

Total-cost frameworks for multimodal and environmentally preferable streets

A Dissertation Presented to
the faculty of the School of Engineering and Applied Science
University of Virginia

In Partial Fulfillment
of the requirements for the degree
Doctor of Philosophy

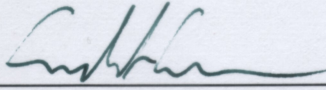
by

Conrad A. Gosse

December 2013

APPROVAL SHEET

The dissertation is submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy



Conrad A. Gosse (Author)

This dissertation has been read and approved by the examining committee:

Dr. Andres Clarens (Dissertation Advisor)

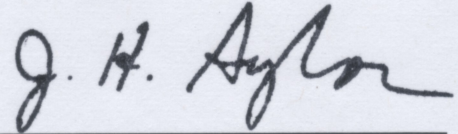
Dr. Brian Smith (Committee Chair)

Dr. Johnathan Goodall

Dr. Randy Cogill

Dr. John Miller

Accepted for the School of Engineering and Applied Science:



Dean, School of Engineering
and Applied Science

December 2013

ABSTRACT

Efforts to reduce the environmental impacts of transportation have often overlooked efficiencies obtained by considering the relevant engineering and economic aspects of roadway infrastructure as a system. Concerns over greenhouse gas emissions (GHGs), degrading infrastructure in the face of limited maintenance budgets, and the declining use of private automobiles all motivate this more comprehensive approach investment prioritization. Here, a three-part framework is presented consisting of a multi-segment discrete pavement management system (PMS), a roadway use-phase microsimulation that evaluates lane allocation between modes, and a spatial-Bayesian bicycle volume model to synthesize necessary but generally unknown street-specific bicycle usage from available data. A more complete assessment of the costs and benefits of various road-use scenarios is quantified with respect to user costs, agency costs, and GHG emissions.

The PMS presented here incorporates GHG emissions using a multi-objective evolutionary algorithm to produce a Pareto-set of discrete long-term maintenance plans. By using deterioration models and cost estimates from the Virginia Department of Transportation (VDOT), comparisons with historical practice are also possible. VDOT has historically relied on corrective maintenance in validation area, but an optimized management plan could achieve the same average pavement condition with 60% of the cost and 50% of the GHG emissions. Solutions from a network-wide optimization also dominated aggregated single-segment solutions, justifying the computational cost of the method. The use phase still accounts for the majority of roadway impacts, however, so the second component of this work employs microsimulation and a vehicle

emissions model to consider the coupled effects of pavement condition and vehicle fuel consumption, as well as travel time costs, by mode.

Both microsimulation and a probabilistic analysis of all feasible combinations of travel lanes, bicycle lanes, and curb parking show mobility reductions on road segments of insufficient width for heavy vehicles to pass bicycles without encroaching on oncoming traffic. This delay is positively correlated with uphill grades and increasing traffic volumes and is inversely proportional to total pavement width. A high bicycle mode share is therefore negatively correlated with total costs and emissions for lane configurations allowing motor-vehicles to safely pass bicycles, while the opposite is true for configurations that inhibit passing. As a result, curb parking exhibits spatial opportunity costs well in excess of feasible hourly use fees when the parking lane could have been devoted to bicycle mobility, even before considering safety benefits often used to justify such conversions. The results are sensitive to street-specific bicycle mode share, however, and these data are not commonly known without dedicated field observation, thus precluding network-wide analysis.

The final component of this work employs Markov-chain Monte Carlo sampling to address the temporal factoring of bicycle count observations into annually representative posterior distributions of critical parameters on direction roadway links that have been sampled. A novel spatial factoring method employs Bayesian updating to combine uncertain volume estimates from a regional travel demand model with the temporally factored intermediate distributions by applying a stochastic edge correlation matrix. For a small city in the United States with some volunteer bicycle counts and no permanent counting infrastructure, the model is able to estimate edge-specific bicycle usage network-wide with large but well-characterized uncertainty.

Overall, the results provide quantitative evidence that efforts to reallocate limited pavement space to bicycles, like those being adopted in several US cities, could appreciably reduce costs for all users and illustrate the value of a total cost approach

to investment optimization in evaluating these decisions. Future work will continue to address the data gap between bicycle and motorized transportation with respect to travel demand and safety.

© Conrad A. Gosse 2013

All Rights Reserved

For Carolyn

ACKNOWLEDGEMENTS

I would like to acknowledge the continual support and sage counsel of my advisor, Prof Andres Clarens, my fellow students, and the faculty of the Civil and Environmental Engineering Department. Many other individuals have contributed to this work and are acknowledged at chapter conclusions, as appropriate.

This work was developed under STAR Fellowship Assistance Agreement no. FP917275 awarded by the U.S. Environmental Protection Agency (EPA). It has not been formally reviewed by EPA. The views expressed are solely those of the authors, and EPA does not endorse any products or commercial services mentioned.

TABLE OF CONTENTS

ABSTRACT	ii
DEDICATION	v
ACKNOWLEDGEMENTS	vi
TABLE OF CONTENTS	vii
LIST OF FIGURES	x
LIST OF TABLES	xiii
CHAPTER	
I. Introduction	1
1.1 Background	4
1.1.1 Environmental Costs	6
1.1.2 Conventional User Costs	9
1.1.3 Health and Safety	10
1.2 Problem Statement	12
1.3 Research Objectives	14
1.3.1 Tool 1	15
1.3.2 Tool 2	15
1.3.3 Tool 3	16
II. Pavement Management	18
2.1 Introduction	18
2.1.1 Algorithm	20
2.1.2 LCA in Pavement Management	22
2.2 Method	25
2.2.1 Pavement Condition	25
2.2.2 Deterioration Model	28
2.2.3 Global warming potential	30

2.2.4	Study Segments	32
2.2.5	Algorithm	35
2.3	Results	37
2.4	Conclusion	41
2.4.1	Acknowledgements	42
2.5	Supporting Information	43
III.	Mobility and Shared Right of Way	47
3.1	Introduction	47
3.2	Method	51
3.3	Results	56
3.4	Sensitivity and Additional Considerations	61
3.5	Implications	63
IV.	Bicycle Volume Model	66
4.1	Introduction	66
4.2	Method	69
4.2.1	Network and Routing	70
4.2.2	Prior AADB	72
4.2.3	Temporal Factoring	73
4.2.4	Spatial Factoring	77
4.3	Results	80
4.4	Implications	84
4.5	Acknowledgements	85
V.	Conclusions and Future Work	86
5.1	MPO Travel Model Integration	88
5.1.1	Non-static mode preference	91
5.1.2	Modeling Additional Modes	92
5.2	Video Data Collection	93
5.2.1	Bicycle crash risk and facility design	94
5.2.2	Continuous counts	98
5.2.3	Dispersed sensing	100
5.3	Optimizing Investment	101
APPENDICES	104
A.	Supporting Information for Pavement Management	104
A.1	Paving LCA Details	104
B.	Justification for Network-wide Pavement Management	108
B.1	Introduction	108
B.1.1	Proposed Framework	110
B.2	Method	111

B.2.1	Modeling Horizon	114
B.3	Results	115
B.4	Conclusions	119
B.5	Acknowledgements	120
C.	Supporting Information for Mobility	121
C.1	Traffic Microsimulation	121
C.2	Vehicle Emissions	122
C.3	Probabilistic Delay	123
C.4	Time and Discount Rate	130
C.5	Pavement Management	131
D.	Publications resulting from this work	136
BIBLIOGRAPHY		137

LIST OF FIGURES

Figure

2.1	Conceptual diagram of the evolution of pavement management from unmanaged to the approach suggested in this work where PMS = Pavement Management System and EPMS = Environmentally Adapted Pavement Management System.	23
2.2	A conceptual diagram of the three most common VDOT treatment intensities in asphalt pavement maintenance. Reconstruction (RC) is not depicted. Data from [38].	26
2.3	Unconstrained maintenance triggers for VDOT Interstate pavements. Data from [38].	27
2.4	VDOT load-related pavement deterioration model for bituminous interstate pavements according to the time elapsed since the last major rehabilitation. Coefficients for this model are given in Table 2.2.	30
2.5	Pareto surface of 15 year solutions for a single segment with $CCI_o = 85$ and initial age of 7 years. Solution id's refer to treatment plans in Table 2.5.	35
2.6	Three sample 15 year plans showing the <i>cumulative</i> performance, GHG, and economic implications of representative treatment plans for a single 1.6km segment with $CCI_o = 85$ and an initial age of 7 years.	44
2.7	Results of model output for 10 years illustrate relationship between CCI, cost, and GHG emissions (large dots). Projections of results between pairs of variables are also shown (small dots).	45
2.8	Pareto front curves for several planning periods. The 3 year solutions are based on VDOT practice from 2007-2010. Arrows represent the improvements that could be achieved using the optimization procedure proposed here.	46
3.1	A reference multi-modal street section and the relative widths of vehicles and lanes.	53
3.2	Agency pavement costs vs. user fuel costs as a result of pavement roughness (upper) and combined costs vs. combined GHG emissions (lower) from non-dominated pavement management plans. The triangle data point indicates the selected plan by minimum GHG. .	55

3.3	Simulated motorized travel time illustrates the delay caused when heavy vehicles are unable to pass bicycles.	58
3.4	The relationship between road width and costs follows a Pareto optimal behavior, shown here for 100m characteristic lengths. Total travel time costs (left) dominate fuel and PMS costs, however when delay costs are used (right) increased cycle mode share can reduce costs for appropriate lane configurations.	59
3.5	Spatial opportunity cost for curb parking, shown here for 100m characteristic length and 4% grade, increases dramatically for narrower roads suggesting that when all the costs are considered, there is a tipping point beyond which curb parking becomes an expensive use of land.	61
4.1	Principle data sources and processing steps.	70
4.2	Hourly factor curve illustrating Equations 4.3b, 4.4, and 4.5 for a directional link with $p_{m_j} = 0.27$	74
4.3	Relative likelihood to bicycle (F_t) vs. effective temperature (T_e) using MCMC-sampled parameter values. The temperature factor F_t is parameterized by effective temperatures T_{mc} and T_m as defined by Equation 4.9.	78
4.4	Network subset showing reference locations A and B. The width of the dark road overlay indicates relative correlation with trips traveling to the left at location A. Circles show nearby count locations scaled by the amount of data.	79
4.5	Posterior distribution for the mean of the morning and evening commute peaks. Kernel density represents a continuous histogram of parameter distribution.	80
4.6	Posterior distribution for beta-distributed global parameters	81
4.7	Hourly aggregated observations and modeled results for a single day at location A (shown in Figure 4.4)	82
4.8	Prior ($AADB_{j_o}$) and posterior ($AADB_j$) distributions and a conventionally factored estimate for location A (see Figure 4.4), a directional link with two weeks of continuous counts.	83
4.9	Prior ($AADB_{j_o}$) and posterior ($AADB_j$) distributions for location B (see Figure 4.4), a site 1km from location A, but that has no data. .	84
5.1	Bidirectional average annual daily bicycles (AADB) where wider lines represent lower variance in the final volume estimate and circles indicate data collection points.	90
5.2	A rendering of the visible portions of a Charlottesville, VA intersection fitted with detection cameras covering all four approach legs. . . .	99
B.1	A general framework for LCA in asset management of networks of public infrastructure. M&R = maintenance and rehabilitation. . . .	111
B.2	Condition trajectories for each segment for a reference plan. Large improvements represent CM, smaller improvements PM.	116
B.3	Non-dominated network solutions comprising the PF. The circled plan is noted in Figures B.2, B.4, and B.5.	116

B.4	Individual PFs for each segment with the network PF normalized by the number of segments. The components of the reference plan from the previous figures are circled.	117
B.5	The network-optimal PF and reference plan from Figures B.2 - B.4 in contrast to two solutions derived from examining the per-segment PFs in Figure B.4.	118
B.6	Per-segment and network-wide relationship between GHG emissions and user cost. Per-segment symbols from Figure B.4.	119
C.1	Predicted travel time aggregated by person trips compared with simulation results	134
C.2	Distribution of stuck and unstuck lane configurations	135

LIST OF TABLES

Table

2.1	Selected reference treatments and VDOT cost estimates for bituminous interstate pavements.	27
2.2	Coefficients of VDOT load-related deterioration model given in Eq. 2.3 for bituminous pavements [162].	29
2.3	Material specifics of the two mixes used by the reference treatments. All percentages by mass.	31
2.4	Summary results of the LCA of reference treatments assuming a 3.66m lane width.	32
2.5	Non-dominated treatment plans for a single segment over 15 years with $CCI_o = 85$. DN not shown for clarity. The id column is an arbitrary identifier that links the treatment plans given here with Figs. 2.5 and 2.6, with boldface rows shown specifically in Fig. 2.6.	34
3.1	Discrete segment parameter space. All unique combinations were evaluated using microsimulation. The first group defines a lane configuration and the second a scenario.	51
3.2	Common parameters and descriptive values for all simulated configurations.	52
4.1	Network-Wide Variables	70
4.2	Directional Link-Specific Variables	71
4.3	Variables Applicable to Count Observations	73
A.1	Model inputs for the LCA developed to estimate the GWP associated with pavement maintenance activities.	106
A.2	Summary totals for each plan for the reference single site.	107

CHAPTER I

Introduction

Electric streetcars enabled the first suburban expansion in the U.S. but eventually fell into disrepair and were abandoned. Their construction was funded by the profits of converting rural and previously inaccessible land to urban use, so when these speculative-turned-critical transportation systems were in need of reconstruction, the land development companies that built them had rewarded their investors and were gone [130]. Private automobiles and the post WWII highway building boom subsequently opened spaces between radial streetcar suburbs to a second larger wave of residential development. These highways were built with public funds, however the source of the funds was ultimately the same. The possibility of new suburban development drove up land values [39], and the tax base of the formerly rural areas expanded rapidly thanks to the new construction enabled by government mortgage subsidies [161]. Fuel tax receipts used to fund highway infrastructure also grew steadily in response to the additional driving induced by these more automobile-oriented developments. Nearly a century after the first federal highway bill was passed in 1916 to “get the farmer out of the mud,” and a half century after the sweeping 1956 highway act that created the interstate system, these funding mechanisms are failing, however, and even basic maintenance is now often deferred [182]. Induced congestion has also illustrated the practical limits of auto-dependant expansion as

traffic delays increasingly counterbalance the economic contributions transportation infrastructure should provide. Fortunately, either as a direct result of these issues or for other reasons, travel demand is now shifting in foundational ways [141] that present new opportunities to transportation planners.

Young adults are not adopting the private automobile as their baby-boomer parents did [106], and the century-long upward trend in vehicle miles of travel (VMT) per capita may have peaked [119]. This shift is further supported by recent dramatic increases in bicycling and public transit use [106] with commuting by bicycle in the 70 largest US cities up more than 50% nationally between 2005 and 2012 [172, 111]. The work to be presented here is predicated on the idea that these changes represent a desire for a sustainable livability, i.e. a consistently high quality of life over the long term, where quality of life assumes a productive environment that can support healthy people and communities and a healthy economy to provide personal amenities. Greater adoption of active transportation modes is a primary means to achieve these objectives by its inherent physical activity and its tendency to build social capital [144] and reduce the financial and environmental burdens of transportation. Realizing this shift will require significant planning over the long term. Automobile use has so displaced other forms of travel in the United States that it will take decades to reestablish a diversity of modes. Highways and other subsidies to automobile use have encouraged sprawling land use patterns that are challenging to serve by any other mode. Ironically, bicycles are better suited to many of these places than motorized public transit since they enjoy many of the freedoms and autonomy of private automobiles without the congestion or parking problems and avoid the high subsidies that would be required to serve low-density residential areas with fixed-route public transit. Travel modes function in the context of networks, however, which means that even a large investment in bicycle facilities may prove ineffective if barriers such as high-speed arterial roadways continue to separate bikeable sections of the network. As density

increases, different modes compete not only for funding, but for right of way as well, further complicating the picture.

Decades of auto-centric investment have created urban environments hostile to people not in cars and correcting this imbalance requires thoughtful and continuous planning over the long term, since the full utility of individual walking and cycling infrastructure will only be realized once coherent networks are once more in place. Bridges are deemed “functionally obsolete” when traffic volumes exceed capacity or for safety reasons [12]. The term could also be applied to maintenance-intensive auto-oriented facilities that function as barriers to more environmentally and fiscally responsible modes of using the street, such as bicycling and walking. The growing concern with [91] and momentum for regulation of [178] greenhouse gas emissions (GHGs) can be viewed as a direct response to our desire to ensure livability over the long term by maintaining the ecological systems to which we have adapted. Since transportation emissions represent 23% of CO₂ emissions globally [89] and 28% in the U.S. [177], lowering the carbon footprint of transportation clearly supports this goal at a global level, however, complex urban-scale changes are necessary to realize these benefits. This shift can only occur as part of a larger shift in expectations of the role and use of roadway infrastructure which presents a modeling task beyond the scope of the present work. The frameworks that will be presented can be viewed as components of such a hypothetical complete model and were selected for study in order to incorporate new metrics like GHG emissions into existing processes like pavement management or to synthesize data for further analysis where no suitable data were available, as is often found when investigating bicycle transport [128]. The sections to follow provide more background information on the need for and state of these metrics and data sets.

1.1 Background

An infrastructure maintenance backlog is not unforeseen. All structures have a finite life, and it is more challenging politically to fund maintenance than bold new projects [142]. Our present situation can also be seen as a reflection of the idea that a transportation investment represents an asset in perpetuity that will not at some point become a liability. Major new bridges can illustrate this idea. A typical cost-benefit analysis would treat the bridge as a factory might a major piece of equipment [118]. A proposed bridge would be estimated to provide a certain vehicle capacity, thereby reducing trip distances and congestion on existing routes. These benefits would be monetized through reduced user costs and considered against amortized capital costs and ongoing maintenance expenditures. When the bridge is first opened, drivers may choose to use it, or not, though many presumably will or it would not have been seen as beneficial. The difficulty comes when maintenance costs become so large that the facility must be replaced. Replacing a now-critical facility under heavy daily use is a much more expensive than new construction of the same facility due to the more constrained construction footprint and the safety and delay costs associated with a prolonged work zone. The replacement of a “successful” bridge must also have a greater capacity than the original to satisfy the demand induced by the original, further expanding costs. Often a parallel bridge will have been built to meet demand long before the original needs replacement, but the aggregate effect is the same.

These costs are not typically included in the original project estimation, though this makes little practical difference. Their net present value after applying typical discount rates over the 50-100 year expected life of the structure would be so small that there would be no reason to invest the time necessary to estimate them. Discounting future cash flows at a rate exceeding inflation means that only the immediate capital costs and user benefits of the bridge are well represented. The early years of operation also tend to show the largest benefits since the new facility reduces congestion but

has not yet induced additional travel. If the factory equipment analogy held, however, planners could evaluate whether the facility should be retired or replaced at the end of its useful life according to conditions at that time.

