



Transit-hub: a smart public transportation decision support system with multi-timescale analytical services

Fangzhou Sun¹ · Abhishek Dubey¹ · Jules White¹ · Aniruddha Gokhale¹

Received: 3 October 2016 / Accepted: 4 January 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

Public transit is a critical component of a smart and connected community. As such, citizens expect and require accurate information about real-time arrival/departures of transportation assets. As transit agencies enable large-scale integration of real-time sensors and support back-end data-driven decision support systems, the dynamic data-driven applications systems (DDDAS) paradigm becomes a promising approach to make the system smarter by providing online model learning and multi-time scale analytics as part of the decision support system that is used in the DDDAS feedback loop. In this paper, we describe a system in use in Nashville and illustrate the analytic methods developed by our team. These methods use both historical as well as real-time streaming data for online bus arrival prediction. The historical data is used to build classifiers that enable us to create expected performance models as well as identify anomalies. These classifiers can be used to provide schedule adjustment feedback to the metro transit authority. We also show how these analytics services can be packaged into modular, distributed and resilient micro-services that can be deployed on both cloud back ends as well as edge computing resources.

1 Introduction

1.1 Emerging trends and challenges

Public transit ridership in the United States increased by 37% from 1995–2015, which is roughly twice as much as the country's population growth (21%) in the same years [1]. In 2013 alone, there were 10.7 billion trips taken on US public transportation [2]. Meanwhile, people in the US have been reducing the use of personal vehicles [3]. Public transportation has become an essential part of communities and cities.

Bus services, which is one of the most important segments of public transportation, are vulnerable to delays and congestion due to traffic congestion, weather conditions, spe-

cial events, etc. Travel and arrival time variation was found to have a substantial impact on commuter satisfaction [4]. Moreover, people's tolerance to errors in bus time predictions is quite low [5]. Providing real-time bus schedules reduces this uncertainty and improves passenger experience and increases ridership. A direct benefit of increased ridership on public transport is the reduced use of personal vehicles and hence reduction in both traffic congestion and greenhouse emissions.

Recently, transit agencies have been integrating real-time sensors into public transit systems. A number of technological systems have been developed by academic researchers and commercial companies to utilize this real-time data. For instance, automatic vehicle location (AVLs) and automatic passenger counter (APCs) can provide real-time data such as vehicle travel time, arrival and departure time, and passenger boarding counts. This data can be used for at-stop displays [6], bus time prediction [7–9], schedule planning optimization [10,11], real-time control strategies [12,13], etc.

However, these sensors have some problems. Accurate real-time bus arrival and departure data that many prediction systems use is not always available. In Nashville, for example, only special bus stops called timepoints are equipped with sensor devices that record exact times. There are over 2700 bus stops all over the city but only 573 timepoints. In

✉ Fangzhou Sun
fangzhou.sun@vanderbilt.edu

Abhishek Dubey
abhishek.dubey@vanderbilt.edu

Jules White
jules.white@vanderbilt.edu

Aniruddha Gokhale
a.gokhale@vanderbilt.edu

¹ Institute for Software Integrated Systems, Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN 37212, USA

addition, the timepoint dataset is not real-time. It is available at the end of each month when the Nashville metropolitan transit authority (MTA) summarizes and analyzes the historical data. Automatic passenger counter can help provide accurate timing of when a bus stops at a transit stop, which can be used in analysis. On the contrary, AVLs do not provide that. Many transit systems, including the city of Nashville, do not have APCs on buses and use automatic vehicle location (AVL) data to estimate the arrival and departure time at bus stops and use the estimated data for bus delay prediction in real-time. The issue with this approach is that the lack of quality data results in worse predictive analytic performance. Even for systems with APCs, real-time sensor systems can have many problems in the real world [14,15], due to reasons, such as low networking bandwidth and delays in uploads. As a result, often GPS position data is noisy.

A typical mechanism for handling noise is to normalize the data. However, normalization requires large data sets, often clustered around transit routes. This is helpful because the transit data of preceding buses may be used to create the models for the current trip on that route. However, if a city does not have high-frequency operations across its routes, then such data is not available.

1.2 Solution approach and contributions

To address the lack of quality data for transit data analytics, yet make effective predictions for bus arrivals, we surmise that the dynamic data driven applications systems (DDDAS) paradigm [16] holds promise as a solution approach. In DDDAS, both real-time and/or historical data is used to learn the model of the system that must be controlled, and subsequently a decision support system uses these learned models to make informed decisions and control the system in a feedback loop. This is the approach we utilize in this paper. It integrates historical and streaming real-time bus location data from multiple routes for short-term delay prediction as well as long-term delay pattern analytics. We also use the data feedback loop to provide results to city planners and end users.

This paper significantly extends our prior work on Transit-Hub [17,18] and provides the following contributions to the study of real-time and predictive analytics for public transportation using DDDAS principles:

- We present a better short-term delay prediction model that combines clustering analysis and Kalman filters and uses real-time data from shared route segments.
- We show the efficacy of our short-term delay prediction model. When predicting the travel time delay of segments 15 min ahead of scheduled time, our model reduced the root-mean-square deviation (RMSD) by about 30 to 65% compared with an SVM-Kalman model [9]. The SVM-

Kalman model that we used for comparison is a dynamic prediction model that combines SVM and Kalman filters, two of the most widely used models in bus delay prediction [7,19–21].

- We provide an algorithm that generates shared bus route segment networks from standard general transit feed specification (GTFS) datasets.
- We illustrate how the analytical algorithms can be packaged into independently deployable and self-contained micro-services.
- We describe how the system's data feedback loop works to provide decision support to city planners by assisting metro transportation authority (MTA) in identifying real-time outliers and optimizing bus timetables to improve bus services and availability.

1.3 Paper organization

The remainder of this paper is organized as follows: Sect. 2 compares the enhanced Transit-Hub system with related work, specifically how we differentiate from and improve on an SVM-Kalman model that also used bus data from multiple routes; Sect. 3 outlines the key challenges faced in realizing a DDDAS-enabled system for accurate prediction of bus schedules; Sect. 4 describes the integrated data sources and potential feedback mechanism to MTA; Sect. 5 describes how we construct the bus delay models and integrate real-time bus data to predict arrival delay in real-time; Sect. 6 presents the system deployment; Sect. 7 describes the performance evaluation of travel time delay in route segments and arrival time delay at bus stops; and finally Sect. 8 presents concluding remarks and future work.

2 Related work

This section compares Transit-Hub with related work on transit data analysis using different models. In the end, we explain the differences between Transit-Hub models and an SVM-Kalman model that also used shared route segment data.

2.1 Statistical models

The basic average models directly use the average delay from historical data as the estimated delay for future and are often constructed for performance comparison purposes. For example, Jeong et al. [22] developed a basic average model and found that the basic average model was outperformed by regression models and artificial neural network (ANN) models for bus arrival time prediction. The reason is that the basic average models only use historical data and perform simple average analysis, the model does not reflect real-time

conditions and is limited by the consistency of route delay patterns.

Many researchers have conducted studies that utilize both historical and real-time bus data. Weigang et al. [23] presented a model to estimate bus arrival time at bus stops using the real-time GTFS data. Their model contains two sub algorithms to determine the bus speed using the historical average speed and the real-time speed information from GPS. Their main algorithm utilizes the calculated real-time speed to predict the arrival time. Sun et al. [24] proposed a prediction algorithm that combines real-time GPS data and average travel speeds of route segments.

Regression models are also used to explain the impact of variables for delay prediction. Since the variables in transit systems are correlated [25], regression models are typically limited to delay prediction. Patnaik et al. [26] presented a set of regression models that predict bus travel times on a route segment. The data they used is real-world data (number of passengers boarding, stops, dwell time and weather) collected by automatic passenger counters (APC) installed on buses. They also found that weather did not have a significant effect on the prediction.

2.2 Kalman filter models

Kalman filters have been used widely for bus delay prediction because of their ability to filter noise and continuously estimate and update actual states from observed real-time data. Chien et al. [21] presented a dynamic travel time prediction model that used real-time and historical data collected on the New York State Thruway (NYST). Shalaby et al. [27] proposed a bus delay prediction model based on two Kalman filter algorithms: one for estimating the running time and another for estimating the dwell time at bus stops. Yang et al. [28] developed a discrete-time Kalman filter model to predict travel time using collected real-time global positioning system (GPS) data. Bai et al. [9] proposed a dynamic travel time prediction model that employed support vector machines to provide a base time estimate and a Kalman filter to adjust the prediction using the most recent bus trips on multiple routes.

