EVALUATION OF A FUZZY LOGIC RAMP METERING ALGORITHM:
A COMPARATIVE STUDY AMONG THREE RAMP METERING ALGORITHMS
USED IN THE GREATER SEATTLE AREA

by

Cynthia Taylor   Deirdre Meldrum
Research Engineer   Associate Professor

Department of Electrical Engineering
University of Washington
Seattle, Washington 98195

Washington State Transportation Center (TRAC)
University of Washington, Box 354802
University District Building
1107 NE 45th Street, Suite 535
Seattle, Washington 98105-4631

Washington State Department of Transportation
Technical Monitor
Dave McCormick
Traffic Systems Manager, Northwest Region

Prepared for

Washington State Transportation Commission
Department of Transportation
and in cooperation with
U.S. Department of Transportation
Federal Highway Administration

February 2000

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the Washington State Transportation Commission, Department of Transportation, or the Federal Highway Administration. This report does not constitute a standard, specification, or regulation

TECHNICAL REPORT STANDARD TITLE PAGE

1. REPORT NO. WA-RD 481.2	2. GOVERNMENT ACCESSION NO.	3. RECIPIENT'S CATALOG NO.
4. TITLE AND SUBTITLE Evaluation of a Fuzzy Logic Ramp Metering Algorithm: A Comparative Study Among Three Ramp Metering Algorithms Used in the Greater Seattle Area		5. REPORT DATE February 2000
		6. PERFORMING ORGANIZATION CODE
7. AUTHOR(S) Cynthia Taylor and Deirdre Meldrum		8. PERFORMING ORGANIZATION REPORT NO.
9. PERFORMING ORGANIZATION NAME AND ADDRESS Washington State Transportation Center (TRAC) University of Washington, Box 354802 University District Building; 1107 NE 45th Street, Suite 535 Seattle, Washington 98105-4631		10. WORK UNIT NO.
		11. CONTRACT OR GRANT NO. Agreement T9903, Task 84
12. SPONSORING AGENCY NAME AND ADDRESS Research Office Washington State Department of Transportation Transportation Building, MS 47370 Olympia, Washington 98504-7370 Dave McCormick, Project Manager, 206-440-4486		13. TYPE OF REPORT AND PERIOD COVERED Draft research report
		14. SPONSORING AGENCY CODE
15. SUPPLEMENTARY NOTES This study was conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.		

16. ABSTRACT

  A Fuzzy Logic Ramp Metering Algorithm was implemented on 126 ramps in the greater Seattle area. Two multiple-ramp study sites were evaluated by comparing the fuzzy logic controller (FLC) to the other two ramp metering algorithms in operation at those sites over a four-month period. At the first study site, the days when the FLC was metering had lower mainline occupancies and higher throughput volumes in comparison to the days when the Local Algorithm was metering. At the second study site, the days when the FLC was metering had mainline occupancies that were similar, queues that were shorter, and throughput that was similar to the days when the Bottleneck Algorithm was metering.

17. KEY WORDS Ramp metering, fuzzy logic control, intelligent transportation systems, freeway operations, transportation management software	18. DISTRIBUTION STATEMENT No restrictions. This document is available to the public through the National Technical Information Service, Springfield, VA 22616
19. SECURITY CLASSIF. (of this report) None	20. SECURITY CLASSIF. (of this page) None
21. NO. OF PAGES	22. PRICE

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

iv

INTRODUCTION

Our region's growth and transportation were ranked number 1 in the concerns of King, Snohomish, and Pierce county residents (Pryne, 1997). This is not surprising considering that the Seattle area tied for first with Los Angeles and San Francisco for the worst peak-hour traffic. Most residents blame increasing traffic congestion on the increasing population, but in actuality, worsening congestion is more a result of increasing mileage driven per motorist than of population growth. Since 1969, the number of vehicles nationally has increased 143 percent, which is much higher than the population increase of 23 percent during this same time period. Local motorists are no exception to the national trend. In King, Pierce, Snohomish, and Kitsap counties, the mileage driven *per motorist* doubled between 1980 and 1996 (Pryne, 1998). In Washington State, 35 million more miles are traveled every day than were traveled 10 years ago, with only 47 miles of new highway added during that time, not counting added lanes (Whitely, 1999).

Adding more roadway is not by itself a solution to solving traffic congestion. A state study concluded that average speeds on I-405 will drop to 26 mph within two decades regardless of what changes are made. Even with 12 superhighway lanes, the roadway will be four times as congested as today. Several studies support a theory called "induced traffic" (Peirce, 1999). The theory is based on complex mathematics, but the idea is very simple: "Build it and they will come." Researchers at the University of California found that for every 10 percent increase in lane-mileage created, there was a 9 percent increase in traffic. The Surface Transportation Policy Project of Washington D.C. found that among 70 cities, those that added extensive new road capacity had no different traffic congestion than those cities that did not.

Over the years, solutions to traffic congestion have shifted from the build-our-way-out mentality to methods that make better use of the existing infrastructure. The solution to traffic congestion must be a multi-faceted one. Mileage driven can be targeted through

managed growth, alternative transportation, carpool incentives, trip reduction, and congestion pricing. Freeway efficiency can be improved through incident response, roadway improvements, driver information, and ramp metering.

Freeway systems are chaotic systems, meaning that a tiny cause can have a huge effect. One driver tapping his or her brakes can cause a shock wave that travels backwards for kilometers. An accident that partially blocks a lane for 10 minutes might cause a backup that takes 45 minutes to clear. Rubbernecking drivers who slow down to gawk at accidents on the roadway in the opposite direction can cause several miles of backup. The more that freeway demand exceeds freeway capacity, the greater the impact of an event.

Just as a minor event can have a huge effect, ramp metering to prevent or delay critical flow breakdown can have a huge benefit, with a relatively inexpensive implementation cost. In 23 urban areas across the U.S., on-line studies where ramp metering was implemented have reported accident rate reductions of 24 to 50 percent, throughput increases of 17 to 25 percent, and mainline speed increases of 16 to 62 percent (Piotrowicz and Robinson, 1995). If the cost of traffic congestion is considered in terms of higher accident rates, lost productivity, and pollution, even moderate gains in freeway system efficiency are worth the cost of ramp metering improvements.

Most drivers are aware that ramp metering smoothes the merge onto the freeway. Ramp metering reduces mainline congestion by reducing the turbulence caused by merging platoons. However, many drivers are not aware that a system-wide ramp metering algorithm provides even more benefit by preventing downstream bottlenecks. The prevention of critical flow breakdown results in higher throughput and faster mainline travel times.

This research project involved the design, on-line implementation, and evaluation of a fuzzy logic ramp metering algorithm for the greater Seattle area. The preliminary stages of the Transportation Systems Management Center (TSMC) software documentation

and the software integration plan were carried out under a previous WSDOT/TransNow research grant. This report describes the research approach, evaluation method, and the results of on-line testing of the Fuzzy Logic Ramp Metering Algorithm in comparison to the Local Algorithm and Bottleneck Algorithm at two different study sites.

For details on the code, see the technical report "A Programmer's Guide to the Fuzzy Logic Ramp Metering Algorithm: Software Design, Integration, Testing, and Evaluation" (Taylor and Meldrum, 2000.) The programmer's guide also contains knowledge gained about the system, recommendations for future software projects, and the algorithm's transferability to other regions. For information regarding the algorithm design and tuning technique, see the technical report "Algorithm Design, User Interface, and Optimization Procedure for a Fuzzy Logic Ramp Metering Algorithm: A Training Manual for Freeway Operations Engineers" (Taylor and Meldrum, 2000). The training manual also contains background on fuzzy logic control, how this algorithm addresses various ramp metering problems, a description of design modifications, and many examples on how to handle special cases.

RESEARCH APPROACH

The research approach included on-line testing to tune the controller for optimal performance, to determine the system-wide parameter defaults, and to compare the behavior of the Fuzzy Logic Ramp Metering with that of the Local and Bottleneck ramp metering algorithms (Jacobson, Henry, and Mehyar, 1988). The scope of the project did not include comprehensive, system-wide testing but rather, preliminary study site testing to determine whether the Fuzzy Logic Ramp Metering Algorithm was beneficial relative to the other ramp metering algorithms. The goal of the on-line testing was to determine whether to proceed with system-wide implementation. Because the test results of this pilot project on the two study sites were promising, managers of the freeway operations group requested system-wide implementation as soon as possible, beginning with the most congested corridors.

Study Sites

The study sites were chosen with the following criteria:

- There was a set of adjacent metered ramps.
- Recurrent congestion was present to provide relatively uniform test conditions for the purpose of comparing different algorithms. Sites with morning metering were preferred because their traffic patterns were more consistent from day-to-day.
- Nonrecurrent congestion and special events were needed to test under a broad range of conditions.
- Adequate surveillance was in place.
- No new construction was planned for these sites for the duration of the study.
- The study site corridor was geographically isolated from other ramp metering algorithms, and congestion cleared upstream and downstream of the site.

The concept was to do preliminary testing with a light to moderately congested corridor, so that if the metering rates were not optimal the impact would be minimal. The purpose of the second study site was to test under heavy congestion and verify that the Fuzzy Logic Ramp Metering Algorithm was transferable from one location to the next. By testing under a variety of circumstances, we could determine whether dissimilar study sites can use similar parameters, and in turn, estimate how much tuning would be necessary for system-wide implementation.

The combination of the following two sites was ideal: 1) westbound I-90 between SR 900 and Eastgate, and 2) southbound I-405 between NE 160[th] St and NE 72[nd] St (Figure 1). With these two sites, we compared the Fuzzy Logic Algorithm to the Local Algorithm in operation on the moderately congested westbound I-90, and we compared the Fuzzy Logic Algorithm to the Bottleneck Algorithm in operation on the heavily congested southbound I-405. These sites met all of the specified criteria and allowed us to test under a broad range of conditions.
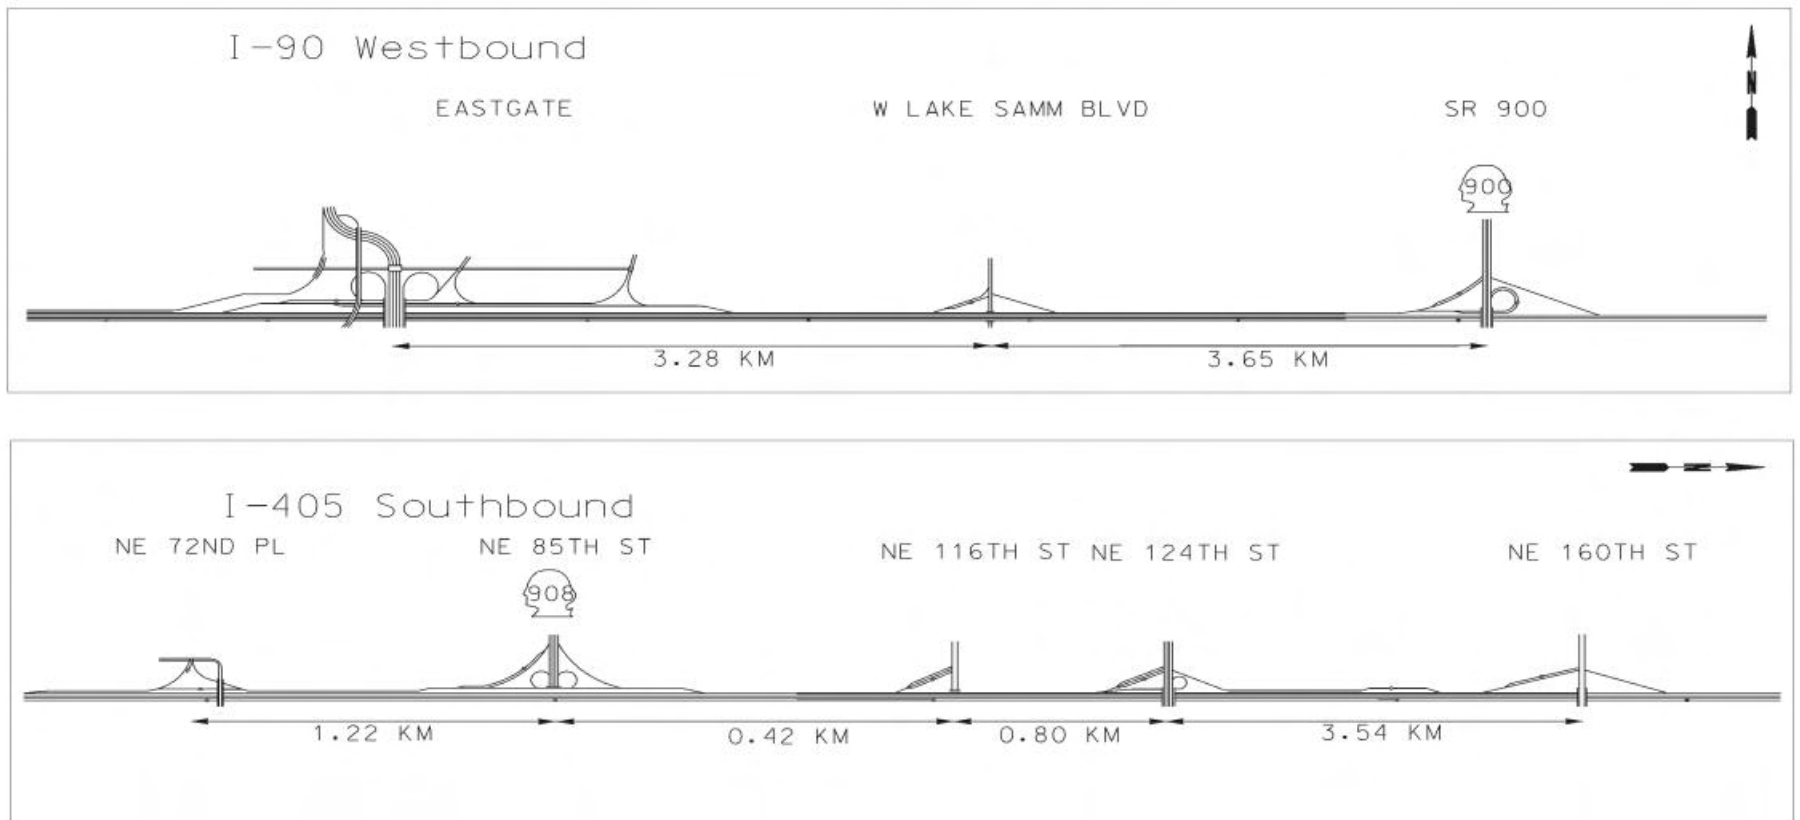
Figure 1: WB I-90 and SB I-405 Study Sites