A major bridge undergoes a transformation that factory equipment does not, however, it becomes critical infrastructure. Transportation links are not geographically fungible and distort connected land uses over time such that they start as amenities but quickly become indispensable. It is politically improbable that a facility will be retired because the user costs of such a move are large. When the facility was first introduced, travel demand gradually evolved to make use of it, but presumably in a user-beneficial way throughout. Closing the bridge does not allow this type of gradual realignment, even if there are clear long-term benefits to doing so. As a result, this strategy is only readily applied to small links with numerous alternative routes [186]. Taken together, these examples suggest that even though transportation infrastructure is designed to maximize benefit to system users, over time it can become an expensive liability for the managing agency which will be left with few options but to maintain the infrastructure, even at suboptimal levels of repair.

When these collected expenses grow to the point of consuming all available funds, there is no money left for alternative investment, and a city is stuck with a deteriorating status quo. Unfortunately this hypothetical situation has largely come to pass [160]. A spatial growth rate well in excess of population growth is almost a definition of suburban expansion itself, largely since spatial expansion on this scale can only be realistically supported by private automobiles. Now that growth in auto-only travel has slowed or stopped in the U.S.[119], it stands to reason that spatial expansion will similarly slow, and therefore future investments in urban infrastructure will occur within already developed areas. While performance-based contracts that include ongoing maintenance and public-private partnerships that exploit novel funding models will somewhat address funding problems in specific cases, much of transportation

planning in the foreseeable future will be a task of repurposing and reconfiguring existing uses.

Urban renewal in the post WWII period produced a lasting wariness over the use of eminent domain to obtain land for new roads [109]. This hesitance, coupled with the high cost of urban land from even willing sellers and the tendency for transportation demand to follow existing corridors, means that most new transportation investment will take place within existing public rights of way e.g. streets. The present degree of automobile dominance means that these reconfigurations will have to address complicated trade-offs between competing travel modes as well as the more fundamental question of how much space should be devoted to mobility vs. other uses. Quantitatively answering these questions demands that facility (re)construction, maintenance, and use be considered in a single multi-objective optimization framework that is able to marry conventional user and agency costs with environmental and liveability measures to capture the full costs and benefits of transportation facilities.

1.1.1 Environmental Costs

Greenhouse gas emissions (GHGs) are perhaps the principal environmental challenge facing transportation planners. Unlike smog, the precursors of which could be addressed through pollution control devices on vehicles, CO₂ is a principal byproduct of the combustion of fossil fuels and cannot be feasibly captured on board vehicles at this time [165]. Reduction of this pollutant is then a question of either shifting its production to a central facility where it could be captured and sequestered, which would require widespread adoption of electric cars, or simply using less energy in vehicles.

Energy intensity of trips can be framed as a combination of trip distance and specific efficiency along the chosen route. Specific efficiency is in turn a function of travel mode (vehicle technology) and operating conditions. Clearly all modes benefit

with respect to energy use from shorter trips, however there are practical length minimums depending on the mode. Trips by public transit will be at least the distance between two stops, and there is always a trade-off between average speed along a route and the density of stops. [168] In the case of automobiles, parking requirements naturally spread out destinations that accommodate them, and the tailpipe emissions per distance of internal combustion autos are much worse for short trips since more of the trip will occur before the engine is warm and operating at peak efficiency [159]. Shorter trips also generally imply a greater density of destinations, however the slower urban streets associated with these places lead to greater specific emissions as vehicles are re-accelerated from each slowdown and idled at stops. In essence, automobiles operate most efficiently on infrastructure associated with longer more inefficient trips and vice versa. Given two trips that serve the same purpose, the longer journey does not create more economic value than the shorter one other than the marginal increase in resource consumption of the vehicle itself. Meaningfully reducing GHG emissions while maintaining economic output therefore presents the challenge of reducing average trip lengths in spite of the limited policy options that influence individual destination choice.

Geographically constrained regions may be able to introduce additional network connectivity, such as a bridge, to reduce trip lengths between existing origins and destinations, but region-wide reductions require more fundamental changes in land use by encouraging a fine mix of use types that distribute particular services and amenities more evenly throughout the region. Dense mixed-use development is not conducive to the efficient movement of motor-vehicles, however, so while such a shift would reduce VMT, if it also introduced more congestion the net reduction in GHGs would be unclear. Fortunately, as development patterns shift to allow shorter trips, additional transportation modes become viable and even preferable alternatives. Walking, bicycling, and public transit are well suited to higher densities since they

are able to move more people with the same infrastructure footprint. For example, a typical 3m wide travel lane has a bicycle-only saturation flow rate of 6,000-7,500 per hour in an urban setting [146], compared with less than 1,000 automobiles per hour using an identical lane [168]. Once parking requirements are included the calculation tips further in favor of the bicycles. Despite this advantage, competition for right of way between the various modes in dense urban neighborhoods is intense, so an objective means of managing this space, including as many categories of costs and benefits as possible, is necessary in order to maximize the net benefit returned by the public right of way. Pavement condition affects vehicle operation and running costs, for example, so while maintenance and mobility are typically part of separate analyses, they need to be considered together to obtain a complete view of a facility.

Differences in rolling resistance exhibited by various pavement types can be quite significant in the aggregate [187, 99], and the emissions associated with maintaining roadways can be as much as 10% of the use-phase totals [151]. While these impacts are still relatively small, they have the distinct advantage of being directly influenced by pavement design and maintenance, unlike user-dependant factors. Pavements also exhibit accelerating rates of deterioration as their condition worsens, such that either periodic maintenance or reversion to an unpaved surface are necessary [137]. In the near term, the urban public is unlikely to accept un-paved streets. Since many cities will not have funds available for any infrastructure projects beyond such basic activity, street reconfigurations that can be efficiently implemented within the context of a paving project, such as converting auto lanes to bus and or bicycle use, show particular promise. Combining the objectives of reconfiguration and basic maintenance will help with very real agency funding challenges, however the net impact of these interventions must still be considered with respect to more typical use-phase metrics, such as travel time and delay.

1.1.2 Conventional User Costs

No aggregate metric can capture the total costs and benefits of competing transportation alternatives since individual stakeholders have diverse priorities, however a handful of measures can capture a great deal more than is typically included in transportation planning. By computing a Pareto surface of non-dominated plans with respect to each metric, decision makers can at least choose from a rich set of sensible alternatives. Defining and estimating these metrics remains far from trivial, however.

Chief among these difficulties are dealing with sunk costs and other inertial forces in personal travel behavior. On a given day, a personal automobile depreciates according to age and insurance premiums regardless of whether the vehicle is driven or not. These sunk costs clearly bias the owner of the vehicle toward use of the automobile since they can dwarf the fuel and maintenance costs incrementally associated with driving. Ownership of an automobile is fixed over short time scales, however it is ultimately sensitive to the availability of other means of transportation. Rather than attempt to model this shift, however, the present work will presume certain mode shares in order to consider the costs and benefits of encouraging shifts in their ratios and so will maintain a more traditional focus on incremental user costs as opposed to a full incorporation of sunk costs and the long-term behavioral modeling necessary to quantify them.

Fuel is the principal incremental cost associated with operating a motor vehicle in the U.S despite low fuel prices [15]. GHGs are also closely correlated with vehicle fuel use. Operating costs are not a perfect proxy for estimating GHG emissions, however, since both fuel use and vehicle maintenance costs increase as pavement condition deteriorates. GHG-intensive pavement maintenance is required to achieve a firm smooth surface which in turn reduces vehicle emissions, so while the correlation between GHG emissions and fuel consumption does narrow the optimal multi-objective design space, it does not simplify the problem to a single optimal point. Ultimately

however, all vehicle-related costs are small in comparison to direct user costs of travel in the form of lost productive time.

Driver’s time is typically valued at half the prevailing wage rate [7]. Since fuel costs are so low in the U.S., and the median hourly wage in 2012 was \$16.71 [30], the time cost of traffic congestion tends to dominate incremental costs, which in turn makes congestion mitigation the central component of typical highway planning. People also experience waiting time, whether in traffic or for a bus, as twice the duration of time spent in motion [121], which further encourages reactive congestion minimization. This approach has the well known outcome of continual highway expansion, though this would not necessarily be the case if additional costs were included. One obvious but difficult to implement idea is to minimize total travel time for all trips by favoring policies that lead to shorter trips. A more holistic approach would be to recognize and incorporate the large and typically externalized health and safety costs associated with transportation with a focus on the differences between the modes.

1.1.3 Health and Safety

Government at all levels in the U.S. presently spends more than four times as much on healthcare as transportation [36], and this still represents only half the total healthcare spending in this country [169]. The other half is contributed by the private sector. With both the total and the government’s share of this expenditure expected to rise in coming years [169], the present transportation funding shortfall can be expected to worsen with time. Clearly some of these costs are endemic to the health system itself and outside the scope of this study, however this spending can also be linked to the present epidemic of overweight and associated chronic diseases which can in turn be traced to the lack of exercise associated with automobile-dependant lifestyles [14]. The projected incremental increase in health care spending in the U.S. due to obesity alone is projected to be between \$22-28 billion annually in 2020 and \$48-66 billion by

2030 [180], though this does not include the lost productivity attributable to obesity which represents an economic burden several times larger [60]. Since this combined 2020 total is perhaps half the total cost of all traffic congestion nation-wide forecast for the same year (\$199billion [156]), investments supporting active transportation that have the potential to improve population fitness while also relieving traffic congestion make clear economic sense.

The area of health and safety also illustrates some of the sharpest distinctions and conflicts between the various modes. Active transportation, such as riding a bicycle or using public transit, is beneficial from the point of view of daily exercise [43, 115]. Bicycles are particularly suited to introducing significant levels of activity in daily routines when substituted for automobiles. Unlike both walking and public transit, bicycles are also able to negotiate the longer trip distances to dispersed locations associated with motor vehicle-dominated land use patterns. Cyclists are also at greater risk of injury than the automobile users sharing the same roadways. These risks, together with noise and airborne particulate exposure, represent negative externalities of automobile use that should be incorporated in any comprehensive planning framework as barrier costs [150] but rarely are. Despite these problems, a Barcelona study found that all-cause mortality was decreased for individuals using a bike share program [148] after considering increased physical activity, particulate inhalation, and crash risk. Bicycle safety is a highly local phenomenon, however. The U.S. has one of the worst bicycle crash risks per distance traveled in the industrialized world [27] which is likely do to the prevalence of urban street designs that favor automobile mobility.

Attempts to estimate crash risk based on facilities and traffic conditions alone are confounded by the so-called safety in numbers effect. Cyclist injuries will generally grow at a slower pace than bicycle volumes, all other things being equal [93]. Cities with the lowest risk also have the highest bicycle mode share [27], which not only

reduces the relative number of cars, but forces drivers to be much more aware of bicycles. The debate over helmet laws illustrates how difficult it is to draw causal relationships, however. Using a helmet reduces the risk of head injuries in crashes somewhat [50], but people may be dissuaded from riding a bicycle if they feel a helmet is necessary [147]. Since increasing the bicycle mode share reduces the risk of *all* injuries per distance cycled, there is a point at which overall safety could be improved if deemphasizing helmet use encouraged enough additional ridership that the safety in numbers effect overcame the slight benefit of helmet use [147]. Basic questions must still be resolved before these more complex analyses can be contemplated. Chief among these are that any safety study requires exposure data with which to normalize observed events, and these data are rarely if ever available for bicycles network-wide [128].

1.2 Problem Statement

The behavior of systems as complex as urban regions can never be completely captured through modeling and aggregate metrics, but since transportation infrastructure is so central to their functioning and evolution over time, all transportation investments should be assessed as objectively and comprehensively as possible with respect to their long-term influence on cities. Holes remain in our ability to perform these comparisons, however, especially with respect to projects targeting different travel modes. In the medium term, maintenance will consume an ever-greater share of infrastructure spending. Without a comprehensive tally of costs and benefits over the long term, it is easy to see how limited funds could be prioritized by facility condition. This “worst first” approach would merely maintain a deteriorating status quo with its associated environmental and health problems.

Global pressure will also continue to mount to reduce or at least minimize GHG emissions. These twin fiscal and environmental pressures demand that maintenance

activities be programmed and executed in the most efficient manner possible, however network-wide long-range discrete asset management algorithms are not yet available to meet these demands. Life cycle assessment (LCA) and life cycle costing (LCC) provide frameworks to assess the long-term costs of infrastructure maintenance, however they are better suited to evaluating direct costs such as paving emissions [152] than incorporating indirect impacts, such as those that effect the use phase of a roadway.

Conventional roadway planning views traffic as an external demand to be accommodated in the most efficient manner possible. The level of service offered by a facility will influence demand for that facility, however, and over several decades land development decisions are shaped by the transportation options afforded to particular areas [77]. Before this macro-level phenomena of integrated transportation and land use can be modeled, however, knowledge gaps in more narrow areas of study must be corrected. The century-long growth in automobile dominance is now reversing, which necessitates more complex questions with regards to allocating existing right of way between the various modes. Transit vehicles and bicycles place very different demands on urban streets both in terms of spatial mobility and structural requirements. Buses are some of the most damaging vehicles to pavements that make regular use of city streets [20] but integrate well with other motor vehicles, while bicycles place no structural burden on pavements yet operate at significantly lower speeds than motorized traffic and may not integrate well with it in shared rights of way. Since lane allocation is typically a matter of lane striping, which is done after paving operations, pavement maintenance planning should also include an assessment of the optimal use of lane width, by mode, but no framework exists combining mobility analysis with environmentally preferable pavement maintenance. Such a framework would also require volume estimates for all modes by time of day in order to analyze peak periods.

Given the dominance of motorized travel in this country, most localities maintain very limited data on bicycle movements when compared with automobile traffic and

public transit ridership. This data shortfall is exacerbated by the inherently variable nature of bicycle use and its potentially large latent demand. As a result, planning efforts to encourage cycling, and thus to reap the numerous environmental and public health benefits it offers, are hindered. Only a handful of US cities have significant networks of bicycle infrastructure, so for most localities, it is difficult to say how many people might bicycle for transportation given investment in safe and attractive facilities. Even setting aside the question of latent demand, network-wide bicycle volume models only exist within the confines of sophisticated activity-based travel demand models developed by a few large cities, such as Portland, OR. The vast majority of urban areas rely on more conventional four-step travel demand models that lack the network-specificity to even attempt a spatially discrete forecast of bicycle travel.

In summary, we lack tools to effectively prioritize roadway funds taking all costs and benefits into account. In practice, this means an integrated suite of tools able to plan and evaluate both the maintenance and use phase of streets.

1.3 Research Objectives

The objective of this work is to provide a suite of complementary tools necessary for comprehensive multi-modal transportation project planning and evaluation. These tools have been selected with the primary goal of integrating typically disparate domains so that a fuller picture of the complex interactions between investments and policies for each mode can be evaluated. Feasible frameworks to assess the total costs and benefits of transportation projects will always be limited. Several notable shortcomings in current practice will be addressed, however, with respect to discrete pavement management optimization, lane allocation in urban corridors, and network-wide bicycle data synthesis.

1.3.1 Tool 1

For computational reasons, pavement management systems (PMS) typically operate under the assumption of a steady state. In practice, however, pavement deterioration cannot be exactly modeled and traffic loads are constantly changing, which violates the assumption that the transition from initial condition to optimal pavement condition cycles is negligible. This is especially true in the likely situation of insufficient funds to implement the optimal steady-state plan. Chapter II proposes a framework to address these PMS shortcomings by using high performance evolutionary algorithms to present solutions to the network-wide discrete pavement management problem over a multi-decade planning horizon. The objective of this framework is to enable pavement managers to select maintenance plans from a Pareto-optimal set that are non-dominated with respect to GHG emissions and costs. This framework also provides an important component of a wider analysis that can consider maintenance together with various scenarios of roadway lane allocation.

Tool 1 identifies paving schedules that jointly minimize GHG emissions and user and agency costs by evaluating direct maintenance actions and incremental vehicle fuel usage due to pavement roughness. This tool helps pavement managers make more optimal management decisions by evaluating economic and environmental factors together and presenting complete actionable maintenance plans. Pavement construction and maintenance GHG emissions are only 10% of use phase emissions [151], however, so major reductions will have to also target travel patterns and the use phase of the roadway itself.

1.3.2 Tool 2

Chapter III builds on the PMS presented in chapter II and presents a higher-level framework with the objective of quantifying the costs and benefits of multiple lane configurations for urban streets. This analysis combines PMS with microsimulation

of car, bus, and bicycle interactions to enable planners to appreciate future shifts in mode share in an integrated way. This tool presents optimal allocations of right of way by mode at the level of a bi-directional link, where optimality is defined by minimizing agency paving costs, user fuel costs, and travel time per passenger-km. The contribution of this tool is to allow cities to quantify the trade-offs between mobility and other costs in order to allocate street space in accordance with larger planning goals. An example of the higher-level questions considered by this framework is the computation of a spatial opportunity cost of curb parking.

The framework can be applied over a network, rather than the single representative street presented in chapter III, however doing so would require network-wide bicycle volume estimates as an input. These are not computed as part of typical travel demand modeling, so further work is necessary to synthesize these data from limited available information.

1.3.3 Tool 3

Network-wide travel patterns for bicycles are difficult to estimate because the minor streets typically excluded from automobile models are often the preferred routes for bicycles. Furthermore, bicycle data collection [funding] substantially lags automobiles efforts. The objective of the bicycle volume model presented in chapter IV is to synthesize network-wide directional link-specific bicycle usage estimates at high temporal resolution from a feasible, which is to say small, amount of count data. Tool 3 will not produce impacts directly visible to the public, but rather serves as the basis for multi-modal planning and evaluation efforts by cities. This volume model can also serve as the necessary basis for critical further studies, most notably by providing exposure values for bicycle safety analyses.

Taken together, these research objectives represent an advancement of the state of the art in environmentally preferable infrastructure planning but also provide the

foundation for equally important future work. Chapter V presents conclusions drawn from the combined development of the three tools and goes on to outline ongoing and proposed work that builds on these tools. This material is inherently more speculative in nature and outlines possibilities rather than completed efforts.

Health care costs continue to rise and the links between a lack of daily physical activity and chronic and expensive health problems are only becoming more clear [163, 43]. Reducing both the perceived and actual risk of bicycling for transportation is vital to achieving the significant mode shift necessary to realize its many societal benefits, and an integrated framework able to quantify health benefits and risks associated with all transport modes is critical to prioritize scarce transportation funds effectively enough that real gains in sustainable livability can be realized within fiscal constraints.

CHAPTER II

Pavement Management

Chapters II, III, and IV each present a relatively self-contained component of an overall framework to prioritize infrastructure investments in order to create more livable and environmentally preferable cities. Because the percentage of total infrastructure spending absorbed by maintenance will only continue to grow over the next several decades, maintenance optimization becomes increasingly critical in order to maximize the societal benefit from these expenditures. Accordingly, this chapter presents a discrete network-wide pavement management system that incorporates agency costs, user costs, and greenhouse gas emissions directly and will form the maintenance component of the more comprehensive analysis of urban streets in chapter III.