2.3 Machine learning models

Artificial neural network (ANN) [20,22,29] and support vector machine (SVM) [7,9,19,20] are two of the most popularly used machine learning techniques in bus delay prediction. For example, Jeong et al. [30] developed an ANN model for bus arrival time prediction using Automatic vehicle location (AVL) data. Mazloumi et al. [31] used real-time traffic flow data to develop ANN models to predict bus travel times. Yu et al. [20] proposed a machine learning model that used bus running times of multiple routes for predicting arrival times

of each bus route and proposed bus arrival time prediction models that include support vector machine (SVM), Artificial Neural Network (ANN), k-nearest neighbors algorithm (k-NN) and linear regression (LR).

2.4 Comparison with our work

Prior work emphasized long-term and short-term transit data analysis and prediction. However, most of them, as mentioned above, focused on a single route and few noticed that many bus routes share segments with other routes. In 2011, Yu et al. [20] recognized that the data from multiple routes can help to improve the delay prediction. In 2015, Bai et al. [9] proposed a dynamic travel time prediction model that combines SVM and Kalman filter using multiple bus routes data. However, when solving the shared-segment prediction problem, they only used the actual travel time of preceding buses and did not consider the scheduled time difference of separate bus routes. Also, their model included the data of all recent preceding buses, which may contain outliers that should be excluded. Transit-Hub extends these concepts, presents a solution to generate shared route segment network (explained later in Sect. 5.2.1) using standard static GTFS dataset, and provides transit data predictive analysis at multiple timescales. The benefit is that the analysis results can be used to provide schedule adjustment feedback to MTA, and real-time delay prediction to commuters.

3 Building multi-timescale analytical services for public transit

In this section, we present the key problems associated with building multi-timescale analytics models for public transportation systems.

3.1 Problem 1: integrating and managing heterogeneous data from multiple sources

Transportation agencies are employing advanced technologies, such as (AVL) and (APC) to monitor and manage bus services to improve service quality. However, the data collected from multiple data sources may require significant effort to be integrated in order to learn a model for the following reasons: (i) Data are collected at different sampling rates: systems such as AVL and APC have different hardware specifications. Data from different sources need to be sampled before being used by the system; (ii) Data may be missing, duplicated or faulty: these issues need to be detected and handled differently before conducting the data analytics.

Furthermore, the scale of the transit system brings its own challenges and requires efficient and reliable data storage management for the following reasons: (i) Data is large-

scale: the real-time transit data, for instance, is accumulated at the scale of several gigabytes per day currently. If Nashville MTA expands its services and updates the devices for faster data rates, more data will be generated and may require more sophisticated management; (ii) Data replication is also required since the system is accessible by the public and needs to be fault-tolerant and reliable.

We address these challenges in Sect. 4 by describing the heterogeneous data sources and how we integrate, store, and prepare the data for use.

3.2 Problem 2: utilizing real-time bus data for multiple routes that sharing similar segments

Delay prediction models rely on training data. The prediction accuracy depends greatly on the quality of training data. However, the data quality varies for the following reasons: (i) The bus timing is vulnerable to various conditions such as accidents, congestion, road constructions, weather conditions, etc. Therefore, the bus travel time of the preceding bus on the same route may just be an outlier and not reflect the future delay trends; (ii) In mid-sized cities there are limited public resources to support public transportation compared to large metropolis. For example, route 3 (one of the busiest bus routes in Nashville) has 37 trips on weekdays in "From Downtown" direction [32], while M15-SBS in New York has 144 trips in one direction [33]. The quantity of data available for historical analysis and future time prediction in mid-sized cities is less than those in their larger counterparts, which makes it harder to learn accurate models of the system; and (iii) Software bugs, hardware malfunctions and wireless communication issues may occur occasionally and result in missing or faulty real-time data. For example, during our experiments, often AVL data was not uploaded in proper sequence and often had repetitions. Curating such data becomes a challenge in itself.

We address these concerns in Sect. 5.2.1 by discussing how our short-term prediction model improves the data quality by dividing all bus routes in the city into shared route segments and utilizing real-time bus data from segments shared by multiple routes.

3.3 Problem 3: providing schedule adjustment feedback to metro transit authority

Improving existing bus schedules is a critical task for metro transportation authorities such as Nashville MTA. MTA regularly examines the historical bus operation reports and updates its bus schedules. Recently real-time sensors are being installed on buses and MTA can track the bus operation in real-time. However, it is still difficult for them to be aware of the actual bus status. For example, by combining real-time bus location feed and static bus schedule, it is not

difficult to tell if a bus has deviated from its schedule or not. However, the capability to differentiate a delay event from a normal delay that fits historical delay patterns and thereby identify outliers that need to be further investigated is still lacking at present.

We present our solution to this challenge in Sect. 4.3 by designing a data feedback loop for metro transportation authority that tracks the real-time status of the bus operating using analytics result from both historical as well real-time prediction models.

3.4 Problem 4: building and deploying the system with high availability and scalability

Traditional applications are often built in a monolithic style where all logic for handling requests runs in a single service process. Even though the monolithic architecture is easy to develop, deploy and it is also easy to scale if a load balancer is used, when the scale of the application increases, it will become too large and complex for developers to understand, improve and conduct continuous deployment [34]. Also, the reliability of the monolithic application will be a problem because a break down in one component has the potential to impact the entire application [35].

To improve the scalability and availability of the system, we adopted a microservice architecture, which is a modular architectural pattern for building and deployment [36,37]. The microservice architecture is well-suited for cloud environments and has many advantages over traditional architectures: (1) Smaller modules are easier to develop and therefore improve the productivity of developers, (2) Services can be developed and deployed independently, (3) The source of faults is more apparent.

However, the microservice pattern is not perfect. It has some unique drawbacks including (1) it is not easy to partition an existing large-scale system into microservices, (2) additional inter-microservice communication mechanism is needed, (3) memory consumption may increase especially if the microservices do not share the same environment. In our current implementation, the microservices are deployed in a single environment to avoid this problem.

Sect. 6 explains how we addressed this challenge by using a microservices architecture to develop and deploy the back-end services.

4 Data management and feedback

In this section, we first present the heterogeneous data sources that the system is using and then describe how we integrate and manage the collected data to address the issues raised in Sect. 3.1.

Table 1 Realtime and static datasets collected and stored in the system

Bus Schedules	
Format	Static GTFS
Source	Nashville MTA
Update	Every public release
Total size	193MB (Version: Mar. 9 2016)
Real-time Transit	
Format	Real-time GTFS
Source	Nashville MTA
Update	Every minute
Total size	278 GB
Time Points	
Format	Excel
Source	Nashville MTA
Update	Every month
Monthly size	300,000 entries/month

4.1 Data sources

We have been collaborating with the Nashville Metropolitan Transit Authority (MTA) for accessing the static and real-time transit data all across the Nashville city. The data sources that we are collecting are as follows (Table 1).

- *Static GTFS data sets* Static bus schedules and associated geographic information in the general transit feed specification (GTFS) [38] are collected. The data sets include routes, trips, stops, stop times and physical layout.
- *Real-time GTFS data feed* Real-time transit fleet feed in GTFS real-time [39] format that contains three types of data: service alerts, trip updates and vehicle positions. The data source of the feed includes streaming AVL data on operating buses.
- *Time point data sets* Time point Datasets are the historical bus data at time points, including route ID, trip ID, drive ID, actual departure and arrival time, etc. This data is not available in real-time and is only made available at the end of the month.
- *Crowd-sourced data feed* Crowd-sourced data feed is collected anonymously from the Transit-Hub mobile app. Anonymous data generated by users is updated to the server when a user uses the app for route planning and navigation. It should be noted that this data set is not being used in the system described in the paper, however, we will exploit the integration of user-supplied data for closing the loop to the users in the future release.

4.2 Data management

4.2.1 Data collection

We have to handle data from each source differently as they have different update rates and formats. For example, (i) Bus schedule data (static GTFS) is updated only when MTA modifies its bus routes or schedules; (ii) Historical time point data set is collected by MTA at the end of the month and is then manually transferred and imported into our MongoDB database. On an average, we collect approximately three hundred thousand entries each month; and (iii) For the real-time transit data, our back-end server requests the data from these real-time feeds every minute and stores the responses in the database (see Table 1).

4.2.2 Data cleaning

Data cleaning is a crucial step for data pre-processing to handle the following issues:

- *Duplicated data* Detecting and eliminating duplicated data is one of the major tasks for data cleaning. We compare and remove data with the same time stamps and key-value pairs.
- *Data with logistic errors* This type of data exists mainly in the real-time bus location data. To deal with it, for example, we remove the records where a bus' distance from a stop changes too fast, or if it moves in the wrong direction. This is done using some custom filters created by us.
- *Missing data* This can happen for various reasons, which are: (a) operational disruptions due to service alerts, (b) hardware failures, or (c) data transmission issues. The missing data is filled in using linear interpolation on the sampled data.

