls. Online fleet-management solutions are simple to implement but impose significant overheads [9] and require considerable infrastructure set-up [10]. The alternative approach is to build mobility models using historical data and use those models to estimate the arrival times by either clustering, extrapolation, regression, Kalman filtering, and machine learning techniques. The majority of the aforementioned works use GNSS traces for this purpose due to the increased availability of this type of data. Of course, other forms of road data including induction loop measurements (as in [11]) have also been used in the past.

However, the current literature decouples the travel-time prediction problem from the vehicle tracking problem in most of the cases. As a result, the efficiency aspect of tracking is overlooked altogether. The issue has first been identified in [1] where the tradeoff between accuracy and efficiency has been investigated. The work presented in this paper takes a more systematic approach to deal with this issue and in the process of doing so develops a new and complementary approach to vehicle tracking based on ET.

A. Mobility Characterization

The literature contains a plethora of studies for characterizing the mobility of public transport vehicles and extensive work has been conducted for inter-arrival models for public transport vehicles. Studies dealing with the latter issue have resulted in the development of four general prediction models. Statistical approaches using historical data have been primarily used to predict the arrival time for a specific time period based on average values at those time periods [12]. However these approaches assume that the traffic conditions remain stationary. Statistical models are divided into two sub-models, those using the average travel time (e.g., [13]) and those using the average travel speed (e.g. [14]). In addition, Kalman filtering has also been employed to model traffic variations as a function of the time-dependent parameters. The main use of Kalman filtering is to offer an estimate of the system's state while providing forecasts of future values or improvements for values that already exist [15].

Further, regression analysis-based models employ mathematical functions to predict the expected arrival time between bus stops, using previously collected data [16], [17]. Unlike statistical prediction models, the latter approach is able to capture dynamic traffic conditions. However, the derived models are only applicable to the particular route at hand. Finally, artificial neural networks (ANNs) have also been employed to express the complex non-linear relationships that exist in traffic scenarios with correlated data and make sense of the collected data [18]. Despite the precision that ANNs have shown to achieve in the general case, they require extensive training and verification for finding the proper structure of the ANN network and the proper set of input parameters.

B. Event Triggering approaches

Due to its simplicity, periodic triggering has primarily been used for vehicle tracking applications in the past [19]. The appealing proposition of this strategy is that there is a definite, predetermined interaction of the local and remote host. Hence, absence of communication results in an indisputable indication of a fault in the system. Nevertheless, periodic triggering can result in an unnecessarily high number of computation and communication actions (that convey no new information or do not indicate a change in state) simply due to the inherent periodicity of the paradigm [20]. At the same time, network scalability issues can arise due to this periodicity.

As exemplified above, ET compensates for these shortfalls by carrying out computation and communication actions only when certain events have taken place and which can potentially change the state of the system at hand [21], [22], [23]. In this way, the resource utilization efficiency is improved (i.e., when no events are triggered, processing can scale down, and communication circuitry can be put to sleep).

Interestingly, most existing ET implementations consider spontaneous events and thus fail to take advantage of the recurrent patterns that may exist in the system. Processing and analyzing streams of data can reveal recurrent patterns which can in turn be used to extract accurate behavior models and eventually more meaningful information. In doing so, all anticipated patterns are incorporated into behavior models and relevant actions take place only when some unanticipated events have occurred.

A plethora of model-based ET strategies have been developed for system *monitoring and control* (a review of such triggering strategies can be found in [24], [25]). The work described in this paper introduces a novel data-driven multimodel ET paradigm and demonstrates its multiple gains in terms of accuracy and efficiency compared to the existing vehicle tracking solutions that assume periodic triggering.

III. System Model

We consider a fleet of buses serving specific routes across a city. Our objective is to track the buses along their routes. For this purpose, each vehicle is equipped with an onboard device that sends messages to a remote central host, providing real-time information about the actual location of the bus along its route. Along the route we consider a set of predetermined measurement locations (i.e., bus stops), at which the time-of-arrival of the bus needs to be reported to the central host. Whereas a conventional solution relies on sending a message to the central host when the bus reaches each one of the measurement locations along the route, we present an ET approach in which both the remote host and the onboard device use a predetermined mobility model to estimate the movement of the bus with respect to time. An event is triggered when a bus reaches a measurement location with a substantial time lag compared to the prediction provided by the mobility model. The onboard device sends a message to the central host only when an event is triggered, that is, only when a measurement location is reached at a different timing than expected.