2.1 Introduction

Few engineered systems in the United States compare with the federal and state highway system in terms of economic impact, materials usage, and environmental burden [81, 88]. In 2007, over \$144 billion were spent on highway activities in the US alone representing 1% of the gross domestic product [171]. These values represent maintenance and new construction, both of which share the same basic techniques. Furthermore, the ratio of new construction to maintenance activities has been steadily decreasing as Department of Transportation (DOT) budgets and

planning priorities necessitate a transition to a maintenance-centric mode of operations [173], underscoring the need to specifically consider the economic and environmental burdens associated with maintenance activities. In general, the emissions associated with maintenance activities trend upward with project cost, though this relationship is complex, particularly once traffic delays are taken into consideration.

GHG emissions have emerged as an important impact associated with industrial activity, and given the magnitude of roadway construction and maintenance activities, it should be no surprise that they represent an appreciable source of emissions for a DOT. Even though DOT paving operations have a variety of environmental impacts, GHGs are of timely concern and have the advantage as a metric in a network-planning context that their burdens are geographically fungible since GHG emissions are equally undesirable regardless of their origin. The US Environmental Protection Agency (EPA) already requires reporting of emissions for facilities generating over 25,000 tonnes CO₂/yr [175], and it is expected that many DOT maintenance and construction projects will be subject to these and future regulations [178]. In light of these factors, DOTs are beginning to evaluate how their engineering activities contribute to emissions as a first step toward identifying reduction strategies [74].

Most efforts to reduce the environmental impacts of maintenance have focused on material and process selection (e.g., cement or asphalt, warm mix or hot mix) rather than overall planning [164, 123, 189]. Life cycle analyses (LCA) with expansive system boundaries and broad but non-specific applicability also exist in the literature [151], as well as detailed analysis of particular situations [192] that cannot be readily generalized. Outside of the academic literature, applied "Green" rating tools at the scope of individual projects are in active development, including the GreenRoads tool [170], Infrastructure Voluntary Evaluation Sustainability Tool (INVEST) from the Federal Highway Administration (FHWA) [19], and the Illinois Livable and Sustainable Transportation System (ILAST) [9]. Minimizing project-level impacts is important,

however confining optimizations within the fixed bounds of individual projects ignores the potential to comprehensively optimize the geographic distribution and intensity of the overall work program, i.e. the boundaries that govern project-level optimizations, thereby reducing *total* economic and environmental impacts.

To illustrate the importance of network-wide planning on emissions, it is useful to consider how pavement decisions are made at most DOTs. Maintenance planning has historically been conducted using a “worst-first” approach, where sections of the highway network showing the most distress are paved until the annual budget is exhausted. Over a multiple year time horizon, this triage approach does not generally produce the most optimal conditions network-wide [185], nor is it likely to ensure the lowest emissions. In response, pavement management systems (PMS) have evolved, to incorporate condition forecasting into a centralized planning framework. They address some of the previous shortcomings in maintenance planning and ideally maximize the condition of the pavement network over the long term within available budgets [4]. In particular, PMS aim to optimize the balance of preventative and restorative treatments over time in order to maximize the return from the maintenance investments. As currently implemented, such systems still often delegate selection of specific treatment programming to the project level, rather than making network-wide optimal decisions in a unified manner [185], complicating efforts to estimate and reduce emissions from these activities. Nevertheless, PMS systems do provide a highly leveraged place to insert environmental goals into infrastructure planning [167].

2.1.1 Algorithm

The quantitative structure underlying most PMS has been under development for many decades and is based, in part, on the principles of life-cycle cost analysis [124, 4]. PMS are typically discrete integer programming problems that lack gradient information and cover sufficiently large solution spaces to preclude exhaustive or

random search. Finding optima in this class of problem is recognized in computer science to be NP-hard, even with an assumption of linearity [100]. Common optimization approaches in the PMS literature include linear optimization [185], dynamic programming [135, 191, 105, 22], genetic algorithms (GA) [63, 58], full enumeration [108], and combination heuristics [136]. Recent work in the PMS area has focused on making the problem computationally tractable [105], including the bottom-up method of *Sathaye and Madanat* [153] which reduces the problem to a single multi-segment steady-state optimum condition threshold vector. Existing frameworks do incorporate one or more of the three specific features that are desirable for design of an environmentally aware PMS: 1) a multi-objective structure that is able to identify a Pareto-set of non-dominated solutions over the desired range of the solution space, 2) direct incorporation of a dynamic environmental LCA to estimate environmental impacts, and 3) solutions at the resolution of individual pavement segments rather than aggregated over the entire network.

Li and Madanat [113] developed a steady-state algorithm to schedule rehabilitation intervals which was expanded by *Ouyang and Madanat* [136] to develop an approximate heuristic solution method to overcome the high dimensionality of the problem by feeding nearly optimal solutions to a branch and bound algorithm. *Wu and Flintsch* [185] used basic weighted linear optimization to provide a quasi multi-objective formulation yielding a Pareto set of non-dominated solutions. The discrete nonlinear nature of the PMS problem does not lend itself to such linear solutions, so dynamic programming, or reducing a non-linear problem through a series of steps, is commonly employed. Non-linear formulations allow versatility in the structure of objective functions and constraints to produce a PMS that accurately reflects the priorities of the funding agency. *Boyles et al.* [22] used dynamic programming in a formulation incorporating non-linear agency preferences. *Ouyang* [135] applied dynamic programming to the question of route selection in light of planned maintenance for a network of two

links, a trivial planning situation, but an important expansion of PMS scope. *Kuhn* [105] sought to overcome the resource constraints of real-world application through approximate dynamic programming (ADP), which can find approximate solutions to combinatorial knapsack problems, such as PMS, in pseudo-polynomial time [18], while the exact problem remains NP-hard. ADP, like dynamic programming, still seeks a specific solution and thus requires parametric iteration of expensive computations to consider a multi-dimensional problem.

Genetic algorithms (GA) can incorporate arbitrary nonlinear factors (agency preferences) while evolving a Pareto set of potential solutions and are readily scalable to employ distributed computing environments. Critically, solutions within several percentage points of the global optimum can be obtained from a GA in a fraction of the time required to achieve true convergence [58]. *Morcoux and Lounis* [122] applied GA for optimization of concrete bridge deck maintenance, and *Fwa et al.* [63] presented a true multi-objective PMS using GA. *Ferreira et al.* [58] further applied a GA-based PMS to a network of specific pavement segments and found GA to be superior to conventional branch and bound techniques in terms of solution quality and speed. *Ferreira et al.* [58] also considered a network of appreciable size (254 segments) and found GA necessary to feasibly obtain solutions. None of these studies included environmental impacts, however. The work described here presents an environmentally adapted PMS or EPMS as depicted in Fig. 2.1 that presents decision makers with a range of non-dominated options to reduce environmental impacts while considering the balance between economic expenditure and resulting network condition.

2.1.2 LCA in Pavement Management

Before environmental constraints can be incorporated into existing PMS algorithms, a dynamic environmental LCA is needed to provide GHG values to the EPMS system. A large number of paving and roadway LCAs have been published over the past

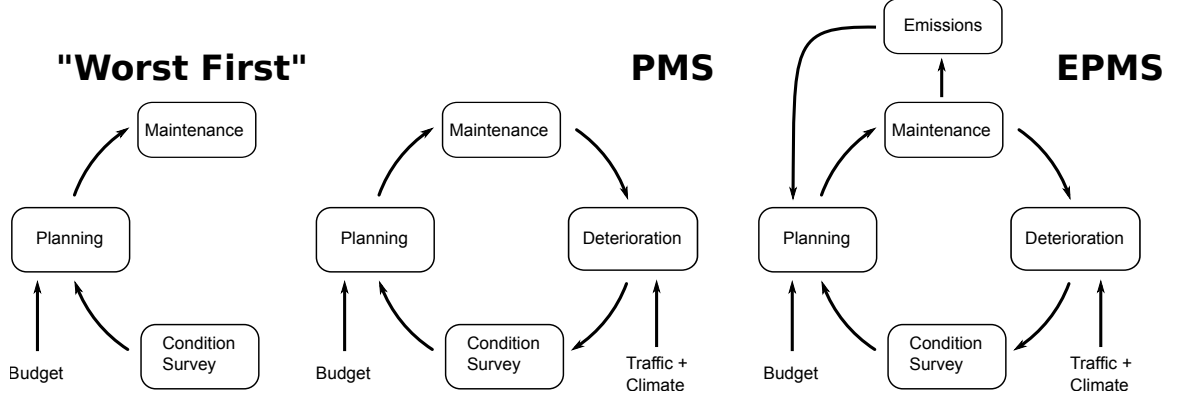


Figure 2.1: Conceptual diagram of the evolution of pavement management from unmanaged to the approach suggested in this work where PMS = Pavement Management System and EPMS = Environmentally Adapted Pavement Management System.

decade, generally with the focus of assessing the trade-offs between reinforced concrete and asphalt pavements. *Hendrickson et al.* [80] considered paving emissions through an economic input-output (EIO) model which captures the systemic implications of paving more thoroughly than a necessarily limited process-based LCA. Available EIO datasets [33], however, contain only a single entry for asphalt paving which does not capture the compositional differences between preventative maintenance and more intensive treatments. Process-based pavement LCAs, which consider the individual unit processes that make up a paving project, have also been published [85, 164, 123, 189, 190, 88, 87, 192]. Such LCAs however, especially in a geographically variable industry such as paving, rely on specific boundaries and assumptions that are not always generalizable [164, 123] or alternatively incorporate a variety of processes and materials rather than focusing on specific actions under consideration in a EPMS context as in *Santero and Horvath* [151], magnifying uncertainty. An LCA specifically tailored to EPMS and the geographic region of interest is required.

Adding an environmental dimension to the design space shifts the mathematics of optimization from a classic cost vs. performance formulation into the multiple-objective domain. Traditional PMS could produce a Pareto, or non-dominated, set

of solutions as in *Wu and Flintsch* [185], where each member of the set achieves the highest possible performance for given expenditure. In practice, it is advantageous for DOTs to spend their entire budget [185], so a single preferred solution is identified from the Pareto set at that budget. In the multiple-objective case created by incorporating an environmental metric, a three dimensional Pareto surface is produced. While an agency may still choose to spend its entire budget, a trade-off between performance and GHG emissions remains at any chosen level of expenditure. This transforms the output provided to decision makers from a single preferred alternative to a region of feasible non-dominated solutions.

Zhang et al. [191] proposed an approach to minimize environmental impacts alongside the conventional objectives of minimizing cost while maximizing performance using multi-objective dynamic programming. This work represents an important step toward integrated environmental decision making, however it only considers a single aggregated pavement segment, and scaling this algorithm up to a network of hundreds of segments would be computationally infeasible. In order to reap the greatest benefits from EPMS, linkage with numerous specific pavement segments is critical in order to reduce subsequent workload and to ensure that the gains achieved by the optimization are accurately implemented.

This work develops the quantitative framework for an EPMS providing segment-specific maintenance plans which combine network and project-level considerations into a single framework. The model includes both a novel GA formulation of PMS and a dynamic LCA for assessing the environmental impacts of a proposed treatment plan. The EPMS is developed and validated using Virginia pavements which provide a compelling case study for several reasons including the relatively large share of Virginia roadways maintained by VDOT rather than localities, and the recent upgrade of VDOT's PMS system. To enable the most effective integration into VDOT's PMS, the model developed here relies on existing elements of the VDOT system, particularly

its use of the critical condition index (CCI) and related deterioration models to quantify and predict pavement condition. By developing an EPMS framework within the constraints of an existing pavement management system, the promise of LCA to achieve environmental gains in real-world application can be realized while leveraging existing investments in computational infrastructure and data collection. The authors stress that the specific components of the framework presented here were chosen to provide a straightforward validation case and to be as consistent with current VDOT practice as possible, rather than as the best possible examples of their type. The framework, as presented here, is intended to illustrate the ease with which GHG emission minimization can be incorporated into DOT PMS systems, enabling objective consideration of environmental performance metrics and therefore evaluate novel GHG-specific maintenance treatments together with more conventional options.

2.2 Method

In order to present pavement managers with a range of non-dominated solutions, each optimal for a given trade-off between cost, network performance, and GHG emissions, a GA framework will be presented that incorporates a well understood metric for pavement condition, means of forecasting future condition, and a model for estimating GHG emissions that result from proposed maintenance activities.

2.2.1 Pavement Condition

Though similar, pavement condition indices typically have underlying differences in methodology and reporting [107, 68], and so are not cross-compatible. VDOT uses a composite measure, the critical condition index (CCI) to represent the worst of either load related or non-load related distresses which are themselves a composite of the data VDOT collects annually through automated surveys of the entire network. Pavement deterioration models have been developed by VDOT as a function of time




Figure 2.2: A conceptual diagram of the three most common VDOT treatment intensities in asphalt pavement maintenance. Reconstruction (RC) is not depicted. Data from [38].

and last treatment applied in units of CCI from 0 (complete failure) to 100 (ideal pavement) (VDOT) for both bituminuous and concrete pavements and is used in interstate, primary, or secondary roadways. VDOT considers a CCI of 60 and below to signify deficient pavement in particular need of attention [38].

Pavement engineers select maintenance treatments appropriate to the pavement type from a discrete list maintained by the DOT. At a planning level, treatments on this list are combined into some small number of planning treatment categories, each with an expected cost and effectiveness over time. VDOT classifies potential treatments into five levels: (0) do nothing (DN); (1) preventative maintenance (PM); (2) corrective maintenance (CM); (3) restorative maintenance (RM); and (4) reconstruction (RC), as shown for asphalt pavements in Fig. 2.2. Concrete, composite, and other pavement types could be included as well, provided appropriate treatments were identified for each planning category and appropriate deterioration models were used. Fig. 2.3 shows the range of conditions acceptable by VDOT for the application of each treatment category. This work made one modification to these thresholds in that DN was acceptable at all treatment intensities to allow \$0 treatment plans.

Each VDOT treatment level covers a range of treatment intensities, so reference treatment designs for each level must be selected in order to quantify costs and impacts




Figure 2.3: Unconstrained maintenance triggers for VDOT Interstate pavements. Data from [38].

Table 2.1: Selected reference treatments and VDOT cost estimates for bituminous interstate pavements.

id	Treatment	Mix	Mill Depth(in)	Pave Depth(in)	Cost per lane km
0	DN	n/a	0.0	0.0	\$0
1	PM	Slurry Seal	0	0.4	\$20,706*
2	CM	HMA	1.5	1.5	\$44,527
3	RM	HMA	3.5	3.5	\$111,992
4	RC	HMA	12.0	12.0	\$314,934

* cost for representative treatment, VDOT cost is \$4,186

due to their specific overlay thicknesses, the particular asphalt mix employed, and the depth of milling performed. The mixes and specific milling and paving depths comprising the reference treatments are given in Table 2.1. PM is a particularly broad category that ranges in intensity from crack sealing to a full-width thin overlay. This wide variation necessitated a more detailed survey of recent bid prices for treatments more closely matching the PM specification of this study, yielding the reference cost in Table 2.1.

Planning is done annually for the entire paving season, in part because ambient temperatures restrict paving to the warmer months. The design space for this EPMS therefore consists of a integer representing one of the VDOT treatment levels for every discrete pavement segment in the roadway network under consideration for every year

in the planning period. Eq. 2.1 gives a conceptual view of one solution in the set.

$$\begin{pmatrix} t_{0,0} & t_{0,y} & \dots & t_{0,Y} \\ t_{s,0} & t_{s,y} & & \vdots \\ \vdots & & \ddots & \vdots \\ t_{S,0} & \dots & \dots & t_{S,Y} \end{pmatrix}, \quad t \in T \quad (2.1)$$

where $t_{s,y}$ is the treatment for segment s at planning year y from the set of all possible treatments T given in Table 2.1. This solution space is determined by the number of segments in the network, S and the years in the planning period, Y , and can be quite large as computed by Eq. 2.2

$$N_{solutions} = (\#T)^{S \cdot Y} \quad (2.2)$$

where $N_{solutions}$ is the total number of unique possible treatment plans and $\#T$ is the number of potential treatments, including DN. A 15-year solution covering just a single segment yields approximately 3×10^{10} possible solutions. Fortunately, the restrictions on feasible solution combinations given in Fig. 2.3 greatly reduce this total. One would not, for example, perform major rehabilitation on the same segment two years in a row. The CCI in the second year will be far above the maximum allowable for major rehabilitation.

2.2.2 Deterioration Model

To evaluate planned maintenance strategies one must forecast pavement rehabilitation needs as a function of time and proposed maintenance. VDOT maintains deterioration models for the five levels of rehabilitation [162] using the base form given by Eq. 2.3 and coefficients from Table 2.2 selected according to the most recent

Table 2.2: Coefficients of VDOT load-related deterioration model given in Eq. 2.3 for bituminous pavements [162].

	CM	RM	RC
a	9.176	9.176	9.176
b	9.180	9.180	9.180
c	1.273	1.251	1.228

rehabilitation performed

$$CCI_t = CCI_o - e^{a+bc^{\ln(1/t)}} \quad (2.3)$$

where t is the time in years since the last major rehabilitation or the effective age, CCI_t is the condition at that time, and CCI_o is the condition immediately after rehabilitation. When a pavement is new, $CCI_o = 100$ since CCI is expressed on a scale from 100 [best] to 0 [worst]. Major rehabilitation, defined as CM and above, is presumed by the model to return the pavement to its initial state, CCI_o and resets the effective age to zero as used in Eq. 2.3. PM does not constitute major rehabilitation and so stand-alone deterioration models following PM have not been developed. Rather, PM is modeled as improving the CCI of a segment 15 points, after which deterioration continues along the CM curve with no reduction in effective age (VDOT). In the case of bituminous interstate pavements, load-related distresses always dominate CCI response. Since CCI is reported as the lesser of load and non-load related distresses, the load-related coefficients are used alone in this work.

The deterioration model is a function of time only, which sacrifices some ability to differentiate pavements with different physical structures or loading. These variations are captured by more complex deterioration models [64], but at the cost of more extensive data inputs, such as the structural capacity of lower pavement and base layers, not consistently maintained by DOTs. The present state of lower pavement and subgrade layers remains in question even with as-built data available, and the




Figure 2.4: VDOT load-related pavement deterioration model for bituminous interstate pavements according to the time elapsed since the last major rehabilitation. Coefficients for this model are given in Table 2.2.

extensive testing needed to complete these data, such as falling weight deflectometer measurements, is infeasible on a frequent network-wide basis. For these reasons, time dependent deterioration models were used to generate average deterioration profiles as shown in Fig. 2.4. The simplistic nature of these models is tempered by the restrictions in Fig. 2.3 which keep treatment applications within the range represented in the data used to derive the model coefficients. Unlike deterioration, GHG emissions result from current actions and so may be directly estimated.