4.2.3 Data storage

The large scale of the historical and real-time transit data that are accumulated over time requires efficient storage and management methods. Also, the stored data must be accessible to multiple clients in the system at the same time. To meet this scale requirement, we employ JSON as the data structure and MongoDB [40] for data storage. MongoDB is a distributed NoSQL database that can efficiently store and query data on the scale of terabytes.

4.3 Opportunities for closing the DDDAS loop

This section shows how data feedback in the transportation decision support system can be used to help (MTA) to identify

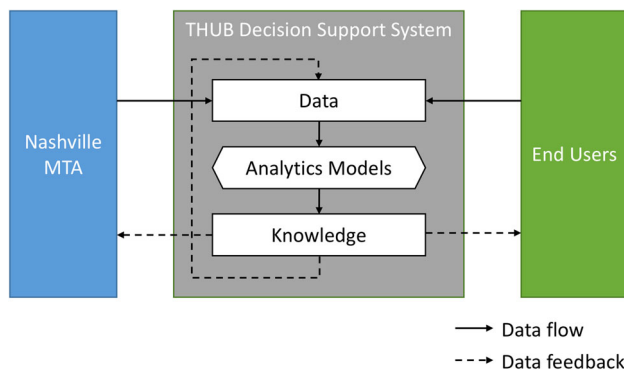


Fig. 1 Proposed DDDAS Loop in Transit-Hub transportation decision support system between MTA, Transit-Hub and end users

real-time outliers and perform long-term delay optimization to improve bus services and availability.

Figure 1 illustrates the data feedback cycle. We utilize the multi-source data from Nashville MTA to conduct real-time and long-term data analytics, and the results can be sent back to them as feedback in different ways:

- *Metro transportation authority (MTA)* By doing long-term bus data analysis, our models can find the delay patterns that are associated with seasons, day of the week, and time of day. This feedback can be used by MTA to identify bottlenecks within routes and adjust the bus timetable or route layout accordingly. Also, by tracking the real-time bus data and comparing it with the historical delay patterns, we are able to find the outlier trips that deviate from the normal ones, which will be used to inform MTA to investigate and avoid these in the future.
- *End Users* We are collecting anonymous usage and location data from application users. This data can be used to provide an alternative real-time data source for buses. If a user plans to take a bus that is full of people, the system can send notifications to advise him/her to take some other bus or routes. In addition, it can also help to optimize the bus route network and reduce rider walking distances as it shows the origins of users to the bus stops and helps MTA to identify areas with low/high transit service availability.

5 Model construction

In this section, we present how we construct the long-term delay model, short-term travel-time model and arrival delay prediction model. In particular, we solve Problem 2 described in Sect. 3.2 by creating a shared route segment network and utilize real-time data from multiple routes.

5.1 Building model for analyzing long-term delay patterns

The section describes a long-term analytics model that constructs historical bus delay patterns at time points. In this model, clustering methods are applied to historical arrival delay and travel delay data.

5.1.1 Clustering analysis

For each weekday, K-means clustering algorithm [41] is used to obtain the cluster of the delay data in accordance with the delay and time of the day by minimizing the within-cluster sum of squares (WCSS).

$$\operatorname{argmin}_S \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2 \quad (1)$$

where μ_i denotes the mean of all points in the cluster S_i .

Silhouette analysis [42] is an approach to measure how close each point is to others within one cluster.

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}} \quad (2)$$

where for each data point i in the cluster, a_i is the average distance between i and the rest of data points in the same cluster, b_i is the smallest average distance between data point i and every other cluster, and $s(i)$ is the silhouette score. We calculate the silhouette scores for 2–5 clusters derived from K-means algorithm to find the optimal number of clusters with the lowest silhouette score.

The normal distribution of the clustered data helps to identify the typical delay patterns of previous buses, which can be given to users when they want an estimate for a future time, or if there is no real-time data available. The time point data is imported into the database at the end of each month. Then the data is stored according to weekday. We subsequently generate the clusters and normal distributions for all the route segments in each group. Meanwhile, the clustered data and normal distributions are cached and stored in the database. Thus, when we have to query the model, there is no need to run clustering analysis again.

Example Consider a time point 'HRWB' on route 3 in Nashville. The historical bus arrival delay data we select is for Wednesday, outbound direction, between June 1, 2016 and June 30, 2016 (for a total of 185 points). Figure 2 displays the delay data for a day during that month. In the figure, there are two obvious groups (yellow and red), one is between 5–2 PM and the other one is between 2 and 12 AM. The two groups reveal that there exist two different delay patterns which happen in the morning and in the afternoon separately.

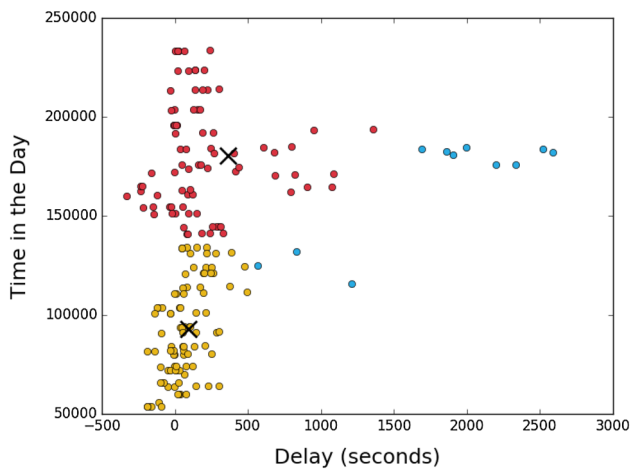


Fig. 2 Cluster historical delay data according to the delay and time in the day at time point “HRWB” on route 3. The figure shows that there are two active delay patterns, one before and one after 2 PM. The blue dots are outliers identified by analysis in Sect. 5.1.3




Fig. 3 Normal distribution of the clustered historical delay data at time point “HRWB” on route 3

This information can be provided to end users to help them plan trips.

5.1.2 Normality test and analysis.

The analytics is based on the assumption that historical delay data has a normal distribution. In order to ensure this, we perform normality test on each cluster that we get in the previous step. We can calculate the confidence interval for long-term delay analysis from the distribution curve.

Example These are the two normal distributions in Fig. 3 that we obtain after performing the normality test on the clusters generated from the data described in the previous example. The cluster for the delay in the afternoon has a higher mean value (92.0 vs. 58.0 s) and a wider normal distribution curve, which indicates that buses on route 3 are more likely to be on

Table 2 Mean value of the delay data distributions for 4 time points on route 3 in morning and afternoon in June

	Timepoints			
	WE23	WE31	HRWB	WHBG
Morning	116.90	127.71	93.14	443.52
Afternoon	121.03	146.28	114.48	545.49

time in the afternoon. In the afternoon, the 95% confidence interval of delay is between -60.4 and 244.4 s while in the morning the 95% confidence interval of delay is between -73.5 and 189.6 s (the negative seconds mean the buses are predicted to arrive earlier than scheduled time).

5.1.3 Outlier analysis

In order to identify outliers from historical bus data, the first step is to generate the normal distribution for each of the clustered data groups described in the former sections. Since for a normal distribution where μ is the mean value and σ is the standard deviation, 95% of all data is within the confidence interval of $[\mu - 2\sigma, \mu + 2\sigma]$, we define that the outliers are the historical data with delay greater than $\mu + 2\sigma$ or less than $\mu - 2\sigma$ in the distribution.

Example For the dataset mentioned in the previous two examples, there exist some outliers (blue points) in Fig. 2. These outliers belong to the two clusters obtained from clustering analysis and are identified by outlier analysis. The outliers mostly emerged during rush hours in the morning and in the evening. One hypothesis is that during rush hours, there are more passengers and more traffic congestion on the route, which will increase the boarding time at stops and travel time on the road. Since our back-end server is monitoring the real-time transit feeds and in the meantime records real-time data, trips that have severe outliers and do not fit in the typical delay pattern can be easily detected and used for further investigation.

5.1.4 Bottleneck identification

After mean delay patterns of all time points and all route segments are derived, we can then identify the bottlenecks along the routes by using those patterns. This also helps so that actions to optimize the route performance can be taken afterwards.

There are 4 time points “WE23”, “WE31”, “HRWB” and “WHBG” on route 3 (traveling away from downtown Nashville). Table 2 shows the findings that the typical arrival delay for “WHBG” is 443.52 s in the morning and 545.49 s in the afternoon. Considering the fact that the typical arrival delays for “WE23”, “WE31” and “HRWB”, timepoints




Fig. 4 Finding shared route segments between two bus routes. The segment that contains the three center points is shared by route 1 and route 2

before “WHBG”, in the morning and afternoon are all below 150 s, we can draw the conclusion that the bus stops between “HRWB” and “WHBG” are the bottlenecks for route 3.

