| 1. REPORT NUMBER | 2. GOVERNMENT ASSOCIATION NUMBER | 3. RECIPIENT'S CATALOG NUMBER |
|---|---|---|
| CA16-2873 | | |

| 4. TITLE AND SUBTITLE | 5. REPORT DATE |
|---|---|
| Coordinating Transit Transfers in Real Time | May 6, 2016 |
| | 6. PERFORMING ORGANIZATION CODE |
| | N/A |

| 7. AUTHOR | 8. PERFORMING ORGANIZATION REPORT NO. |
|---|---|
| Carlos Daganzo, Paul Anderson | N/A |

| 9. PERFORMING ORGANIZATION NAME AND ADDRESS | 10. WORK UNIT NUMBER |
|---|---|
| University of California at Berkeley Institute of Transportation Studies Berkeley, CA 94720 | N/A |
| | 11. CONTRACT OR GRANT NUMBER |

| 12. SPONSORING AGENCY AND ADDRESS | |
|---|---|
| California Department of Transportation Division of Research, Innovation and System Information P.O. Box 942873 Sacramento, CA 94273-0001 | 65A0529 |
| | 13. TYPE OF REPORT AND PERIOD COVERED Final Report, 3/1/15 – 3/31/16 |
| University of California Center on Economic Competitiveness in Transportation (UCCONNECT) 2616 Dwight Way, Berkeley, CA 94720-1782 | 14. SPONSORING AGENCY CODE N/A |

15. SUPPLEMENTARY NOTES

16. ABSTRACT

Transfers are a major source of travel time variability for transit passengers. Coordinating transfers between transit routes in real time can reduce passenger waiting times and travel time variability, but these benefits need to be contrasted with the delays to on-board and downstream passengers, as well as the potential for bus bunching created by holding buses for transfers. We developed a dynamic holding strategy for transfer coordination based on control theory. We then obtained the optimal control strategy, where maximum holding time is a function of real-time estimates of bus arrivals and passengers and the uncertainty in these estimates. Total travel time (waiting plus in-vehicle) with the optimal control is found to be globally less than or equal to total travel time without control when uncertainty is bounded. The time savings from transfer coordination increase with the ratio of transferring to through passengers but diminish as uncertainty in the real-time estimates of bus arrivals increases. Field observations at a multimodal transfer point in Oakland show that the proposed control strategy could reduce net transfer delay by 30-39% in a real-world scenario. The data collected also confirm that the upper bound on uncertainty in bus arrivals can be satisfied with existing bus location technology. We conclude with a discussion of complementary measures, such as the provision of real-time information at transfer points and conditional signal priority, which could allow coordination to be applied in more cases.

| 17. KEY WORDS | 18. DISTRIBUTION STATEMENT |
|---|---|
| Public transportation; interlining; transfers; coordination. | No restrictions. This document is available to the public through the National Technical Information Service, Springfield, VA, 22161. |

| 19. SECURITY CLASSIFICATION (of this report) | 20. NUMBER OF PAGES | 21. COST OF REPORT CHARGED |
|---|---|---|
| Unclassified | 26 | |

Reproduction of completed page authorized.

**DISCLAIMER STATEMENT**

This document is disseminated in the interest of information exchange. The contents of this report reflect the views of the authors who are responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the State of California or the Federal Highway Administration. This publication does not constitute a standard, specification or regulation. This report does not constitute an endorsement by the Department of any product described herein.

For individuals with sensory disabilities, this document is available in alternate formats. For information, call (916) 654-8899, TTY 711, or write to California Department of Transportation, Division of Research, Innovation and System Information, MS-83, P.O. Box 942873, Sacramento, CA 94273-0001.

UCCONNECT

# Coordinating Transit Transfers in Real Time

## Final Report
### UCCONNECT 2016 TO 021 - 65A0529

Carlos Daganzo, Chancellor's Professor
University of California, Berkeley

Paul Anderson Ph.D. Candidate
University of California, Berkeley

# Coordinating Transit Transfers in Real Time

Caltrans contract no. 65A0529

Paul Anderson Carlos Daganzo (PI)
University of California Berkeley

May 6, 2016

# Contents

## Figures

## Tables

## Abstract

Transfers are a major source of travel time variability for transit passengers. Coordinating transfers between transit routes in real time can reduce passenger waiting times and travel time variability, but these benefits need to be contrasted with the delays to on-board and downstream passengers, as well as the potential for bus bunching created by holding buses for transfers. We developed a dynamic holding strategy for transfer coordination based on control theory. We then obtained the optimal control strategy, where maximum holding time is a function of real-time estimates of bus arrivals and passengers and the uncertainty in these estimates. Total travel time (waiting plus in-vehicle) with the optimal control is found to be globally less than or equal to total travel time without control when uncertainty is bounded. The time savings from transfer coordination increase with the ratio of transferring to through passengers but diminish as uncertainty in the real-time estimates of bus arrivals increases. Field observations at a multimodal transfer point in Oakland show that the proposed control strategy could reduce net transfer delay by 30-39% in a real-world scenario. The data collected also confirm that the upper bound on uncertainty in bus arrivals can be satisfied with existing bus location technology. We conclude with a discussion of complementary measures, such as the provision of real-time information at transfer points and conditional signal priority, which could allow coordination to be applied in more cases.

## 1  Introduction

Public transit networks usually resemble one of several common patterns including ring-radial, grid, hierarchical (e.g. feeder-trunk), etc. These designs are closely tied to the historical development of cities and to the layout of the street network. They also tend to be efficient methods for achieving spatial coverage. Transfers are essential for reaching destinations in grid and hierarchical systems and increasingly in ring-radial systems too because travel demand has become much more polycentric in recent years.

Transfers take on particular importance in California transit systems. California metropolitan statistical areas (MSAs) developed much after the automobile, and this has led to considerable spatial growth (sometimes referred to as sprawl). Constraints such as inland mountain ranges have led to increases in density in parts of these MSAs. The major cities in California now have enough density in many areas to support good transit coverage, but in most cases not enough for high frequency service. Many other North American cities face a similar problem.

To illustrate this qualitative conclusion, we consider the whole region served by AC Transit, which provides bus service in Alameda and Contra Costa Counties. Of the 90 possible trips between the randomly selected points shown on the map in Figure 1, 38% of the trips did not require a transfer, 51% required one transfer, and 11% required two transfers. All of these points were served by BART or AC Transit routes, which shows that the East Bay can indeed support transit, albeit many of the routes run with large scheduled headways (20, 30, 40, or 60 minutes), as we anticipated. Headways like these[1] mean that there is a huge penalty for missing a transfer, which is likely to discourage people from riding the bus if missed transfers are a regular occurrence. With 62% of the trips in this example requiring at least one transfer, it is clear that transfers are critically important for transit accessibility in California. Transfer waiting time could be greatly reduced with real-time coordination.