2.2.3 Global warming potential

Global warming potential (GWP) measured as CO_2 equivalent emitted per area of pavement treated, is selected as the representative environmental impact in the EPMS model. Generalized LCAs, most notably the PaLATE tool [84], allow detailed LCA of individual paving projects, however these tools are unsuited to direct integration within the GA and rely on a level of input detail unavailable in the network-planning context. Instead, a process-based LCA was constructed consistent with the work of

Table 2.3: Material specifics of the two mixes used by the reference treatments. All percentages by mass.

Mix	Coarse Aggregate	Fine Aggregate	Asphalt Binder
HMA	57%	38%	5%
Slurry Seal	0%	90%	10%

Huang et al. [88] and *Zhang et al.* [192].

Two types of materials, slurry seal and conventional hot-mix asphalt (HMA), were modeled as shown in Table 2.3 to represent current VDOT practice in consultation with a VDOT Research Scientist. Each of the treatments listed in Table 2.2 can be generated using combinations of these two types of materials. A slurry seal, used by the representative PM treatment, is a non-structural overlay that uses only fine aggregate and approximately twice the binder content of the conventional hot mix asphalt (HMA) used by the more intensive treatments. The other treatments are differentiated by their depth of milling of distressed pavement and depth of new asphalt placed. Paving emissions are dominated by the energy required to heat aggregate and asphalt binder from ambient temperature to 180°C [10], which was modeled as discrete heating steps assuming all moisture entrained in the aggregate was boiled away at 100°C. Fine and coarse aggregates were treated separately since fine material tends to trap more moisture, and boiling this water requires an appreciable proportion of the total heating energy. The remaining emissions are the result of operating the mixing plant itself, transportation of material, and the equipment used to place the material on site. Milling old pavement and transportation of that material back to a storage facility was also included for treatment levels above PM. Additional details are available in supporting materials online.

Warm-mix asphalt (WMA) has received a great deal of attention recently for its lower GHG emissions as compared with HMA, due to lower mixing and compaction

Table 2.4: Summary results of the LCA of reference treatments assuming a 3.66m lane width.

Treatment	tonne CO ₂ eqv./lane-km	tonne CO ₂ eqv. / \$
DN	0	n/a
PM	9.8	4.73E-4
CM	32.4	7.28E-4
RM	75.7	6.76E-4
RC	259.5	8.24E-4

temperatures. Many industry leaders believe that foamed binder WMA will become the de-facto standard as asphalt producers invest in new equipment due to the large fuel savings such equipment provides [59]. However, since considerable process variability remains in industry usage of WMA, only conventional HMA is represented in this work for sake of simplicity.

Table 2.4 reports the results of the LCA in the form of GHG emissions and GHG intensity per dollar for each treatment level. The supporting text online contains more details on this model with specific values in Table S1. Not surprisingly, the emissions increase with the intensity of the treatment. The GHG emissions normalized to dollars spent are more constant. The normalized emissions vary primarily because different treatments have disproportionately higher binder content or use larger amounts of high moisture fine aggregates. PM is also relatively inexpensive so its GHG intensity per dollar is lower than some of the other treatments even though it has a high binder content.

2.2.4 Study Segments

VDOT divides the state of Virginia into nine administrative districts. District 8 was selected to validate the model in this work because it is primarily rural with several hundred miles of entirely bituminous interstate pavement. Interstate pavements were considered exclusively in this study because more complete and higher quality data is

available for these pavements, and because they have consistent properties and cross-sections over long distances, reducing un-modeled variability. Only Virginia data were considered because the deterioration models used are calibrated for Virginia conditions, however the generic nature of these pavements suggests a more general applicability of the results to other bituminous high-volume facilities. Several sources of pavement data are maintained by VDOT that facilitate this EPMS, however some pre-processing is required to obtain a suitable unified data-set. CCI data was available from an automated condition survey carried out annually by VDOT, with results reported every 0.06km, while pavement structure is recorded at arbitrary points according to the extents of past paving jobs. DOTs typically try to maintain these management sections of homogeneous structure over time in order to minimize the size of the PMS solution space. In this work, segment data were aggregated into 430 segments each approximately 1.6km in length. Maintenance plans (Eq. 2.1) therefore consist of a particular activity to perform on each of these segments for each year in the planning period. Restricting plan definitions to a pre-determined inventory of discrete segments is not theoretically necessary, however this formulation was chosen over a continuous network approach to maintain consistency with VDOT PMS practice. For a 7-year plan, this is a 430×7 matrix with 5 possible values for each cell corresponding to the 5 VDOT treatment levels.

In practice, the size of the solution space is greatly reduced by pre-determining suitable courses of treatment for each segment individually over all years in the plan. Further consolidation is achieved by retaining only the non-dominated plans for each segment, reducing the computation burden by orders of magnitude. A reference single segment case with $CCI_o = 85$, the \overline{CCI} of the 430 total segments in the study, and initial age of 7 years, has a total solution space of 5^{15} . The treatment restrictions in Fig. 2.3 reduce this to approximately 7,500 feasible plans, of which only 19 are non-dominated. As in *Ferreira et al.* [58], pre-computation of summary values for the

Table 2.5: Non-dominated treatment plans for a single segment over 15 years with $CCI_o = 85$. DN not shown for clarity. The id column is an arbitrary identifier that links the treatment plans given here with Figs. 2.5 and 2.6, with boldface rows shown specifically in Fig. 2.6.

plan id	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0															
1							PM								
2							CM								
3						CM					PM				
4				PM	PM										
5				PM	PM		PM								
6			PM	PM			PM		PM						
7			CM							CM					
8			CM						CM						
9		PM		PM		PM				CM					
10		PM		PM		PM			PM		PM				
11		CM						CM					PM		
12		CM					PM			CM					
13	PM					CM					PM				
14	PM				CM					PM			PM		
15	PM			PM		PM			CM						PM
16	PM			PM		PM		PM			PM		PM		
17	PM			CM					PM			PM		PM	
18	PM			CM					PM			CM			

non-dominated plans for each segment greatly speeds final computation, in addition to the gains resulting from the smaller solution space. Table 2.5 gives the specific actions for those 19 plans with summary values for each plan in Table S2. In order to give the reader a sense of the variations possible between non-dominated solutions that cover the full range of the solution space, Fig. 2.5 shows the Pareto surface for this reference single segment case. The nearly linear correspondence between cost and GHG emissions can be seen in contrast to the asymptotic response of CCI to either variable. Fig. 2.6 shows the year-by-year progress of three representative plans from the non-dominated set. These three plans are shown in bold in Table 2.5 and circled in Fig. 2.5.




Figure 2.5: Pareto surface of 15 year solutions for a single segment with $CCI_o = 85$ and initial age of 7 years. Solution id's refer to treatment plans in Table 2.5.

2.2.5 Algorithm

A multi-objective unconstrained GA patterned after MOGADES [82] was used to perform the primary optimization in this EPMS. Lane-km weighted average CCI (\overline{CCI}) is the performance metric as described by Eq. 2.4

$$\overline{CCI} = \frac{\sum^s \sum^y CCI_{s,y} L_s}{\sum^s \sum^y CCI_{s,y}} \quad (2.4)$$

where $CCI_{s,y}$ is the CCI of segment s at year y of the plan and L_s is the lane-km of segment s . Depending on agency priorities for evaluating network condition, other

factors could be included in the weighting of performance results, such as the traffic volume experienced at each segment (AADT), or a segment's relative importance in the network.

After each iteration of the GA, a Pareto-front operation removes dominated solutions, i.e. plans that are worse than some other plan in the population on all three metrics of cost, performance, and environmental impact, from this group of child individuals combined with the non-dominated individuals of the previous generation. If this set of non-dominated solutions then exceeds the target size, the most densely populated regions of this Pareto-front are thinned. This set is retained for use in the next iteration and is ultimately the result of the optimization. Pair-wise comparisons between individual solutions are used exclusively in order to avoid the problems associated with normalizing disparate metrics to produce a single rank score.

In order to cover the entire solution space efficiently, the overall population is broken up into a number of sub-populations, islands in evolutionary terminology, with different relative weighting w_i between environment (GHG) and performance (CCI). A small number of individual maintenance plans, some with comparatively high and others with low GHG values, are periodically migrated to adjacent islands in order to maintain the relative distinctions between islands' GHG values and thus overall solution diversity. Island weight, w_i , is uniformly distributed over the islands initially and is adjusted during each migration interval according to Eq. 2.5

$$w_i^{new} = w_{i+1} \frac{d_{i+1,i}}{d_{i,i-1} + d_{i+1,i}} + w_{i-1} \frac{d_{i,i-1}}{d_{i,i-1} + d_{i+1,i}} \quad (2.5)$$

where $d_{i,i+1}$ is the normalized Pythagorean distance between the expected value of island i and the adjacent island $i + 1$. Shifting weights in this manner avoids the natural tendency for the solution population to concentrate in one area of the solution space.

In the tournament and thinning operations, individual plans are compared according to Eq. 2.6

$$P_\alpha < P_\beta \Leftrightarrow (1 - w_i) \frac{\overline{CCI}_\beta - \overline{CCI}_\alpha}{\overline{CCI}} < w_i \frac{GHG_\alpha - GHG_\beta}{\overline{GHG}} \quad (2.6)$$

where w_i again determines the relative importance of GHG emissions versus \overline{CCI} .

The GA results presented here took approximately 60 hours (10 year case, 400 individuals on each of 4 islands) on an 8-core commodity server running multi-threaded C++ code. The population size was halved after one third of the iterations had elapsed to speed further refinement.

2.3 Results

The GA model successfully identified multi-year segment-specific pavement maintenance plans distributed over a range of budgets, \overline{CCI} , and GHG emissions forming a Pareto surface of non-dominated plans for the Virginia interstate highways considered. As expected, increases in annual expenditures result in increases in the maximum achievable \overline{CCI} with diminishing marginal returns at higher budget values. PM and CM comprise the majority of the treatments applied with roughly equal expenditure on each. PM is applied to more lane-km given its lower cost. The preponderance of PM and CM relative to the other treatments follows from the characteristics of the deterioration curves used by VDOT and common to many transportation agencies. These models predict similar deterioration between the three levels of major rehabilitation, yet the costs of RM and RC are much greater than CM (Table 2.1). Fig. 2.6 illustrates the value of PM treatments in between major rehabilitation, as is accepted practice. It is also illustrative that for the sample segment represented by Table 2.5, only PM and CM appear in the non-dominated plans. RC does appear in small but significant quantities for network-wide solutions at higher budget levels since it is

costly but applicable to segments with poor initial condition.

Fig. 2.7 shows the results of a 4 million iteration optimization for a 10 year planning period over all 430 segments, comprising the interstate pavements of VDOT District 8. GHG emissions are well correlated with economic activity as previously shown [80]. The Pareto-front produced by the EPMS is in principle three dimensional, however the strong correlation between economic expenditure and GHG emissions means that in practice, the surface representing non-dominated solutions is nearly two dimensional. The Pareto fronts shown in Fig. 2.7 exhibit smooth curvature in contrast to the Pareto front for a single segment (Fig. 2.5), owing to the much smaller potential intervals between solutions when hundreds of segments are considered together. Pavement managers are therefore still presented with a variety of potential solutions within a narrow range of budget values, despite the inherently discrete nature of segment-specific maintenance plans.

The duration of the planning horizon was found to be quite important in evaluating GA results. Fig. 2.8 shows GA performance for a range of planning durations. \overline{CCI} is computed over the finite planning period, and only modest deterioration is predicted by the model over the duration of typical, e.g. 6 year, plans. More appreciable levels of deterioration occur during the longer plans under low budget conditions. As budgets increase, plan performance slowly converges to a presumed steady state value. The hollow points in Fig. 2.8 represent the most expensive treatment plan considered in the optimization for each planning horizon, as constrained by the CCI restrictions in Fig. 2.3 and the process of excluding dominated plans from the set of treatment plans covering each individual segment. As a result, the solutions are prevented from continuing into the region of very high annual budgets and eventually converging. The locations of the maximum budget points further illustrates the effect planning periods have in light of the discrete nature of maintenance programming both in terms of annual cycles and defined treatment levels. The Pareto-fronts do not extend

to cover even these maximum allowable budgets because either the most expensive single-segment treatments together produce a globally dominated network-wide plan, or because the GA was unable to reach those upper regions within the computational time allowed. The algorithm is designed to focus on the high-value plans found near the knuckle in the Pareto curve, the point of diminishing returns. All planning periods produced solutions up to approximately triple the annual VDOT budget, however, so the region of practical interest is well covered. The maximum allowable budgets under the treatment restrictions imposed on this work are approximately four times the modeled VDOT expenditures.

Computing \overline{CCI} as a plain average also tends to discount the value of treatments applied near the end of the planning period. This is especially true of more intensive treatments that one expects to provide many years of service. For the 19 non-dominated plans for a single segment presented in Table 2.5, none includes a treatment in the final year of the planning period. This suggests the importance of either using sufficiently long planning periods to approximate steady state performance (computationally infeasible given 20+ year typical pavement life), or the inclusion of a salvage value to capture the long-term value of treatments applied late in the planning period. However, since maintenance planning occurs annually, the treatments applied in the first year have the most practical significance. Considering management based on a modeled long term performance average may focus too little on this actionable portion of the plan. Conversely, weighting early performance too heavily may also lead to sub-optimal long term performance. A plain average was used in this work to allow direct comparison with VDOT actions which follow that method.

Fig. 2.8 shows the GA results for a three year planning period and the VDOT actions over this same period. VDOT initiated comprehensive automated data collection in 2007, so a three year planning time frame can be used to compare GA results to actual VDOT practice from 2007 to 2010. Because GA results are

necessarily based on predicted pavement deterioration rather than observed data, this comparison is between past actions with modeled deterioration and hypothetical GA maintenance plans also evaluated using modeled deterioration. Discrepancies exist between modeled and observed behavior such as those caused by the underlying structural capacity of a segment or the amount of heavy truck traffic a segment experiences. Perfect information about these factors is rarely available and since PMS must rely on forecast deterioration, a comparison with respect to modeled performance is appropriate here. The curves in Fig. 2.8 show the Pareto front as it is constrained by practical restrictions on back to back treatments. Without these constraints, the maximum cost plans indicated by the point symbols in the figure would be far more expensive.

For a three year planning period, the results in Fig. 2.8 represent solutions that improve on VDOT practice in all three of the metrics modeled here. Solutions can be found that achieve a 3% improvement in \overline{CCI} for the same GHG emissions, or 2% for the same budget. More interestingly the same \overline{CCI} can be achieved with a 40% reduction in cost and a 50% reduction in GHG emissions. The greater relative reductions in costs and emissions vs. increased performance are likely attributable to a more precise targeting of maintenance actions, since a conservative approach would otherwise involve re-paving earlier than necessary in order to preempt extensive damage in the lower pavement layers.

During the three study years 2007-2009, VDOT applied only CM in the study area. Given the relatively short time horizon, it is not possible to draw any broad conclusions from this trend, but the focus on CM could have been due to limited budgets and/or a strategy in which PM occurs cyclically, e.g., in 2010. This analysis demonstrates that use of one treatment type is not optimal in the long term. From Table 2.4 we see that CM has relatively high specific GHG emissions. Over a 15-year period, however, we see from Fig. 2.6 that it provides better performance per dollar,

though with higher GHG emissions, than an equivalent course of PM treatments. GA solutions use a mix of treatment intensities, primarily PM and CM at roughly equal levels, which explains some of the differences between GA and current practice.

2.4 Conclusion

A GA-based EPMS was developed capable of producing a Pareto-set of segment-specific maintenance plans for a range of budget, performance, and GHG emission values. This EPMS was validated for the case of rural interstate highways in Virginia. The authors believe this case to be generally representative, however these conclusions have not been verified in other states. GHG emissions were found to be strongly coupled to economic expenditures, so minimization of embodied GHG also improves economic performance. Since GHG emissions are tied to the underlying physical processes rather than market forces, a low-GHG maintenance plan may also prove more resilient to price shocks such as spikes in oil price, than a plan optimized on costs alone. The EPMS system was able to identify solutions which either improved on network performance for the same budget and GHG emissions, or greatly reduced those values while maintaining equivalent performance over a three year planning period. Solutions were found to be highly dependent on LCA results and deterioration models, which are themselves both highly dependent on available data. Thus, this approach is most readily implemented on well-characterized asset classes, such as interstate highways. Further refinement of both these components will also improve the ability of the EPMS to forecast performance and emissions over many years. For example, expanding the LCA system boundary to include the use phase of the roadway, and including a greater diversity of treatment options, especially those with higher costs coupled with long term benefits like reduced rolling resistance, could more comprehensively optimize overall emissions. This work is focused on the DOT maintenance office itself, however, as the relevant decision-making body, and so the

system boundaries are limited to reflect only modest expansions of DOT environmental mandates and the selected treatment options represent standard practice, in order to illustrate the benefits of policy-level optimization. A more complete systems dynamic model, or something similar, could address maintenance planning in its larger economic and political context.

Solution performance was also found to be sensitive to the planning period used, which is intrinsically linked with the chosen network performance metric. Consolidating the performance of an entire pavement network over multiple years to a single value is complex and a matter of agency preference, so significant attention must be given to the selection of this metric in practice. High performance maintenance plans incorporate a variety of treatment intensities, however moving to this more complex approach requires managers to face exponentially larger decision spaces. DOTs must work comprehensively to reduce emissions. Paving emissions are small compared to the use phase of roadways, however paving dominates the activities under direct DOT control, and maintenance is an ever-growing share of that work. Optimization of DOT maintenance actions is central to their success in reducing emissions, and an EPMS provides a framework to make these comprehensive multi-objective decisions in the face of shrinking budgets and coming GHG constraints.

2.4.1 Acknowledgements

The authors would like to thank Dr. Brian Diefenderfer (Virginia Center for Transportation Innovation and Research) and Kevin Chisnell (VDOT District 8) for data and technical expertise as well as Dr. Raja Shekharan (VDOT Maintenance Division) for manuscript review and deterioration and cost models.

2.5 Supporting Information

Tables A.1 and A.2 and additional details on the LCA employed in this chapter are provided in appendix A. Appendix B provides additional justification for a network-wide approach from an LCA perspective.