5.2 Real-time data integration

This section describes a short-term bus arrival delay prediction model that we have developed to address the challenges presented in Sect. 3. The model integrates real-time bus location data of shared route segments and combines clustering analysis and Kalman filters for delay prediction.

5.2.1 Utilizing shared route segment data

Problem 2 from Sect. 3.2 describes the issue that real-time bus data is not always available due to infrequency of buses. To address this challenge, the short-term delay prediction model in Transit-Hub creates a shared bus route segment network, and uses the real-time data from shared route segments for short-term predictive analysis.

Our prior work [17] was based on shared route segments, but at that time we used shared segments that were manually selected and we did not provide a solution to automatically identify shared route segments. In this paper we present an algorithm to create a shared bus route segment network for all the existing routes in the city [43]. Also, the data that the algorithm uses is in standard GTFS format, so the algorithm can be easily applied to other cities that use the same data format.

A Route segment is defined as a maximal part of bus route that is shared by a set of bus routes. In GTFS format, the physical path of bus routes is described using a sequence of coordinate points (in the *shapes.txt* file) on the map. If there are two segments from two bus routes that share the same sequence of coordinate points, then we can assume that the routes share that road segment. The outline of the algorithm to generate the shared route segment network is described as below (The key steps are illustrated in Fig. 4):

Input: *Static GTFS dataset* Static bus and associated geographic information are loaded from database.

Output: *Shared route segment network* Segment layout for each bus route is saved in the database.

Step 1: *Map grid initialization* The Nashville map is divided into map grids of squares. The length of each square is about 8.97 meters, so each grid cell covers about 80.51 square meters on the map.

Step 2: *Route path re-sampling and smoothing* The sequences of points in all bus routes are re-sampled to the centers of grid cells if the point is covered by the cell. Also, if the distance between adjacent points in the sequences is larger than the width of a grid cell, points will be interpolated to fill the cells that are missing points. The re-sampled points of each route are cached in the database for determining the shared route segments in the later step. As shown in Fig. 4, the paths of route 1 and 2 are re-sampled to the center points of grid cells.

Step 3: *Calculating segments for bus routes* Each cell is tagged by every route that uses that cell. If a cell contains tags from multiple routes then it becomes part of a new shared segment. For example, the three-point segments in Fig. 4 are shared by route 1 and 2, so this segment is marked as a shared route segment. New segments are checked to make sure no duplicated segments are generated.

Step 4: *Segment length limitation* Any segment that has a length that is greater than 1 mile is divided into smaller segments. because our model is based on the assumption that the travel delay within each segment is equally distributed, and hence the division of larger segments into smaller ones will satisfy this assumption and reduce prediction error.

Using Nashville’s static GTFS (version of March 9, 2016), we generated a shared route segment network shown in Fig. 5. The 57 bus routes in Nashville city were divided into 5139 segments. The lines in different colors show different route segments. Since the static bus schedules are updated regularly by MTA, the shared route segment network should be updated when new schedules are released.

There are many benefits to using real-time data of shared route segments, such as: (1) Utilizing the real-time data from other routes can greatly increase the volume of data that are available for short-term delay prediction analysis. For example, the route 3 in Nashville from White Bridge to Downtown has a schedule interval of 40 min at holiday and weekends. Only using route 3 data means the most recent data is at least 40 min old, which is not recent enough to predict the currently delay on route 3. (2) The length of each segment in the network can be controlled by the one-mile limitation mentioned in the last step of the algorithm. Since the delay pattern varies along a bus route, segments with longer length are divided by the algorithm to produce more accurate analytics results. (3) By creating a shared route segment model, the divide and conquer design pattern is used. Individual and self-maintained microservice model can then run for each of the segments concurrently.

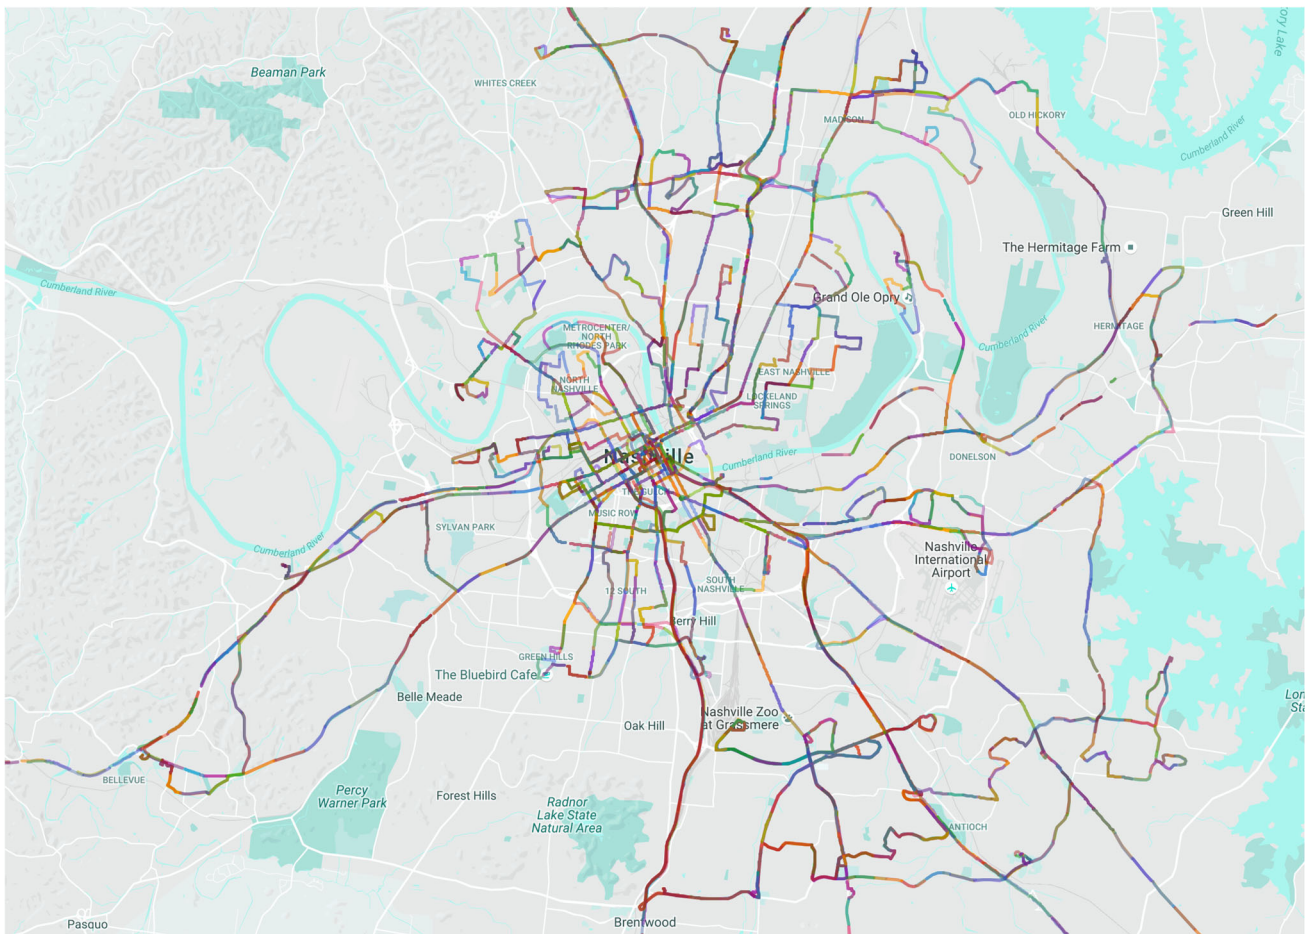


Fig. 5 Generated shared bus route segment network in Nashville. The lines with different colors represent the 5139 shared route segments in all 57 bus routes in the network. The length of the segments are limited to less than 1 mile

5.2.2 Estimating the arrival time at bus stops

Since the actual arrival time at bus stops are not included in the real-time GTFS feed in Nashville, we integrate the real-time bus location data and the static bus stop locations to estimate the arrival time of buses.

From the real-time bus location feed, we can get the bus location and timestamps in the following array format: $[(t_1, d_1), \dots, (t_k, d_k), \dots]$. Because the update rate of the original data varies from seconds to minutes, we first aggregate the collected data into 1-min average data using sliding time windows. Then, we assume that bus speed is approximately the average of the two adjacent data points and apply the following equation to calculate the bus arrival time at stops:

$$t_{stop} = t_{k-1} + (t_k - t_{k-1}) \frac{d_{stop} - d_{k-1}}{d_k - d_{k-1}} \quad (3)$$

where t_{stop} denotes the estimated arrival time, d_k is the bus's distance from the current location to the first bus stop of the

route along the route path at time t_k . Also, $d_{k-1} \leq d_{stop} < d_k$.