Test Plan

The software testing was done in a way that would minimize impacts to drivers. Implementation and testing of the new algorithm did not cause any downtime for the TMSC. Testing of the new software was done off-line on a microVAX, which replicated the real system software. Regression testing was performed to ensure that all old functionality of the code still existed. Tests were performed to verify that the new code worked properly. Off-line testing was so thorough that on-line testing proceeded without any problems.

The on-line tests were structured in incremental steps toward progressively more realistic conditions in order to mitigate the risks of using the new software. The main risks of using the new software were bugs, non-ideal operation, and improper operator usage. Because of extensive off-line diagnostics, an incremental on-line test plan, the hierarchical software design, training provided to the operators, and the presence of the algorithm designer during all test periods, none of these problems occurred.

The first step of real-time, on-line testing was to test the fuzzy logic controller on a test 170 rack. The test 170 rack is identical to those in the field, except that the user can specify the loop data, and the generated metering rates have no impact on field operations. Through this test bed, the software quality was verified.

The next step of on-line testing was to generate metering rates from real-time field data, but to not actually implement the rates, so that there was no impact on operations. We verified proper algorithm behavior by comparing the Fuzzy Metering rates to those produced by the Local Metering Algorithm, given the same data from the I-90 study site. Of course, without actually implementing the Fuzzy Metering rates, the actual effect of the controller behavior could not be determined. However, this observation mode (a software compile option of Fuzzy Metering discussed in the programmer's guide, Taylor and Meldrum, 2000), in which the rates were calculated with real data but not implemented, proved very useful for preliminary tuning. Using the observation mode, the algorithm was

tuned to behave similarly to the Local Ramp Metering Algorithm upon initial deployment. The idea was to minimize impact to drivers and the risk of unknown metering behavior.

At this time, a baseline study was done to gain familiarity with the test sites. "Before" performance measures were calculated during observation mode. Intimate knowledge of the study site prepared the algorithm tester to understand the effects of tuning the fuzzy controller.

After preliminary tuning and verification of reasonable rates, the new algorithm was ready for actual field testing. Field testing began with the Eastgate ramp on the I-90 study site. After verification of proper controller behavior, the implementation expanded to two more ramps on the I-90 study site. Then the fuzzy control parameters were retuned for optimal control, diverging from the behavior of the Local Ramp Metering algorithm. (See the Training Manual for the tuning procedure to achieve the control objectives.)

With a noticeable improvement in mainline congestion while reasonable queue lengths were maintained, the algorithm was ready for implementation on the second study site, I-405. Again, the metering rates were first generated in observation mode before actual implementation, this time to see how the algorithm would behave under heavy congestion. The generated metering rates were compared to those of the Bottleneck Metering Algorithm in operation at that site, and the fuzzy parameters were further tuned. Upon verification of the algorithm's desired behavior, the fuzzy metering rates were deployed on the I-405 study site. The fuzzy parameters were fine-tuned for optimal control.

Performance Objectives

An unusual aspect of the Fuzzy Logic Ramp Metering Algorithm in comparison to other ramp metering algorithms is its flexibility in balancing performance objectives. The performance objectives are flexible in two respects: 1) the operator can dynamically alter the performance objectives, and 2) the operator can specify a *weighted cost function* of

multiple performance objectives. The algorithm design allows tunable performance objectives to accommodate local variations (including politics) and long-term changes in traffic patterns. Most ramp metering algorithms have a single, static performance objective, which is embedded into the control logic. Alternatively, the fuzzy logic controller allows the operator to indicate the relative importance of multiple control objectives, specified through the rule weights. In this way, the metering rates are determined by the simultaneous, weighted consideration of relevant factors, rather than utilizing only one objective or oscillating between objectives.

We balanced several objectives at the implementation sites:

- to minimize mainline congestion
- to maximize throughput volumes
- to minimize overall travel times

while meeting two constraints:

- to prevent a secondary queue (this is when vehicles cannot merge onto the freeway as fast as the metered rate)
- to prevent excessive queue formation (this definition varies from ramp to ramp, depending on devoted arterial turn lanes, the safety of the left hand turn, peak volume, available storage, and local politics).

Inherently, ramp metering has conflicting demands. The objective to reduce mainline congestion produces restrictive metering rates, and the constraint to reduce ramp queues limits how slow these metering rates can be. The constraint to prevent a secondary queue produces minimal metering rates during heavy local congestion, and this conflicts directly with the constraint to reduce ramp queues. The objective of minimizing travel times may conflict with the objective of maximizing throughput. The objective of maximizing throughput may require more vehicles on the freeway at the expense of lower travel times, providing that flow does not break down. In tuning the Fuzzy Logic Ramp Metering Algorithm, the balance point between these objectives was found, with variations

to handle special cases. For details on performance objectives and how to achieve them, see the training manual (Taylor and Meldrum, 2000).

EVALUATION METHOD

Evaluation of the on-line performance of a ramp metering algorithm is complicated by two facts: 1) traffic is not uniform from one day to the next, and 2) performance measures are limited to those that can be measured or estimated reliably.

Controlled Experiment

Because traffic patterns vary with the season, weather, holidays, special events, and the day of week, these factors must explicitly be taken into account during the study. The difficulty lies in distinguishing whether an improvement in freeway system performance is the result of the metering algorithm or other factors such as lower demand, better weather, or perhaps an incident upstream that decreased the volume to the study site. Although it is important to evaluate the ramp metering algorithm under all of these circumstances, it is difficult to *compare* different algorithms under all situations without several months of data, during which traffic patterns may change.

Performing a controlled experiment involves minimizing the variance caused by factors outside of the metering algorithm used. To determine what subsets of similar conditions could be compared, we examined how various factors outside of ramp metering affected traffic patterns. Each day of testing, the operators classified the weather as bad (heavy rain), typical (some rain), or good (no rain). By examining occupancy contour maps and throughput histograms (of similar ramp metering, similar day of week, no special events, and no incidents), we found that there was no discernable difference between traffic patterns of typical and good weather, but that bad weather had a noticeable effect on traffic patterns. Bad weather produced dramatic outlier traffic patterns in the same way that incidents do. To reduce variance in the data set used to compare the metering algorithms, days on which bad weather occurred were not used in the comparative study between algorithms.

The operators also attempted to classify demand. The demand to a site cannot be directly measured. Throughput summed over the metering period and overall congestion provide an idea of what the demand is, but many vehicles may divert or cancel their trip if freeway congestion is high. By itself, throughput is not an accurate gauge of demand, because throughput may be low following freeway breakdown or incidents. Nor is heavier congestion necessarily caused by higher demand. For morning metering periods, demand appeared similar for Monday through Thursday, but Fridays were often lighter. Most Mondays were similar to Tuesday through Thursday, but some Mondays were lighter. When we compared the way in which the operators classified demand with the total throughput volumes of the study sites during the metering period, there was not a reliable relationship. For this reason, we could not accurately classify the data sets into subsets of demand level in order to further reduce the variance caused by demand. Our timeline did not allow us to gather enough data to compare statistically significant subsets of Mondays to Mondays, Tuesdays to Tuesdays, and so forth. Given the data that we had collected, the best solution was to use all of the demand levels in the same data set, provided that the data sets contained similar days of the week for each type of metering.

Each incident is unique, and produces traffic patterns that can be quite different from non-incident data. No two incidents can be compared to each other. For any days on which incidents affected the study site, these days were not used in the comparative study.

The comparison between the Local Metering and Fuzzy Logic Metering Algorithms on the I-90 study site took place between March 15 and June 22, 1999. To reduce the effect of seasonal variations in traffic, alternation between the metering algorithms took place during the study. Both data sets were composed of 28 metered days, containing 3 Mondays, 22 Tuesdays/Wednesdays/Thursdays, and 3 Fridays. The data sets excluded days on which heavy rainfall, incidents, or special events affected the study site.

The comparison between the Bottleneck Metering and Fuzzy Logic Metering Algorithms on the I-405 study site took place between March 15 and July 26, 1999. With

the heavier congestion on I-405, accidents occurred more frequently and with greater effect, particularly on Tuesdays through Thursdays. For this reason, a longer study period was required to gather sufficient incident-free data, and the data sets contained more Mondays and Fridays. We were not able to match the data sets exactly with regard to day of the week. Both data sets comprised 27 metered days. The Bottleneck Metering data set contained 4 Mondays, 19 Tuesdays/Wednesdays/Thursdays, and 4 Fridays. The Fuzzy Logic Metering data set contained 4 Mondays, 21 Tuesdays/Wednesdays/ Thursdays, and 6 Fridays.

Performance Measures

Desired performance measure include the total distance traveled by all vehicles in the system, total travel time of all vehicles in the system, and queue delays, but we were not able to accurately measure or estimate these performance measures for the duration of the study. We were limited to performance measures that we could accurately estimate through a combination of hardware and software processing, because this project did not have the resources required to gather data by hand.

Mainline Performance

Of the mainline performance measures, the one that was the best representation of mainline congestion was occupancy. Using CDR (Compact Disk Retrieval Software) and CD Analyst, a new software package developed under the FLOW project (Ishimaru and Hallenbeck, 1999), we were able to process 5-minute data to create occupancy contour maps in an efficient, standard, and reproducible methodology. The CD Analyst software estimates bad or missing loop detector data. During the study, good data availability was not a problem.

Mainline speeds are also a good barometer of congestion. However, the speeds available to us were not very accurate. The speeds estimated from the loop detector data assume a constant headway between vehicles. With heavier congestion, the actual

headway drops, and the real speeds are slower than the estimated speeds. With no congestion, this headway is greater than the assumed constant, and the real speeds are faster than the estimated speeds. Thus, a change in congestion would not be reflected in the estimated speeds to the extent that it would be represented by the occupancy data. Due to the inaccuracy of the speeds when we most needed them (heavier congestion), the speeds were not as good of a barometer of mainline congestion as occupancy.

In turn, mainline travel time estimates that rely on these estimated speeds were not very accurate either, and were not used. Travel times available through TrafficView (Microsoft's traffic page) were considered, but they were no more accurate than the travel times produced by CD Analyst (both used the same data). Within CD Analyst, the Kalman filter option of calculating speeds might have produced more reliable travel times, because it does not assume a constant headway. However, only the 'Normal' option (where a fixed headway is assumed) was available for our beta test version of the software. Nor could these software packages calculate a travel time that included queue delay, which would have been of great use. The reasons for these limitations are described below.