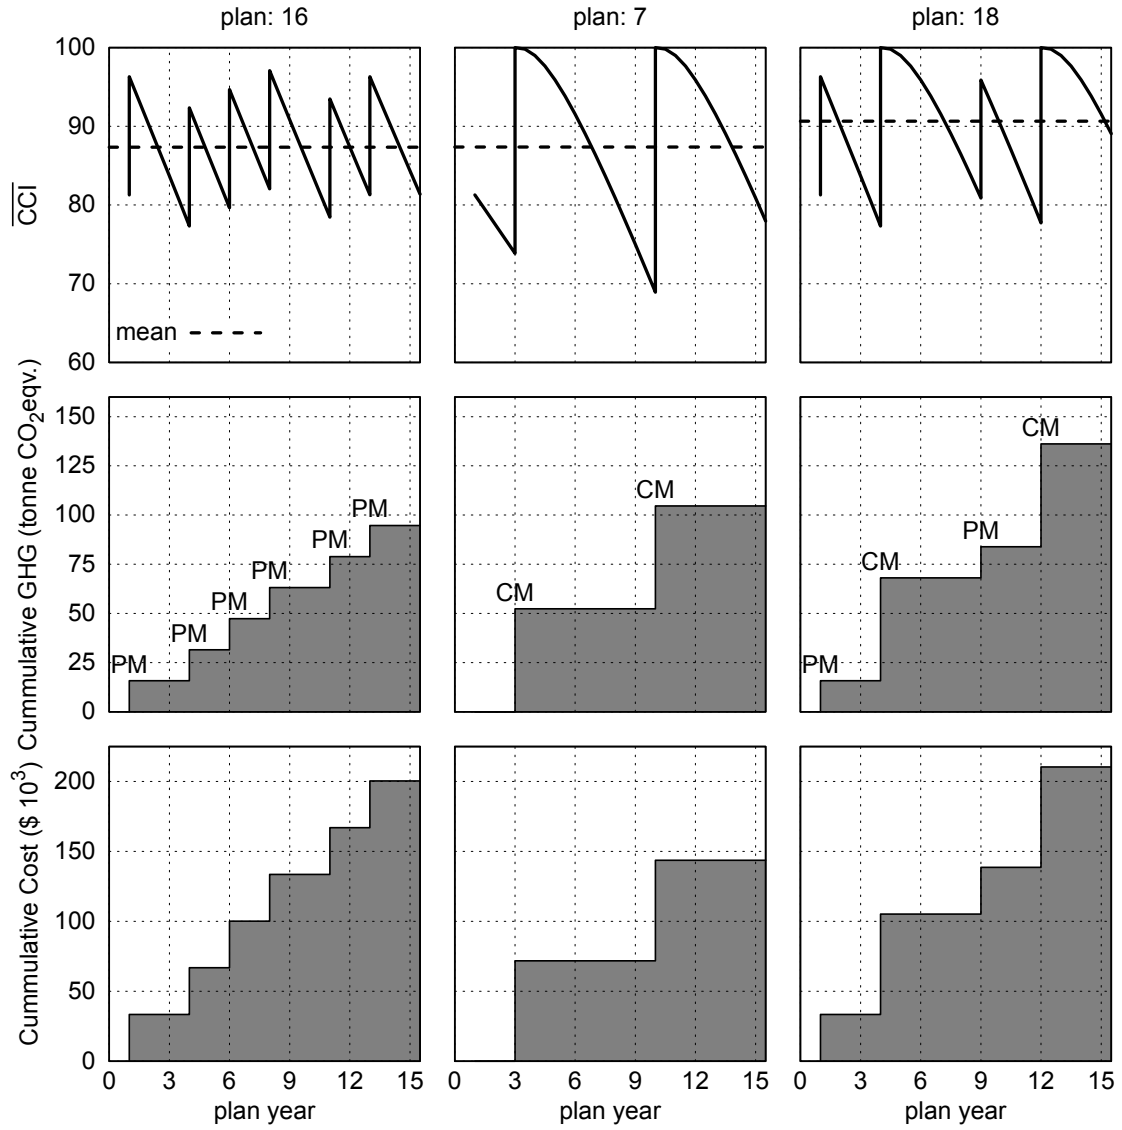


Figure 2.6: Three sample 15 year plans showing the *cumulative* performance, GHG, and economic implications of representative treatment plans for a single 1.6km segment with $CCI_o = 85$ and an initial age of 7 years.




Figure 2.7: Results of model output for 10 years illustrate relationship between CCI, cost, and GHG emissions (large dots). Projections of results between pairs of variables are also shown (small dots).




Figure 2.8: Pareto front curves for several planning periods. The 3 year solutions are based on VDOT practice from 2007-2010. Arrows represent the improvements that could be achieved using the optimization procedure proposed here.

CHAPTER III

Mobility and Shared Right of Way

Chapter II quantified pavement maintenance costs and GHG emissions and provided a set of multi-year paving plans that were non-dominated with respect to agency costs, and a subset of use-phase impacts. This chapter expands on the user cost and greenhouse gas emission components by considering the use phase explicitly through microsimulation and a coupled vehicle emissions model. By considering a range of street configuration and traffic scenarios, the question of when to pave is expanded into a total-cost analysis of different lane configurations and the relative costs of mobility, vehicle efficiency, and maintenance.

3.1 Introduction

Ground transportation is responsible for nearly 30% of the primary energy consumption and 27% of the greenhouse gas (GHG) emissions in the United States [183]. Related infrastructure also results in significant material movement — every \$1 million investment in roadway construction requires 9×10^4 tonnes of aggregate and 3.3×10^3 tonnes of cement [101]. In an effort to try and reduce these burdens, numerous life cycle assessments (LCAs) have been performed over the past decade to understand how specific technological choices contribute to emissions, energy consumption, and materials use. These studies have tended to focus on either the design or the use phase

of the road [81]. Design generally involves the selection of a material, e.g., concrete or asphalt, or specification of roadway width and configuration [166]. Use entails a variety of other processes including vehicle selection or roadway maintenance [11]. Even though these studies have identified many obvious opportunities for environmental improvement, there is little evidence to suggest that they have provided deep reductions in material use or emissions.

A principal limitation of many published analyses is that they consider technological options for ground transportation narrowly and evaluate only specific elements of the design or use of roadways at once. An asphalt road may have a lower life cycle emission profile than a concrete road, for example, but that difference is small when compared to the overall emissions from the use phase of the road [80]. Similarly, the conclusion that a greater bicycle mode share will reduce the carbon emissions of a roadway is not useful if it is not considered along with other factors discouraging bicycle use and the impact of more bicycles on overall traffic flow. To date, Life cycle assessment has been employed as a method for environmental bean counting that considers problems removed from the broader system within which they exist. Consequently, even though LCA has been actively pursued in academic circles, it has had only limited impact in policy circles.

At the same time that conventional LCA tools have been insufficient for solving many of the existing problems faced by transportation managers, emergent challenges make the need for new tools even more pressing [89]. Conventional development patterns have led to widespread congestion in urban and suburban areas around the world. Shrinking maintenance budgets at a time when facilities built in the post WWII boom period are reaching their design life span are making it ever harder to maintain the level of service that was envisioned for roadways during design. Declining pavement quality is also exacerbating the emissions[151] and safety costs associated with the use of these aging facilities, which only compound the impacts of unchecked

growth in vehicle kilometers of travel (VKT) worldwide. Efforts to address these problems and provide meaningful improvements will require systems thinking that considers life cycle impacts, personal choice, and policy realities.

In practice, ground transportation is constrained by a few overarching factors. The most obvious is budget. Road construction is expensive, though less so than alternatives like public transport, because much of the cost is borne by users in the form of vehicles and fuel. Roadways are expensive to maintain, and so many exhibit condition ratings below their design values. In many urban areas, space is also constrained and a limiting factor in terms of enabling more mobility. Where space is available, additional road and parking facilities relieve congestion in the near term but only further separate typical destinations, increasing trip lengths over the long term[129]. The carbon emissions from ground transportation are appreciable and growing as more and more developed nations move toward car ownership levels on par with the United States. Paradoxically, the convenience afforded by automobiles has contributed to significant and consistent traffic-related delays in almost all of the worlds' urban centers. These delays translate into appreciable costs to the users of the transportation systems [87], in addition to increased crash and health risks [70], and overall environmental impacts.

Planning that tackles these challenges involves both near-term adaptive strategies and long-term improvement projects. Existing facilities, including the functionally obsolete, have significant embodied emissions and sunk costs that preclude their immediate replacement, even when sufficient funds are available. This lag between identifying changing needs and building new infrastructure results in increased total public costs with respect to design projections. Interim adaptive strategies can be considered, however, to minimize the monetary, environmental, and safety impacts of a sub-optimal design still in the middle of its useful life until it is time to replace the facility [138]. Adapting existing facilities to new use patterns and goals also provides

a bridge between generational shifts in infrastructure planning objectives.

In the case of transportation, planning has historically focused on mobility, with the outcomes of ever-increasing VKT and sprawling development patterns that discourage alternatives to the private automobile [52, 120]. In the near term, however, adapting existing facilities to maximize mobility in light of an increased bicycle and transit mode share is an appropriate measure to maximize the value provided by these facilities. Adaptive strategies comply with existing system constraints and involve lower cost measures to capture the remaining value of past infrastructure investments until such time as a major change consistent with a lower-impact vision is warranted. Evaluation of these transitional actions is complicated, however, given the absence of any sort of steady state and the number of analyses that must be integrated.

Most of the individual elements constraining ground transportation systems have been studied in isolation but few examples of integrated multi-criterion analysis of roadway use have been published. Pavement management systems (PMS) have been developed to help maintenance managers maintain large systems of pavements under budget constraints[76, 114, 193, 154]. Separately, traffic engineers have developed micro simulation tools of vehicle dynamics to understand the effect of different road configurations[57, 54, 17] and work-zone traffic management [87]. Economic analyses rely in part on the engineering analysis of road utilization and make the case between more infrastructure and induced economic activity. The interests of pedestrians and bikers are also considered in the context of safety[125, 96, 139] and congestion[46] but rarely in terms of providing viable alternatives to automobile transportation.

Efforts to identify significant reductions in the environmental burden of transportation will need to consider these tools together to inform optimal use of roadways under multiple constraints. Here we present a method for combining these analyses with conventional LCA of roadways. We consider the results in the context of adaptive roadway lane (re)configuration, such as converting curb parking to bicycle facilities,

that many cities, including Washington, DC and New York City, are currently undertaking in an effort to reduce congestion directly and indirectly by supporting alternatives to automobile travel that make more efficient use of the public right of way and incentivize reduced trip distances.

3.2 Method

A total cost minimization approach is proposed in order to identify preferable lane configurations for two lane urban roadways, given the physical parameters of the site, available right of way width, and traffic volumes by mode. A lane configuration is defined by the number and width (or presence) of parking, bicycle, and conventional travel lanes for each direction. Costs include: annualized pavement maintenance, motor vehicle fuel costs [13], and travel time at half the prevailing wage rate [7]. GHG emissions are also calculated. Computationally, the proposed framework is structured as a series of distinct codes.

Table 3.1: Discrete segment parameter space. All unique combinations were evaluated using microsimulation. The first group defines a lane configuration and the second a scenario.

Travel Lane Width	3.4m, 4.3m
Bicycle Lane	None, 1.22m ^a
Parking Lane	None, 2.5
Characteristic Length Between Passing Zones	50m, 100m, 200m
Grade	0%, $\pm 4\%$
Bicycle Mode Share	1%, 10%

^a AASHTO guidelines call for a 1.52m bicycle lane adjacent to curb parking which are applied here as appropriate.

Micro simulation of idealized roadway segments was carried out using VISSIM 5.4 for all parameter combinations given in Table 3.1 using common values in Table 3.2.

Table 3.2: Common parameters and descriptive values for all simulated configurations.

Total Width	6.8m - 16.6m
Motorized Speed	16.7m/s
Bicycle Speed	varies by grade
Car Occupancy	1.2
Bus Occupancy	20
Bicycle Occupancy	1.0
Potential Curb Parking	6.1m spaces covering 80% of segment length

The first section of Table 3.1 lists the parameters that define a lane configuration. The remaining parameters in Table 3.1, taken together, will be referred to as the scenario. The key dynamic explored in this work is the use of a single lane by motor vehicles and bicycles simultaneously, and whether the former is able to safely pass the latter within the lane. VISSIM is able to model lateral behavior within lanes, in addition to more conventional vehicle following and lane changing behaviors, and so is able to consider this question. Peak and off peak traffic volumes were simulated separately and combined using 12 hours of each to arrive at daily totals, which were then inflated to annual values. Additional detail is provided in appendix C.

A low and high bicycle mode share were explored, with 1% representing typical urban mode share in the US and 10% representing a target that leading cities, such as Portland, OR, could achieve in the next decade with sufficient investment. Total person trips are constant between the bicycle volume scenarios, with the difference made up by automobiles with an occupancy of 1.2. Figure 3.1 shows the relative sizes of vehicles and lanes considered in this work. The design standard for lane width in the US is 3.6m [168], however it is both permissible and likely that narrower lanes are used in the width constrained urban corridors considered in this study, so we




Figure 3.1: A reference multi-modal street section and the relative widths of vehicles and lanes.

assume a 3.4m base case. Assuming a 1m passing buffer, 1.75m wide automobile, and 0.5m wide bicycle, an automobile is able to pass a bicycle within the lane, however a 2.6m wide transit bus is not. To evaluate the benefit of alleviating this "stuck" condition, a wider lane is also considered that allows buses to pass cyclists without either vehicle departing the lane. "Dooring" accidents, crashes between bicycles and opening automobile doors, may also justify additional buffer width between parked cars and lanes used by bicycles, however the present work does not model crashes, so this effect is not represented in the analysis. The impact of these passing conditions is dependent upon the likelihood that buses will encounter a bicycle and become stuck behind it, thereby delaying itself and following motor vehicles.

Realistic urban corridors vary in width, so it is overly conservative to assume that a given lane width will restrict passing movements indefinitely. Here, we adopt the concept of a characteristic length between passing zones, as given in Table 3.1 to determine the likelihood that a bicycle and bus will be present and the delay expected to result from the encounter. VISSIM simulates these interactions directly. Since Poisson vehicle arrivals are assumed, however, in order to make the results as general as possible, a probabilistic analysis can be carried out to compute expected delay. This analysis is presented in appendix C and agrees with the results of the microsimulation. Motor-vehicle speeds were assumed constant for each lane width

given the considerable variation that exists in the literature on the effect of lane width of motorist speed choice [125, 149, 168], however, this behavior could be readily altered in the microsimulation parameters where local data is available. Bicycle speeds for each grade were computed according to first principles formulas[194].

Vehicle fuel use and emissions are affected by pavement roughness though not consistently between various operational regimes [15, 187, 99] due to the varying contribution of rolling resistance to required power. A power-based vehicle emissions model, CMEM [16], was used to post-process the microsimulation vehicle data at 1hz and two roughnesses using a lookup table computed at a reference international roughness index (IRI) of 1.0 m/km, and for a rough case with an IRI of 4.0 m/km by inflating vehicle rolling resistance after Karlsson et al. [99]. Only automobile emissions were affected due to the inconclusive results of that study for heavy vehicles. Final emissions and fuel consumption were computed by the pavement management module through linear interpolation of the two roughness cases.

Pavement management activity and emissions were computed based on the previous work of the authors [71] with additional dynamic pavement loading due to roughness [56]. Explicit treatment of heavy vehicles is important in the comprehensive analysis of a roadway given their disproportionate impact [42]. Pavement maintenance plans were computed for both directions of travel lanes, bicycle lanes, and parking lanes independently using aggregated annual vehicle volumes and emissions from the appropriate microsimulation trials. A network Pareto front for all the lanes was then computed, with a representative example given in Figure 3.2. For this analysis, the non-dominated plan with the minimum total GHG emissions, subject to agency constraints, was selected. Further detail on this method, and the larger issues of discounting and temporal variation, can be found in appendix C.

In combing the separate models of pavement management, vehicle microsimulation, and vehicle emissions computation, a hierarchy exists according to the sensitivity of




Figure 3.2: Agency pavement costs vs. user fuel costs as a result of pavement roughness (upper) and combined costs vs. combined GHG emissions (lower) from non-dominated pavement management plans. The triangle data point indicates the selected plan by minimum GHG.

one to another. Pavement roughness influences vehicle fuel use and emissions as well as dynamic pavement loading. Loading affects pavement durability, and maintenance investment determines the resulting pavement condition. Vehicle behavior is assumed to be insensitive to pavement condition within the specified limits, however, which makes travel time cost insensitive to changes in fuel and agency costs. Critically, this means that the PMS optimization can be performed after microsimulation. Otherwise, the task would be computationally intractable with existing microsimulation tools, since a microsimulation would have to be run at each iteration of the PMS genetic algorithm.

A Pareto front of lane configurations can be identified for each scenario that are non-dominated with respect to total costs and width. Lane configurations that include curb parking will naturally not appear in this set since they incur pavement maintenance costs but provide no counterbalancing benefit as computed. A spatial opportunity cost of curb parking can be computed, however, by computing the cost differential between configurations with parking and a point interpolated on the Pareto cost curve at the same total width. This opportunity cost allows decision makers to quantify the potential mobility value of public right of way allocated to parking.

3.3 Results

The effect of bicycle mode share on average travel time for a particular segment can be significant for specific conditions as shown by the results in Figure 3.3. Differentiation between the cases occurs when heavy [wide] vehicles, such as transit buses encounter a bicycle and have insufficient room to safely pass resulting in significant delays for themselves and the motor vehicles behind them. In these graphs, individual data points represent microsimulations of specific cases and are plotted with random jitter on the length axis for legibility. As with other results presented here, they are normalized to a kilometer of travel. The trend lines are second order polynomials

used to illustrate the relationships of interest. The effect of grade, and whether or not vehicles are traveling uphill or down, have important effects on the results presented in Figure 3.3. The data are grouped based on whether or not buses are stuck behind bikers in the different configurations. For the level ground segment, the results are equivalent. For the inclined segment, the uphill travel time is always considerably higher than downhill travel time if trucks get stuck behind buses. The impact of the stuck condition is proportional to the relative speed difference between bicycles and motor vehicles, which comes from roadway grade, and the likelihood of a heavy vehicle encountering a bicycle within the characteristic distance between passing zones. This is determined by modal volumes and headway distribution.