Since traffic changes dynamically over the course of a day and also between different days, multiple mobility models need to be defined as the trip of a vehicle along the route is expected to exhibit multiple expected behaviors. For example, the traveling times between measurement locations along the route are likely to differ significantly on weekdays and on weekends at particular time instances. Therefore, we need to define a set of models and a switching algorithm to determine which model to use.

A. Problem Formulation

Without loss of generality we focus on a particular bus route. We specify a set of K measurement locations along the route and a required tracking accuracy α in units of time. The parameter α is the time lag above which an event is triggered. The density of the measurement locations as well as the value of α are set accordingly to meet the desired tracking accuracy of the system.

We define a sequence of K random variables, $\{S_1, S_2, \cdots, S_k, \cdots, S_K\}$, where the random variable S_k corresponds to the time needed to travel from measurement location $k-1$ to k and S_1 to the time needed to travel from the beginning of the route to the first measurement location. A trace, $\{s_1, s_2, \cdots, s_k, \cdots, s_K\}$, is a realization of the sequence of random variables $\{S_k, k = 1, \cdots, K\}$. Because the trip of a vehicle along the route is expected to exhibit multiple expected behaviors, we define a set of M models. At each time instant, the system uses one of the models. A model, m_i for the particular problem is defined as a K-dimensional vector, $\{v_{i,1}, v_{i,2}, \cdots, v_{i,k}, \cdots, v_{i,K}\}$, that includes the times needed to travel between consecutive measurement locations.

When reaching measurement location k, the onboard device generates an event if there is substantial time lag between the prediction and the actual vehicle movement, that is, if:

$$\left| \sum_{j=n}^{k} s_j - \sum_{j=n}^{k} v_{i,j} \right| > \alpha \tag{1}$$

where n is the last measurement location where a triggering event occurred and m_i is the model currently in use since measurement location n. When an event trigger occurs, the value of the accumulated time lag has exceeded α, therefore the onboard device sends a message to the remote host for changing the currently used model.

The objective is to minimize the number of control messages that need to be sent by the onboard device to the remote host during a trip, i.e., the event triggers that occur. This is equivalent to maximizing the probability that the next expected measurement will not create an event (and thus implicitly build models of increased tracking accuracy).

$$\max P\left(|s_k - v_{i,k}| \leq (\alpha - \epsilon) \mid s_n, s_{n+1}, \cdots, s_{k-1}\right) \tag{2}$$

where

$$\epsilon = \sum_{j=n}^{k-1} s_l - \sum_{j=n}^{k-1} v_{i,j} \tag{3}$$

B. Behavior Modeling

Given a set of collected measured traces, we apply a clustering algorithm in order to arrange the training traces into M clusters. Each cluster corresponds to a model and each model is defined by averaging the values of the traces that have been classified in the corresponding cluster.

We need to determine the number of models, M, so as to minimize the number of event triggers, i.e., maximize the optimization objective of Eq. 2. The optimal number of models increases with the dependencies between the random variables $\{S_k, k = 1, \cdots, K\}$, i.e., between the values measured at consecutive measurement locations.

The benefit of defining multiple models stems from the fact that if consecutive measurements are not independent, a measured value will provide information about the next expected measurement, thus will provide information about which one out of several available models will more likely provide a sufficiently accurate prediction. On the other hand, if the values are independent, there is no prior information about which model shall be better than others for the measurement that will follow. Therefore, the selection of the model cannot be any better than random and the use of several models cannot increase the probability of making an accurate prediction. In fact, under some highly realistic assumptions, it can be proven that the optimal number of models is 1 if consecutive measurements are independent (please see the proof in Appendix I).

Considering that there is a certain level of dependency between consecutive measurements, measurement s_{k-1} already provides some information about the following measurement. In mathematical terms, the optimal value that should be used for the prediction of s_k is not simply the value that coincides with the maximum of the probability density function $f_k(x)$, but with the maximum of the conditional probability function $f_{S_k}(x|S_{k-1} = s_{k-1})$. When measurements are independent this conditional probability function is equal to the probability density function of S_k for all possible realizations of S_{k-1}. Therefore, a single value should be used for the prediction at measurement location S_k, independently of previous measurements. As the dependencies between measurements increase, the conditional probability function starts becoming different for different realizations of S_{k-1} and shows a peak for different values, thus providing an indication of the need to use alternative prediction values based on the value of S_{k-1}. Therefore, the use of several models increases the probability that a value is predicted accurately enough. Nevertheless, defining too many models will again worsen the performance, as we practically try to be more precise than what we are allowed by the information provided by the dependencies. Therefore, depending on the degree of dependencies between the measurements there is an optimal number of models, M, which provides the maximum probability that a value is predicted accurately enough. A smaller number of models will worsen the performance due to the fact that we have information that remains unused, and a larger number of models will also worsen the performance because we are trying to become more precise than what can be extracted from the information provided by the dependencies between measurements. To the other extreme, if consecutive measurements exhibit a correlation coefficient equal to 1, then the value of the next expected measurement s_k is known given the measurement s_{k-1}.