5.2.3 Updating the travel delay prediction using K-means algorithm and smoothing filter

Excluding the outliers If the travel time of a preceding bus differs greatly from other preceding buses, we consider this point an outlier and exclude it from the model computation.

To identify the outliers from the data, we employ K-means algorithm to cluster the preceding bus data according to travel time and time in the day. The Silhouette analysis that was introduced in Eq. (2) is also used here to find the optimal number of clusters. We choose the cluster whose time of day is closest to the current time. The data points from that cluster are smoothed through the filter described in the next section and used as an estimate for the current travel time on that segment.

Smoothing the preceding bus data By comparing the travel time of preceding buses and the scheduled travel time within the route segment, we compute the travel delay of the pre-

ceding buses in the segment. The travel delay data is then through a filter to eliminate noise and predict the segment's current travel delay. The state transition equation is:

$$x_k = x_{k-1} + \omega_{k-1} \quad (4)$$

where the state variable x_k denotes the time step for which the travel delay needs to be predicted, ω_k denotes the zero mean normal distribution noise with covariance Q_k .

The observation equation used is:

$$z_k = x_k + \nu_k \quad (5)$$

where variable z_k represents the observation of delay at time step k . ν_k represents the zero mean Gaussian distribution observation noise with covariance R_k . ω_k and ν_k are assumed to be independent.

5.2.4 Example

In this section we use an example to explain the workflow of Transit-Hub multi-timescale analysis services. Figure 6 illustrate a common scenario where a bus b_1 is running along a bus route r_1 and the system needs to predict on request, the expected delay for a bus at stop s_i :

1. *Creating shared route segment network* From the figure we can see that routes r_1 and r_2 are divided into 5 segments: seg_1 , seg_2 , seg_3 , seg_4 , seg_5 . The segment seg_2 is shared by the routes.
2. *Getting preceding buses using static bus schedules* From the static bus schedules we find that there are many buses (b_2 , b_3 , etc.) from route r_1 and r_2 that have passed through segment seg_3
3. *Estimating travel time of the buses in segments* Preceding buses' travel time can be estimated using the collected real-time bus location data.
4. *Predict travel delay in segments* The data from recent buses are clustered by travel time and time in the day. The group of data whose mean value (time in the day) is closed to the current time will be smoothed with a Kalman filter.
5. *Getting arrival delay at bus stop* The sum of the delays for each segment between the current bus position and the target stop s_n is the model's prediction for arrival.

6 Deployed architecture

In this section, we describe the implementation architecture for the short term online delay prediction service, which




Fig. 6 Use case: example of using shared route segment data to predict a bus's delay at a bus stop

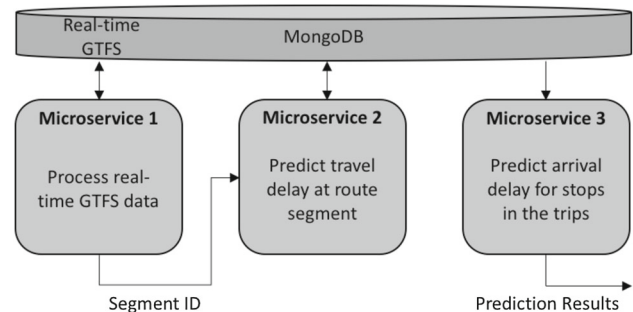


Fig. 7 Microservice architecture of Transit Hub back-end analytics services

addresses problem 4 described in Sect. 3.4 concerning scalability and availability.

Section 3.4 compared two deployment patterns: traditional monolithic style and microservice style. Microservice deployment is an application architectural pattern where independently deployable and self-contained services can work together, which may be more suitable for complicated web applications [34–36]. Microservices communicate with each other via lightweight network mechanisms, such as using REpresentational State Transfer (REST) API, message broker, etc. Fig. 7 illustrates the overall architecture of the Transit-Hub analytics.

Microservice 1: Smoothing the real-time GTFS data Microservice 1 first cleans the raw real-time GTFS data by removing the duplicate and missing data, and then re-sample it to estimate the bus arrival time on bus routes. This microservice tracks real-time bus location and when a new bus travels through a route segment, it will inform Microservice 2 which updates the travel time delay for this route segment. Microservice 1 is activated by a scheduler every 5 min.

Microservice 2: Predicting arrival delay at segments Microservice 2 collects the data processed by Microservice 1 and employs short-term delay prediction (Sect. 5.2) to update the estimated delay for the route segments. When Microservice 2 receives a prediction update request for a route segment, it wakes up and runs the prediction process to update the travel delay prediction for that route segment.

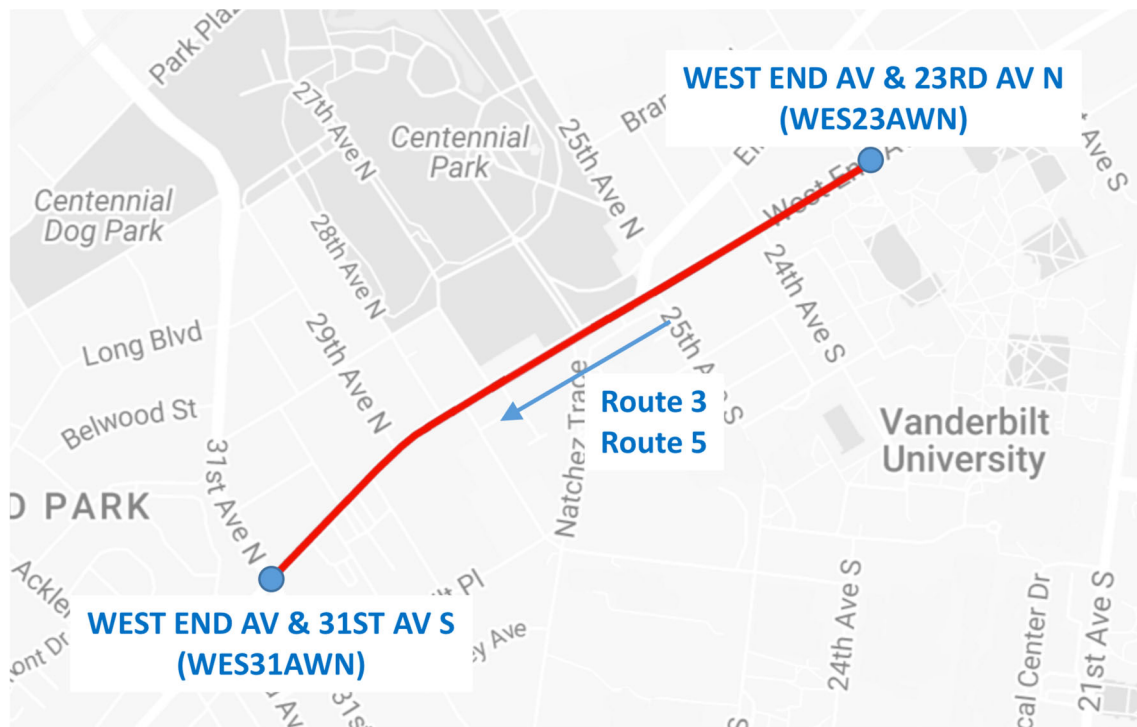


Fig. 8 Studied road segment shared by route 3 and 5

Microservice 3: Predicting arrival time at bus stops
 Microservice 3 combines the current delay of all buses and the predicted travel delay for all route segments to produce the arrival delay prediction at all bus stops for all routes. This microservice is activated every minute and stores the prediction results in the database. Note that Microservice 2 runs per route segment whereas service 3 runs to update the arrival time for all routes.

Representational state transfer (REST) API and message broker are two of the popular approaches for providing a communication mechanism between microservices. The REST approach is synchronous by default and uses DNS or a registry for service discovery, and supports load balancing by using software like Ribbon [44]. The message broker is an asynchronous mechanism, which uses queues to manage message queues and can achieve load balancing very easily. Asynchronous message passing is a better choice for microservices because: (1) the individual microservice that sends a message will not be blocked before the other microservice responds; (2) using asynchronous communication can help to reduce unnecessary duplicate computation. For example, in our architecture, microservice 1 is continuously sending the IDs of route segments that need to update prediction to microservice 2. If we find that there are two identical segment IDs in the message queue, then the duplicate can be removed to avoid duplication of work. Based on these considerations, we use RabbitMQ [45], which is a message broker that provides asynchronous messaging.