For lane widths more narrow than those considered here, all motor vehicles with more than two wheels would be unable to pass bicycles within their lane, with the result that the expected speed of all traffic would approach that of bicycles as characteristic length and bicycle volume increased. These cases are not presented in order to focus on the more typical but less intuitive stuck condition, and because very narrow lanes are likely to have an effect on driver speed decisions according to the particular characteristics of the site, such as sight distance, land use, number of driveways, and other factors. This is not to say, however, that the framework presented here is not suitable for 3.1m lane widths, only that the results would not be transferrable to other situations. For wider lane widths, the differentiation observed here disappears as heavy vehicles are able to pass bicycles. All four stuck groups include multiple lane configurations which reveal more subtle differences in annual costs with respect to width (distance between curbs) when considered individually. In Figure 3.4, this relationship is seen as a distinct Pareto optimal frontier for each characteristic length between passing zones, grade, and bicycle mode share considered, with non-dominated lane configurations shown in bold. The influence of the stuck condition can be seen in the abrupt transition in the Pareto curve around 8.6m, or the minimum width of

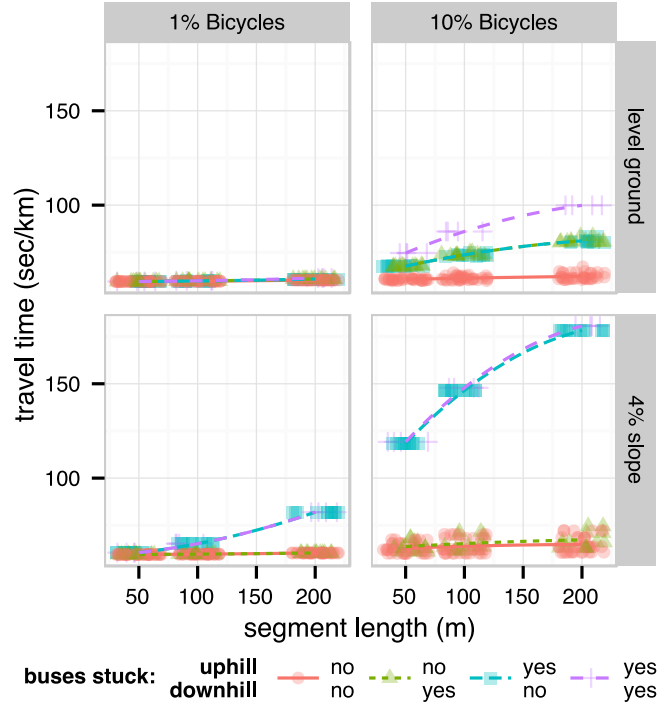


Figure 3.3: Simulated motorized travel time illustrates the delay caused when heavy vehicles are unable to pass bicycles.

a configuration not stuck in both directions. The initial drop between the first two non-dominated points is larger for the 4% case, as compared to level ground, since alleviating the stuck condition on the uphill segment is considerably more important than in the downhill direction given the dramatic difference in expected bicycle speeds.

For the urban situations considered in this work, both travel time and fuel use are positively correlated with vehicle delay. These costs are also considerably larger than agency expenditures for non-dominated pavement maintenance plans, which can be seen in Figure 3.2. This relationship supports the decision to select the PMS plan with the minimum total GHG emissions, since a comparatively small agency investment provides a larger reduction in user costs. It follows that minimizing total costs also minimizes GHG emissions since total costs are dominated by time and fuel and are sensitive only to delay once pavement condition has been established. This is encouraging since typical planning processes do not explicitly quantify GHG emissions.

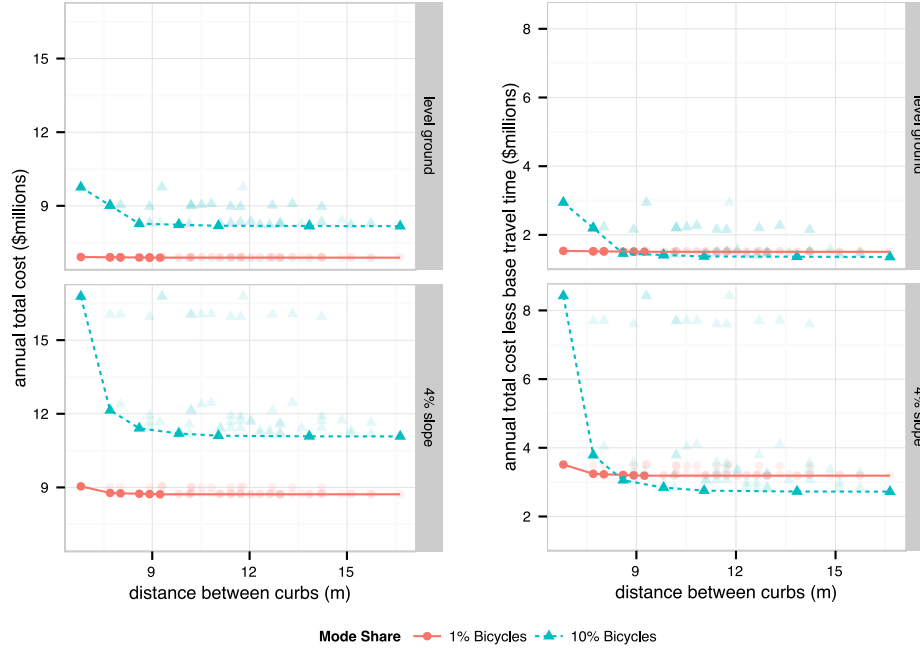


Figure 3.4: The relationship between road width and costs follows a Pareto optimal behavior, shown here for 100m characteristic lengths. Total travel time costs (left) dominate fuel and PMS costs, however when delay costs are used (right) increased cycle mode share can reduce costs for appropriate lane configurations.

Figure 3.4 presents the Pareto front with respect to total costs and then using delay rather than total travel time costs. Neither approach is strictly more accurate, however since bicycle travel is generally more time consuming over the same roadway segment, increasing bicycle mode share tends to dramatically increase time costs. Using the time cost of delay only, assumes that travelers had already accounted for this cost externally, which is not unreasonable, and is typical in traffic analysis. Under this assumption, Figure 3.4 reveals a tipping point where increased bicycle mode share lowers total costs, given sufficient roadway width. For cases where significant increases in bike ridership are not paired with enhanced facilities like wider roads, this will increase travel time for all users.

Parking is not valued in the total cost reported in Figure 3.4. As a result, lane configurations with parking bays incur pavement maintenance costs without corresponding

negative costs from parking’s value as a service, rendering these configurations sub-optimal. This also explains why costs reach their minimum in Figure 3.4 approximately 5m before the maximum width since that is the width occupied by two 2.5m parking bays. Curb parking does have a site-specific value both in terms of vehicular accessibility and in the broader economic sense by supporting value capture along the street through increased business patronage. Because the magnitude of this value capture is so site-specific, Figure 3.5 presents a spatial opportunity cost of curb parking as the difference between the total cost of a lane configuration that includes parking and a linear interpolation of the Pareto front at the same width. Visually, this is the vertical distance between the lighter data points in Figure 3.4 and the Pareto front, normalized to a daily value per parking space. The narrowest lane configuration that can include parking is 9.3m wide, so opportunity cost is reported from this value up to the maximum configuration width.

A spatial opportunity cost of parking is shown in Figure 3.5 for the inclined case with a 100m characteristic distance between segments. The likelihood of buses’ getting stuck behind a bicycle is related to characteristic length, so other scenarios look similar to this plot but shifted in magnitude accordingly by length. Level ground also exhibits consistent behavior, albeit with a different curve shape that can be inferred from Figure 3.4. Since the maximum hourly volumes considered in this study are below saturation levels, delay comes almost entirely from the bicycle / motor vehicle interactions. Therefore, it is not surprising that parking opportunity costs are highly sensitive to the percentage of bicycles in the traffic stream. For a given width, parking and bicycle infrastructure essentially compete for the same space which results in large opportunity costs for parking in narrow configurations, as might be seen in a traditional urban neighborhood of places like New York City and Washington, DC. These larger values are also well in excess of typical parking meter returns, which is consistent with the pervasive subsidization of automobile parking in the US [158], though adjacent

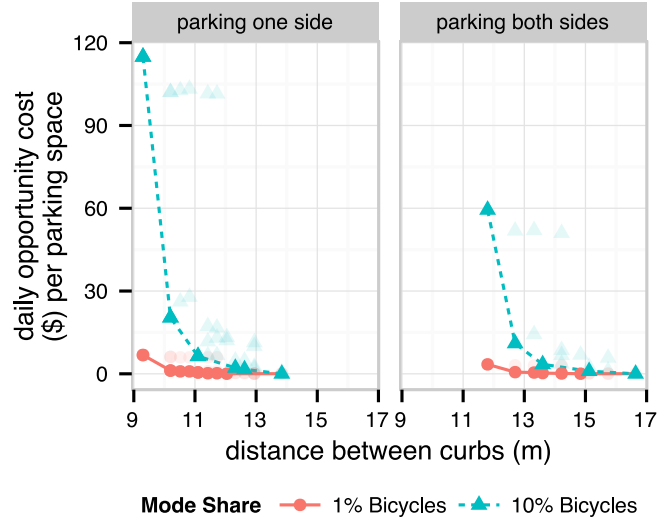


Figure 3.5: Spatial opportunity cost for curb parking, shown here for 100m characteristic length and 4% grade, increases dramatically for narrower roads suggesting that when all the costs are considered, there is a tipping point beyond which curb parking becomes an expensive use of land.

businesses may benefit by attracting auto-dependant customers. These values could also be considered conservative since they do not include any opportunity costs for land in the pedestrian zone on either side of the roadway proper that could be put to other uses, such as pedestrian mobility or restaurant seating.

3.4 Sensitivity and Additional Considerations

To develop a more complete understanding of how urban roadway design can be informed by a model like this, it is useful to consider both those factors that had little impact on the results and those factors that were not included in the analysis for one reason or another. Pavement roughness and its effects on fuel consumption are considered here even though in the context of urban driving, roughness is far less important than acceleration cycles in determining fuel efficiency. However, since travel time is independent of roughness, pavement management costs for a given configuration are optimized against marginal vehicle fuel consumption which is on the

same order of magnitude. Additionally, the fiscal reality of the pavement manager may be such that maintaining serviceable pavement is a constant struggle, and so minimization of agency costs may replace minimization of GHG emissions in selecting the best plan. The impact pavement condition has on a given commuter’s likelihood to select the bicycle or automobile mode is difficult to quantify, but also more significant to the ultimate makeup of the Pareto set of lane configurations by determining relative volumes between the modes as can be seen in Figures 3.4 and 3.5.

Similarly, only select measures of the health and safety factors associated with modal shift are captured by this model since many of these are difficult to quantify with confidence even though they are often cited as an important driver and/or obstacle in mode shift toward more active forms of transportation [94]. Crashes, though not considered in this work, are an important factor in selecting bicycle transport and even though AASHTO guidelines dictate an additional 0.3m of width for bicycle lanes adjacent to curb parking, that width remains insufficient to prevent dooring crashes between bicycles and parked cars [2, 41]. The cost of curbside parking goes up fractionally for this additional 0.3m of pavement to address the risk of dooring to bikers, however the actuarial cost of a potential fatality is on the same order of magnitude as total annual travel cost per kilometer in this study. Crashes are difficult to model generally, and difficult to even estimate based on past reports involving bicycles due to pervasive under-reporting of non-fatal encounters. Never the less, the potential health and safety costs are potentially large enough to influence the results presented here.

Another limitation of the model is that it does not consider roadway segments within the larger roadway network. In many urban areas, vehicle arrival is not Poisson distributed but rather appears as a decaying platoon progressing between controlled intersections. In addition, driveways play an important roll in platoon migration and overall capacity in urban settings. Finally, buses stop periodically for passengers and

this presents an opportunity for the traffic to clear and the biker and bus to separate. These factors were not considered here in order to make the analysis as general as possible but could be readily incorporated into the microsimulation for a specific site. Many of these factors would need to be included in a similar model to derive site specific estimates for the opportunity cost of parking even though we expect that the general trends discussed here would hold for most urban roadways.

3.5 Implications

A major challenge for urban areas around the world is to improve livability, which is often achieved by reducing reliance on automobile transport [126, 145]. A shift away from the automobile also results in significant reductions in energy use and GHG emissions, both of which are increasingly relevant policy objectives. For countries with a legacy of auto-mobility-dominated planning and policy decisions, such as the US, these goals are especially daunting. Reducing VKT through more compact development and alternative transportation modes is a long term objective, but previous studies have observed that many of our projected future emissions are 'locked in' by virtue of the inefficient nature of our existing infrastructure stock and the relatively slow rate at which this is replaced [112]. Short-term strategies are needed to achieve some of these gains without major changes to our existing infrastructure. Improved bicycle facilities is one such change that could encourage higher mode share without significant investment. In cities like Washington, DC and New York City, where many of the trips are short and well suited to bicycle transport, efforts have been underway for several years to provide such facilities. Consequently, bike mode share is increasing significantly [3], but the ridership rates in these US cities are still an order of magnitude lower than in many European cities [1]. Additionally, these changes are not without controversy in light of their direct costs and use of limited space [103].

Tools like the one developed here can be used to help resolve the apparent conflicts that could inhibit progress toward more environmentally sustainable infrastructure. We find that increasing bicycle mode share can have a significant impact on motor vehicle delay, and indeed greatly increase total costs where sufficient right of way is not provided. Conversely, with sufficient space to allow wide vehicles to pass bicycles, a reduction in total costs for all users is obtained through an increased bicycle mode share. Parking, specifically curb parking, emerges as space that is potentially most eligible for reallocation to bicycles in width constraints urban corridors because of its ideal position and the relatively minimal expense of lane reconfiguration. The spatial opportunity cost of parking quantifies this trade-off and exhibits sharp transitions between realistic monetary amounts for wider streets and infeasible high expected returns for more narrow areas. These results provide a clear guide to traffic engineers and urban policy makers with respect to optimal allocation of limited pavement, although such a transformation is often politically fraught.

Even though life cycle assessment and other tools are useful for beginning to understand pieces of this problem, a systems-based approach like the one proposed here is needed to directly support policy decisions. By considering infrastructure systems in an integrated, yet quantitative manner using existing modeling frameworks, short-term low cost opportunities emerge for efficiency improvements that would not be obvious using other tools alone. The general framework proposed here could be applied in a variety of infrastructure contexts. For example, the green building industry standard, leadership in energy and environmental design (LEED), has already recognized that buildings cannot be analyzed outside of the larger infrastructure context within which they exist and is making changes to consider the community within which the building exists [40]. Efforts to quantify and minimize road trucking emissions exist, such as the EPA SmartWay program [131], however this approach is still confined within the silo of the trucking industry and excludes the very significant

interdependence between pavement and their durability and vehicle fuel use. Other types of infrastructure have been less studied and offer hereto untapped efficiency improvements. Water treatment and distribution, electric power grids, and wireless communications, and others could all benefit from considering engineering performance criteria, agency and user cost, and human factors for informing improved policy.

CHAPTER IV

Bicycle Volume Model

By incorporating the use phase of a street, chapter III added a critical aspect to the total-cost evaluation of road infrastructure, however it also greatly increased the scope of necessary input data. For illustrative purposes, the method was presented a range of possible scenarios, however this approach quickly becomes dimensionally impractical when multiple streets are considered. Instead, the analysis can be narrowed to the actual topography, available right of way, and traffic volumes of the particular network under study to make the analysis feasible. Network-wide bicycle volumes are not computed in typical travel demand modeling, however, so the final component of the dissertation is an empirical bicycle volume model that can estimate directional volumes for all streets from comparatively sparse count observations.

4.1 Introduction

Bicycles offer a compelling alternative to the congestion, air quality, obesity, and infrastructure funding problems associated with private automobile travel [73, 29]. Much of automobiles' dominance in the United States is the direct result of decades of targeted planning and investment [53]. By the same token, the negligible and essentially flat bicycle mode share over the past several decades [55] can be traced to a lack of corresponding investments in attractive and safe [55, 184] cycling infrastructure.

Computational advances, combined with the modest scope and cost of typical bicycle facilities, means that comparable bicycle infrastructure planning and investment could be realized relatively quickly, except that relevant data on bicycle usage patterns at the street level is lacking, especially in the United States. Specifically, planners need spatially comprehensive estimates of bicycle usage to perform comprehensive safety and cost / benefit analyses [140, 104, 72]. Without reliable volume estimates, bicycle projects are at a disadvantage when competing with motorized projects, despite the enormous potential for long-term societal cost savings associated with increased bicycle use [34].

Existing volume estimation methods can be roughly divided into two groups: forecast and empirical. Forecast volumes can be produced by either regional travel models that consider trip generators and attractors to model network usage by travel mode or by developing localized regressions of surrounding land use metrics, topography, and facilities. Forecasting methods share the strengths of being applicable to future scenarios that necessarily lack direct empirical data and of being able to assess latent cycling demand which is not, by definition, revealed by direct observation of bicycle counts. Localized regression methods are helpful in very specific planning studies and to analyze general determinants of cycling, such as proximity to a major university [75], however they are not suitable to estimate internally-consistent network-wide volumes. Regional travel models are designed to estimate network-wide volumes, though typically on only a subset of roads and with a focus on automobiles [118].

Automobile volumes are relatively insensitive to changes in weather, the subtleties of facility condition, or topography. Bicycles, however, are quite sensitive to steep grades, poor pavement quality, presence of gravel, inclement weather, and a host of other highly disaggregate factors [44]. As a result, bicycle route choice is only somewhat determined by travel time, and since bicycle congestion is rarely a factor in the contemporary United States, traditional delay-based traffic assignment tools are of

little relevance. Treatment of bicycles in even very sophisticated travel models remains limited [24], which motivates the direct incorporation of observed data into the volume estimation process. Unfortunately, these count observations are not immediately policy-relevant. They must first be factored into long-term representative values, such as average annual daily bicycles (AADB), and this task remains an active area of research [128].

Cities with extensive counting infrastructure, such as Arlington, VA, Portland, OR, or Boulder, CO can address the temporal factoring problem more or less directly by examining multiple continuous volume histories over a period of years. Permanent counting hardware will still produce data anomalies over the span of a year, however, and from a statistical point of view, even a year’s data is only completely representative of that historical year. Nation-wide, continuous bicycle count stations like these are also quite rare, which precludes such direct and potentially non-parametric factor development in most localities. This equipment can also only be feasibly installed in a few locations, and the additional counts typically undertaken to improve spatial coverage will still present a temporal factoring challenge as a result of their limited duration.

Generic temporal factors can be used in the absence of continuous data and have been developed [8], however such factors are still unable to capture highly localized behavior such as differing morning and evening peaks between the two directional links of a single street. Temporal factoring at one or more count locations is really a special case of the larger task of transferring count data from a particular location and time to a different location and time, i.e. combined temporal and spatial factoring, that is required to estimate empirically derived bicycle usage at a location some distance away from direct observations. Route choice studies and travel demand models are the closest analogs to spatial factoring in the cycling literature, though each has only a limited ability to forecast bicycle usage patterns with enough specificity to prioritize

infrastructure investments.

Bicycle route selection has been studied, producing sophisticated though typically locality-specific algorithms [157, 116, 25, 83, 143, 132]. User preferences, as revealed through GPS traces over the study period provide a particularly rich data set for the local calibration of route selection codes. By incorporating these routing algorithms into regional metropolitan planning organization (MPO) travel models, network-wide volumes can be estimated, however these models typically operate at the relatively coarse spatial resolution of a traffic analysis zone (TAZ) [75], and estimate bicycle mode share in a cursory fashion at best.

Unobservable long-term characteristics of both individual directional links and the network as a whole are the policy-relevant values sought by transportation decision makers. MPO models can provide a logical starting point for estimating this information by treating directional link-specific volumes as weakly informative Bayesian prior estimates that can be updated by empirical data. The following section presents a parameterized multi-level Bayesian model of bicycle usage designed for a typical small to medium-size city with a basic four-step travel model and sparse bicycle count data. The proposed framework leverages all existing, if limited, investments in bicycle data collection and volume estimation into a unified framework to estimate posterior distributions of relevant parameters, even on the majority of links that lack any direct observational data. Sample results will follow based on a small set of count observations taken over the last three years in the Charlottesville, VA MPO region, home to approximately 200,000 people.