The literature on travel behavior, summarized in the following section, has established that transit users do not value all components of their trip equally. This finding sets the stage for our work by suggesting that a control strategy which exchanges out of vehicle time for in vehicle time is beneficial for the system.

Figure 1: Location of random OD points
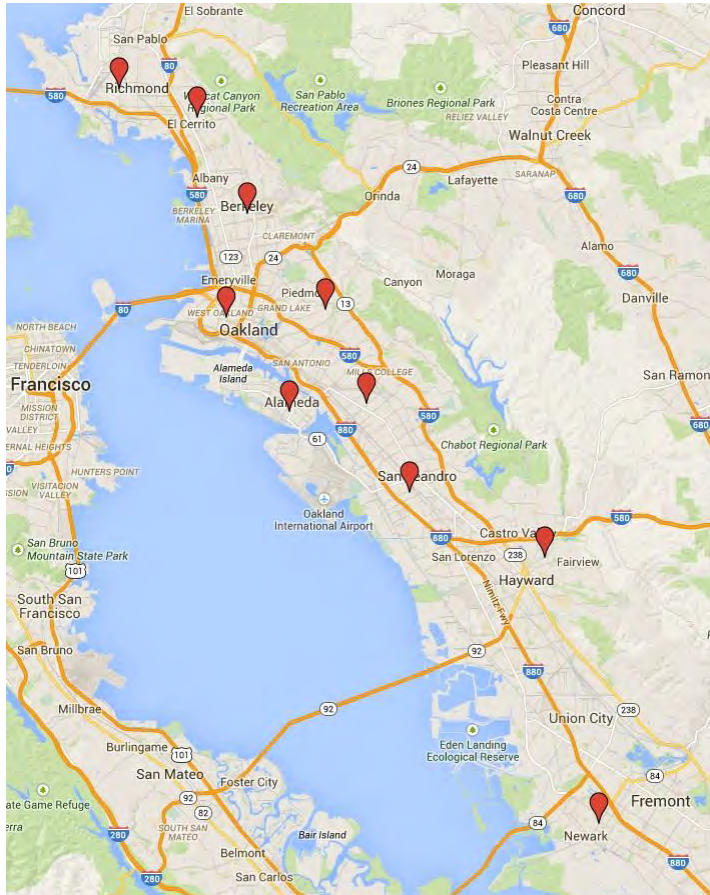


Unfortunately there are real world phenomena that stand in the way of real-time coordination as we propose to investigate. The main phenomenon is generically known as "bus bunching". This is the tendency of buses to form clusters as they move along their routes, especially where they travel in traffic. Bus bunching is a common problem on public transportation systems worldwide. The trigger is usually traffic, which continually knocks buses off schedule. Once a bus is sufficiently delayed so there is a large gap between it and the bus in front, a positive feedback loop that makes matters worse kicks in. The bus encounters more passengers, which delay it further and compound the problem. Conversely, if the bus runs ahead of schedule so its gap is small, it encounters fewer passengers and tends to catch up with the bus in front. This is why bunches form.

The literature on control measures for bus bunching is described in detail in the following section. There are three types of approaches: optimization strategies that use real-time information and "rolling horizon" heuristics, informal strategies that are reactive and use measures like boarding limits and stop skipping, and control theory which uses simple preventative principles. To our knowledge, no works exist at present which have extended this third approach to a network of routes including transfers, or even to two connecting routes, although the issue was identified as a direction for future research in Daganzo and Pilachowski (2011). Therefore, this is the main task of the proposed research. We believe that the analysis methods used in the control approach can and should be extended to account for transfers and improve their reliability.

2

Transfer coordination is a logical extension of control measures for bus bunching; transfer coordination is concerned with interactions between buses on connecting routes while bus bunching control strategies focus on interactions between sequential buses on the same route. Having a control strategy in place for bus bunching is not explicitly required for transfer coordination, but would keep buses closer to their scheduled headways, which makes arrivals more predictable.

This work looks at ways of controlling buses at transfer stations, and at traffic signals while en route, in order to improve the reliability of passenger connections between transit routes. A literature review is given (section 2) which describes relevant work in travel behavior and bus control strategies as well as other papers that have treated transfers. We first developed a mathematical framework for optimal holding control at transfer stations (section 3). We verified the analytical findings through simulation (section 4), and then conducted experiments to evaluate the performance of the proposed control strategy in real world situations (section 5). We conclude with a discussion of complimentary measures that could shorten the required holding time or reduce the downstream effects and therefore allow coordination in more cases (section 6). The overarching goal of these control measures is to improve the reliability of transfers and reduce out-of-vehicle passenger delay, the type of delay that is most annoying.

## 2   Literature Review

A first set of works illustrate the importance of efficient transfers for the traveling public. There is an extensive literature on travel behavior, including the way that users perceive the time spent on various components of their trip (Hickman and Wilson, 1995; Dziekan and Kottenho, 2007; Watkins et al., 2011), and transfer time is always valued very highly. Two findings from the literature support this conclusion. First, users place a higher value on out-of-vehicle waiting time, such as when transferring, than in-vehicle time (Ben-Akiva and Lerman, 1985; Dube et al., 1991). Second, travel time reliability is consistently found to be amongst the most important components of transit utility, even more important than the total travel time (de Palma and Picard, 2005; Bhat and Sardesai, 2006; Perk et al., 2008). Since transfers are an important source of travel time variability, and may be the dominant source of variability in bus systems with long headways, coordinating transfers in real time has the potential for improving matters considerably and increasing ridership. Various control strategies have been developed to mitigate bus bunching. Newell and Potts (1964) was the first work to study the bus bunching phenomenon and showed the inherent instability of bus systems. Osuna and Newell (1972) proposed to build slack time into the schedule, which would allow buses to get back on schedule at control points. This strategy and others of the same vintage do not involve any system coordination; drivers merely check the time against a schedule at control points. Strategies of this type are commonly used today but are largely in effective because the slack time required to prevent bunching is impractically large, and therefore the slack time provided is much smaller.

More recent papers have examined bus control strategies that incorporate coordination, using real-time information and "rolling horizon" heuristic optimization methods; see e.g., Eberlein et al. (2001). Other papers have proposed to deal with bunching less formally, using passenger boarding limits (Delgado et al., 2009) and skipping stops (Sun and Hickman, 2005; Liu et al., 2013) as control mechanisms to react to bunching. Unfortunately, the optimization strategies increase in complexity with the size of the system and cannot be scaled easily; and the informal approaches only act after bunching has taken hold. A third type of approach that overcomes these two drawbacks is based on control theory (Daganzo, 2009; Daganzo and Pilachowski, 2011; Xuan et al., 2011). The approach uses simple preventive principles rather than detailed optimization. As such it is scalable and averts bunching before it happens.