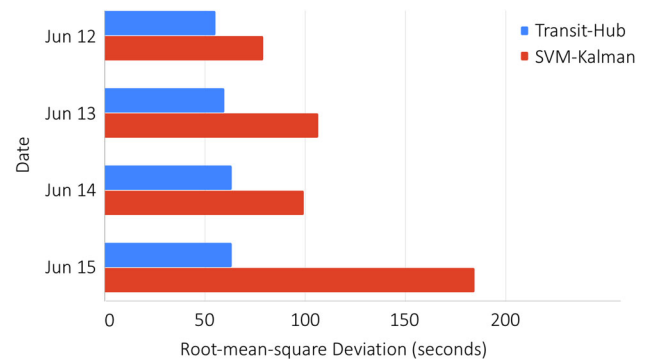


Fig. 9 RMSD of travel time delay prediction for each day when comparing the Transit-Hub model with the SVM Kalman model proposed in 2015. Transit-Hub model outperforms the SVM-Kalman model: (1) RMSD values are smaller (2) it shows less variation on different days

The microservices are deployed on an OpenStack [46] cloud operating system. We created an *m1.large* nova computing instance for the microservices which has 4 virtual CPUs, 8GB RAM and runs Ubuntu 14.04 (LTS). The microservices all together use 10.9% CPU resources and 28% RAM on average. The performance and resource consumption of the microservices will not be affected by user interactions. They run separately and repeatedly in the back end and store analysis results for later use. When an end user sends a prediction request for a route, an independent ser-

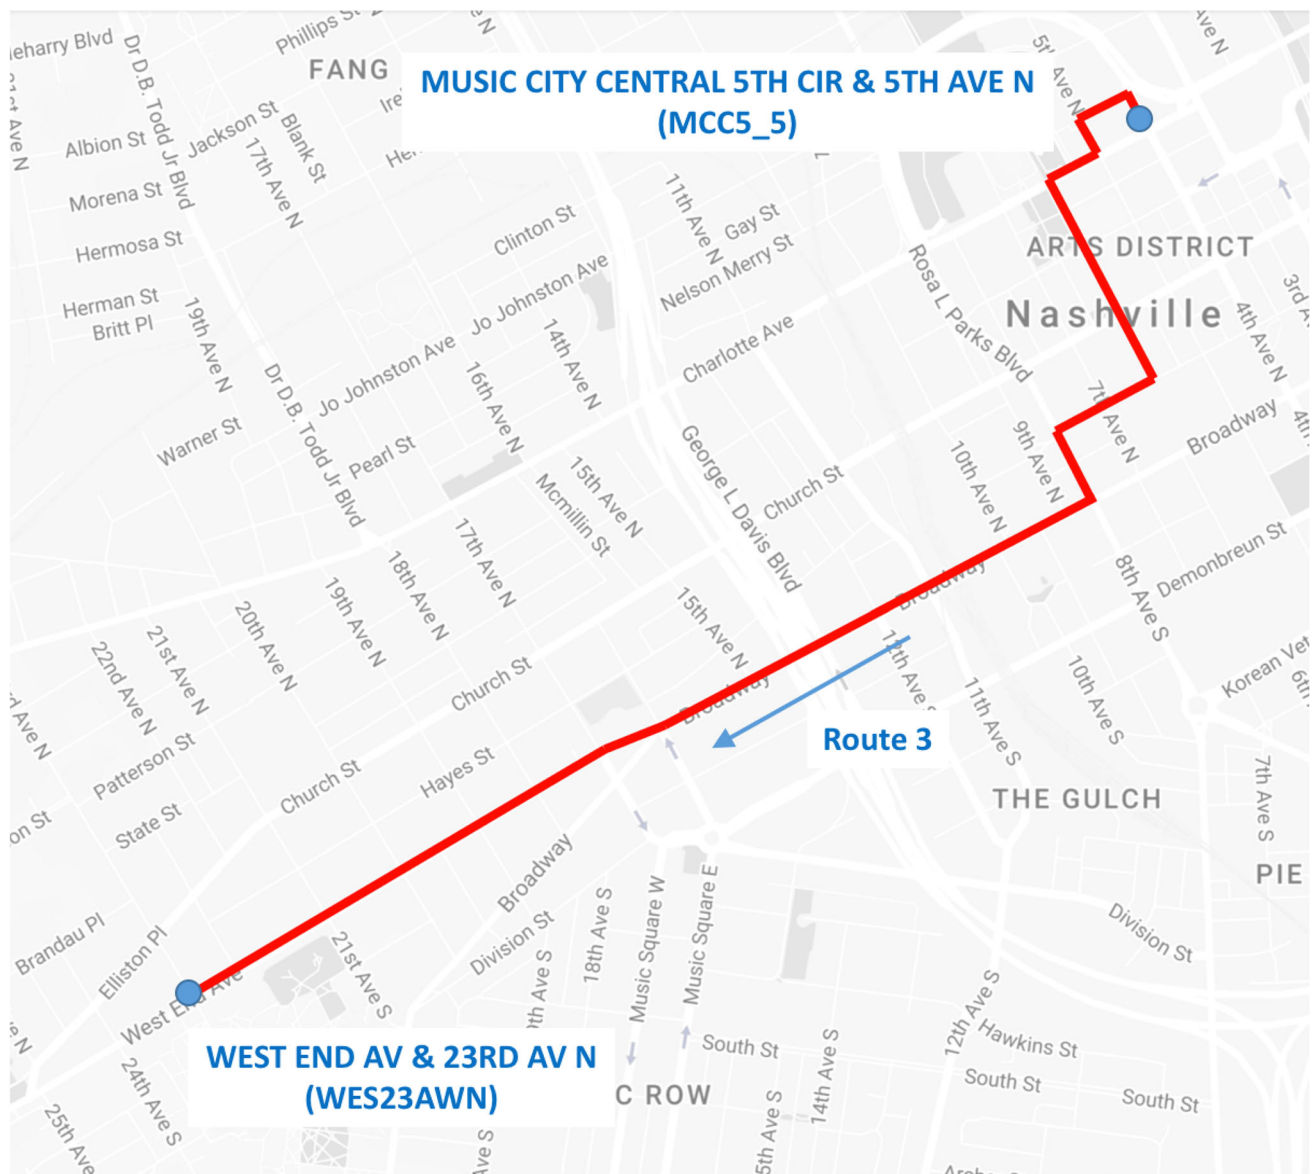


Fig. 10 Arrival time delay prediction for a bus stop of a trip: (1) actual arrival delay, (2) predicted mean value – standard deviation, (3) predicted mean, (4) predicted mean value + standard deviation

vice in the system will fetch the prediction results from the database and provide the information to the end user.

7 Prediction performance evaluation

This section presents experimental results from Transit-Hub's real-time delay prediction model. These results empirically evaluate Transit-Hub's bus travel time delay prediction ability against an SVM-Kalman model [9] using real-time data collected in Nashville. Compared to the SVM-Kalman model, our model takes the scheduled time of preceding buses into consideration, and since we are clustering the data of pre-

ceding buses according to time of day and delay, only clusters with an average time of day close to the current time of day will be used. We also evaluate how well our model predicts arrival delay comparing it against real-world data.

7.1 Experiment 1: evaluating the travel time delay prediction

The first experiment is designed to evaluate Transit-Hub's ability to predict travel time delay, using its prediction model and comparing against other prediction models using the same real-world data.

7.1.1 Experiment setup

Routes 3 and 5 are two of the major bus routes in Nashville. As shown in Fig. 8, they share the same route segment between time point WES23AWN and time point WES31AWN along West End Avenue. We select this route segment of route 3 and 5 towards WHITE BRIDGE to test our proposed model.

The data used in this experiment is the real-time and static GTFS data for routes 3 and 5 that we collected from Nashville MTA in June 2016. We divide the data into two parts: a training dataset and a validation dataset. The training dataset contains bus data from June 6th to June 12nd and the validation dataset contains data from Jun 13rd to Jun 15th. Our model and the SVM-Kalman are evaluated using the same validation dataset. From our previous paper [17] we learned that only data 120 min old or newer is important for real-time delay prediction. Therefore, in this experiment we use the data for buses in the past 2 h.

7.1.2 Comparing with a SVM-Kalman model

In order to evaluate the performance of the proposed short-term delay prediction model, we chose and implemented a dynamic SVM-Kalman model that was proposed by Bai, et al. in 2015 [9]. The dynamic model consists of a support vector machines (SVM) model that uses historical data to estimate the current travel time as a baseline prediction, and a Kalman filter model that uses real-time preceding bus data to adjust the base time. The features that they use in the SVM model include: (1) time of the day, (2) road segment ID, (3) weighted average bus travel time of preceding buses, and (4) the travel time of the preceding buses on the same route.