Because mainline travel times are dependent on mainline occupancies and volumes, the combination of occupancy contour maps in conjunction with throughput volumes was a sufficient measure of mainline performance. The volume of the mainline station which best represented the throughput of the study site was summed from 5 to 10 AM, in order to encompass the demand throughout the morning commute. (Metering typically begins and ends well within this period.) The percentage change and distribution of the throughput volume between the two algorithms were compared. Accident rates were not used as a performance measure because they were not statistically significant with less than two months of each type of metering algorithm.

Ramp Queues

To determine whether the ramp metering algorithm had a beneficial effect overall, it was important to examine the queue delay in addition to mainline congestion. Ramp

metering must trade off between mainline congestion and queue delay. We wanted to know whether an improvement in mainline performance was achieved at the expense of further queue delay and the route diversion that goes hand-in-hand with excessive queue delay. Ideally, we would measure total travel times that included queue delays, with the objective of improving overall travel times. However, this was not possible because we were unable to accurately estimate queue delay with the available loop detector data, nor did this project have the resources required to continuously gather travel time data by driving the study sites, or to measure queue delays by hand.

Poor loop detector placement was the limitation in accurately estimating queue delay. If the loop detectors had been better placed, we could calculate the number of vehicles added to the queue each sample as the total volume into the queue minus the total volume out of the queue, including HOVs. Aggregating the queue storage rate would give us the queue size. (The queue calculation has an initial condition of zero queue size because the calculation begins before metering.) From the passage rate (which is the realized metering rate) and the queue size, we could estimate the queue delay. Using this method, we wrote software to process 20-second data, estimate missing 20-second data, verify data quality, and calculate queue performance measures. (See the software manual by Taylor and Meldrum, 2000).

However, the number of vehicles in the queue was far from accurate for three reasons: 1) The queue often continues far past the last detector, so the queue calculation would not encompass all of the vehicles in the queue. 2) Vehicles are often counted by more than one detector because of weaving patterns or overly sensitive detectors. If the detectors for adjacent lanes are not located adjacently, a single weaving vehicle could be counted by both. Even when the detectors are adjacent, a vehicle changing lanes over adjacent detectors can be counted twice. 3) Vehicles are often not counted because of poor detector placement or weaving patterns. If a vehicle changes to an adjacent lane but the detectors are non-adjacent, it may elude detection. Some detectors are located too far left

or right to capture most of the vehicles that pass near it. (The Appendix provides a list of ramps that need better loop detector placement for the purposes of control.)

For over half of the metered ramps on the study sites, calculations of the vehicles in the queue were not at all accurate because of miscounted vehicles. When the queue estimate was inaccurate, there was a definite trend for each ramp. Of the ramps with poor detector placement, about half of them erred in the direction of positive queue size, and the other half erred in the direction of negative queue size. Even when the error was only one vehicle per minute, by the end of a three hour metering period, the total error of the vehicles in the queue would reach 180 vehicles. Commonly, the calculated queue size would be positive or negative by hundreds of vehicles by the end of the metering period. Using the camera, we could easily verify that these volume counts and queue sizes were not accurate. The queue calculation was reliable for single metered lanes with no HOV bypass (where weaving is not possible), but few of the metered ramps are of this design.

If the loop detector data are so inaccurate, how can the ramp metering algorithms possibly function so well? Because the loop detector occupancies are more reliable than volumes for indicating a queue presence, and this is what we use. The ramp queue occupancies are more reliable than calculations of queue delay for both the purpose of real-time ramp metering control and evaluation of queue characteristics. In particular, it is the aggregated storage rate that is problematic. Storage rates have a poor signal-to-noise ratio by nature, because they have all of the error but little of the volume. When we aggregate this storage rate, we propagate and build the error over time, while the queue magnitude remains small. The ramp metering algorithms do rely on volume to some extent – to indicate when the vehicle waiting at the meter has passed the meter. For the ramps where the vehicles consistently miss the passage loop, the amplifier of the passage loop was turned up to increase its sensitivity, which was able to solve the problem for the most part. In general, the demand and passage loops are located so close to the meter that the vehicles usually hit them as intended. Near the queue and advance queue loops, the

vehicles have more freedom of movement, and vehicle miscounts abound where there are multiple adjacent lanes.

Interpreting queue performance measures can be tricky because there are situations in which a ramp queue is desirable, such as to prevent a secondary queue from forming at the merge. If we do not prevent a secondary queue, there is no benefit to metering. If a secondary queue forms, the metered vehicles are contributing to a mainline bottleneck at the merge. Because all mainline vehicles through this section are affected by this bottleneck, there is system-wide benefit to preventing a local bottleneck. For these reasons, preventing a secondary queue takes precedence over maintaining a reasonable queue when the local mainline merge is highly congested. Drivers obey this restrictive rate because they understand that there is no point to metering faster than the vehicles can merge onto the mainline.

Similarly, a high occupancy at the *queue* detector is not considered problematic when we want to utilize the available storage at the ramp to prevent mainline bottlenecks. However, when the *advance queue* detector frequently reads high occupancy, this queue may block the arterial. This is acceptable only if preventing a secondary queue and if local politics allow an excessive queue.

How could we come up with meaningful queue performance measures when there are times that a queue is desirable and data quality is limited? To answer this question, we evaluated 11 queue performance measures on 14 metered on-ramps (the ones at the study sites) for several days of metering. Table 1 shows the results of this study, with the desired characteristic that we wanted to measure and the usefulness of the performance measure. All performance measures were calculated from 20-second data of good quality, where any missing data points were estimated.

Of the 11 queue performance measures investigated, two of them proved useful: the number of minutes that the queue activated the queue detector, and the number of minutes that the queue activated the advance queue detector. These are the only

performance measures that were consistently accurate and meaningful for all ramps. Although this performance measure did not tell us how far the queue extended beyond the detector, a reduction in the number of minutes that the queue activated the detector implied that the queue was shorter because it dissipated faster. For the advance queue detector, we wanted the minimal time that the queue exceeded its available storage, except in the case of a secondary queue. For the queue detector, this measure was more ambiguous, depending on the queue's relationship to mainline events, but using it in conjunction with the mainline occupancy contour plots helped explain any ambiguities in its performance.

Table 1: Queue Performance Measures

Performance Measure	Desired Characteristic	Usefulness
Maximum # of vehicles in the queue during the metering period	Does not exceed maximum allowable storage, except in the case of a secondary queue	Could not estimate accurately --unusable.
% change in maximum number of vehicles in the queue	Reduction in maximum queue size, except if preventing secondary queue or mainline bottleneck.	Although the queue calculation itself is inaccurate, a consistent bias in the data would mean that the % change is usable. However, this measure was of limited usefulness without knowing the relationship between it and the mainline. It was not practical to write software that checked the ramp queue size with the history of the mainline bottlenecks, because if the metering algorithm prevented a bottleneck from forming, the mainline data would not contain absolute evidence of it. For some ramps, this measure was usable, but for others, it was not meaningful.

Performance Measure	Desired Characteristic	Usefulness
# of Minutes that the # of vehicles in the queue exceeds the available storage	Does not exceed maximum allowable storage, except in the case of a secondary queue	Could not estimate accurately the # of vehicles in the queue --unusable.
Aggregate the # of minutes that the occupancy of the *queue detector* exceeds 35% during the metering period, and calculate the average/day for each ramp. In the case of multiple queue detectors on adjacent ramps, the maximum occupancy was used.	We assume that the queue had reached the detector if the occupancy exceeded 35%, and that otherwise the queue had not reached the detector. A reduction in minutes would imply that the queue dissipated faster because it was shorter. Generally, a reduction is desired because this correlates with less ramp delay, but a high number of minutes for this detector is not necessarily a penalty because the algorithm should utilize available storage when mitigating bottlenecks.	These queue loops tend to read either very low occupancy if the queue has not yet reached the detector (less than 8%), or very high after the queue has reached the detector (typically greater than 60%). The queue occupancy rarely reads mid-range, so the 35% threshold effectively classified whether or not the queue had reached the detector. This was one of the few performance measures that we could accurately obtain for all ramps. Although it was not perfect in that we didn't know anything about timing in relation to mainline events or actual queue size, it was a useful gauge for whether or not the ramp delay had increased or decreased.

Performance Measure	Desired Characteristic	Usefulness
Aggregate the # of minutes that the occupancy of the *advance queue* detector exceeds 35% during the metering period, and calculate the average/day for each ramp. In the case where there were multiple advance queue detectors, the maximum occupancy was used.	We assume that the queue had reached the advance queue detector if the occupancy exceeded 35%, and that otherwise the queue had not reached the detector. A reduction is highly desired, because there is typically no more storage available beyond this detector.	These queue loops tend to read either very low occupancy if the queue has not yet reached the detector (less than 8%), or very high after the queue has reached the detector (typically greater than 60%). The advance queue occupancy infrequently reads mid-range, so the 35% threshold effectively classified whether or not the queue had reached the detector. This was one of the few performance measures that we could accurately obtain for all ramps. Although it was not perfect in that we didn't know anything about timing in relation to mainline events or actual queue size, it told us how long the queue exceeded its available storage, and if the ramp delay had increased or decreased.

Performance Measure	Desired Characteristic	Usefulness
Initial time that the occupancy of *the queue detector* exceeds 35%	One problem with this measure is that the desired characteristic depends so much on what the mainline conditions are. A late start time for the queue would be desirable to reduce queue delay, but not desired if it exacerbated mainline bottlenecks.	Not meaningful because of inconsistency of queue size. The demand peaks at certain times, but the queue does not persist between peaks.
Initial time that the occupancy of the *advance queue detector* exceeds 35%	A late start time for the queue would imply that the algorithm was able to maintain a reasonable queue for longer, and queue delay is shorter.	Not meaningful because of inconsistency of queue size. The demand peaks at certain times, but the queue does not persist between peaks. For most ramps, the advance queue occupancy rarely exceeds 35%.

Performance Measure	Desired Characteristic	Usefulness
End time that the occupancy of the *queue detector* no longer exceeds 35%	An earlier end time for the queue would imply less queue delay.	Not meaningful because of the inconsistency of the queue. The demand peaks at certain times, but the queue does not persist between all peaks. For many ramps, there was a late morning spike, although the queue had dissipated much earlier.
End time that the occupancy of the *advance queue detector* no longer exceeds 35%	An earlier end time for the queue would imply less queue delay.	Not meaningful because of the inconsistency of the queue. The demand peaks at certain times, but the queue does not persist between all peaks. For many ramps, there was a late morning spike, although the queue had dissipated much earlier.

Performance Measure	Desired Characteristic	Usefulness
# of positive transitions between no queue and queue, where queue presence is defined as *queue occupancy* > 35% for at least two out of three consecutive samples	The idea of this measure is to see how oscillatory the queue is, because this is one of the pre-existing problems that we are trying to address. Less oscillation implies smoother control.	This measure was not meaningful. Peaks in the queue size were frequently a result of peaks in demand, platooning caused by the signal, rather than a function of the metering algorithm.
# of positive transitions between no advance queue and advance queue, where advance queue presence is defined as *advance queue* occupancy > 35% for at least two out of three consecutive samples	The idea of this measure is to see how oscillatory the queue is. Less oscillation implies smoother control.	This measure was not meaningful, and for many ramps, this measure was zero because the queue never reached the advance queue detector. Peaks in the queue size were frequently a result of peaks in demand, platooning caused by the signal, not a function of the metering algorithm.

RESULTS

At the first study site, on fuzzy logic metered days mainline occupancies were smaller, throughput volumes were larger, and queues were slightly longer in comparison to the days when the Local Algorithm was metering. At the second study site, on fuzzy logic metered days mainline occupancies were similar, throughput volumes were similar, and queues were significantly shorter in comparison to days when the Bottleneck Algorithm was metering.

I-90

At the I-90 study site, the Fuzzy Logic Algorithm resulted in a reduction in mainline occupancies relative to the Local Algorithm (figures 2 and 3). This 8.2 percent change in mainline congestion was significant enough that it was noticeable on a day-to-day basis with the closed-circuit television (CCTV) and FLOW map. Most importantly, the Fuzzy Logic Algorithm prevented the bottleneck near the Eastgate on-ramp, while the Local Algorithm did not. Because the Fuzzy Logic Algorithm uses downstream inputs and the Local Algorithm does not, these results were expected.