There are relatively few papers about transfers. Several works have proposed ways of designing bus schedules to reduce expected transfer waiting times (Hall, 1985; Abkowitz et al., 1987; Lee and Schonfeld, 1991; Bookbinder and Desilets, 1992; Knoppers and Muller, 1995; Hall et al., 2001; Hadas and Ceder, 2010). In addition to schedule design, Hadas and Ceder (2010) propose using informal strategies to help buses reach transfer points on time. Dessouky et al. (2003) develop holding strategies to preserve timed transfers. Lo and Chang (2012) propose a real time fuzzy bus holding system. Delgado et al. (2013) consider transfer coordination on a heavy rail system using a rolling horizon approach. An innovation of this paper is that it considers a trapezoidal distribution of passenger walking time, so it is possible for some to make the transfer while others miss it. Most recently, Nesheli and Ceder (2014, 2015) propose a combination of holding and segment skipping (i.e. skipping several stops in a row) to effectuate transfer coordination.

## 3 Development of Control Strategy

The control strategy that we propose for transfer coordination is distributed. This means that there is no centralized controller. Instead, each bus becomes a decision maker (subscript d in the notation) when it arrives at a transfer point. The basic procedure is that the decision-making bus arrives at the transfer point, consults the real-time information, calculates its maximum hold time, and then looks to see if there is a connecting bus (subscript c) arriving at or before that maximum hold time. If there is a connecting trip, the decision maker holds until the transferring passengers from that connecting bus arrive (even if it is late), and then departs. If there is no connecting bus expected before the maximum hold time, the decision maker departs as soon as it is done boarding and alighting passengers and any transferring passengers must wait for the next bus on the decision maker's route.

The information needed to implement this control strategy consists of estimates of upcoming bus arrivals and counts of transferring and affected passengers. The decision making bus may already have some of the real-time information it needs, like the number of passengers on board and the backward headway (part of some bus bunching control strategies). The arrival times of connecting buses can be requested when the decision making bus is approaching the transfer point. Other parameters, like the standard deviations, transferring passengers, and recovery, are difficult to estimate in real time so historical values can be used and stored locally.

The notation is as follows:

### 3.1 Parameters

**Arrival time** $A_i$ [s]: arrival time of the next bus on route $i$ at the transfer point.

**Headway** $H_i$ [s]: real time headway on route $i$

**Standard Deviation** $\sigma_x$ [s]: standard deviation of variable $x$

**Affected Passengers** $P_a$ [persons]: passengers on route $i$ who would be affected by a hold

**Transferring Passengers** $P_t$ [persons]: passengers transferring from route $i$ to $j$ at the transfer point

**Recovery** $\rho$ []: reflects the portion of holding time experienced by the average affected passenger. $\rho = 1$ means no holding time is recovered en route, $\rho < 1$ means some of the holding time is recovered before affected passengers get off the bus

### 3.1.1 Decision Variable

**Maximum Hold Time** $a_{max}$ [s]: maximum time that a bus should be willing to hold to receive transferring passengers

*For all variables, the notation $A_i$ refers to the true value, while the notation $\hat{A}_i$ refers to the estimate available at the time of decision.*

### 3.2 Deterministic Form

The simplest case to describe is the deterministic one, where all real time estimates are equal to the true value. Each instance of this control strategy begins when the decision-making bus arrives at the transfer point and consults the real-time information. As such, we call its arrival time $A_d = 0$. We mentioned above that the decision-making bus should calculate the maximum time that it is willing to hold and then compare the arrival times of connecting buses to this value. This simple rule forms the basis of our control strategy. We call the cost of our control strategy the net transfer delay, which can be expressed in terms of the maximum holding time:

$$E(C|A_c, \hat{A}_c, P_a, P_t, a_{max}, H_d) = \begin{cases} \rho P_a A_c & \hat{A}_c \le a_{max} \\ P_t(H_d - A_c) & \hat{A}_c > a_{max} \end{cases} \quad (1)$$

We are interested in finding the optimal control strategy, which is the expression for maximum holding time that would minimize the expectation of net transfer delay. In the deterministic case, this is easy to find because the cost of coordinating, $\rho P_a A_c$, and the cost of not coordinating, $P_t(H_d - A_c)$, both scale linearly with $A_c$. We can set these two equal to find the point where coordinating and not coordinating give us the same cost:

$$a_{max} = \frac{P_t H_d}{\rho P_a + P_t} \quad (2)$$

This $a_{max}$ is the optimal control strategy. We can see that it is simply the ratio of transferring to affected passengers multiplied by the headway on the deciding bus route. The affected passengers are adjusted by the recovery factor because of our assumption that passengers only experience the schedule deviation when they get off the bus. Therefore, if most of the holding time can be made up before the on-board passengers get off, transfer coordination becomes very attractive.

### 3.3 Uncertainty

However, real bus systems experience perturbations from traffic signals, mixed traffic, variable dwell times, and so on. We know that the information available at the time of decision consists of several estimates. The optimal control strategy could be improved by taking the uncertainty in these estimates into account. We assume the following distributions for the parameters:

$A_c \leftarrow U(\hat{A}c, \hat{A}c + a_a{}^p\overline{12})$

$A_c \leftarrow U(0, H_d)$

$\hat{P}_a = E(P_a)$ (distribution and variance do not matter if the estimate is unbiased)

$\hat{P}_t = E(P_t)$ (distribution and variance do not matter if the estimate is unbiased)

$H_d \sim U(\hat{H}d, \hat{H}d + a_H\sigma 12)$

The distributions for bus arrivals are biased: buses may be later than the estimate, but not earlier. We think that this construction is reasonable because buses are unlikely to be significantly earlier than a near-term arrival estimate but could be late due to some unexpected delay. Early arrivals would also be beneficial for the control (lower than expected cost of the selected action), so it is more useful to test potentially harmful disturbances.

The entire derivation of cost with uncertainty, incorporating uncertainty one parameter at a time, is contained in the appendix. The resulting cost with uncertainty in all parameters is as follows:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{1}{4\sqrt{3}\hat{H}_d + 12\sigma_H}\Big(2\sqrt{3}a_{max}^2(\rho\hat{P}_a + \hat{P}_t) - a_{max}(4\sqrt{3}\hat{H}_d\hat{P}_t$$
$$- 3\rho\hat{P}_a\sigma_a - 12\hat{P}_t\sigma_a + 12\hat{P}_t\sigma_H) + 2\hat{P}_t(\sqrt{3}\hat{H}_d^2 + 6\hat{H}_d\sigma_H \quad (3)$$
$$+ 4\sqrt{3}\sigma_H^2)\Big)$$