Since the objective is to minimize the probability that an event is triggered, the optimal number of models, M, can be estimated by determining the maximum value of the conditional probability that the value s_k is predicted with certain accuracy by model m_j, given that this particular model was the most accurate for the previous measurement, s_{k-1}. That is, let us use the notation C_{k-1} to denote the model closest to s_{k-1}, defined as follows:

$$\left|s_{k-1} - v_{C_{k-1},k-1}\right| \leq |s_{k-1} - v_{i,k-1}| \quad \forall i \in \{1, 2, \cdots, M\} \setminus \{M\} \tag{4}$$

Then, we are looking for the value of M that maximizes the following conditional probability:

$$\max_{M} P\left(|s_k - v_{j,k} \leq \alpha| \, \middle| C_{k-1} = j\right). \tag{5}$$

The value of α is not critical, since it does not affect the shape of the conditional probability with respect to M.

C. Model Selection

Let $\rho_{i,j}$ denote the correlation coefficient between the random variables S_i and S_j. Let us consider that an event trigger occurred upon receiving measurement s_k, thus a new model needs to be chosen. Under the assumption that the correlation coefficient between consecutive measurement locations is the highest, i.e., $\rho_{k,k+1} > \rho_{j,k+1}, \forall j = 1, 2, \cdots, k-1$, it is optimal to choose the new model based solely on the measurement s_k and not on previous measurements. Under this assumption, considering the previous measurements will only negatively affect the prediction for the following measurement. Therefore, the model m_i closest to the last received measurement s_k should be chosen, that is:

$$\min_{i}\{| \; s_k - v_{i,k} \; | : i = 1, 2, \cdots, M\} \tag{6}$$

IV. Numerical Investigation

In order to illustrate the proposed ET approach we perform in this section, an extensive numerical investigation using artificial, randomly generated data. In particular, we generated random traces to represent the collected measured traces (training set), and the proposed classification algorithm was used in order to generate the desired mobility-tracking models. Then, a second set of random traces were generated (test data) to evaluate the proposed ET signaling performance.

A trace consists of K time intervals between consecutive measurement locations. Vectors of K elements were drawn from a multivariate normal distribution with an arbitrarily selected vector of mean values. The following covariance matrix is used in order to represent realistic dependencies between the measurements in a trace,

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 & \rho^{2f}\sigma_1\sigma_3 & \dots & \rho^{(K-1)f}\sigma_1\sigma_K \\ \rho\sigma_1\sigma_2 & \sigma_2^2 & \rho\sigma_2\sigma_3 & \dots & \rho^{(K-2)f}\sigma_2\sigma_K \\ \rho^{2f}\sigma_1\sigma_3 & \rho\sigma_2\sigma_3 & \sigma_3^2 & \dots & \rho^{(K-3)f}\sigma_3\sigma_K \\ \vdots & \vdots & \vdots & \dots & \vdots \\ \rho^{(K-1)f}\sigma_1\sigma_K & \dots & \dots & \dots & \sigma_K^2 \end{pmatrix} \tag{7}$$

Consecutive measurements are correlated; two consecutive measurements exhibit the highest correlation represented by the correlation coefficient ρ. Then, the dependency, i.e., the correlation coefficient drops from measurement to measurement at a rate controlled by parameter f.

In addition, in a realistic situation it is expected to have single outlier measurements or a sequence of outliers due to unexpected incidents such as accidents on the road, road works, local traffic congestion due to an organized event, etc. Therefore, to a percentage of the traces we create transitions from one randomly drawn vector to another following the Markov model of Figure 1. State I in the state-transition diagram represents the initial randomly generated trace before adding any outlier bursts. With probability p_{IO} the next measurement will jump to another randomly generated vector and afterwards will return to the initial vector with probability p_{OI}.

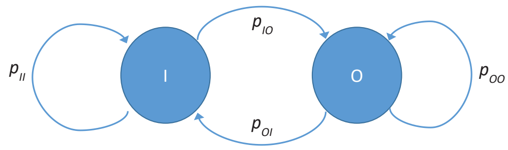

Fig. 1. The Markov model used to simulate sequence of outliers in a trace.