The bottleneck upstream near SR 900 was slightly more congested with the Fuzzy Logic Algorithm. There are two possible explanations for this. Most significantly, there was higher throughput during the days metered with the Fuzzy Logic Algorithm, and because most of that additional volume originated at the SR 900 on-ramp, it is not surprising to see more congestion at this merge. There also may be a minor trade-off effect between preventing the downstream bottleneck at Eastgate and additional congestion upstream at SR 900. If so, this trade-off is worthwhile, given that more vehicles can get through the corridor without experiencing freeway breakdown when the congestion is further upstream.

The volume histogram indicates that more high-flow days occurred during fuzzy logic metering than during local metering (Figure 4). Overall, there was a 4.9 percent increase in throughput. Because occupancy is inversely proportional to speed, and volume is proportional to speed, the duo of lower occupancies and higher volumes means that the mainline speeds increased as well. With the combination of lower average mainline occupancies and higher average throughput, the data supports that the Fuzzy Logic Algorithm utilized the mainline more efficiently than did the Local Algorithm.

Figure 5 shows the number of minutes that the queue reached the queue detector and advance queue occupancy detector, averaged per day for each type of metering. If the advance queue data is not shown, that is because there were no instances when the queue reached the advance queue detector.

At the Eastgate on-ramp, the queue detector read 5 minutes more of high occupancy fuzzy metering than for local metering. However, this ramp has plenty of storage, and the queue never reached the advance queue detector for either metering algorithm. Considering that this merge is a bottleneck, a slightly longer queue at this ramp is acceptable and desirable for system-wide benefit.

West Lake Sammamish Way data showed that fuzzy metering reduced the number of minutes of high queue occupancy by roughly a third. This reduction is probably due to the preventative nature of the Fuzzy Algorithm's queue control in comparison to the Local Algorithm's threshold activated queue control. As at Eastgate, this ramp has adequate storage, and neither ramp metering algorithm exceeded its storage capacity.

The ramp volumes for the SR 900 slip ramp far exceeded the storage of the ramp. Neither algorithm was able to meter the vehicles quickly enough to avoid an excessive queue. Traffic is already moving slowly because of a bend in the mainline just upstream of the merge. Between the bend in the freeway and the high volumes merging, this location tends to be a bottleneck, with a difficult merge onto the mainline for SR 900. Fuzzy metering has more minutes than local metering at this ramp for two reasons. 1) Fuzzy

metering restricts the metering rate during heavy mainline congestion more so than local metering to prevent a secondary queue. The congestion on the contour maps reflects the difficulty in this local merge during higher demand days and explains why fuzzy metering restricted the metering rates at this ramp. Preventing the bottleneck at this merge preserves mainline efficiency despite the higher volumes during fuzzy metering. 2) Much of the higher flow that occurred during the fuzzy metered days originated at this ramp and is reflected in the queue.

The SR 900 loop ramp is low volume, and the queue is reasonable. Fuzzy metering restricted this metering rate more than local metering in order to better utilize the storage, especially because this ramp merges with the overloaded slip ramp. By metering the SR 900 loop ramp more restrictively, some of the burden was taken off of the slip ramp. Part of the reason that fuzzy metering produced a higher queue here than local metering is that these algorithms handle the HOV bypass differently. This ramp is an example of how the fuzzy metering rate reduction caused by the HOV bypass is distributed among ramps that merge onto the mainline with those HOV vehicles. While the local metering does all of the HOV reduction on the slip ramp adjacent to the HOV bypass, the Fuzzy Metering Algorithm distributes the HOV reduction on both the slip ramp and the loop ramp, because both of these ramps are affected by the HOVs at the merge with the mainline.

Since this study was done, two events have occurred to improve the queue problem at SR 900: 1) Fuzzy metering was retuned to be more responsive to an excessive queue at the SR 900 loop ramp, and 2) a new road was added to access the W. Lake Sammamish on-ramp more easily. With this new road access, some drivers now use W. Lake Sammamish way instead of the SR 900 on-ramp. This alternative route is desirable because the W. Lake Sammamish on-ramp was under capacity.
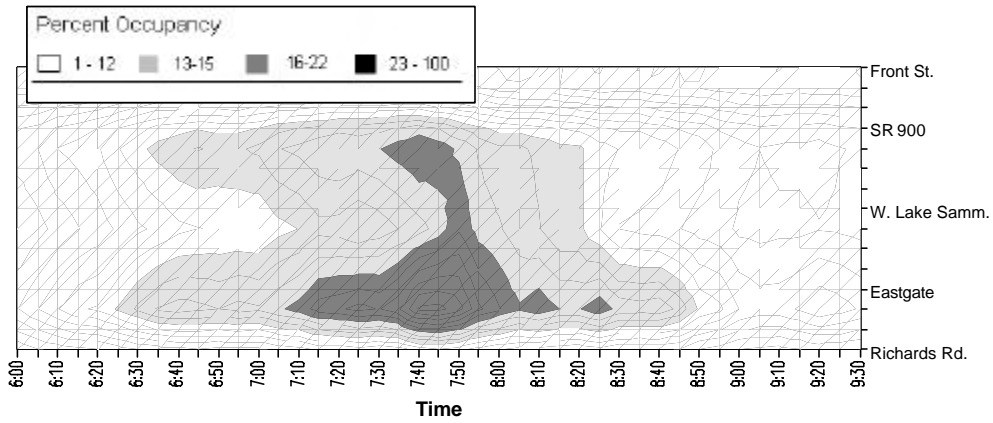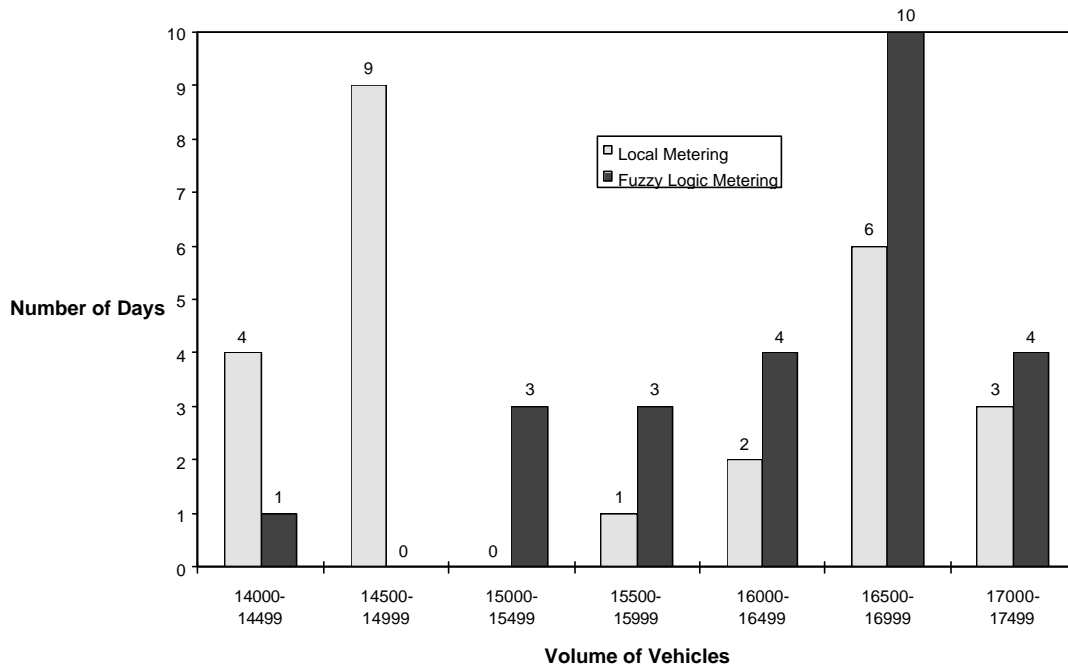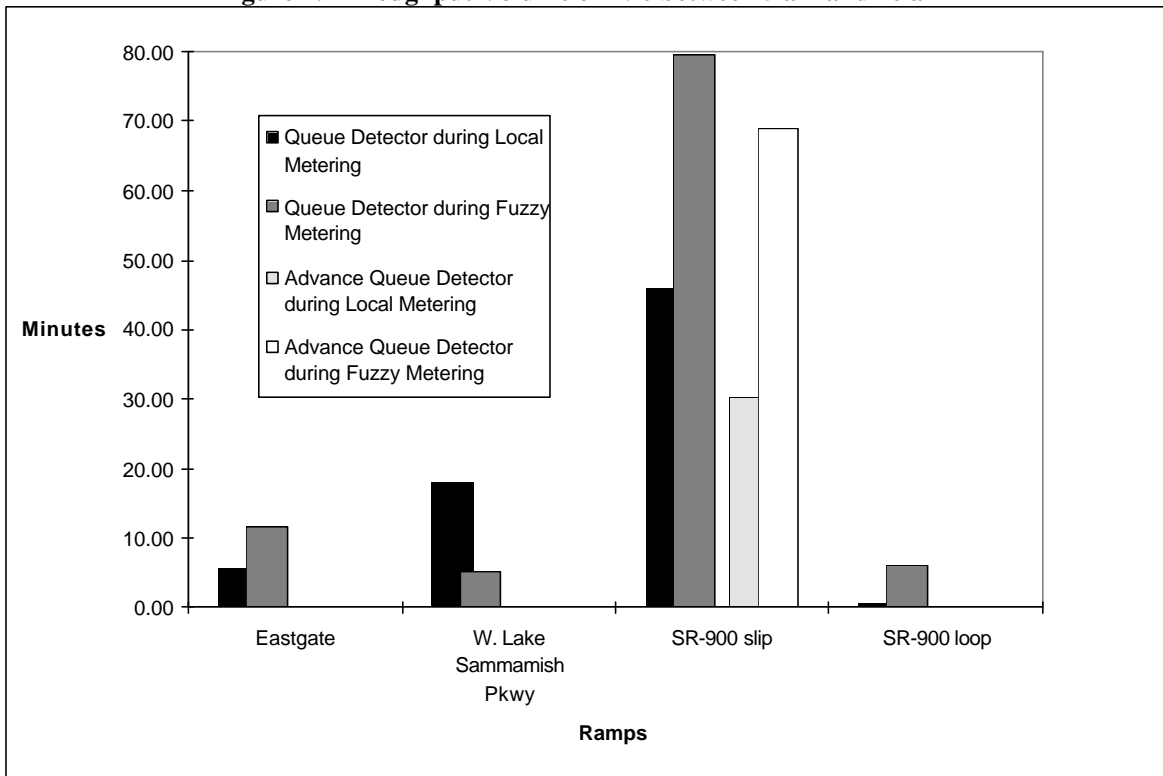
Figure 2: Occupancy Contour of Local Metering on I-90
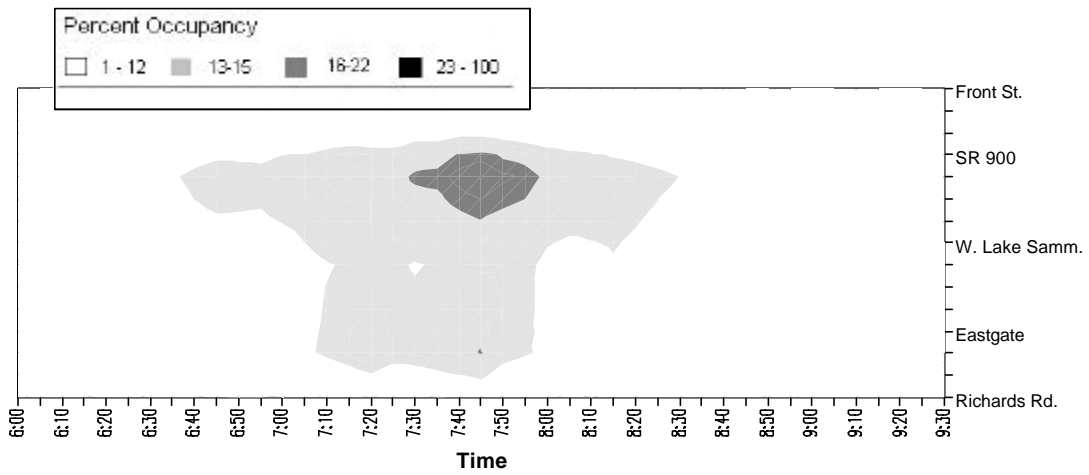
Figure 3. Occupancy Contour of Fuzzy Logic Metering on I-90

28

Figure 4: Throughput Volume of I-90 between 5 am and 10 am

Figure 5: Average Minutes/Day that Queue Reaches Detectors on I-90 Ramps

Overall, the effect of the Fuzzy Metering Algorithm in comparison to that of the Local Algorithm appears beneficial. Mainline efficiency was improved, with lower mainline occupancies and higher throughput. Although the queue occupancy (average/day for the study site as a whole) was high 32 minutes more during fuzzy metering than during local metering, this increase is justifiable given that some ramps were underutilized and that secondary queue prevention occurred at the SR 900 merge. The advance queue occupancy (average/day for the study site) was high for 6 minutes more during fuzzy metering than during local metering, but this excessive queue was largely the result of higher demand, and the overall system benefits outweigh the slightly higher queues.