Now that the true parameter values have been replaced with real-time estimates, we can solve for the value of $a_{max}$ that minimizes expected cost. This value will represent the optimal control action, based on the information available. We take a derivative with respect to $a_{max}$:

$$\frac{d}{da_{max}}E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{2\sqrt{3}a_{max}(\rho\hat{P}_a + \hat{P}_t)}{2(\sqrt{3}\hat{H}_d + 3\sigma_H)} - \frac{2\sqrt{3}(\hat{H}_d\hat{P}_t + \sqrt{3}\hat{P}_t\sigma_H - \sqrt{3}\rho\hat{P}_a\sigma_a - \sqrt{3}\hat{P}_t\sigma_a)}{2(\sqrt{3}\hat{H}_d + 3\sigma_H)}$$

By setting the derivative equal to 0 and solving for $a_{max}$, we obtain an expression for the maximum holding time in terms of real-time information:

$$a_{max} = \frac{\hat{P}_t(\hat{H}_d + \sqrt{3}\sigma_H) - (\rho\hat{P}_a + \hat{P}_t)\sqrt{3}\sigma_a}{\rho\hat{P}_a + \hat{P}_t} \quad (4)$$

### 3.4 Key Assumptions

We make two key assumptions to reduce the complexity of the analytical expressions. Their validity and impact on real-world applications will be discussed later. The first is a constraint placed on the uncertainty in the connecting bus' arrival time:

$$\sigma_a\sqrt{12} \le H_d - a_{max} \quad (5)$$

This assumption means that a bus whose expected arrival time is within the maximum holding time of the decision-making bus will actually arrive within one headway.
The second assumption is:

$$\hat{A}_c \sim U(0, \hat{H}_d + \frac{\sqrt{12}}{2}\sigma_H) \quad (6)$$

In effect, we replace the true value of $H_d$ in the upper bound of the distribution with its expected value, $E(H_d)$, to avoid complicated integration.

# 4 Verification and Evaluation through Simulation

We use simulation for two purposes. The first is to confirm the analytical findings, and the second is to explore some reasonable combinations of parameter values to see what kind of results can be expected from the proposed control scheme.

In the first case, we hold all of the parameters constant except for $\hat{A}_c$ and $A_c$, which are drawn from the specified distributions. Simulations of two different sets of parameter values are plotted in Figure 3. Each point in the scatter represents one simulation run. One simulation run consists of a bus arriving at the transfer point, analyzing the situation represented by the constant parameters and a randomly drawn $\hat{A}_c$, making its control decision, and calculating the cost of that decision based on the true value $A_c$. The optimal $a_{max}$ from the analytical model is plotted with a dashed line. We can see that, in both plots, the expected costs of holding for a coordinated transfer (the ascending green points) and not holding (the red scatter) are equal at the analytical $a_{max}$. This is what we expect, as $a_{max}$ should represent an equilibrium point between the two control actions. To the left of $a_{max}$, coordination is the better control action as a short hold can eliminate a long wait for transferring passengers. To the right of $a_{max}$, the best action is not to wait because the transferring passengers do not have as long to wait for the next bus, and the cost of holding is high. The slope of the two lines and their intersection depends on the numbers of transferring and through passengers.

The second purpose of simulation is to explore how the control scheme performs under different parameter combinations. Figure 2 includes three dimensionless plots of the cost in passenger delay relative to the no control case vs. $x$, the ratio of transferring to through passengers. The difference between the three plots is the uncertainty parameters $a_a$ and $a_H$, which are given in the captions.

The interesting result in these figures is that our proposed control scheme is always better than or equal to no control, with or without uncertainty in the bus arrival estimates, and for any value of $x$. The percent reduction in passenger travel time with control varies widely, but very large reductions on the order of 60-70% are possible if transferring passengers significantly outnumber through passengers and the uncertainty in bus arrival estimates is not too large. We see that there is a high cost of uncertainty, particularly uncertainty in $A_c$, the connecting bus' arrival. However, uncertainty does not undermine the control strategy as it remains superior to no control. What is lost are the potential benefits that could be realized with perfect information.

Figure 4 compares deterministic and uncertainty control under more parameter combinations. Each point in these scatter plots is the average of 10,000 simulations. One parameter varies per plot, the other parameters are held constant with the values given in the plot title. Like the other simulations, $\hat{A}_c$ and $A_c$, are drawn from the specified distributions. As we can see, deterministic and uncertainty control have virtually the same net benefit per transferring passenger for most parameter combinations. The exceptions are when the ratio of transferring to affected passengers is small ($P_t/P_a < 0.1$) and when uncertainty in bus arrivals is high ($a_a > 150$s, $a_H > 250$s).

The intuition for the first case is that when there are relatively few transfers, the cost of holding increases much faster than the cost of leaving the transferring passengers to wait for the next bus, so a late arrival of the connecting bus is costly. For the second case, obviously when uncertainty is high it pays to take it into account.

(a) Simulation with $a_a$ = 0.10 and $a_H$ = 0.10



(b) Simulation with $a_a$ = 0.10 and $a_H$ = 0.15



(c) Simulation with $a_a$ = 0.15 and $a_H$ = 0.05

8

Figure 2: Simulations over a range of parameter values



(a) Passenger delay vs. arrival time of connecting bus with low uncertainty



(b) Passenger delay vs. arrival time of connecting bus with high uncertainty

Figure 3: Confirmation of Analytical Result

(a) Simulations with different connecting route headway

(b) Simulations varying the ratio of transferring to affected passengers

(c) Simulations with different values of $a_a$

(d) Simulations with different values of $a_H$

Figure 4: Comparison of deterministic and uncertainty control

# 5   Verification and Evaluation with Real Data: A Case Study

In order to test the proposed control scheme in a real-world scenario, we collected data at a subway- to-bus transfer point in Oakland, California. The Rockridge BART station is elevated and the main exit leads directly to the westbound bus stop. This bus stop is the starting point for AC Transit's 51B route, which heads west to the UC Berkeley campus and west Berkeley. We used two observers, one at the exit to the BART station and one at the bus stop. The observer at the exit started shooting video when a train arrived (which could be observed because the tracks are elevated) and stopped when all the passengers had exited the station. The observer at the bus stop counted transferring passengers and shot video of the bus boarding process.  At the same time, we ran    a script to request information on upcoming trips from the BART and AC Transit APIs once a minute. This experiment obtained the time that every BART passenger exited the station, the numbers of transferring and total passengers boarding each bus, and forecasted arrival times for each BART and bus trip over time.

   These data were used to associate BART passengers to trains and determine the walking time from train platform to bus stop. The station exit and bus boarding videos were compared to confirm the number of transferring passengers and determine their arrival time at the bus stop. A histogram of train platform to bus stop walking times is shown in Figure 5.  The mean walking time is 00:01:33 (HH:MM:SS), with a standard deviation of 00:00:43.  In 160 observations, 82% fall between 00:00:30-00:02:00, and 95% of the observations are contained between 00:00:30 and 00:02:30.