For the clustering of the traces, and in order to extract the mobility models, we used an agglomerative clustering algorithm. For each set of parameters the illustrated results are an average obtained by running 20 random realizations. The Markov model to create outlier bursts was applied to half of the traces with $p_{IO} = 0.1$ and $p_{OI} = 0.8$.

Our objective is to minimize the number of messages sent from the onboard device to the central host. The benefit of reducing the number of messages sent from the onboard device to the central host is not only related to the reduction of the communication overhead. It is furthermore linked to the capability of estimating future arrival times along the route as fewer messages indicate the the proposed tracking scheme is generally efficient at predicting times-of-arrival. Figure 2 shows the number of messages sent for routes of different sizes. Using a conventional approach (periodic triggering), the number of messages would be equal to the number of measurement locations along the route. An ET approach with a single mobility model would cut the number of messages approximately in half and the utilization of multiple models achieves a further reduction.

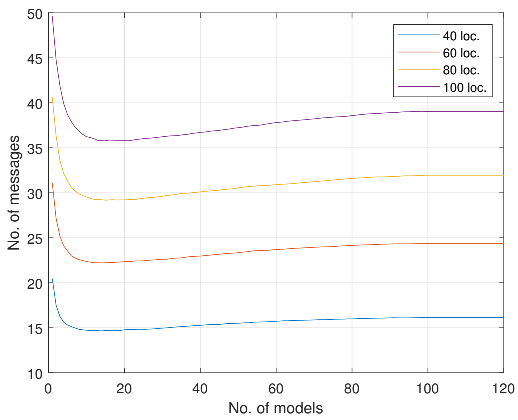

Fig. 2. The number of messages sent along routes of different sizes ($\rho = 0.9,\ f = 2$).

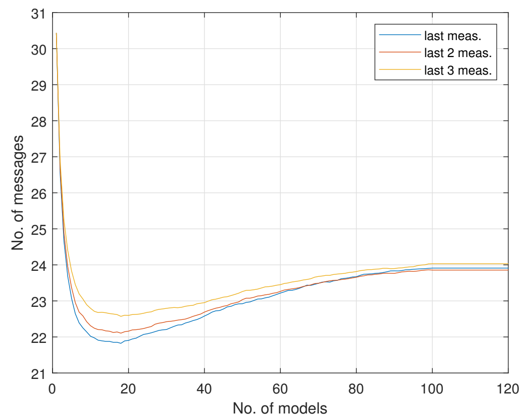

Fig. 4. The number of messages sent for a route with 60 measurement locations and model selection based on different number of past measurements ($\rho = 0.9,\ f = 2$)

Figure 3 illustrates the effect of the dependency between consecutive measurement locations and verifies the discussion in Section III-B. The higher the correlation coefficient of Eq. 7, the higher the potential benefits from employing the proposed multimodel ET approach.

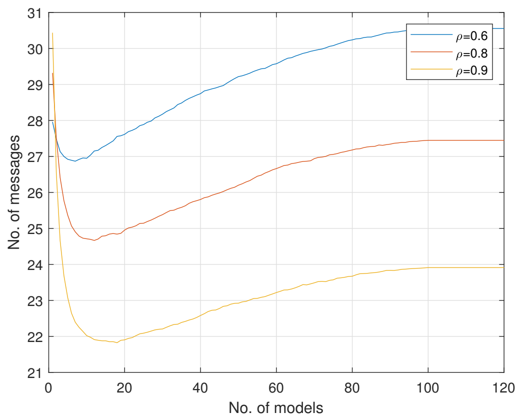

Fig. 3. The number of messages sent for a route with 60 measurement locations for different values of the correlation coefficient ρ of Eq. 7 ($f = 2$)

Figure 4 shows the number of messages when the last one, two, or three measurements are taken into account for selecting the next mobility model when an event is triggered. As discussed in Section III-C, given that the correlation is stronger between consecutive measurement locations and drops as we move further, the optimal choice is to use only the last measurement in order to determine the next model to be used.

Finally, the simulation results show that, as explained in Section III-B, there is an optimal number of models for which the number of messages obtains a minimum value. Figure 5 illustrates the conditional probability of a measurement being predicted accurately enough by a certain model, given that the particular model was the closest to the previous measurement. The conditional probabilities in Figure 5 correspond to the scenarios of Figure 3. We can observe that the maximum value of the conditional probability indeed coincides with the number of models that achieve the minimum number of messages.