I-405

The I-405 study site is much more congested than the I-90 site, often experiencing freeway breakdown for hours a day. The bottleneck begins around the 116th Street merge, and congestion continues up to 160th Street. The merges at 124th and 160th Street can be problematic as well. The occupancy contour maps indicate that the mainline congestion was slightly worse with fuzzy metering than with bottleneck metering, with an average increase of 1.2 percent occupancy (figures 6 and 7). The Bottleneck Algorithm almost always meters at the minimum allowable metering rates. The lower mainline congestion with Bottleneck Metering presents an argument for metering restrictively in general, and for using downstream inputs to mitigate bottlenecks.

Vehicle throughput was more distributed for days during which fuzzy metering was applied (Figure 8). Fuzzy metering took place during more high end days and more low end days, while bottleneck metering throughput was more consistent day to day. There is no obvious explanation for this distribution, except that the I-405 study site is more chaotic than I-90, with a more dramatic breakdown phenomenon. With more days in the study, it would be expected that the throughput distributions would be more similar. Averaged over

all days, the flows were nearly identical between the two algorithms, with fuzzy metering producing a 0.8 percent increase over bottleneck metering.

The Bottleneck Algorithm achieved a slightly less congested mainline with slightly less throughput at the expense of much longer queues (Figure 9). Although the CCTVs frequently displayed ramp queues that were politically unacceptable, the engineers were unable to easily tune the Bottleneck Algorithm to meter at mid-range rates. (Reasons for the difficulty in obtaining non-minimal metering rates from the Bottleneck Algorithm are described in the training manual by Taylor and Meldrum, 2000.)

At the 72nd ramp, bottleneck metering produced a high queue occupancy for 43 minutes in comparison to the Fuzzy Algorithm's 3 minutes. This ramp is downstream of any appreciable mainline bottleneck, so there is no noticeable benefit to a restrictive rate at this ramp, as shown in the occupancy contour maps. The lightly congested mainline at this ramp's merge does not justify the 8 minutes of excessive queue formation during bottleneck metering, which compares to 2 minutes of excessive queue formation with fuzzy metering.

Likewise, at the 85th St. on-ramps, the merge is downstream of the huge bottleneck at 116th, and therefore fuzzy metering did not restrict it as much as bottleneck metering, with 18 fewer minutes of high queue occupancy for both the loop and slip ramps. The occurrence of excessive queue formation at the loop ramp was low and nearly identical between the two algorithms. At the slip ramp, an excessive ramp queue occurred twice as often during bottleneck metering than during fuzzy metering.

The ramp volumes at 116th St. are high relative to the available ramp storage, and the local merge is very congested. Both algorithms attempted to mitigate this mainline bottleneck with restrictive metering rates. The advance queue detector data indicates that the Bottleneck Algorithm exceeded allowable storage for almost twice as long as the
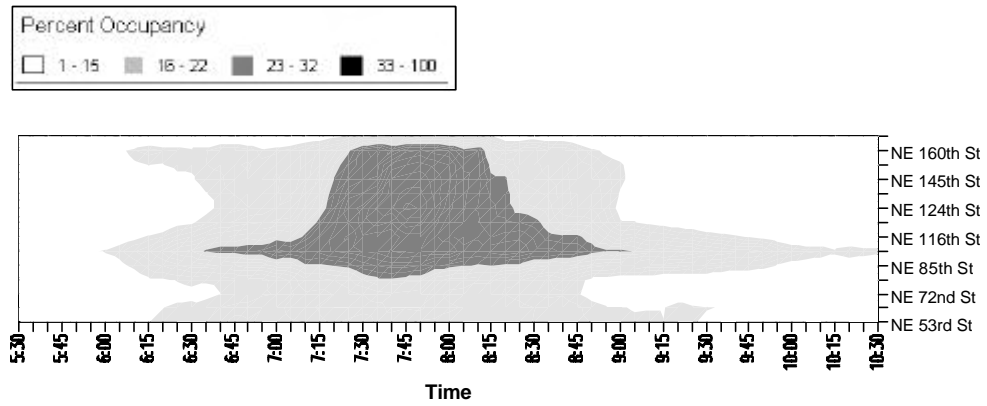
Figure 6: Occupancy Contour Map of Bottleneck Metering on I-405
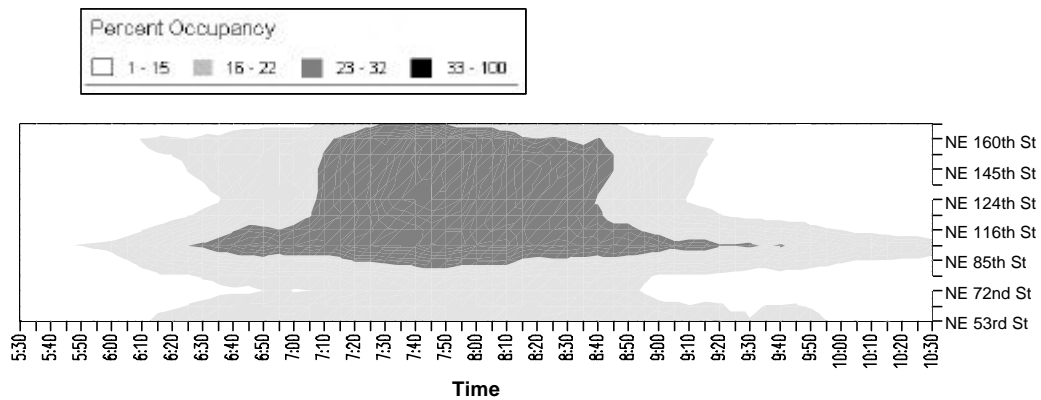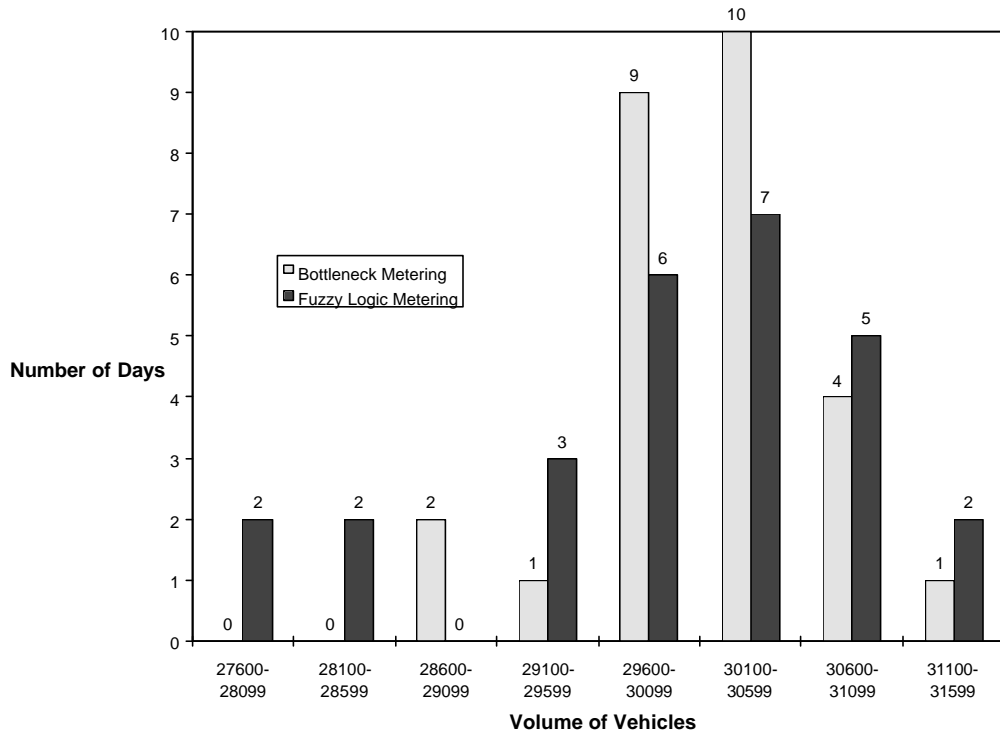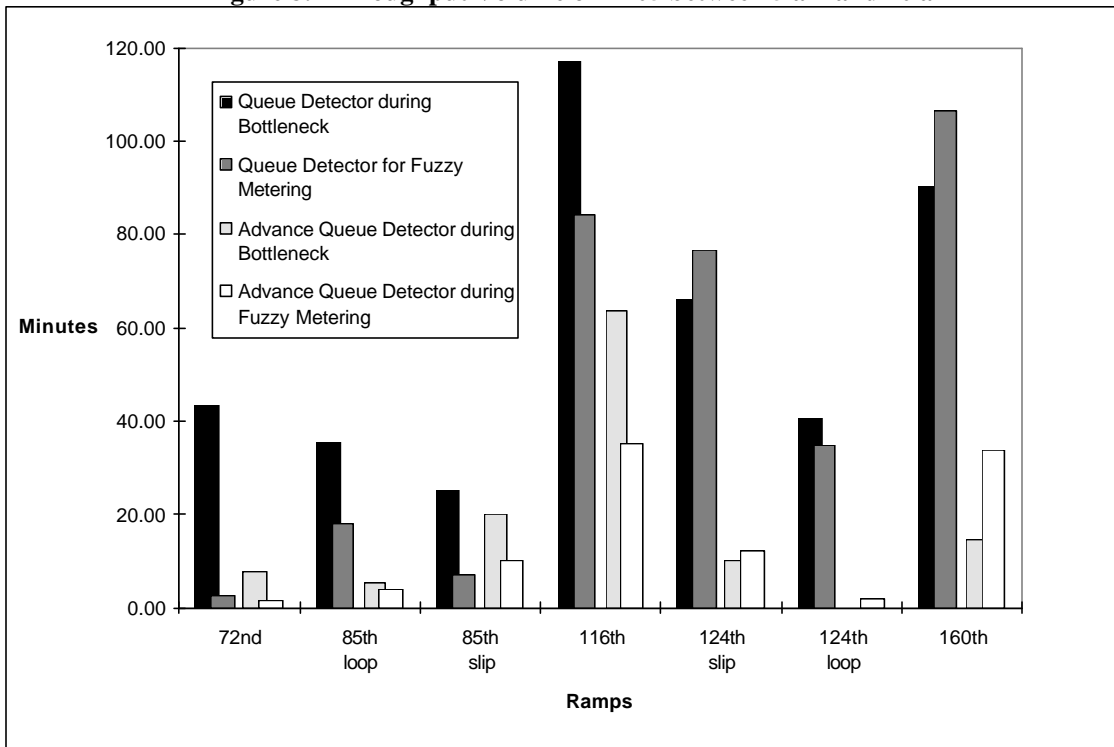

Figure 7: Occupancy Contour Map of Fuzzy Metering on I-405

**Figure 8: Throughput Volume of I-405 between 5 am and 10 am**

Figure 9: Average Minutes/Day that Queue Reaches Detectors on I-405 Ramps

33

Fuzzy Algorithm, creating a political situation evidenced by the complaints received while the Bottleneck Algorithm was metering this location.
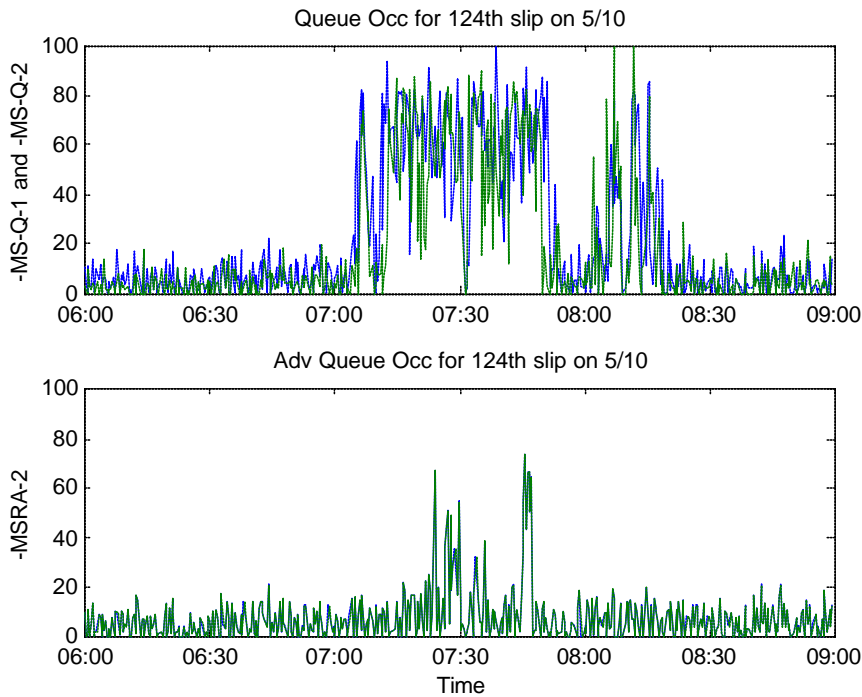
At 124[th] Street, fuzzy metering and bottleneck metering produced similar ramp queues for both the slip and loop ramps. Fuzzy metering produced a queue that was slightly higher at the slip ramp and slightly lower at the loop ramp. The excessive queue formation of 10 minutes during fuzzy metering and 12 minutes during bottleneck metering is justifiable because a restrictive rate is necessary to prevent a secondary queue at this location.