7.1.3 Results

Figure 9 shows the root-mean-square deviation (RMSD) of the travel time delay prediction results for three days in June. The RMSD of travel time delay is calculated using the following equation:

$$t_{ij}^{act_tra} = t_j^{act_arr} - t_i^{act_dep} \quad (6)$$

$$RMSD = \sqrt{\frac{\sum_i^n (t_{ij}^{act_tra} - t_{ij}^{pred_tra})^2}{n}} \quad (7)$$

where i and j are indexes of the timepoints along the route, and $i < j$. Variable $t_i^{act_arr}$ and $t_i^{act_dep}$ represent the actual arrival and departure time at timepoint i , $t_{ij}^{act_tra}$ and $t_{ij}^{pred_tra}$ represent the actual and predicted travel time at the segment between timepoint i and j , respectively. n is the number of bus trips in the dataset.

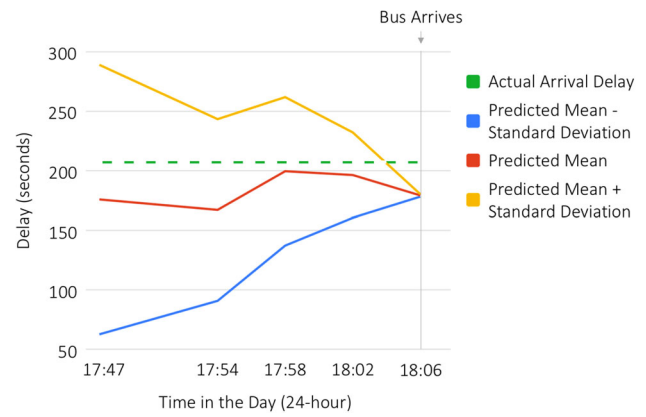


Fig. 11 Studied segment of route 3 that starts from first bus stop (MCC5_5) to the 15th bus stop (WES23AWN)

Since the SVM model ignores the differences that exist in the scheduled travel of preceding buses and the model does not exclude outliers, we expect our model to outperform the SVM model from [9]. The experimental results validate our hypothesis. When predicting the travel time delay 15 min ahead using collected data from Jun 12 to Jun 15, the RSMD of the our model is about 30–65% lower compared to the SVM-Kalman model.

7.2 Experiment 2: evaluating the arrival time delay prediction

The second experiment is designed to evaluate the short-term prediction model's performance when the prediction horizon changes. For this experiment, we choose a trip from route 3 on June 14th 2016. The studied segment is shown in Fig. 10.

Results Figure 11 shows the actual delay and the predicted arrival delay with confidence interval as the prediction horizon decreases from 19 to 0 min (the time just before the bus arrived).

Since our model integrates the predicted the travel delay in route segments and estimated arrival delay at the most recent bus stop that the bus passed, we expect the predicted confidence interval will become smaller and the error will decrease as the prediction interval reduces, i.e., as we make the prediction closer to the scheduled time of arrival.

This example shows that when predicting 19 min before the actual arrival time, the confidence interval is 226.5 s and the interval decreases to 2.1 s. We notice a 27.8 s difference between predicted delay and actual delay when the bus arrives, we attribute this to normal system variance.

8 Conclusion and future work

In this paper, we presented research on a DDDAS-enabled smart public transportation decision support system that sig-

Table 3 Summary of architectural decisions in Transit-Hub

Challenge	Design principle	Approach	Section
Learning the historical delay patterns	Long-term analytics model	Clustering analysis, normality test and outlier analysis	Sec 5.1
Accurate bus delay prediction	Better usage of the real-time bus data	Prediction model that combines clustering and Kalman filter	Sec 5.2
Lack of quality real-time data	Integrating shared route segment data	Shared route segment network	Sec 5.2.1
Improve scalability and fault isolation	Separation into independent modules	Microservice architecture	Sec 6
Optimize the bus service	Data feedback loop	Provide the analytics results via feedback loop to MTA	Sec 4.3

nificantly extends our prior work on Transit-Hub [17] by illustrating and validating the methods developed for long-term and short-term predictive analytics services. Table 3 summarizes the work by presenting the challenges we resolved, the corresponding design principle used, and approaches when developing the system. Our long-term delay analysis service excludes the noise of outliers in the historical dataset and identifies the delay patterns of time points and route segments that are associated with different times of day, day of the week and seasons. The city planners can utilize the feedback data to optimize the bus schedules and improve rider satisfaction. Residents and travelers in cities like Nashville can also benefit from our short-term delay prediction services.

In the future, the work presented in the paper can be extended in the following ways: (1) We want to integrate more data sources into the analysis and prediction models. New data sources, such as traffic flow, weather conditions, special events, can impact public transportation and can be used as new feature vectors to improve the current services. The crowd-sourced data is being collected and the integration of user data will be explored in the future. (2) The services can be deployed further in edge devices using tools like Apache Edgent [47] to reduce the data transmission between edge nodes and a central analytics engine. (3) The system can fit into a smart city platform called Cyber-pHysical Application Architecture with Objective-based reconfiguratiOn (CHAR-IOT) [48], which will improve the system's resilience and communication heterogeneity. (4) Storm is a scalable, fast and distributed computation system. In order to scale the Transit-Hub system to serve multiple cities in the future, Storm can be integrated into the system to consume the distributed streaming real-time data feeds, run the multi-timescale analytics and then make the results available to all users.


Acknowledgements This work is supported by The National Science Foundation under the award numbers CNS-1528799 and CNS-1647015 and a TIPS grant from Vanderbilt University. We acknowledge the sup-

port provided by our partners from Nashville Metropolitan Transport Authority.

References

1. APTA: Americans took 10.6 billion trips on public transportation in 2015. (2016)
2. APTA: Record 10.7 billion trips taken on U.S. public transportation in 2013. (2014)
3. Federal Highway Administration: Travel monitoring and traffic volume (2014)
4. Bates, J., Polak, J., Jones, P., Cook, A.: The valuation of reliability for personal travel. *Transp. Res. Part E* **37**(2), 191–229 (2001)
5. Gooze, A., Watkins, K., Borning, A.: Benefits of real-time transit information and impacts of data accuracy on rider experience. *Transp. Res. Record* **2351**, 95–103 (2013)
6. Dziekan, K., Kottenhoff, K.: Dynamic at-stop real-time information displays for public transport: effects on customers. *Transp. Res. Part A* **41**(6), 489–501 (2007)
7. Bin, Y., Zhongzhen, Y., Baozhen, Y.: Bus arrival time prediction using support vector machines. *J. Intell. Transp. Syst.* **10**(4), 151–158 (2006)
8. Zhang, C., Teng, J.: Bus dwell time estimation and prediction: a study case in shanghai-china. *Procedia-Soc. Behav. Sci.* **96**, 1329–1340 (2013)
9. Bai, C., Peng, Z.R., Lu, Q.C., Sun, J.: Dynamic bus travel time prediction models on road with multiple bus routes. *Comput. Intell. Neurosci.* **2015**, 63 (2015)
10. Mazloumi, E., Mesbah, M., Ceder, A., Moridpour, S., Currie, G.: Efficient transit schedule design of timing points: a comparison of ant colony and genetic algorithms. *Transp. Res. Part B* **46**(1), 217–234 (2012)
11. Dessouky, M., Hall, R., Nowroozi, A., Mourikas, K.: Bus dispatching at timed transfer transit stations using bus tracking technology. *Transp. Res. Part C* **7**(4), 187–208 (1999)
12. Fu, L., Liu, Q., Calamai, P.: Real-time optimization model for dynamic scheduling of transit operations. *Transp. Res. Record* **1857**, 48–55 (2003)
13. Sun, A., Hickman, M.: The real-time stop-skipping problem. *J. Intell. Transp. Syst.* **9**(2), 91–109 (2005)
14. Real-time port authority bus tracking system not always real. <http://www.post-gazette.com/news/transportation/2014/10/16/Real-time-Port-Authority-tracking-not-always-real/stories/201410160155> (2014) Accessed 30 Sept 2016
15. Cota says its real-time bus-tracking system doesn't work. <http://www.dispatch.com/content/stories/local/2014/07/23/COTA->