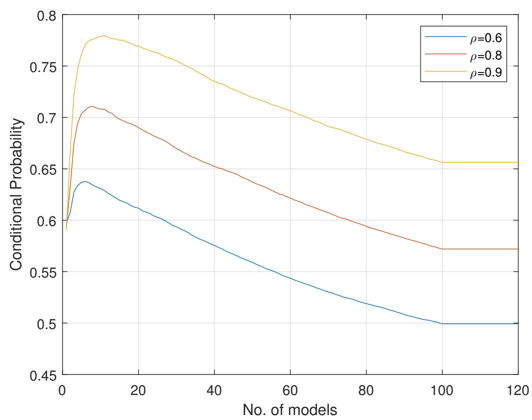

Fig. 5. The conditional probability for the scenarios depicted in Figure 3.

V. Conclusions

In this paper we have presented a data-driven ET approach for public transport tracking systems. The proposed technique uses vehicle traces to create a set of mobility models that are used to predict the time-of-arrival at certain locations along the route.

The aim of this work is to introduce a method that is more efficient in terms of signaling traffic while improving the tracking accuracy compared to the conventional periodic triggering approaches. By means of numerical analysis, we verified that our technique decreases significantly the number of messages that need to be sent by the onboard device to the central host compared to periodic signaling.

Finally, we would like to note that the data-driven ET technique illustrated in this paper is a highly promising architecture for IoT applications in general. Therefore, the analysis and methodology presented here can be useful in the context of a plethora of other applications as well which experience recurrent patterns in their behavior.

ACKNOWLEDGMENTS

This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 739551(KIOS CoE) and from the Government of the Republic of Cyprus through the Directorate General for European Programmes, Coordination and Development.

REFERENCES

[1] J. Biagioni, et al., "Thrifty Tracking: Online GPS Tracking with Low Data Uplink Usage", *ACM SIGSPATIAL*, 2013.

[2] Q. Wu, G. Ding, Y. Xu, S. Feng, Z. Du, J. Wang, and K. Long, "Cognitive Internet of Things: A New Paradigm Beyond Connection", *IEEE Internet of Things Journal*, 1(2):129-143, Apr. 2014.

[3] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos, "Context Aware Computing for The Internet of Things: A Survey", *IEEE Communications Surveys & Tutorials*, 16(1):414-454, 2014.

[4] M. Lemmon, "Event-Triggered Feedback in Control, Estimation, and Optimization", *Networked Control Systems - Lecture Notes in Control and Information Sciences*, 406:293-358, 2010.

[5] R.K. Megalingam, N. Raj, A. L. Soman, L. Prakash, N. Satheesh, and D. Vijay, "Smart, Public Buses Information System", *International Conference on Communication and Signal Processing*, Apr. 2014.

[6] S. Sankarananrayanan, and P. Hamilton, "Mobile Enabled Bus Tracking and Ticketing System", *International Conference on Information and Communication Technology*, Aug. 2014.

[7] F. Abdessemed, "An Integrated System for Tracking and Control Pilgrims Shuttle Buses", *IEEE International Conference on Intelligent Transportation Systems*, Oct. 2011.

[8] M. Berlingerio, F. Calabrese, G.Di Lorenzo, R. Nair, F. Pinelli, and M.L. Sbodio, "AllAboard: A System for Exploring Mobility and Optimizing Transport in Developing Countries using Cellphone Data", *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, Sept. 2013.

[9] IBM Ireland Report, "Dublin City Council", *Travel & Transport*, July 2013.

[10] E. Bouillet, M. Feblowitz, Z. Liu, A. Ranganathan, A. Riabov, F. Yen, S. Shao, and D. Schlosnagle, "Data Stream Processing Infrastructure for Intelligent Transport Systems", *IEEE Vehicular Technology Conference*, Sept. 2007.

[11] J. Kwon , B. Coifman, P. Bickel, "Day-to-day Travel Time Trends and Travel Time Prediction from Loop Detector Data", *Transportation Research Record: Journal of the Transportation Research Board*, 1717(1554):120–129, 2000.

[12] D. Sun, H. Luo, L. Fu, W. Liu, X. Liao, M. Zhao, "Predicting Bus Arrival Time on the Basis of Global Positioning System Data", *Transportation Research Record*, 2007.

[13] R.H. Jeong, "The Prediction of Bus Arrival Time Using Automatic Vehicle Location Systems Data", *Ph.D. Dissertation*, Texas A&M University, 2004.