   The API data was used to evaluate the quality of forecasts and to estimate the parameters $a_a$ and $a_H$.  The standard deviation of bus arrival forecasts depends on the prediction horizon.  If we use all the data to calculate standard deviation (with respect to the actual arrival time), we obtain $a_H$ = 1.47 min.  However, the AC Transit API includes forecasts up to 90 minutes in advance of arrival (see Figure 6).  As we can see in the figure, the long-range forecasts are probably not real time, and are just quoting the schedule.   We can see that two estimates are revised backwards at around 30 minutes before arrival, so perhaps this is the limit of real-time forecasting. Using all forecasts within 33 minutes of arrival, we obtain a standard deviation $a_H$ = 1.10 min. The route's scheduled headway is 10 minutes, and if we consider only forecasts within 10 minutes of arrival, the standard deviation drops to $a_H$ = 0.519 min.  The BART forecasts do not show much deviation from schedule, probably because rail systems have fewer disturbances than buses. The BART API does not give forecasts as far in advance, as the most long-range estimate for all but one trip is 15-20 minutes before arrival. We simply use all the data to calculate standard deviation and obtain $a_a$ = 0.50 min. Recalling Equation 5, we need to check that this value of $a_a$ satisfies our assumption. Using the smallest observed headway, $H_d$ = 7 min, we need $a_a$ < 1.35 min. The BART forecasts easily satisfy the condition. The condition does not apply to $a_H$, but we can see that the 10 min and 33 min horizon values for $a_H$ satisfy the condition. This finding means that the upper bound on $a_a$ can be satisfied with existing bus location technology, which is necessary to apply the control scheme to bus-to-bus transfers.

   We then used the data collected in this experiment to virtually apply our control scheme. The uncertainty parameters are those obtained from the API data. We use the 33 minute horizon for the bus route.  The estimated arrival times, $\hat{A}_c$ and $\hat{H}_d$, are directly observed in  the API data. Because this stop is the beginning of the bus route, there are no through passengers on board. Instead, we counted the passengers who were waiting at the stop when the bus arrived as the affected passengers, $P_a$.  The remaining unknown is the number of transferring passengers, $\hat{P}_t$. In the data, each BART train contains between 0-4 transferring passengers, with 9 out of 12 trains (and 7 out of 8 westbound trains) containing at least one transferring passenger.

Figure 5: Histogram of walking times

(a) Time-space diagram of 51B route from API requests



(b) Time-space diagram of Rockridge BART from API requests

Figure 6: API data

| bus | through | transfers | amax | passenger | time | train | time | action |
|---|---|---|---|---|---|---|---|---|
| 08:14:56 | 14 | 1 | 0.25 | 08:16:38 | 1.70 | Daly City | 08:15:33 | depart |
| | | 2 | 1.11 | 08:16:45 | 1.82 | Daly City | 08:15:33 | depart |
| | | 3 | 1.81 | 08:21:28 | 6.53 | Pittsburg/Bay Point | 08:20:21 | depart |
| 08:21:55 | 10 | 1 | 1.28 | 08:22:04 | 0.15 | Pittsburg/Bay Point | 08:20:21 | wait |
| | | 2 | 2.82 | 08:22:53 | 0.97 | SF Airport | 08:21:10 | wait |
| | | 3 | 3.97 | 08:23:07 | 1.20 | SF Airport | 08:21:10 | wait |
| | | 4 | 4.87 | 08:23:22 | 1.45 | SF Airport | 08:21:10 | wait |
| | | 5 | 5.59 | 08:30:59 | 9.07 | Daly City | 08:29:41 | depart |
| | | 6 | 6.17 | 08:31:01 | 9.10 | Daly City | 08:29:41 | depart |
| | | 7 | 6.66 | 08:31:01 | 9.10 | Daly City | 08:29:41 | depart |
| | | 8 | 7.08 | 08:31:23 | 9.47 | Daly City | 08:29:41 | |
| 08:33:09 | 12 | 1 | 0.98 | 08:36:46 | 3.62 | SF Airport | 08:35:23 | depart |
| | | 2 | 2.36 | 08:36:48 | 3.65 | SF Airport | 08:35:23 | depart |
| | | 3 | 3.44 | 08:43:45 | 10.6 | Daly City | 08:42:22 | depart |
| 08:45:02 | 5 | 1 | 2.25 | 08:51:12 | 6.17 | SF Airport | 08:50:13 | depart |
| 08:55:07 | 7 | | | | | | | |

Table 1 shows the data on bus trips and transferring passengers. The left two columns show the actual bus departure time and $P_a$, the passengers waiting at the stop when the bus arrived. This $P_a$ figure includes transferring passengers who did not get a coordinated transfer and had to wait for the next trip. The number of transfers is unknown, so we present in the fourth column the $a_{max}$ (in minutes) calculated from the other real-time estimates for that trip and the $\hat{P}_t$ value in the third column. The fifth through eight columns show the passenger's arrival time, the relative time (in minutes) after the bus available for coordination, the BART train the passenger came from, and that train's arrival time. The rightmost column shows the correct control action if we knew the number of transferring passengers in advance.

We can see that the first bus trip is close to a Daly City train, but would not be able to coordinate if fewer than 3 passengers are expected because it already has 14 people on board and has a short headway behind it. The second trip departed in real life while Pittsburg/Bay Point passengers were still exiting the station and after the SF Airport train had already arrived. It has fewer passengers on board, so both the Pittsburg/Bay Point and SF Airport trains look like good possibilities for coordination if at least 1 transfer is expected. The third trip has a long headway behind it, but many onboard passengers.   An average passenger from the Daly City train would not be expected for nearly 4 minutes, so it is not a likely target for coordination unless at least 4 transfers are expected. The fourth bus trip has relatively few onboard passengers, but the SF Airport train is too far away to coordinate. We can see that if we use the average value for $\hat{P}_t$, which is 2 transferring passengers per BART train, it would produce the correct control actions. That might not be true in general because $a_{max}$ is very sensitive to $\hat{P}_t$. However, we might be able to come up with estimates on the number of transfers to expect from individual trains if we collected data over a longer period of time.

Table 2:  Out of vehicle passenger delay (min)

| | no control | | control (0% recovery) | | control (50% recovery) | |
|---|---|---|---|---|---|---|
| | through | transfer | through | transfer | through | transfer |
| Trip  1 | 0.0 | 10.9 | 0.0 | 10.9 | 0.0 | 10.9 |
| Trip  2 | 0.0 | 49.4 | 14.5 | 10.2 | 7.2 | 10.2 |
| Trip  3 | 0.0 | 17.8 | 0.0 | 17.8 | 0.0 | 17.8 |
| Trip  4 | 0.0 | 3.9 | 0.0 | 3.9 | 0.0 | 3.9 |
| Total | | 82.0 | | 57.3 | | 50.1 |
| Savings | | | | 30% | | 39% |

## 6   Complementary Measures

We observed in the preceding results that the ratio of transferring to affected passengers is the single most important variable for determining whether transfer coordination can be applied in a particular instance. Therefore, we expect that transfer coordination could be enhanced by complementary measures that increase this ratio, either by improving our knowledge of incoming transfers, or by reducing the number of affected passengers.