- [says-its-GPS-system-doesnt-work.html](#) (2014). Accessed 30 Sept 2016
16. Darema, F.: Dynamic data driven applications systems: a new paradigm for application simulations and measurements. *Computat. Sci.-ICCS* **2004**, 662–669 (2004)
 17. Sun, F., Pan, Y., White, J., Dubey, A.: Real-time and predictive analytics for smart public transportation decision support system. In: 2016 IEEE International Conference on Smart Computing (SMARTCOMP), pp 1–8, <https://doi.org/10.1109/SMARTCOMP.2016.7501714> (2016)
 18. Shekhar, S., Sun, F., Dubey, A., Gokhale, A., Neema, H., Lehofer, M., Freudberg, D.: Transit hub. In: Geng, H. (ed.) *Internet of Things and Data Analytics Handbook*, pp. 597–612. Wiley, Hoboken (2016)
 19. Wu, C.H., Ho, J.M., Lee, D.T.: Travel-time prediction with support vector regression. *IEEE Trans. Intell. Transp. Syst.* **5**(4), 276–281 (2004)
 20. Yu, B., Lam, W.H., Tam, M.L.: Bus arrival time prediction at bus stop with multiple routes. *Transp. Res. Part C* **19**(6), 1157–1170 (2011)
 21. Chien, S.I.J., Kuchipudi, C.M.: Dynamic travel time prediction with real-time and historic data. *J. Transp. Eng.* **129**(6), 608–616 (2003)
 22. Jeong, R.H.: The prediction of bus arrival time using automatic vehicle location systems data. PhD thesis, Texas A&M University (2005)
 23. Weigang, L., Koendjibiharie, W., de M Juca, R., Yamashita, Y., MacIver, A.: Algorithms for estimating bus arrival times using gps data. In: *Intelligent Transportation Systems, 2002. Proceedings of the IEEE 5th International Conference on*, IEEE, pp 868–873 (2002)
 24. Sun, D., Luo, H., Fu, L., Liu, W., Liao, X., Zhao, M.: Predicting bus arrival time on the basis of global positioning system data. *Transp. Res. Record* **2034**, 62–72 (2007)
 25. Chien, S.I.J., Ding, Y., Wei, C.: Dynamic bus arrival time prediction with artificial neural networks. *J. Transp. Eng.* **128**(5), 429–438 (2002)
 26. Patnaik, J., Chien, S., Bladikas, A.: Estimation of bus arrival times using apc data. *J. Public Transp.* **7**(1), 1 (2004)
 27. Shalaby, A., Farhan, A.: Bus travel time prediction model for dynamic operations control and passenger information systems. *Transp. Res. Board* **2** (2003)
 28. Yang, J.S.: Travel time prediction using the gps test vehicle and kalman filtering techniques. In: *Proceedings of the 2005, American Control Conference, 2005.*, IEEE, pp 2128–2133 (2005)
 29. Chen, M., Liu, X., Xia, J., Chien, S.I.: A dynamic bus-arrival time prediction model based on apc data. *Comput.-Aided Civil Infrastruct. Eng.* **19**(5), 364–376 (2004)
 30. Jeong, R., Rilett, R.: Bus arrival time prediction using artificial neural network model. In: *Intelligent Transportation Systems, 2004. Proceedings of the 7th International IEEE Conference on*, IEEE, pp 988–993 (2004)
 31. Mazloumi, E., Moridpour, S., Currie, G., Rose, G.: Exploring the value of traffic flow data in bus travel time prediction. *J. Transp. Eng.* **138**(4), 436–446 (2011)
 32. Nashville mta maps and schedules. <http://www.nashvillemta.org/Nashville-MTA-Maps-and-Schedules.asp> (2016). Accessed 26 Sept 2016
 33. M15 service between east harlem and south ferry. <http://web.mta.info/nyct/bus/schedule/manh/m015scur.pdf> (2016). Accessed 26 Sept 2016
 34. Rama, G.M., Patel, N.: Software modularization operators. In: *Software Maintenance (ICSM), 2010 IEEE International Conference on*, IEEE, pp 1–10 (2010)
 35. Sarkar, S., Ramachandran, S., Kumar, G.S., Iyengar, M.K., Rangarajan, K., Sivagnanam, S.: Modularization of a large-scale business application: a case study. *IEEE Softw.* **26**(2), 28–35 (2009)
 36. Newman, S.: *Building Microservices*. O'Reilly Media, Inc., Newton (2015)
 37. Thönes, J.: Microservices. *IEEE Softw.* **32**(1), 116–116 (2015)
 38. General transit feed specification (gtfs) static overview. <https://developers.google.com/transit/gtfs/> (2016). Accessed 18 Sept 2016
 39. General transit feed specification (gtfs) real-time overview. <https://developers.google.com/transit/gtfs-realtime/> (2016) Accessed 18 Sept 2016
 40. The mongodb 3.2 manual. <https://docs.mongodb.com/manual/> (2016). Accessed 25 Sept 2016
 41. Lloyd, S.P.: Least squares quantization in pcm. *IEEE Trans. Inf. Theor.* **28**(2), 129–137 (1982)
 42. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math* **20**, 53–65 (1987)
 43. Sun, F.: Transit hub—shared route segment network generation algorithm. <https://github.com/visor-vu/thub-shared-route-segment-network> (2016)
 44. Ribbon, a inter process communication (remote procedure calls) library. <https://github.com/Netflix/ribbon> (2016). Accessed 29 Sept 2016
 45. Rabbitmq. <https://www.rabbitmq.com/> (2016). Accessed 24 Sept 2016
 46. Openstack documentation. <http://docs.openstack.org/> (2016). Accessed 30 Sept 2016
 47. Apache edgent documentation. <http://edgent.apache.org/docs/home> (2016) Accessed 30 Sept 2016
 48. Pradhan, S.M., Dubey, A., Gokhale, A., Lehofer, M.: Chariot: A domain specific language for extensible cyber-physical systems. In: *Proceedings of the Workshop on Domain-Specific Modeling*, ACM, pp 9–16 (2015)



Fangzhou Sun is currently a Ph.D. student in computer science at Vanderbilt University. He received his M.S. degree in computer science from Vanderbilt University in 2015 and completed his undergraduate studies in computer science from Nanjing University, China in 2013. His main research topics include (1) developing data mining and Machine Learning techniques to solve Smart City and cybersecurity problems; (2) developing and managing applications, analytics tool boxes and platforms

for Smart City. He is also an active iOS app developer and web developer.



Abhishek Dubey is an Assistant Professor in the Department of Computer Science and Computer Engineering at Vanderbilt University. He is also a senior research scientist at the Institute for Software Integrated Systems at Vanderbilt University. His research area is resilient cyber-physical systems, including fault diagnostics and prognostics and performance management algorithms. He is particularly interested in applying his work to smart and connected communities. His

current projects are related to transportation, smart grid and emergency response domains. Abhishek completed his Ph.D. in Electrical Engineering from Vanderbilt University in 2009. He received his M.S. in Electrical Engineering from Vanderbilt University in August 2005 and completed his undergraduate studies in electrical engineering from the Indian Institute of Technology, Banaras Hindu University, India in May 2001. He is a senior member of IEEE.



Aniruddha Gokhale is an Associate Professor in the Department of Electrical Engineering and Computer Science, and Senior Research Scientist at the Institute for Software Integrated Systems (ISIS) both at Vanderbilt University, Nashville, TN, USA. His current research focuses on developing novel solutions to emerging challenges in edge-to-cloud computing, real-time stream processing, and publish/subscribe systems as applied to cyber physical systems including smart transportation and

smart cities. He is also working on using cloud computing technologies for STEM education. Dr. Gokhale obtained his B.E. (Computer Engineering) from University of Pune, India, 1989; M.S. (Computer Science) from Arizona State University, 1992; and D.Sc (Computer Science) from Washington University in St. Louis, 1998. Prior to joining Vanderbilt, Dr. Gokhale was a member of technical staff at Lucent Bell Laboratories, NJ. Dr. Gokhale is a Senior member of both IEEE and ACM, and a member of ASEE. His research has been funded over the years by DARPA, DoD, industry and NSF including a NSF CAREER award in 2009.



Jules White is an Assistant Professor of Computer Science in the Department of Electrical Engineering and Computer Science at Vanderbilt University. He was previously a faculty member in Electrical and Computer Engineering at Virginia Tech and won the Outstanding New Assistant Professor Award at Virginia Tech. His research has won 5 Best Paper and Best Student Paper Awards. He has also published over 85 papers. Dr. White's research focuses on

securing, optimizing, and leveraging data from mobile cyber-physical systems. His mobile cyber-physical systems research spans four key focus areas: (1) mobile security and data collection, (2) high-precision mobile augmented reality, (3) mobile device and supporting cloud infrastructure power and configuration optimization, and (4) applications of mobile cyber-physical systems in multi-disciplinary domains, including energy-optimized cloud computing, smart grid systems, healthcare/manufacturing security, next-generation construction technologies, and citizen science.