[14] Li Weigang, et al., "Algorithms for Estimating Bus Arrival Times Using GPS Data", *IEEE International Conference on Intelligent Transportation Systems*, pp.868 - 873, 2002.

[15] J.S. Yang, "Travel Time Prediction Using the GPS Test Vehicle and Kalman Filtering Techniques", *American Control Conference*, 2005.

[16] S.I.J. Chien, Y. Ding, and C. Wei, "Dynamic Bus Arrival Time Prediction with Artificial Neural Networks", *J. Transportation Engineering*, 128(5):429-438, 2002.

[17] J. Patnaik, S. Chein, and A.Bladihas, "Estimation of Bus Arrival Times Using APC Data", *J. Public Transportation*, 7(1):1-20, 2004.

[18] T. Park, S. Lee, and Y.-J. Moon, "Real Time Estimation of Bus Arrival Time Under Mobile Environment", *International Conference in Computational Science and its Applications*, 3043:1088-1096, 2004.

[19] O. Vermesan and P. Freiss, "Internet of Things - From Research and Innovation to Market Deployment", *River Publishers Series in Communications*, 2014.

[20] M. Mahmoud and M. Sabihb, "Networked Event-triggered Control: An Introduction and Research Trends", *International Journal of General Systems*, 43(8):810-827, Apr. 2014.

[21] W. Heemels, K.H. Johansson, and P. Tabuada, "An Introduction to Event-triggered and Self-triggered Control", *IEEE Conference on Decision and Control*, Dec. 2012.

[22] O. Vermesan and P. Freiss, "Internet of Things - Converging Technologies for Smart Environments and Integrated Ecosystems", *River Publishers Series in Communications*, 2014.

[23] P. Huang, C. Wang, and L. Xiao, "RC-MAC: A Receiver-Centric MAC Protocol for Event-Driven Wireless Sensor Networks", *IEEE Transactions on Computers*, 64(4):1149–1161, Feb. 2014.

[24] R. Isermann, *Fault-Diagnosis Systems: An Introduction from Fault Detection to Fault Tolerance*, Springer, 2006.

[25] C. Cassandras, and S. Lafortune, *Introduction to Discrete Event Systems*, Springer, 2nd Ed., 2008.

APPENDIX

Lemma 1: If the measurements in a trace are independent from each other then the optimal number of models to use is 1.

Proof: Consider an arbitrarily selected trace. The next expected measurement is s_k and the last event was triggered at measurement location n. Without loss of generality, let us assume that the model m_i was a sufficiently good choice for the measured values from location n to $k-1$ (i.e., with $\epsilon < \alpha$) and the one that was actually used from measurement location n onwards. In order to minimize the number of event triggers, we want to define the value $v_{i,k}$ so as to maximize the probability that an event will not be triggered upon reception of value s_k. Thus, we consider the optimization objective of Eq. 2 with respect to the value $v_{i,k}$,

$$\max_{v_{i,k}} P\left(|s_k - v_{i,k}| \leq (\alpha - \epsilon) \mid s_n, s_{n+1}, \cdots, s_{k-1}\right) \quad (8)$$

If measurements at consecutive measurement locations are independent, then the random variables $\{S_k, k = 1, \cdots, K\}$ are independent. In particular, the following holds:

$$P\left(S_k = s_k \mid S_n, \cdots, S_{k-1}\right) = P\left(S_k = s_k\right) \quad (9)$$

Then, Eq. 8 becomes equivalent to:

$$\max P\left(|s_k - v_{ik}| \leq (\alpha - \epsilon)\right) \quad (10)$$

Let $f_k(x)$ be the probability density function of the random variable S_k. Then, the following is to be maximized

$$\max_{v_{ik}} \int_{v_k-(\alpha-\epsilon)}^{v_k+(\alpha-\epsilon)} f_k(x)\mathrm{d}x \quad (11)$$

The optimal value of $v_{i,k}$ depends on the value of $\alpha-\epsilon$ which ranges from 0 to α. However, in reality the value of α needs to be very small compared to the range of possible values that random variable S_k takes, otherwise the usefulness of using a tracking application is questioned. However, under this assumption the range of values that will maximize Eq. 11 is $v^*_{i,k}+\mathrm{d}v$, where $\mathrm{d}v$ is very small and can be neglected. The optimal value $v^*_{i,k}$ is the maximum of the probability density function $f_k(x)$, which concludes that since a single optimal value maximizes the above expression, a single model is to be defined. ∎