4.2 Method

The framework estimates network-wide and directional link-specific parameters, given in Tables 4.1 and 4.2 respectively, in three distinct phases: stochastic routing, temporal factoring using MCMC, and spatial factoring. Figure 4.1 outlines the flow of




Figure 4.1: Principle data sources and processing steps.

Table 4.1: Network-Wide Variables

symbol	definition
p_c	Proportion of a week's trips on an directional link that occur during commute days (weekdays)
p_{bc}	Proportion of commute-day trips explained by the base volume sinusoid rather than morning and evening peaks
p_{bnc}	Proportion of non-commute-day trips explained by the base volume sinusoid rather than morning and evening peaks
T_m	Effective temperature where $f_t = P_m$; parameterizes the temperature factor F_t
T_{mc}	Effective temperature where $f_t = 0.5$; parameterizes the temperature factor F_t
s_p	Scaling value for the precipitation factor
h_m	Peak decimal hour of the morning commute
h_{ev}	Peak decimal hour of the evening commute
h_{min}	decimal hour when the base demand curve is zero, defined as 3.75
τ_p	Precision of the morning and evening Gaussian peaks

data between components of the framework. This section will first describe how the network is assembled from open-source data and the bicycle-specific routing algorithm used, since these are necessary to compute $AADB_{j_o}$. Next the temporal factors and complete MCMC model will be described, followed by the spatial factoring process to produce network-wide results.

4.2.1 Network and Routing

Open street map (OSM) [133] provides freely available global transport network data that is continuously updated by millions of volunteers and includes relevant metadata, such as the type of bicycle facilities present on a network link. SQL commands within an open-source PostGIS database are used to split and reassemble

Table 4.2: Directional Link-Specific Variables

symbol	definition
$AADB_j$	Average Annual Daily Bicycle Trips, assumed gamma distributed
$AADB_{j_o}$	Initial AADB estimate from the MPO travel model and stochastic routing
α_{AADB_j}	Shape parameter of the $AADB_j$ distribution
β_{AADB_j}	Rate parameter of the $AADB_j$ distribution
p_{m_j}	Proportion of total peak area contained in the morning peak for directional link j
α_{pm_j}	α parameter of the beta distribution of p_{m_j}
β_{pm_j}	β parameter of the beta distribution of p_{m_j}
$\rho_{j,k}$	correlation between directional links j and k
$\gamma_{j,k}$	AADB ratio between local directional link j and remote directional link k
l_j	Length of link j
l_{e_j}	Effective length of link j

the raw GIS lines from OSM into a topology suitable for route prediction. The resulting network has a vertex representing every approach street to intersections, and directional links connecting adjacent intersections and valid movements within intersections. Links are directional and indicated with the subscript j , such that a typical two-way street contains two links. This topology creates a larger computational burden than the more basic approach using a single vertex to represent an intersection, but has the benefit that intersection-specific bikeability can be captured independently of the nature of the approach streets. For example, a difficult left turn can be modeled independently of relatively simple straight and right turn movements.

Route selection is central to estimating $AADB_{j_o}$ as well as the correlation between directional link pairs $\gamma_{j,k}$. Bicycle-specific routes are computed using a least-cost algorithm, where cost is given as an effective link length l_{e_j}

$$f_{elevation} = \begin{cases} (\frac{\partial z}{\partial x} - m_o)^2 & \text{if } \frac{\partial z}{\partial x} > m_o \\ 0 & \text{if } \frac{\partial z}{\partial x} \leq m_o \end{cases} \quad (4.1a)$$

$$l_e = F_{facility} F_o \left(l + m_o^{-2} \int_0^l f_{elevation} dx \right) \quad (4.1b)$$

where $F_{facility} = 0.67$ for links with bicycle facilities, to reflect cyclist preference for dedicated infrastructure, and unity otherwise [132], z is the instantaneous elevation in meters along a link from $x = 0$ to $x = l$ at the destination vertex, and $F_o = 0.5$ is the minimum ratio of effective length to geometric length. Equation 4.1 uses a quadratic increase in effective length as the instantaneous up-slope increases past -1.4%, defined here as m_o , the point at which a typical cyclist on a city bike can maintain 15kph without pedaling [194]. Effective length is normalized such that a level road has an effective length equal to its geometric length. The integral in Equation 4.1a is evaluated every 10m, which is the pixel size of the National Elevation Dataset (NED) used here [66]. These data are freely available for the entire United States, and similar files are available for most of the world.

4.2.2 Prior AADB

MPOs are required to maintain travel demand models which, at a minimum, can provide estimates of daily bicycle trips between pairs of traffic analysis zones (TAZs). Using the previously described routing algorithm, sets of links R used to travel between randomly chosen origin and destination (OD) vertices in each TAZ can be estimated, with more routes estimated between nearby TAZs to capture the greater variation in route that can be expected between randomly chosen points in nearby TAZs. By weighting these stochastic trips by the MPO expected AADB for their OD pair, a prior estimate of AADB by directional link ($AADB_{jo}$) is computed. The Charlottesville travel model contains 264 TAZs, and the number of routes per OD pair ranged from 1 to 47, leading to 452,972 total stochastic routes in the set R .

Table 4.3: Variables Applicable to Count Observations

symbol	definition
μ_i	expected Poisson mean for count i
T_{e_i}	Effective temperature in °C during count i
F_{t_i}	Temperature factor for observation i
F_{p_i}	Precipitation factor for observation i
F_{h_i}	Hourly factor for observation i
F_{c_i}	Commute-day factor for obseravation i

4.2.3 Temporal Factoring

Sufficient model complexity to estimate the set of unobserved parameters from sparse data can be achieved using a multi-level approach. Table 4.1 describes the parameters informed by count data from any directional link. Parameters $AADB_j$ and p_{m_j} are link-specific to capture directional street volumes and peak commute patterns by time of day. Despite this multi-level structure, fitting such a large number of parameters remains a computational challenge in closed form.

Markov-chain Monte Carlo (MCMC) sampling is a statistical technique to estimate complex multi-level models by sampling from the posterior distribution of the fitted model, without explicitly solving for its parameters. After an initial number of “burn in” iterations are discarded, the Markov chain(s) are considered to have converged on the desired Bayesian posteriors at which point distribution parameters can be fit to the resulting histograms. MCMC sampling is used to solve the temporal factoring problem on the subset of network links with original data by modeling the expected values of those counts using a directional link AADB ($AADB_j$) modified by a series of factors with an expected value of 1. The selected factors are a function of whether the count was on a commute day or not, the weather conditions at the time of the observation, and a parameterized hourly variation curve

$$\mu_i = AADB_{j(i)} F_{p_i} F_{t_i} F_{c_i} F_{h_i} \quad (4.2)$$




Figure 4.2: Hourly factor curve illustrating Equations 4.3b, 4.4, and 4.5 for a directional link with $p_{m_j} = 0.27$.

where $AADB_j$ is the unobserved annualized average daily volume for directional link j , the location of count observation i . Tables 4.2 and 4.3 name the remaining variables in Equation 4.2 with their functional forms defined in the subsections to follow.

The MCMC model relates expected volumes to observed counts assuming that the expected counts are Poisson distributed with mean μ_i . To reflect the relatively minor attention given to bicycles in typical MPO travel models, the prior volume estimates are assumed to follow a gamma distribution with mean $AADB_{j_o}$ and a sizable variance of four times that value. 1812 hours of automated 15-minute counts from a tube counter moved between three locations (6 directional links) and 789 hours of manual observation at 23 intersections (232 directional links) were used to estimate the results given in this work.

4.2.3.1 Hourly and Commute-Day Factors

Hourly variation is modeled as morning and evening Gaussian peaks overlaid on a base volume sinusoid with a 24-hour period. This curve, represented as the sum of the three components given by Equations 4.3b, 4.4, and 4.5, and shown in Figure 4.2, is integrated over the duration of a particular count observation to obtain the proportion of the daily total for that count window:

$$p_b = \begin{cases} p_{bc} & \text{if commute day} \\ p_{bnc} & \text{if non-commute day} \end{cases} \quad (4.3a)$$

$$f_{hb_i} = \frac{p_b}{24} \left(1 - \sin \left(\frac{\pi}{12} (h_{min} + t) \right) \right) \quad (4.3b)$$

$$f_{hm_i} = p_{m_j} (1 - p_b) dnorm(h_m, \tau_p) \quad (4.4)$$

$$f_{hev_i} = (1 - p_{m_j}) (1 - p_b) dnorm(h_{ev}, \tau_p) \quad (4.5)$$

$$f_{h_i} = f_{hm_i} + f_{hev_i} + f_{hb_i} \quad (4.6a)$$

$$F_{h_i} = \int_{t_i}^{t_i+d_i} f_{h_i} dt \quad (4.6b)$$

where F_{h_i} is the resulting hourly factor, t_i is the start time of count session i with duration d_i , and the remaining inputs are named in Table 4.1. Commute and non-commute days exhibit quite different hourly variations [75], so hourly parameters are fit for commute and non-commute days separately.

The hourly factor does not take into account the different daily average volumes

between commute and non-commute days, so an additional term, p_c is necessary to describe the proportion of trips during the work week. The term is beta-distributed, and thus bounded (0,1), and so must be converted to a daily factor using Equation 4.7

$$F_{c_i} = \begin{cases} 7p_c/(2 + 3p_c) & \text{if commute day} \\ (7 - 7p_c)/(2 + 3p_c) & \text{if non-commute day} \end{cases} \quad (4.7)$$

4.2.3.2 Weather Factors

Expected bicycle usage is sensitive to weather effects such as precipitation and temperature [78, 61], and while there may be seasonal effects driven solely by date, such as university schedules, in this work we propose a normalized pair of weather factors to simultaneously account for instantaneous conditions and seasonal variation. Additional factors, such as a purely seasonal component, may be appropriate for inclusion given sufficient data to support the additional model complexity.

Equation 4.8b gives the functional form of the precipitation factor F_{p_i} for observation i

$$f_{p_i} = \frac{1}{1 + s_p P_i} \quad (4.8a)$$

$$F_{p_i} = \overline{f_p}^{-1} f_{p_i} \quad (4.8b)$$

where s_p is the network-wide scaling value for hourly precipitation P_i in mm/hr. Daylight weather conditions for the past two years are used to compute the factor normalization $\overline{f_p}$ in order to give the precipitation factor an expected value of 1. Weather data were taken from the freely available North American Regional Reanalysis (NARR) [117] which retrospectively estimates weather conditions every three hours nation-wide for the past several decades on a 32km grid.

Temperature, relative humidity, and surface wind velocity are combined into an

effective temperature T_e , computed as a piecewise continuous function of the wind chill [31] below 10 °C, its maximum value, the heat index [32] above 26.7 °C, its minimum value, and a spline fit over the intervening range matched to the slope at each boundary. The effect of this temperature modification is that cold days are viewed as even colder when it's windy and hot days are hotter when it's humid. Despite collapsing these three metrics into a single effective temperature, the effect on cycling likelihood remains less straightforward than an obvious deterrent, such as precipitation.

Brandenburg et. al. [23] found that at very warm temperatures, cycling volumes remained at the peak levels of more moderate weather and retained a small but non-zero tail on the coldest days. This trend can be described by a sigmoid function

$$s_t = \frac{\ln(1/P_m - 1)}{T_m - T_{mc}} \quad (4.9a)$$

$$f_t = \frac{1}{1 + e^{s_t(T_e - T_{mc})}} \quad (4.9b)$$

where T_{mc} is the effective temperature when $f_t = 0.5$ and T_m is the effective temperature where $f_t = P_m = 0.95$. The final temperature factor F_{t_i} for observation i is normalized in the same way as precipitation using NARR data and can be seen in Figure 4.3.

$$F_{t_i} = \overline{f_t}^{-1} f_{t_i} \quad (4.10)$$

4.2.4 Spatial Factoring

The MCMC model defined by Equation 4.2 provides a means to address the temporal factoring of individual observations into representative parameters, however greater insight can be offered by using these values to update prior assumptions about AADB on all network links. The posterior distributions of $AADB_j$ and p_{m_j} , the two values estimated by directional link rather than globally, can be represented by




Figure 4.3: Relative likelihood to bicycle (F_t) vs. effective temperature (T_e) using MCMC-sampled parameter values. The temperature factor F_t is parameterized by effective temperatures T_{mc} and T_m as defined by Equation 4.9.

gamma and beta distributions respectively. The final step of the framework spatially extends available count data by updating the prior $AADB_j$ and p_{m_j} distribution parameter estimates ($\alpha_{AADB_j}, \beta_{AADB_j}, \alpha_{p_{m_j}}, \beta_{p_{m_j}}$) for all directional links. In this Bayesian updating step, prior values are either MCMC sampled posteriors fit to their respective distributional forms, or the original prior estimates from the MPO travel model ($AADB_{j_o}$) for the majority of links that lack direct observations.

Values used for the actual updating are drawn from the MCMC posteriors, since only these contain additional information above and beyond $AADB_{j_o}$. In the case of gamma and beta distributions, updating is a simple sum of each parameter [21], however a correction must be made for the expected ratio of AADB between the “local” directional link j being updated and the “remote” directional link k with temporally factored data, as well as for the estimated route correlation between the two directional links.

Directional link volume ratios ($\gamma_{j,k}$) are computed as a simple comparison of



Figure 4.4: Network subset showing reference locations A and B. The width of the dark road overlay indicates relative correlation with trips traveling to the left at location A. Circles show nearby count locations scaled by the amount of data.

$AADB_{j_o}$ between local and remote directional links.

$$\gamma_{j,k} = \frac{AADB_{j_o}}{AADB_{k_o}} \quad (4.11)$$

Correlation between directional links ($\rho_{j,k}$) is computed as the AADB-weighted percentage of stochastic routes R that utilize both links

$$\rho_{j,k} = \frac{1}{AADB_{j_o}} \sum_{r \in R_j \cap R_k} AADB_{r_o} \quad (4.12)$$

where R_j is the set of stochastic routes that pass through the local directional link and R_k the routes that use a remote directional link k . Figure 4.4 graphically demonstrates directional link correlation with respect to a single network link. Computing Equations 4.11 and 4.12 is the most expensive step in the framework, however it need only be done once. Only a change to the network, routing algorithm, or MPO OD estimates necessitates a recomputation.

Equations 4.13 and 4.14 describe the actual updating step, which does not present

computational challenges:

$$\alpha_{AADB_j} = \sum_{k \in E} \rho_{j,k} \gamma_{j,k} \alpha_{AADB_k} \quad (4.13a)$$

$$\beta_{AADB_j} = \sum_{k \in E} \rho_{j,k} \beta_{AADB_k} \quad (4.13b)$$

$$\alpha_{pm_j} = \sum_{k \in E} \rho_{j,k} \alpha_{pm_k} \quad (4.14a)$$

$$\beta_{pm_j} = \sum_{k \in E} \rho_{j,k} \beta_{pm_k} \quad (4.14b)$$

where E is the set of all directional links.

4.3 Results




Figure 4.5: Posterior distribution for the mean of the morning and evening commute peaks. Kernel density represents a continuous histogram of parameter distribution.

The framework was able to leverage the modest amount of bicycle available for the Charlottesville area into reasonable posterior distributions for the parameters

outlined in Tables 4.1, 4.2, and 4.3. Figure 4.5 shows the distribution of the peak hour parameters h_m and h_{ev} only since the minimum volume hour h_{min} was fixed at 3:45 AM to ensure that the sinusoid represented a base daily demand and not some other component of the volume signal. As expected, some parameters were identified quite precisely, such as h_m and h_{ev} , while others, such as p_{bnc} and p_c , exhibit greater posterior variance.




Figure 4.6: Posterior distribution for beta-distributed global parameters

Figure 4.6 shows the posterior distribution of the beta-distributed parameters p_c , p_{bc} , and p_{bnc} . The base sinusoidal curve represented a greater portion of trips on non-commute days rather than during the week, as evidenced by the greater expected value of p_{bnc} compared with p_{bc} . The preponderance of trips also occurred during commute days, illustrated by the expected value of p_c that is much greater than 0.5.

Figure 4.7 shows 15min tube counter observations for a single day at location A, identified in Figure 4.4, aggregated into 1 hour counts. The violin plots labeled MCMC show the posterior distribution of hourly-aggregated μ_i for the same conditions. This result is typical of the agreement between the parametric estimate μ_i and observed data. The principal signals are well represented, though unmodeled variances will always exist. Location A is between an area of student housing and the University of Virginia, so the notable discrepancy in Figure 4.7 between the observed counts




Figure 4.7: Hourly aggregated observations and modeled results for a single day at location A (shown in Figure 4.4)

and μ_i at 11:00 PM can be explained by the unique travel behavior of undergraduate students, which is not seen over the region as a whole.

Figure 4.8 shows the prior estimate of AADB and the final posterior distribution for location A from Figure 4.4. Since several weeks of continuous observations were available for this location, the posterior variance is quite small, as would be expected.

Figure 4.4 identifies a second directional link, location B, approximately 1km from A and on the same street heading the same direction. Figure 4.9 shows the prior and posterior AADB distributions for location B plotted on the same vertical scale as Figure 4.8. As expected, there is a larger variance in the posterior estimate at B since there are no direct count observations in the immediate vicinity of B and the trip correlation between A and B is only 17%. There are count locations beyond the bounds of Figure 4.4 which also influenced the updated values at B, though the counts at A were the most influential.

A conventional volume study typically involves two-hour counts at the same time




Figure 4.8: Prior ($AADB_{j_o}$) and posterior ($AADB_j$) distributions and a conventionally factored estimate for location A (see Figure 4.4), a directional link with two weeks of continuous counts.

of day over several days of the week. By taking the average volume observed between 4:00 and 6:00 PM at location A for the Wednesday depicted in Figure 4.7 as well as the Monday and Friday of the same week, conventional factors [8] can be applied to evaluate the agreement of this more basic method with the framework of this study. As indicated by the arrow labeled “NBPD Extrapolation” in Figure 4.8, the conventionally factored AADB estimate is nearly twice the posterior mean at location A. Applying the same Mon-Wed-Fri 4-6pm count window to a location on Main St. outside the bounds of Figure 4.4, the NBPD factors yielded directional AADBs 122% greater than the posterior mean headed downtown and 20% less than the posterior mean in the opposite direction, despite an only 6% difference between the respective posterior AADBs at this location. While these marked discrepancies highlight the performance of the MCMC temporal factoring, the primary contribution of this work is that outputs of comparable specificity are also provided for uncounted network links where conventional methods are unable to provide any information.




Figure 4.9: Prior ($AADB_{j_o}$) and posterior ($AADB_j$) distributions for location B (see Figure 4.4), a site 1km from location A, but that has no data.