One idea for a complementary measure is a feature in smartphone apps that allows users to announce their travel plans to the transit agency. In current practice, the flow of real time information is unidirectional: transit agencies make their estimates of bus arrivals available, and users can modify their route or departure time accordingly. If users could communicate their travel plans to the transit agency (perhaps by "accepting" a set of directions calculated by an app), this knowledge could be factored into the estimate of transferring passengers at a particular stop. The impact that this would have on transfer coordination depends on the transit system and the day-to- day variability that it experiences.  By default, we estimate the number of transferring passengers   at a particular stop using historical data. Knowing users' travel plans would set a floor for the number of transfers to expect and would make coordination more likely when this floor is greater than the long-term average.

A second complementary measure is conditional signal priority (CSP). Transit signal priority (TSP) allows a bus to send a priority request when it is approaching a traffic signal. When a signal receives a priority request, it adjust the phase timing subject to feasibility constraints, using a green extension or an early green to pass the bus through the intersection with zero (or at least reduced) delay. In current practice, buses typically request priority at every equipped signal. The idea of conditional signal priority is to establish a virtual schedule and only allow buses which are behind schedule to request priority. Conventional TSP slightly improves regularity by reducing signal delay; the main benefit is an increase in commercial speed. CSP also improves commercial speed, but not as much since buses do not request priority every time.  The main benefit of CSP is that it dramatically improves regularity and appears to bound deviations from schedule: early buses are left to experience the full signal delay and soon return to schedule, while late buses can use priority at every signal until they have caught up.

CSP would aid transfer coordination by reducing the number of affected passengers and by reducing the uncertainty on the connecting bus arrival and the real time headway. After holding for transferring passengers, the deciding bus will likely be behind its virtual schedule. CSP would allow it to request priority at the next few intersections until it has caught up to schedule. Since our assumption is that passengers do not care about schedule delay that is made up before they disembark, CSP would reduce the number of affected passengers by helping the bus return to schedule more quickly. If both the deciding and connecting routes adopt CSP, we would expect the uncertainties in the bus arrival estimates to decrease, which also decreases the cost of transfer coordination.

So far we have only mentioned the synergies of using transfer coordination and conditional signal priority on the transit system. The two control strategies could be more explicitly linked by adjusting virtual schedules to benefit transfer coordination. For example, a connecting bus could be allowed to use priority until it reaches the transfer point to reduce the time that the deciding bus has to hold.

## 7  Conclusion

In this work, we developed a dynamic holding strategy for coordinating transit transfers in real time that is based on control theory. The strategy is simple and scalable: when a bus arrives at a transfer point, it looks at the real-time information and immediately makes a decision on whether to depart   or to hold for a connecting route. We started with an expression for expected cost of the control strategy given full knowledge of the system. True parameter values are replaced with real-time estimates one by one to obtain an expression for expected cost that depends only on parameters known at the time of decision. We then obtained an expression for the maximum holding time that minimizes expected cost.  A bus using the control strategy calculates its maximum holding time and then checks to see if any connecting vehicles are expected within this time horizon.

We used simulation to confirm the analytical results and to explore the expected performance of the control strategy over a range of parameter values. Two simulations with all fixed parameters except for random arrivals of the connecting bus confirm that the optimal $a_{max}$ obtained in the analysis is indeed the equilibrium point where the expected costs of holding and not holding intersect. Dimensionless plots of relative cost vs. the ratio of transferring to through passengers show that the proposed control strategy is universally better than or equivalent to no control. Very large reductions in passenger delay are possible when transfer volumes are high relative to through passengers. We showed that uncertainty increases costs significantly, but that the control strategy remains superior to no control.

We then conducted field observations at a multimodal transfer point in Oakland. All of the parameters for the control strategy were either measured directly or estimated from the data. The observations showed that real time estimates of bus arrivals are accurate enough to satisfy one of the key assumptions in the analytical model, an upper bound on uncertainty in bus arrivals. This finding suggests that our control scheme can be used with existing bus location technology. We showed that if our control strategy were used in the four real-world scenarios contained in the data, we would hold one bus for a coordinated transfer. This control action would reduce net transfer delay by 30%, even assuming that the bus is not able to make up the schedule delay incurred. If, through complementary measures, the bus were able to make up the schedule delay so that the average through passenger experiences only half of the hold time, the reduction in net transfer delay increases to 39%.

A promising complementary measure is conditional signal priority. Conditional priority can help to speed up late buses by allowing them to use signal priority while denying priority to buses that are ahead of schedule. Priority would be granted to all late buses, not just those recovering from holding for transfer coordination, but it could greatly reduce the downstream effects of transfer coordination. If small schedule delays could be made up in a few blocks, transfer coordination would be less costly and could be used in more cases.

We acknowledge that the analytical model does not fully account for downstream effects and the walking time of transferring passengers. The model currently compares the schedule delay experienced by through passengers at the time they alight with the out of vehicle waiting time transferring passengers experience if they are left to wait for the next bus.  However, holding for a transfer would increase out of vehicle time for passengers waiting at downstream stops.  It is   very difficult to measure passengers who are not yet in the transit system, but we could estimate a demand rate for downstream stops. Accounting for these passengers would further underline the importance of using complementary measures to make up schedule delay.   The analytical model is best suited for a case where all of the transferring passengers from a connecting vehicle arrive simultaneously.  In our real world data, the train platform and bus stop are not adjacent, and there is some walking time in between. We collected data on walking time and showed that there is a relatively tight distribution, but have not fully accounted for the uncertainty that variable walking time introduces. There is also potential for complementary measures, like directional signage and screens with real-time information that could encourage transferring passengers to walk faster. Future work will extend the model to consider dispersed passenger arrivals and to fully account for downstream effects of holding for transfers.

# A. Appendix: Derivation of Optimal Control

We start by writing an expression for the expected value of the cost of the control strategy, assuming the true values of all parameters are known. This expression is:

$$E(C|A_c, \hat{A}_c, P_a, P_t, a_{max}, H_d) =$$

$$\begin{cases} \rho P_a A_c & \hat{A}_c \leq a_{max} \\ P_t(H_d - A_c) & \hat{A}_c > a_{max} \end{cases}$$

We point out that the two conditional expressions are linear functions of the parameters $P_a$ and $P_t$, respectively. As a result, we can simply replace the true values of these parameters with their estimates, provided the estimates are unbiased.