The 124[th] slip ramp provides an interesting opportunity to compare the queue characteristics of the two algorithms, given similar performance measures. For instance, Figure 10 shows the representative queue occupancy and advance queue occupancy (using 20-second data) during a day of bottleneck metering, April 22nd, and Figure 11 shows representative queue and advance queue occupancy during a day of fuzzy metering, May 10th. Although the queue performance measures produced by these occupancy plots are similar, there are some noticeable differences in the queue characteristics. During bottleneck metering, the queue length oscillated much more than it did during fuzzy metering. This pattern was consistent from day to day, from ramp to ramp. Differences between these algorithms explains the difference in queue characteristics. Bottleneck metering uses threshold activation to indicate when to administer a queue override and advance queue override. The queue is allowed to build to a certain extent, and only then is action taken, up until the queue dissipates. Fuzzy metering, on the other hand, does not wait for an excessive queue but instead provides smooth, continuous, preventative control.

At 160[th] St., fuzzy metering produced a bit longer queue to help with this difficult merge and the downstream bottleneck at 116[th]. An excessive queue occurred more frequently with fuzzy metering because of the fact that fuzzy metering specifically addresses the secondary queue formation.
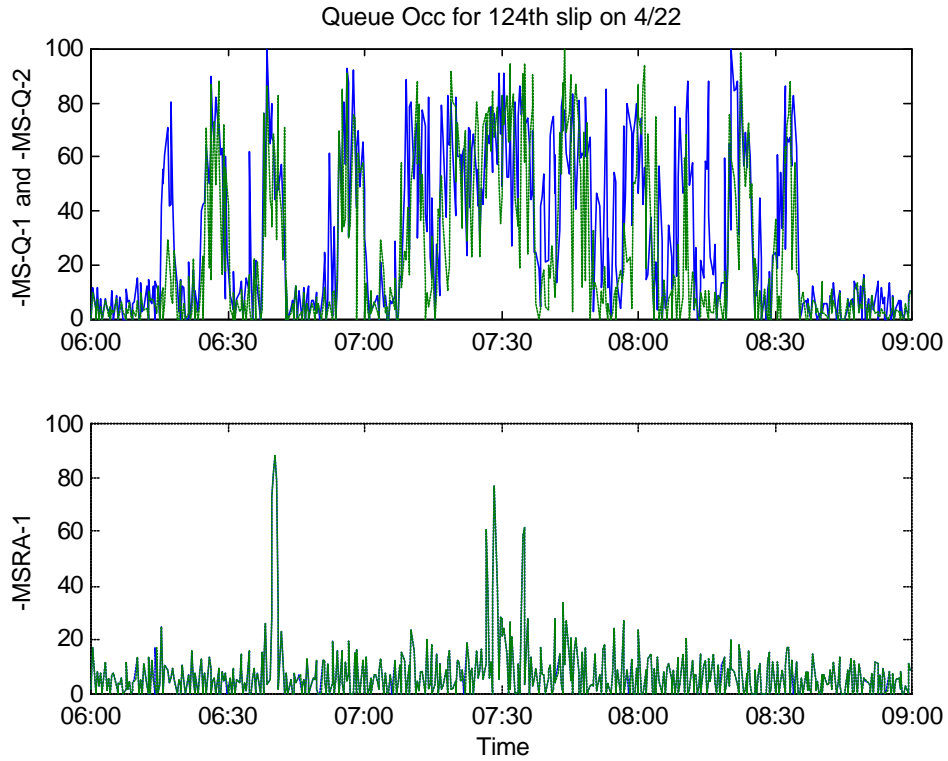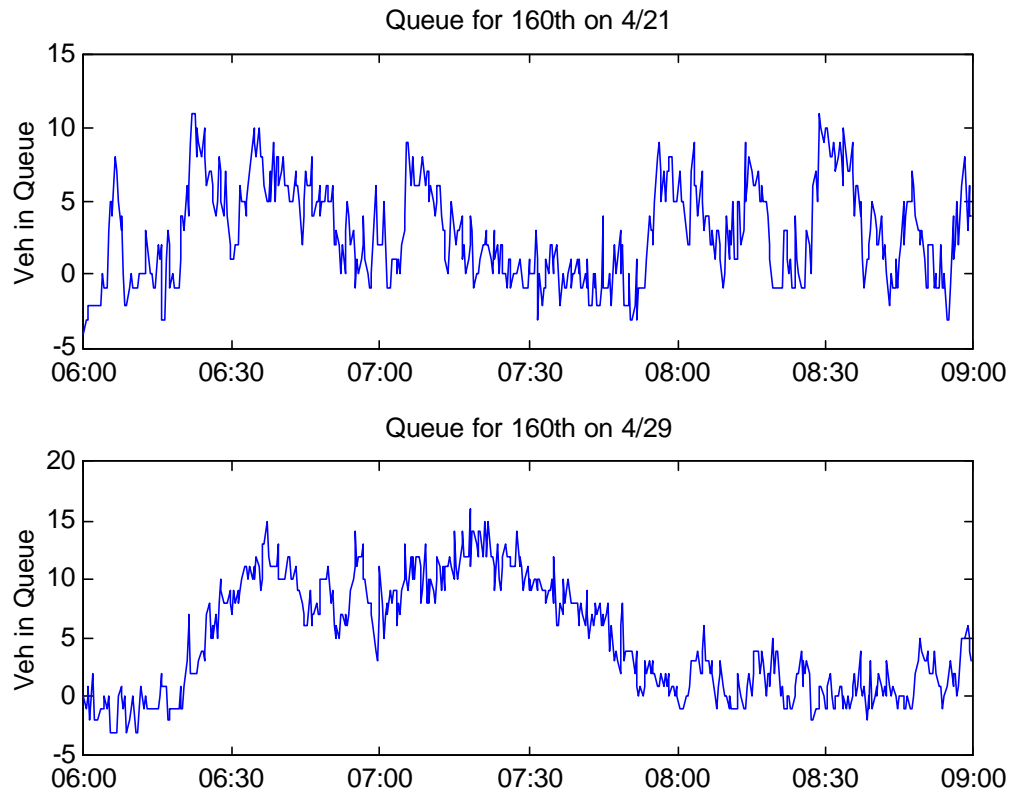
Figure 10: Queue and Advance Queue Occupancies during Bottleneck Metering at 124th

Figure 11: Queue and Advance Queue Occupancies during Fuzzy Metering at 124th

Figure 12: Vehicles in Queue at 160th for Bottleneck and Fuzzy Metering

For 160th Street, the estimates of vehicles in the queue were accurate and could be used to examine queue characteristics. Figure 12 shows the vehicles in the queue (using 20-second data) during a representative day of bottleneck metering (April 21st) and fuzzy metering (April 29th). During bottleneck metering, there is a distinctive sawtooth pattern as the queue suddenly builds, then slowly dissipates. The pattern for each type of metering was repeated consistently for various days and ramps. Again, the queue oscillation during bottleneck metering is explained by threshold activation in contrast to the graduated control of fuzzy metering. The Bottleneck Metering Algorithm gets into an oscillatory cycle between alleviating mainline congestion and dispersing the ramp queue. When the queue

36

exceeds certain thresholds, the queue override turns on, dissipating the queue. This dump of vehicles onto the mainline exacerbates the mainline bottleneck at this location. The metering rate is then restricted to mitigate this mainline bottleneck. The queue builds, and the sawtooth pattern repeats. Fuzzy metering is more consistent regarding when a queue is acceptable. A sustained queue is acceptable to prevent a bottleneck at the merge and the resulting secondary queue formation.

On average per day for the I-405 study site, the Fuzzy Metering Algorithm achieved 133 fewer minutes of high queue occupancy than the Bottleneck Algorithm, and 2 fewer minutes of advance queue occupancy. The few ramps where the queue did increase during fuzzy metering were locations already prone to secondary queue formation. These data show that queue delay decreased overall with fuzzy metering, without decreasing the efficiency of the mainline. The mainline performance was similar for both algorithms, with slightly higher occupancies and volumes for fuzzy metering. Given that the mainline performance did not change appreciably but the queues did, there appears to be an overall benefit to using the Fuzzy Metering Algorithm.

As a whole, the Fuzzy Metering Algorithm did not simply meter faster or slower than Local or Bottleneck Algorithm. The results depended on the conditions. The Fuzzy Logic Algorithm metered more restrictively than the Local Algorithm or Bottleneck Algorithm when preventing a mainline bottleneck, secondary queue formation, or an excessive queue. In the cases where there were not tangible system-wide benefits to metering more restrictively, the metering rates during fuzzy metering were higher than those of the Local or Bottleneck algorithms in order to increase the throughput.

SYSTEM-WIDE IMPLEMENTATION

On the basis of the success of Fuzzy Metering Algorithm's successful performance at the study sites, ease of tuning, and handling of incidents, we proceeded to implement fuzzy metering system-wide. At the request of the freeway operations management, we began implementation on the most congested corridors first.

We categorized corridors into segments according to which ones were turned on/off at the same time, and which meters could mitigate the same bottlenecks because of similar destinations. By categorizing the ramp meters into the subsets shown in Table 2, we were able to tune the relative metering rates within a group to achieve system-wide objectives. We cannot optimize a single ramp by itself because the conditions produced by the metering rates at one ramp affect the metering rates at other ramps within the group. Thus, local optimization does not necessarily produce optimization of the group. However, we believe that the metering rates between different groups were independent enough that we could optimize each group to approximate system optimization. (See the training manual for the tuning method, Taylor and Meldrum, 2000.)

The implementation of the Renton ramps was unique in the sense that fuzzy metering was the original algorithm used at this site, and these ramps were highly political. For years, WSDOT had been negotiating with the City of Renton to meter the ramps on northbound I-405 from SR 169 to NE 44th, and southbound I-405 from Coal Creek to NE Park Drive. The City of Renton finally agreed reluctantly that ramp metering in Renton would begin July 19 of 1999, provided that queues were not excessive.

Because of the results of fuzzy metering at the study sites and on southbound I-5, the WSDOT managers requested fuzzy metering as the default algorithm for these new meters. With so many meters going on-line at once, it would be difficult to observe and tune them all simultaneously. To ensure proper metering behavior for the high profile

Table 2: Ramp Metering Groups in Order of Implementation

GROUP	LOCATION	FROM	TO	Typical Metering Period
1	WB I-90	Front St	Eastgate	AM
2	SB I-405	160th	72nd	AM
3	SB I-5	128th	Boylston	AM/PM
4	SB I-405	8th	4th	AM/PM
5	NB I-405	4th	8th	PM
6	NB I-405	Interurban	44th	AM/PM
7	SB I-405	112th	Interurban	AM/PM
8	NB I-405	60th	160th	PM
9	NB SR-167	277th	84th	AM
10	SB SR-167	84th	277th	PM
11	NB I-5	Dearborn	220th	PM
12	WB 520	Montlake	Lake WA Blvd	PM
13	SB I-5	Dearborn	South Center	PM
14	NB I-5	South Center	University	AM/PM
15	WB I-90	Richards	W. Mercer	AM/PM
16	EB I-90	Rainier	E. Mercer	AM

debut of these ramp meters, we estimated optimal control parameters before implementation. We needed a technique to do this that neither relied on the accuracy of the queue estimate nor required prior knowledge of metering conditions, and with our limited timeline, could be completed in under a week.