Now we aim to write expected cost without knowing $A_c$. This can be accomplished by integrating out $A_c$ in both conditional expressions. The interval $(a_max, H_d]$ is subdivided into cases where the entire distribution of $A_c$ falls before $H_d$ and cases where $A_c$ can fall before or after $H_d$. The integration is as follows:

$$E(C|\hat{A}_c, \sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, H_d) =$$

$$\begin{cases} \int_0^{a_{max}} \rho \hat{P}_a A_c (\frac{1}{\sigma_a \sqrt{12}}) & \hat{A}_c \leq a_{max} \\ \int_{a_{max}}^{H_d - \sigma_a \sqrt{12}} \hat{P}_t(H_d - A_c)(\frac{1}{\sigma_a \sqrt{12}}) & a_{max} < \hat{A}_c < H_d - \sigma_a \sqrt{12} \\ \int_{H_d - \sigma_a \sqrt{12}}^{H_d + \sigma_a \sqrt{12}} \hat{P}_t(\frac{3}{2}H_d - A_c)(\frac{1}{\sigma_a \sqrt{12}}) & a_{max} < H_d - \sigma_a \sqrt{12} \leq \hat{A}_c \leq H_d \end{cases}$$

and results in the following expected cost:

$$E(C|\hat{A}_c, \gamma_a, \hat{P}_a, \hat{P}_t, a_{max}, H_d) =$$

$$\begin{cases} \rho \hat{A}_c \hat{P}_a + \sqrt{3}\rho \hat{P}_a \sigma_a & \hat{A}_c \leq a_{max} \\ -\hat{A}_c \hat{P}_t + H_d \hat{P}_t - \sqrt{3}\hat{P}_t \sigma_a & a_{max} < \hat{A}_c < H_d - \sigma_a \sqrt{12} \\ -\hat{A}_c \hat{P}_t + \frac{3}{2}H_d \hat{P}_t - \sqrt{3}\hat{P}_t \sigma_a & a_{max} < H_d - \sigma_a \sqrt{12} \leq \hat{A}_c \leq H_d \end{cases}$$

The next step is to eliminate $\hat{A}_c$. We have assumed that the two routes are completely uncoordinated, and that the connecting bus could arrive at any time within one headway. We now integrate out $\hat{A}_c$:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, H_d, \hat{H}_d) = \frac{1}{\hat{H}_d + \sigma_H \frac{\sqrt{12}}{2}} \left( \int_0^{a_{max}} (\rho \hat{A}_c \hat{P}_a + \sqrt{3} \rho \hat{P}_a \sigma_a) d\hat{A}_c + \int_{a_{max}}^{H_d - \sigma_a \sqrt{12}} (-\hat{A}_c \hat{P}_t + \right.$$

$$\left. H_d \hat{P}_t - \sqrt{3} \hat{P}_t \sigma_a) d\hat{A}_c + \int_{H_d - \sigma_a \sqrt{12}}^{H_d} (-\hat{A}_c \hat{P}_t + \frac{3}{2} H_d \hat{P}_t - \sqrt{3} \hat{P}_t \sigma_a) d\hat{A}_c \right)$$

and obtain the following expected cost:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, H_d, \hat{H}_d) = \frac{1}{\hat{H}_d + \sqrt{3}\sigma_H} \left( \frac{1}{2} \rho a_{max}^2 \hat{P}_a + \frac{1}{2} a_{max}^2 \hat{P}_t - a_{max} H_d \hat{P}_t + \frac{1}{2} H_d^2 \hat{P}_t + \sqrt{3} \rho a_{max} \hat{P}_a \sigma_a + \right.$$

$$\left. \sqrt{3} a_{max} \hat{P}_t \sigma_a \right)$$

Now we can eliminate $H_d$, leaving us with only parameters that will be known at the time of decision. $H_d$ is assumed to follow a one-sided uniform distribution, with all probability density falling at or after the estimate. Integrating over the probability density of $H_d$:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \int_{\hat{H}_d}^{\hat{H}_d + \sigma_H \sqrt{12}} \frac{1}{\hat{H}_d + \sqrt{3}\sigma_H} \left( \frac{1}{2} \rho a_{max}^2 \hat{P}_a + \frac{1}{2} a_{max}^2 \hat{P}_t - a_{max} H_d \hat{P}_t + \frac{1}{2} H_d^2 \hat{P}_t + \right.$$

$$\left. \sqrt{3} a_{max} \hat{P}_a \sigma_a + \sqrt{3} a_{max} \hat{P}_t \sigma_a \right) \left( \frac{1}{\sigma_H \sqrt{12}} \right) dH_d$$

we obtain the expected cost:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{1}{4\sqrt{3}\hat{H}_d + 12\sigma_H} \left( 2\sqrt{3} a_{max}^2 (\rho \hat{P}_a + \hat{P}_t) - a_{max}(4\sqrt{3}\hat{H}_d \hat{P}_t - 3\rho \hat{P}_a \sigma_a - \right.$$

$$\left. 12\hat{P}_t \sigma_a + 12\hat{P}_t \sigma_H) + 2\hat{P}_t (\sqrt{3}\hat{H}_d^2 + 6\hat{H}_d \sigma_H + 4\sqrt{3}\sigma_H^2) \right)$$

Now that the true parameter values have been replaced with real-time estimates, we can solve for the value of $a_{max}$ that minimizes expected cost. This value will represent the optimal control action, based on the information available. We take a derivative with respect to $a_{max}$:

$$\frac{d}{da_{max}} E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{2\sqrt{3} a_{max}(\rho \hat{P}_a + \hat{P}_t)}{2(\sqrt{3}\hat{H}_d + 3\sigma_H)} - \frac{2\sqrt{3}(\hat{H}_d \hat{P}_t + \sqrt{3}\hat{P}_t \sigma_H - \sqrt{3}\rho \hat{P}_a \sigma_a - \sqrt{3}\hat{P}_t \sigma_a)}{2(\sqrt{3}\hat{H}_d + 3\sigma_H)}$$

By setting the derivative equal to 0 and solving for $a_{max}$, we obtain an expression for the maximum holding time in terms of real-time information:

$$a_{max} = \frac{\hat{P}_t(\hat{H}_d + \sqrt{3}\sigma_H) - (\rho \hat{P}_a + \hat{P}_t)\sqrt{3}\sigma_a}{\rho \hat{P}_a + \hat{P}_t} \tag{7}$$

By substituting this expression for $a_{max}$ back into expected cost, we can obtain an expression for expected cost in terms of the real time estimates alone:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{1}{8(\rho \hat{P}_a + \hat{P}_t)(\sqrt{3}\hat{H}_d + 3\sigma_H)} \left( 4\sqrt{3}\hat{H}_d^2 \rho \hat{P}_a \hat{P}_t + 12\hat{H}_d \hat{P}_t(2\hat{P}_t \sigma_a + 2\rho \hat{P}_a(\sigma_a + \right.$$

$$\left. \sigma_H)) + \sqrt{3}(-12\rho^2 \hat{P}_a^2 \sigma_a^2 + 4\hat{P}_t^2(-3\sigma_a^2 + 6\sigma_a \sigma_H + \sigma_H^2) + 8\rho \hat{P}_a \hat{P}_t(-3\sigma_a^2 + 3\sigma_a \sigma_H + 2\sigma_H^2)) \right)$$

This expression can be rearranged:

$$E(C|\sigma_a, \hat{P}_a, \hat{P}_t, a_{max}, \hat{H}_d, \sigma_H) = \frac{\hat{H}_d \rho \hat{P}_a \hat{P}_t}{4(\rho \hat{P}_1 + \hat{P}_t)}$$

$$+ \frac{1}{8(\rho \hat{P}_a + \hat{P}_t)(\sqrt{3}\hat{H}_d + 3\sigma_H)} \left( \hat{H}_d \hat{P}_t(2\hat{P}_t \sigma_H \right.$$

$$+ 2\rho \hat{P}_a(\sigma_a + \frac{1}{2}\sigma_H)) + \sqrt{3}(-12\rho^2 \hat{P}_a^2 \sigma_a^2$$

$$+ 4\hat{P}_t^2(-3\sigma_a^2 + 6\sigma_a \sigma_H + \sigma_H^2)$$

$$\left. + 8\rho \hat{P}_a \hat{P}_t(-3\sigma_a^2 + 3\sigma_a \sigma_H + 2\sigma_H^2)) \right) \tag{8}$$

The fraction on the first line is the deterministic component while the rest represents the added cost due to uncertainty.

# References

Abkowitz, M., Josef, R., Tozzi, J., and Driscoll, M. K. (1987). Operational feasibility of timed transfer in transit systems. *Journal of transportation engineering*, 113(2):168–177.

Ben-Akiva, M. E. and Lerman, S. R. (1985). *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press.

Bhat, C. R. and Sardesai, R. (2006). The impact of stop-making and travel time reliability on commute mode choice. *Transportation Research Part B: Methodological*, 40(9):709–730.

Bookbinder, J. H. and Desilets, A. (1992). Transfer optimization in a transit network. *Transportation science*, 26(2):106–118.

Daganzo, C. F. (2009). A headway-based approach to eliminate bus bunching: Systematic analysis and comparisons. *Transportation Research Part B: Methodological*, 43(10):913–921.

Daganzo, C. F. and Pilachowski, J. (2011). Reducing bunching with bus-to-bus cooperation. *Transportation Research Part B: Methodological*, 45(1):267–277.

de Palma, A. and Picard, N. (2005). Route choice decision under travel time uncertainty. *Transportation Research Part A: Policy and Practice*, 39(4):295–324.

Delgado, F., Contreras, N., and Munoz, J. C. (2013). Holding for transfers. In *Transportation Research Board 92nd Annual Meeting*, number 13-2913.

Delgado, F., Mun˜oz, J. C., Giesen, R., and Cipriano, A. (2009). Real-time control of buses in a transit corridor based on vehicle holding and boarding limits. *Transportation Research Record: Journal of the Transportation Research Board*, 2090(1):59–67.

Dessouky, M., Hall, R., Zhang, L., and Singh, A. (2003). Real-time control of buses for schedule coordination at a terminal. *Transportation Research Part A: Policy and Practice*, 37(2):145–164.

Dube, L., Schmitt, B. H., and Leclerc, F. (1991). Consumers' affective response to delays at different phases of a service delivery1. *Journal of Applied Social Psychology*, 21(10):810–820.

Dziekan, K. and Kottenho↵, K. (2007). Dynamic at-stop real-time information displays for public transport: e↵ects on customers. *Transportation Research Part A: Policy and Practice*, 41(6):489– 501.

Eberlein, X. J., Wilson, N. H., and Bernstein, D. (2001). The holding problem with real–time information available. *Transportation science*, 35(1):1–18.

Hadas, Y. and Ceder, A. A. (2010). Optimal coordination of public-transit vehicles using operational tactics examined by simulation. *Transportation Research Part C: Emerging Technologies*, 18(6):879–895.

Hall, R., Dessouky, M., and Lu, Q. (2001). Optimal holding times at transfer stations. *Computers & industrial engineering*, 40(4):379–397.

Hall, R. W. (1985). Vehicle scheduling at a transportation terminal with random delay en route. *Transportation Science*, 19(3):308–320.

Hickman, M. D. and Wilson, N. H. (1995). Passenger travel time and path choice implications of real-time transit information. *Transportation Research Part C: Emerging Technologies*, 3(4):211– 226.

Knoppers, P. and Muller, T. (1995). Optimized transfer opportunities in public transport. *Transportation Science*, 29(1):101–105.

Lee, K. K. and Schonfeld, P. (1991). Optimal slack time for timed transfers at a transit terminal. *Journal of Advanced Transportation*, 25(3):281–308.

Liu, Z., Yan, Y., Qu, X., and Zhang, Y. (2013). Bus stop-skipping scheme with random travel time. *Transportation Research Part C: Emerging Technologies*, 35:46–56.

Lo, S - C and Chang, W.-J. (2012). Design of real-time fuzzy bus holding system for the mass rapid transit transfer system. *Expert Systems with Applications*, 39(2):1718–1724.

Nesheli, M. M. and Ceder, A. A. (2014). Optimal combinations of selected tactics for public-transport transfer synchronization. *Transportation Research Part C: Emerging Technologies*, 48:491–504.

Nesheli, M. M. and Ceder, A. A. (2015). A robust, tactic-based, real-time framework for public-transport transfer synchronization. *Transportation Research Part C: Emerging Technologies*, 60:105–123.

Newell, G. F. and Potts, R. B. (1964). Maintaining a bus schedule. In *Australian Road Research Board (ARRB) Conference, 2nd, 1964, Melbourne*, volume 2.

Osuna, E. and Newell, G. (1972). Control strategies for an idealized public transportation system. *Transportation Science*, 6(1):52–72.

Perk, V., Flynn, J., and Volinski, J. M. (2008). Transit ridership, reliability, and retention. Technical report, National Center for Transit Research and Research and Special Programs Administration and Florida Department of Transportation.

Sun, A. and Hickman, M. (2005). The real–time stop–skipping problem. *Journal of Intelligent Transportation Systems*, 9(2):91–109.

Watkins, K. E., Ferris, B., Borning, A., Rutherford, G. S., and Layton, D. (2011). Where is my bus? Impact of mobile real-time information on the perceived and actual wait time of transit riders. *Transportation Research Part A: Policy and Practice*, 45(8):839–848.

Xuan, Y., Argote, J., and Daganzo, C. F. (2011). Dynamic bus holding strategies for schedule reliability: Optimal linear control and performance analysis. *Transportation Research Part B: Methodological*, 45(10):1831–1845.