We developed what we called a Q-factor, which indicated the factor of the queue response needed relative to a similar ramp for which fuzzy metering had been optimized. We compared the Renton ramp to fuzzy metered ramps with similar mainline conditions and similar geometry (in terms of how many adjacently metered lanes and if there was an HOV bypass), so that we expected similar activation of the mainline rules. Because the queue rules balance the mainline rules within the fuzzy controller, tuning is a matter of determining the appropriate rule weights for the queue rules relative to the mainline rules. (See the training manual for a description of the rule base, Taylor and Meldrum, 2000.) We calculated the Q-factor below, where volume is the peak 5-minute volume during the metering period, and storage is the allowable number of vehicles in the queue between the stop bar and undevoted arterial:

$$Q_{factor} = \frac{VOLUME_{\text{Re}nton} * STORAGE_{optimal}}{VOLUME_{optimal} * STORAGE_{\text{Re}nton}}$$

The Q-factor simply indicates how much queue response we needed in comparison to the optimal ramp, based on relative peak volume and available storage. We multiplied the Q-factor by the rule weights of the optimal ramp to estimate the rule weights of the Renton ramp. We compared each Renton ramp to more than one optimal ramp. The higher the Q-factor, the greater the volume relative to the available storage, so the more attention we gave that ramp when it went on-line. Although this method certainly was not perfect, it provided accurate initial estimates of the control parameters and indicated which ramps might be problematic. The implementation of the Renton ramp meters was considered a success by the City of Renton and WSDOT.

Full scale implementation of the Fuzzy Logic Ramp Metering Algorithm was completed in September 1999. At this time, 126 ramps use the new algorithm as the default metering algorithm. For all implementation sites, the algorithm performance was observed and tuned for optimal performance. Although the scope of this project did not include system-wide evaluation, the freeway operations engineers' response to the algorithm has been very favorable. They claim that they could see an improvement in metering behavior, congestion, and queues at several locations after fuzzy metering implementation.

One of the most noticeable operational differences from the engineers' standpoint is that they no longer need to continually adjust the metering rates. With the Bottleneck and Local Metering Algorithms, the engineers frequently adjusted the minimum and maximum rates to force a particular metering rate. In fact, monitoring and adjusting the metering rates used to fully occupy one engineer. With fuzzy metering, there is no need to baby-sit the rates produced. The rule-based design allows the fuzzy controller to perform well under a wide range of conditions. This feature is vital, because more often than not, incidents, special events, poor data, and unusual weather occur. When properly tuned, the algorithm expertly handles both the recurrent congestion under which it was optimized and nonrecurrent congestion, without any need to modify the control parameters.

From the controller standpoint, incidents are effectively treated like bottlenecks. If the incident is downstream of the ramp meter merge, the metering rate will be reduced to mitigate bottleneck formation, or if there is heavy local congestion from the incident, the metering rate will be low enough to prevent secondary queue formation. If the incident is upstream of the merge, the resulting reduction in mainline congestion both locally and downstream of the ramp meter produces higher metering rates to increase throughput.

One concern of the engineers was that the fuzzy logic controller required on-line tuning for implementation. This is the way in which we escape modeling the system. It turned out that tuning the fuzzy logic controller was a much easier task than anticipated.

Once we determined the system-wide defaults for the control parameters, we were able to use the system defaults for initial deployment at all implementation sites. Then we adjusted the parameters if necessary. It turned out that 80 percent of the 126 ramp meters performed best using the system-wide defaults. Where tuning was necessary, it was to handle special cases, such as poor detection, inadequate ramp storage, and secondary queues. This consistency in control parameters between sites of widely varying geometry and congestion is another indication of the controller's robustness, that is, the ability of the controller to adeptly handle a wide range conditions.

During the process of implementation, observation, and tuning, the ramp metering algorithms were found to be quite sensitive to detector location, and for many ramps, detector placement was poor. The reason that ramp metering algorithms are sensitive to detector placement is that ramp queue occupancy data are of a binary nature. Occupancy is very low (typically less than 8 percent) if the queue has not yet reached the detector, and very high (typically greater than 60 percent) when the queue has reached the detector.

There are two common situations in which the advance queue occupancy is not indicative of the long wait time in the queue: 1) Advance queue detectors that are located at the very entrance to a ramp where a signal is located tend to read very low unless a vehicle is blocking the arterial. A surprising number of drivers are willing to take this risk on the left hand turn movement, but the frequency of occurrence is not consistent. This blocking only takes place immediately after the left hand turn movement. For the remainder of the cycle, the advance queue occupancy reads very low and does not reflect the long queues that continue on the arterial. In this case, the advance queue detector data are misrepresentative, located at the only consistent gap in the queue. 2) Advance queue detectors that are located far beyond where we would like the queue to extend are of limited usefulness. Although this placement of the advance queue detector serves as an ultimatum during excessive queues, for the majority of the time, we could prevent that excessive queue from ever forming if we had intermediate detection.

We can compensate for poor advance queue detector placement by reacting more strongly to the queue detector. However, over reacting to the queue detector may result in a queue that ends before the queue detector. For most ramps, the queue detectors are short of where we want the queue to end, and a strong queue detector response would underutilize the ramp storage between the queue and advance queue detectors. For many long, high-volume ramps, the region of the ramp where we most need detection is nebulous. With a high queue occupancy and a very low advance queue occupancy reading, we may have anywhere from 5 to 40 vehicles in the queue.

An ideal situation is where we have an intermediate queue detector located where we want the ramp queue to end for most of the metering period. This ideal detector would be located far enough downstream of the ramp entrance to allow room for a platoon dump from the arterial signal that feeds the ramp, and to avoid the gap in the queue. This way, when the metering algorithm maintains the queue just short of the intermediate detector, there would still be sufficient room for the platoon dump from the left hand turn of the arterial signal, and the detector data would not be so oscillatory or misrepresentative.

The Fuzzy Logic Ramp Metering Algorithm can be tuned to adequately compensate for poor detector placement. (See "Compensating for Poor Detector Placement" in the training manual, Taylor and Meldrum, 2000.) To some extent though, our ramp queues are only as good as our detector placement and detector data. There are some locations where the addition of intermediate queue detectors on high demand ramps would improve ramp metering performance. Proper placement of intermediate queue loops will give us more preventative control to maintain acceptable ramp queues, allow room for platoon dumps from signals, and reduce oscillation. The Appendix lists ramps that could benefit from better detector placement for more precise queue control.

SUMMARY

This research project had several products. Some of the benefits, such as improving the efficiency of the freeway system, were anticipated at the start of the project. Other benefits, such as improving the state of the TSMC software, were unanticipated spin-off effects.

At the beginning of this project, the TSMC VAX software was not in a user-friendly state. There was little documentation or understanding of pre-existing TSMC software. There was no method for maintaining the proper configuration of the highly specialized and complex TSMC software following a code modification. Severe bugs were hindering previous operation. There was no methodology for making changes to the TSMC software because no major revisions had been made before this project.

With the team work of the WSDOT TSMC software group, these problems were addressed during the implementation of the Fuzzy Logic Ramp Metering Algorithm. This collaborative effort resulted in several software products (details in the software manual by Taylor and Meldrum, 2000):

- documentation of pre-existing TSMC software. (Taylor and Meldrum, 1997a)
- better in-house knowledge of TSMC software
- creation of makefiles and installation of a code manager to maintain proper configuration following a modification
- development of protocol for future modifications of TSMC software
- software integration of the Fuzzy Logic Ramp Metering Algorithm
- software testing of all code
- detection and fixing of bugs that previously hindered TSMC Software
- documentation of new software (both psuedo code and in-code)
- study of methods to process 20-second data, and creation of code to automatically gather specified data

- study of techniques to analyze queue data, and determination of which performance measures are the most meaningful in terms of accuracy and relevant information

- creation and documentation of queue analysis software for 20-second data.

The benefits of software improvements include better operation, more user-friendly software development, new tools for evaluation, and easier software maintenance.

The success of this project required the cooperation of many people, including freeway operations engineers, programmers, system administrators, and hardware maintenance persons. With frequent communication between multiple programmers working on separate projects, we avoided improper software configuration. Sharing of knowledge allowed for thorough software testing and prevented any problems upon initial deployment. With thorough documentation of the algorithm usage and a training workshop for freeway operations engineers, in-house knowledge of the algorithm should remain high despite the turnover rate at WSDOT. Several tasks were completed to coordinate test activities, convey on-line test results, and document knowledge gained about the system:

- presented software design, integration plan, and evaluation plan to WSDOT freeway engineers and software engineers

- presented on-line testing results of the new algorithm to WSDOT freeway engineers and software engineers

- wrote a training manual, complete with many examples on how to handle various situations (Taylor and Meldrum, 2000)

- gave a training workshop for TSMC personnel on how to use, tune, evaluate, and maintain the new algorithm

- presented on-line testing results to the Annual Meeting of the Transportation Research Board and published an article (Taylor and Meldrum, 2000).

The most obvious and far-reaching benefit of this research was the implementation of the algorithm and the resulting improvement in freeway system efficiency. The following tasks were completed:

- implementation of the new algorithm on all 126 metered on-ramps in the greater Seattle area,
- determination of the system-wide parameter defaults of the fuzzy logic controller
- optimization of the Fuzzy Logic Ramp Metering Algorithm at all implementation sites.

On-line testing of ramp metering algorithms is challenging because traffic is not uniform from one day to the next and because performance measures are limited to those that we can accurately measure. Despite the ambiguities of on-line testing, we found on-line testing to be very valuable in comparison to simulation testing. Because of limitations in freeway models, the fluctuations in real traffic data are sharper than those produced by simulations (Taylor and Meldrum, 1995). Because ramp metering can smooth some of that oscillation, the difference between ramp metering algorithms is more pronounced on-line than it is in simulation.

The performance of fuzzy metering was evaluated by comparing it to that of local metering at the I-90 study site and to bottleneck metering at the I-405 study site. To reduce variance in the data set used to compare the metering algorithms, days on which bad weather, special events, or incidents affected the study site were not used in the comparative study. The metering algorithms were also alternated to reduce the effects of seasonal variations in weather and traffic patterns. On the I-90 study site, days on which fuzzy metering took place had lower mainline occupancies, higher throughput volumes, and slightly higher queues than days of local metering. On the I-405 study site, days on which fuzzy metering took place had slightly higher mainline occupancies, slightly higher throughput volumes, and significantly reduced queues. With the combination of similar or better mainline efficiency, and similar or reduced queues, there appears to be a system-

wide benefit to the Fuzzy Logic Ramp Metering Algorithm in terms of improved travel times and higher throughput.

Because we compared fuzzy metering to both the Local Algorithm (which does not use downstream inputs) and the Bottleneck Algorithm (which does use downstream inputs), we were able to assess whether the improvements were due to the downstream inputs or to the nature of the algorithm. Aside from the downstream inputs, the Local Algorithm and the Fuzzy Algorithm use identical detector data. Consequently, it is no surprise that the benefits in mainline efficiency were more pronounced during fuzzy metering. Fuzzy metering was able to prevent the downstream bottleneck from forming at the I-90 study site. Fuzzy metering and bottleneck metering used identical detector data at the I-405 site, so this comparison allowed us to assess the nature of fuzzy logic for the ramp metering application. The fact that bottleneck metering did as well as fuzzy metering in preserving mainline efficiency, while local metering did not, is an argument for including downstream detector data to prevent mainline bottlenecks. Fuzzy metering did a far better job than bottleneck metering at maintaining reasonable queues. This benefit is attributed to the fuzzy logic controller's ability to balance conflicting objectives and to use continuous preventative control rather than threshold activation.

In short, the benefits of the Fuzzy Logic Ramp Metering Algorithm are due to both the inclusion of downstream inputs and to the fuzzy controller's use of smooth, graduated control in a preventative manner. Downstream inputs are essential to achieve system-wide optimization. In regions where public support for ramp metering is a factor of importance, the inclusion of ramp queue inputs is strongly recommended for preventing excessive queues. The fuzzy logic control provides a format for achieving the desired performance, which for highly congested locations, is inherently a balance between multiple, conflicting objectives for the mainline, queue, and secondary queue.

The long-term success of a ramp metering algorithm requires more than decent metering rates. It requires the ability to handle a broad range of conditions aside from

"typical" operation, and it requires the ability to tune the algorithm easily to accommodate changes in the system. Because incidents, missing data, special events, and bad weather are the norm in Seattle, the algorithm's ability to perform well under these situations is a key feature. The Fuzzy Logic Ramp Metering Algorithm is specifically designed to handle practical situations, without the need to modify the control parameters. Because traffic patterns and performance objectives vary from location to location and from year to year, it is important that the algorithm is tunable. With linguistic variables and rule-based logic, the fuzzy logic controller uses a format similar to human reasoning, and for that reason, is easy to tune.

RECOMMENDATIONS

While there are many innovative ideas for ramp metering algorithms, not many algorithms make it all the way to on-line implementation. The key to successful deployment is perseverance, software investment, and frequent communication.

The time and funds required for software development are typically greatly underestimated. For this project, over 90 percent of the budget was spent on software, compared to the 10 percent of the budget needed for design, hardware, implementation, and evaluation. In particular, the software integration was time-consuming: over 77,000 lines of pre-existing C code were modified for the integration. The controller code itself is compact and was relatively effortless to write. When dealing with software integration and testing, particularly for a complex system that does not permit downtime, allow ample funds to develop, test, integrate, and evaluate software in a quality manner. For instance, the Bottleneck Algorithm never worked properly for years after its implementation and was only recently deployed successfully when its 170 code was debugged through this project. With high turnover rates in employees, documentation is important for the long-term success of software applications (see Lessons Learned in Piotrowicz and J. Robinson, 1995).

At the onset of this project, there was a lack of software support that made this implementation more time-consuming. Investment in proper software infrastructure and support is recommended for improving operations and risk management. WSDOT is heading in the right direction regarding software infrastructure. It now has a knowledgeable system administrator for the TSMC VAX. It has greater in-house knowledge of how to make modifications to the TSMC VAX code, PC code, and 170 microprocessors. This is very helpful when fixing bugs, making improvements, or expanding the system. WSDOT has improved its file maintenance, including backups, security, makefiles, and code management software.

Like software infrastructure, the importance of communication cannot be overrated. This implementation required extensive coordination among programmers, system administrators, freeway operations engineers, software maintenance persons, and other researchers. When communication between software engineers did not take place on a daily basis, there were problems with shared resources, incompatible integration schedules for different projects, and software configuration issues. With frequent feedback from the freeway operations and software engineers, the quality of the design was improved. Testing had to be scheduled carefully to avoid affecting other projects and events. Software status needed to be communicated to hardware personnel to prevent incompatibilities with field devices. Correspondence with commuters at the study sites allowed us to fine-tune the metering performance. Progress and results needed to be communicated to managers to build support for the project. Although it seems excessive to send out daily or weekly emails regarding the project status, schedule, and anticipated needs, this was found to be necessary to coordinate activities among all individuals who were affected or who could affect the project.

It is recommended that WSDOT develop a new protocol for determining optimal loop detector placement on metered on-ramps and institute a plan for testing their effectiveness. While WSDOT is a step ahead of most departments of transportation in that it has more than one detector for each on-ramp, still, the importance of detector placement on ramps and its effect on queue control have been underrated. There are wide-spread problems with the current methodology of determining detector placement on metered on-ramps (see System-wide Implementation). For many ramps, the control of ramp queues could be improved with the inclusion of intermediate loop detectors or better placed advance queue detectors. Ramps with problematic detector placement are described in the Appendix.

Even before publication of the test results, we received many requests regarding the applicability of this algorithm to other regions. The concepts behind this algorithm are

certainly transferable, but the algorithm may need modification depending on detector types, detector placement, sampling frequency, and system objectives. Although the controller code itself is relatively simple, the interface between the system software, controller, field devices, and user interface may need considerable customizing. Successful implementation requires knowledge of the site specifics, with controller inputs determined as described in the training manual (Taylor and Meldrum, 2000). Regions that will see the most benefit from this type of logic are those that have over saturation, ramp queue detection, and the need to balance mainline objectives with queue objectives.

We anticipated that during the course of on-line testing we would discover further improvements that would be beneficial, such as a global hierarchical controller. Given the detector data available at this time, it does not appear that this modification would have substantial benefit. Because there is no limit to how far downstream a ramp meter can "look," the ramp metering algorithm is as systemic as we want to make it regarding the mainline. We found that including upstream inputs degraded the performance of the ramp meters (see Design Changes in the training manual), particularly during incident handling.

We expect that there is some benefit to using arterial data at some ramps, and the best means of doing this would require further study. Anticipation of a large platoon would allow the controller to provide more room in the queue. However, we do not want to be so queue responsive that we no longer smooth the turbulence of platoons impacting the mainline. The graduated, preventative control of the fuzzy logic controller accomplishes some of the same goals that the inclusion of arterial inputs would do, that is, preventing an excessive queue from forming in the first place (as long as there is not a secondary queue) and allowing room for the next platoon. We already know which ramps commonly have platoons that nearly exceed the storage capacity of the ramp, and we have tuned the controller to accommodate these platoons to their best ability. Another goal of including arterial data would be to better handle incidents. To some extent, the queue

variables that we currently use accomplishes this because these variables indirectly take into account incidents, weather, demand, and special events.

Even with the foresight provided through arterial data, we still must work within the confines of available ramp queue storage and reasonable ramp queue delay: Over saturation (both on the ramps and mainline) is often what dictates the metering rates in Seattle. Metropolitan areas that are not over saturated would be able to utilize arterial data to a greater extent and with more noticeable benefits. However, the first hurdle for both unsaturated and over saturated facilities is adequate detection on the ramps and on the arterial segments devoted to queue storage, which we expect would confer noticeable benefit.

While the inclusion of arterial data would certainly have some advantages, we do not expect it to radically improve the ramp metering performance for Seattle, except in the instances where arterial segments are devoted to ramp metering queue storage. The usefulness of arterial data applied to ramp metering in Seattle is tempered by the need to smooth turbulence, the adequacy of the Fuzzy Logic Algorithm in preventing problems and responding to atypical conditions, and the limitations of our over saturated facility.

ACKNOWLEDGMENTS

REFERENCES

J. Ishimaru and M. Hallenbeck, 1999. "Flow Evaluation Design," Technical Report, WA-RD 466.2.

L. Jacobson, K. Henry, and O. Mehyar, 1988. "Real-Time Metering Algorithm for Centralized Control," *Transportation Research Record* 1232, National Research Council, Washington, D.C., pp. 17-26.

N. Peirce, 24 January 1999. "The more roads we build, the more traffic we create," Editorials, Seattle Times.

G. Piotrowicz and J. Robinson, 1995. "Ramp Metering Status in North," Office of Traffic Operations, Federal Highway Administration, U. S. Department of Transportation, Washington, D.C.

E. Pryne, 29 June 1997. "Front Porch Forum," Local News, Seattle Times.

E. Pryne, 17 May 1998. "Front Porch Forum: Growth – Enough Already?" Local news, Seattle Times.

C. Taylor and D. Meldrum, 1997a. "Documentation of TSMC Software that Interfaces with Traffic Analysis Programs," Final Technical Report. Washington State Department of Transportation, National Technical Information Service, WA-RD 442.2.

C. Taylor and D. Meldrum, 2000. "Algorithm Design, User Interface, and Optimization Procedure for a Fuzzy Logic Ramp Metering Algorithm: A Training Manual for Freeway Operations Engineers," WA-RD Technical Report to be published, Washington State Department of Transportation, National Technical Information Service.

C. Taylor and D. Meldrum, 1997b. "On-line Implementation of a Fuzzy Neural Ramp Metering Algorithm," Final Technical Report. Washington State Department of Transportation, National Technical Information Service, WA-RD 442.1.

C. Taylor and D. Meldrum, 2000. "A Programmer's Guide to the Fuzzy Logic Ramp Metering Algorithm: Software Design, Integration, Testing, and Evaluation," WA-RD Technical Report to be published, Washington State Department of Transportation, National Technical Information Service.

C. Taylor and D. Meldrum, 1995. "Simulation Testing of a Fuzzy Neural Ramp Metering Algorithm," Final Technical Report. Washington State Department of Transportation, National Technical Information Service, WA-RD 395.1.

P. Whitely, 1 January 1999. "Buried in Traffic? There's more cars on the Road," Local News, Seattle Times.

APPENDIX: Detector Placement Problems

164th To I-5 SB   The advance queue detection at this ramp is inadequate. It is located just below the signal on the ramp, which is rarely occupied unless a driver turning left blocks the intersection. It reads low, oscillatory occupancies, which do not reflect the enormous queue that continues onto 164th. We need an intermediate queue detector just downstream of the advance queue detector, which will read more accurately, consistently, and provide more preventative control. We also need to add advance queue detection on the arterial itself for devoted right hand turns and left hand turns. This ramp is scheduled for a redesign, where a second metered lane will be added. Although this will double the storage, the ramp will still utilize all of its storage and still need better advance queue detection.

NE 45th to I-5 SB Queue detection here is poor. There is no advance queue to indicate that the queue continues far along 45th in either direction. The one detector is located at the entrance of the ramp, but so far to the left side that vehicles turning right do not activate it. The only vehicles that activate it are the ones turning left. Only when drivers turning left *block* the intersection does the detector go high. While this problem would be urgent on most ramps, the Ship Canal is such a significant bottleneck that we do not want to be very responsive to the queue at 45th. If the ramp storage capacity, queue detection, and queue response at this ramp are increased, more vehicles will use it instead of NE 50th, which will make the mainline much worse. For the option of a queue response when the mainline is not so bad, a queue detector added just downstream of the current detector is needed. Advance queue detectors on the arterial are needed for both the right hand turn and left hand turn.

NE 50th to I-5 SB There is no advance queue detection. It should be added, particularly to the right hand turn on the arterial.

128th to I-5 SB   The advance queue detectors are located right at the entrance of the ramp just below a signal. Consequently, the detector tends to read low, oscillatory occupancies that misrepresent the large queues. Intermediate queue detection between the queue and advance queue detectors should be added for smoother, more preventative control based on more representative occupancy data.

SE 8th This queue continues far past the detectors onto the arterials. The merge is particularly congested here because of the mainline curve and tunnel entrance, which prevents faster metering. Advance queue detectors on the arterial would be helpful.

NE 4th/8th to I-405 SB   This ramp has a significant safety hazard. The vehicles exiting must cross over this ramp queue when it is extensive, which occurs all too frequently. The drivers exiting often speed off the exit, and vehicles frequently block their path. To use the substantial storage (which is needed) beyond the dangerous crossover weave, the exit would have to be redesigned so it does not cross over the on-ramp queue. Another problem is poor placement of the advance queue detectors. In this case, both the queue detectors and advance queue detectors are located too close to the stop bar, reducing the usage of the available storage. To remedy this, the advance queue detectors could be renamed to be intermediate queue detectors, while advance queue detectors could be added farther back. If the exit is not redesigned, these new detectors should be placed downstream of the dangerous crossover weave. Otherwise, the new detectors should be placed near the beginning of the on-ramp to utilize the significant additional storage.

WB NE 8th to I-405 NB   This advance queue detector's placement is somewhat poor. The occupancy tends to read low and oscillatory, despite a queue beyond the detector. It would be helpful to add an intermediate queue detector halfway between the queue and advance detector for more accurate occupancy readings and more preventative control.

Swamp Creek to I-5 SB   Although storage here is great, intermediate queue detectors are definitely needed here. The advance queue detector is so far back that the

wait times are much longer than we would like even though the queue does not reach the advance detectors. We highly recommend adding intermediate detectors here for smoother control.

Eastgate to I-90 WB The advance queue detectors are located farther back than where the queue should end. Intermediate queue detectors would provide smoother control.

205th lanes 1 and 2, and 236th, lane 1 to I-5 SB With the combination of very low metering rates (to reduce the secondary queue at the merge) and advance queue detection, which is quite far back, it would be helpful to have intermediate queue detectors for better management of wait times here. The advance queue detection is not adequate representation of the wait times at these ramps because the metering rate is so low.

SR-516 to SB SR-167 The queue detector is too close to the arterial, and a strong response to this detector would underutilize the storage of the ramp. For this reason, we depend entirely on the advance queue detector for control purposes. The advance queue detector placement is very useful, but it should be renamed the intermediate detector, while advance queue detectors should be added to the arterial. This ramp has the combination of a very difficult merge and high demand, so better detection would make it easier to maintain the proper balance between secondary queue prevention and excessive queue prevention.