4.4 Implications

As transportation planning becomes ever more data-driven, the bicycle mode suffers from a significant information gap. While the framework presented here could be applied to automobile or even pedestrian movements, it was designed to address the particular challenges of bicycles, which have neither the spatial constraints of pedestrian movements, nor the volume and predictability of automobiles. Toward this end, the framework is able to produce directional link-specific estimates of bicycle usage patterns over an entire region from a modest dataset collected by volunteers and a single tube counter deployed for several months.

In addition to providing more accurate temporal factoring at counted locations, the framework supports systematic project evaluation by enabling direct comparisons between these locations and those for which no count data exists. Since the final outputs are represented by distributions, planners and software systems are also able to move beyond the typical assumption that all transportation model outputs are equally valid. The framework will be used, as presented, to enhance the long range planning process for bicycle infrastructure in the study area. For regions with more

complete data, several avenues exist for further enhancement.

Regions with long bicycle count histories would presumably be interested in an inter-annual growth term, both to improve model fit and as an important policy benchmark that is difficult to estimate by other means. Areas with very small bicycle mode share or short duration analyses could feasibly employ a constant growth rate, while longer analysis periods or more established mode shares would necessitate a more complex formulation, such as a sigmoid to model market saturation. An additional factor level can also be added for parameters with a scope between link-specific and network-wide. These parameters would typically be associated with a special trip generator, such as a major employer or university, that has a significant but localized impact on bicycle usage patterns.

Looking forward, the complete nature of the results obtained from this method facilitate data-driven investment optimization schemes that would otherwise be infeasible. This is especially important for bicycle travel, since the small streets typically excluded from regional travel models often serve as significant bicycle thoroughfares. Safety is one of the most-cited reasons more people do not travel by bicycle. Optimizing scarce infrastructure funding using empirical Bayes methods based on this comprehensive exposure data is critical not just to individual safety, but for society as a whole to realize the compounding benefits of a significant mode shift toward bicycles.

4.5 Acknowledgements

The authors would like to thank Sarah Rhodes at the Thomas Jefferson Planning District Commission and Austin Angulo for their integral involvement in this project.

CHAPTER V

Conclusions and Future Work

Evolving travel demand and mode preferences in the U.S., combined with unresolved funding challenges, are ushering in a period of change in transportation infrastructure. Reconfiguring roadways to facilitate changing travel behaviors while keeping existing facilities in acceptable condition calls for a more comprehensive assessment of the costs and benefits of infrastructure projects than has been practiced to date. The work presented contributes several components to this ideal framework: network-wide multi-objective pavement management, right of way allocation between competing modes on a total cost basis, and a network-wide empirical bicycle volume synthesis. Taken together, these components enable roadway maintenance and use-phase costs and GHG emissions to be considered in an integrated manner revealing otherwise unknown quantities, such as the spatial opportunity cost of curb parking relative to bicycle use which quantifies the auto subsidy such parking represents. No modeling exercise can be truly comprehensive, however, and the work presented in chapters II, III, and IV is focused on adding specificity and GHG emissions to more established transportation planning methods based on travel time, vehicle operating, and maintenance agency costs.

These contributions are predicated on the idea that sustainable livability is the ultimate goal of infrastructure planning, i.e. providing the highest quality of life by

maintaining environmental and economic productivity over the long term. When considering total costs and benefits, health and safety costs are the metrics most notably lacking in existing frameworks. The average annualized personal costs of being obese, including value of lost life, are \$8,365 for obese women and \$6,518 for obese men [45], and 50% of the US population is expected to be obese by 2030 [45]. The magnitudes of these costs are such that when the exercise benefits of increased bicycle and transit use and the comparative risks of injury for people in and out of cars¹ are included in a comprehensive framework, a very different prioritization of street use is likely to emerge. Since non-motorized infrastructure is also much more affordable to build and maintain than its motorized equivalent, significant improvements to both fiscal and personal health could result from such a shift.

While the health benefits of active transport modes likely dwarf mobility-related metrics in magnitude, the fortunately rare nature of crashes makes them difficult to study, and physical activity represents only one of many conflated influences on personal health. Unfortunately, these research challenges are only exacerbated in the case of bicycles due to the pervasive underreporting of minor crashes and the relationship between cycling and overall lifestyle choice. Despite these challenges, since safety concerns are likely the most significant single deterrent to a greater bicycle mode share [94], an investment policy to address this problem is needed to realize the potential benefits of a significant mode shift. Chapters II, III, and IV present several necessary steps toward including crash impacts in a total cost analysis, however further refinement and additional data are also needed to realize this goal.

Work is ongoing at present in two major areas. The first is a trial implementation of the bicycle volume model described in chapter IV for the greater Charlottesville, VA area in partnership with the local metropolitan planning organization (MPO), the Thomas Jefferson Planning District Commission. The second is a more fundamental

¹The US Environmental Protection Agency currently values a human life at \$7.4 million in 2006 US dollars [176].

investigation into analyzing vehicle movements using video sensors. The first goal of this work is to collect turning movements by vehicle type, including bicycles, using existing traffic signal detection cameras at intersections and over a larger area using a dispersed network of video sensors. Related work at a finer scale will precisely quantify vehicle trajectories and interactions with the goal of identifying conflict situations even when no actual crashes are observed.

5.1 MPO Travel Model Integration

The empirical volume model described in chapter IV has been implemented for the Charlottesville area as a trial case. Charlottesville, like most small to medium size urban areas, relies on a travel demand model largely supported by its state Department of Transportation (DOT), VDOT. State DOTs are concerned with state-wide mobility, and since bicycle transportation is primarily confined to the city scale, these models are typically not robust with respect to non-motorized travel. Another reason for this limitation is that the algorithmic structure of such traditional zonal travel models does not readily accommodate the more nuanced factors that influence non-motorized mode choice such as weather, time of day, crime, and myriad other influences [184]. The model presented was designed with these limitations specifically in mind, however, and so can still enhance the data available to MPOs in several ways.

Zonal models typically only include streets classified as collectors or larger. This simplification is appropriate for estimating auto volumes, especially with the understanding that traffic calming will be applied to residential streets that experience high volumes in order to divert traffic back to larger facilities. This assumption is inappropriate for bicycles, however, since they often prefer low-volume streets and their more efficient use of space precludes bicycle-only congestion. The result is that a practical bicycle network must include all street types below limited-access freeways as well as off-street paths and other connections inaccessible to motor vehicles. The model

presented in chapter IV allows the estimated bicycle trips generated by traffic analysis zone (TAZ) to be routed on a complete bicycle network, and thus to present estimated volumes for bicycles on all streets, as shown in Figure 5.1. These estimates represent a significant advancement in the region’s knowledge of bicycle movements, however they are still only based on a few months of observations, and so more informative results will have to wait on the collection of additional count data. Along with using observations taken over a longer time period and covering multiple seasons, several as yet unexplored structural enhancements could also improve model results by more directly integrating the bicycle model with the zonal travel demand model.

Feedback loops are critical in the 4-step modeling process to account for phenomena such as induced demand. In the case of the bicycle model, incorporating a feedback loop with the zonal model would allow the final empirically-derived bicycle volume estimates to calibrate the bicycle origin-destination (OD) matrix of estimated daily trips between TAZ pairs the MPO model provides as an input to the bicycle model. These OD trip volumes cannot be directly estimated from feasible amounts of count data, so they are instead computed in two distinct steps: trip generation and trip distribution. Trip generation rates are derived from socio-economic and land use data for each TAZ, while trip distribution is a function of the expected difficulty in traveling between a pair of TAZs, called the impedance. This process is far from exact, however, so incorporating feedback based on empirical data could greatly improve the final OD estimates, and perhaps more importantly, the estimated bicycle volumes for areas of the network for which there is little or no count data which therefore rely heavily on this prior estimate.

Trip distribution can be improved by using the stochastic routing code developed for the bicycle model to produce a zonal impedance matrix specific to bicycles. This will not change the total number of bicycle trips the model estimates, however it will improve the allocation of those trips between appropriate origins and destinations,

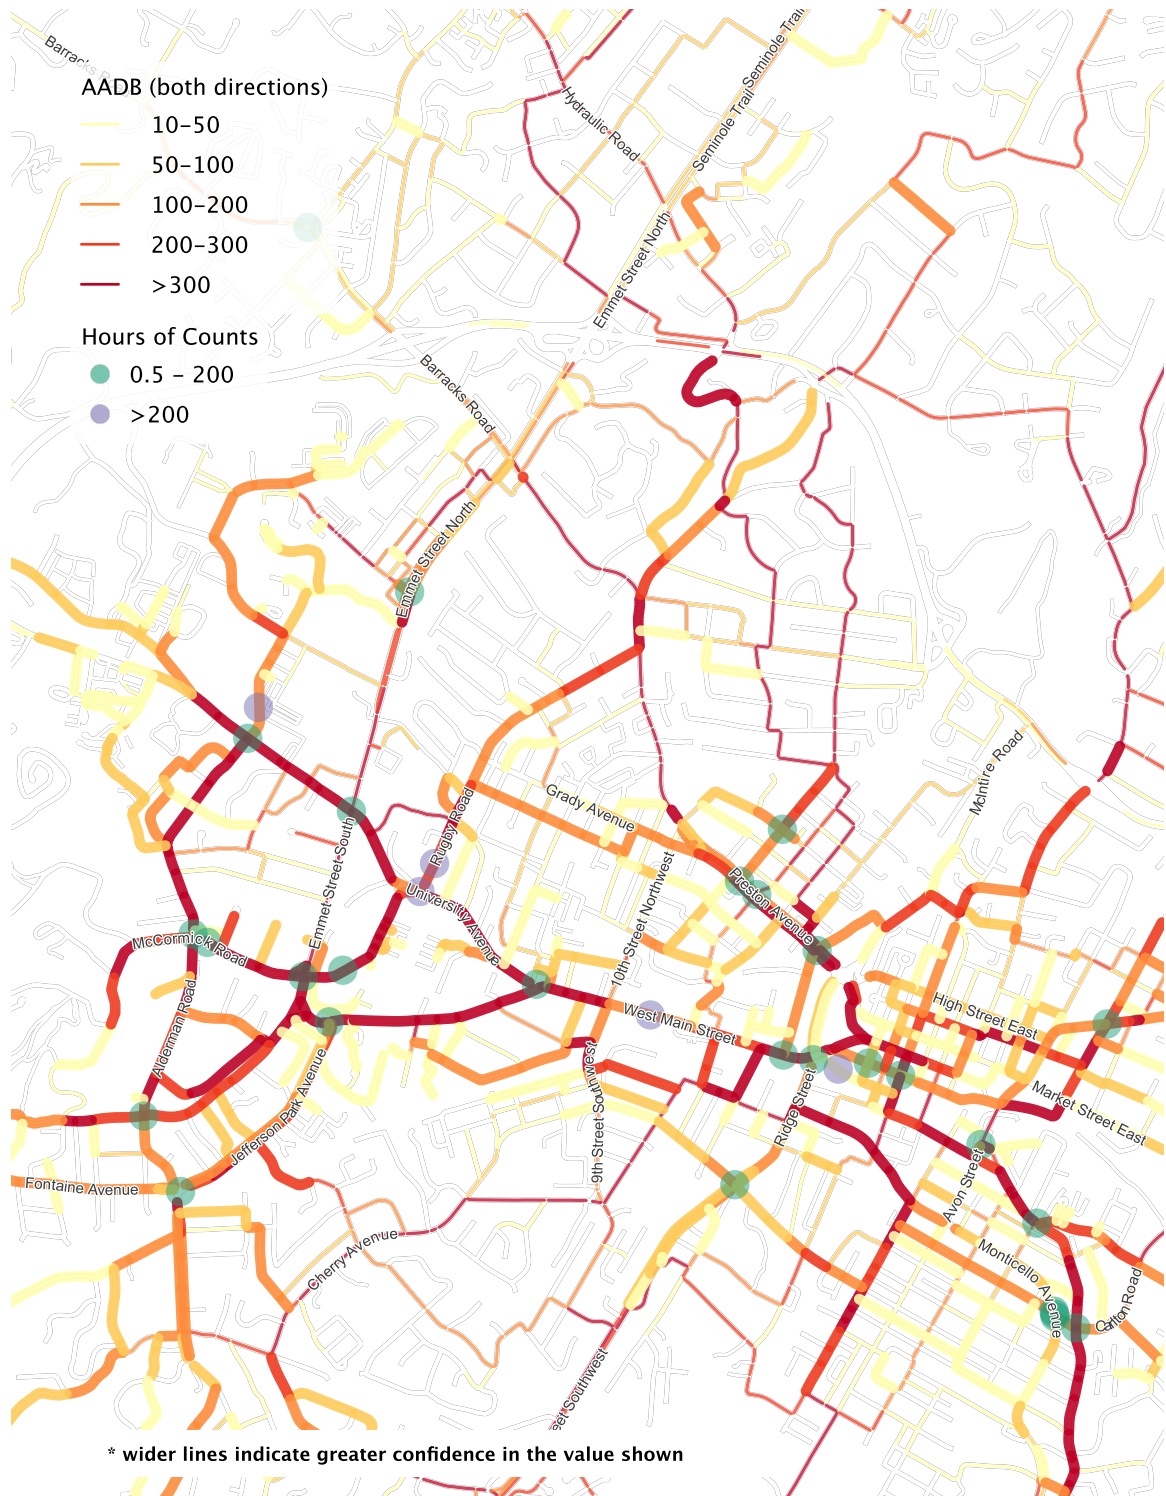


Figure 5.1: Bidirectional average annual daily bicycles (AADB) where wider lines represent lower variance in the final volume estimate and circles indicate data collection points.

thereby leading to a more realistic prior estimate of bicycle volumes. Once this more bicycle-aware distribution component is in place, the full feedback loop can be applied in order to calibrate the mode split model that is responsible for bicycle trip generation rates. An additional benefit of this step is that more accurate estimation of bicycle mode split will improve the estimates for other modes, thereby enhancing the value of the entire travel demand model. Integration with the full travel demand model is also essential from a planning standpoint because as an empirical model, the framework presented in chapter IV is necessarily constrained to the present day. Long-range planning is done with entirely forecast data, so the contribution of the bicycle volume model in this context is its ability to calibrate more basic models in a data-poor environment such that future scenarios can be evaluated as accurately as possible. Forecasting travel demand decades into the future also means considering changing public attitudes toward the various transportation modes.

5.1.1 Non-static mode preference

Bicycle mode share is increasing in the US after a long period of relatively flat or declining values [172], and this is not captured by current mode split models that have been calibrated to reflect static user attitudes. Even a modest annual increase in bicycle mode share will lead to quite significant values at the typical 30-year modeling horizon for long range transportation plans. The decision to go by bicycle is informed by myriad factors [184] beyond the simple cost and travel time values typically included in travel demand models. Along with transient influences such as trip purpose, weather, and time of day, there are other less variable factors that influence this choice. Some of these influences can be included in travel demand modeling, such as access to a private motor vehicle, but even that fails to capture the habitual and cultural factors that influence the choice, such as the prevalence of bicycle commuting in one's peer group, and whether one has access to a bicycle in

good repair. All of this suggests that modeling travel mode split at the typical 30 modeling horizon of a long range transportation plan requires a careful evaluation of these trends. Given several years of observations, an annual growth term can be included in the bicycle volume model to capture a non-static city-wide trend in bicycle volumes. Assuming anecdotal reports of increasing cycling are true, such a term actually becomes necessary to preserve the temporal model fit over multiple years of data since this long-term temporal variation would otherwise be perceived as noise in the many other parameters fit by the temporal factoring operation.

By parameterizing this model component as an annual growth rate and explicitly incorporating it into the mode split model, scenarios for arbitrary future years can be more realistically considered. Depending on how large the term is and the length of the forecast period, a more sophisticated functional form could also be employed. There will certainly be saturation at some bicycle mode split well below 100%. In either case, census data can inform a reasonable Bayesian prior assumption of annual growth, though empirical data will again be necessary for its final computation since the journey to work information available in the census is not necessarily representative of total cycling [28]. The National Household Travel Survey [55] provides much more relevant information, however the infrequent nature of this study makes it inappropriate for quantifying recent trends.

5.1.2 Modeling Additional Modes

While the empirical volume model presented here was developed specifically to overcome the limitations of typical travel demand models in estimating bicycle volumes, many of the key concepts apply equally to other travel modes. Charlottesville’s travel model only estimates motor vehicle volumes as an annualized average value for a 12-hour day. Moving to an hourly model able to more accurately assess peak periods requires significantly more input data than a daily model, and this data gap is quite

similar to the bicycle data gap addressed by the model in chapter IV. While that model was developed to address the challenge of bicycle data synthesis, the only significant difference between it and a motorized model is the trip distribution and routing steps. As with bicycles, these results will require a smaller but still significant amount of count observations. To reduce the cost of obtaining these data for all modes, sensors other than traditional tube counters could be deployed. More advanced sensor networks could also collect more detailed vehicle movement information than simple counts, which would be difficult to analyze in traditional frameworks but could be directly incorporated in a multi-modal version of the chapter IV volume model. The following section proposes several such data collection efforts based on low-cost or existing video sensors.

5.2 Video Data Collection

Work is ongoing to use video sensors to track the movements of individual bicycles and motor vehicles. This work aims to develop a unified software architecture to address two separate problems. The first is a lack of continuous counts of vehicle movements by type. Continuous information is critical with respect to highly variable travel modes, such as walking and bicycling, however installing dedicated sensors city-wide to collect this information is impractical. The second problem is the lack of accurate crash reports for incidents involving bicycles and pedestrians, and the statistical difficulty inherent in the scarce nature of these events, even with perfect records.

Installing and maintaining dedicated sensors in the roadway is expensive and even purpose-built hardware can have significant inaccuracy [127]. Video sensors can be placed out of harm's way and are quite flexible in their detection capabilities. In contrast to dedicated sensors, cameras are primarily software rather than hardware limited, which presents a significant and low-cost opportunity to apply additional

processing to existing video feeds. Cameras are necessarily limited in their spatial coverage and available mounting locations, however, so multiple units may be required to view even a single intersection.

The task of matching moving objects across multiple video sensors varies in nature according to the separation of the sensors. Vehicles, including bicycles, traveling between nearly or partially overlapping camera views can be matched through trajectory analysis, since there is relatively little ambiguity when considering orderly traffic flows. As the sensors are placed further apart the travel time between sensors becomes more uncertain, and a vehicle may only appear in a subset of the total cameras, as is the case at an intersection with a camera for each approach leg. As these sources of variability grow, the matching task shifts from one of trajectory analysis to matching the visible properties of the vehicle itself. The ideal sensor array ultimately depends on the type of analysis desired. Dense overlapping camera views can be used to establish precise vehicle extents and trajectories to analyze intersection safety by computing safety surrogate measures of potential crashes, while matching vehicles over a large dispersed network of cameras can provide detailed travel patterns beyond the scale of an individual intersection.

5.2.1 Bicycle crash risk and facility design

Increased physical activity is one of the primary benefits of a transportation mode shift away from the private automobile, however its direct inclusion in an integrated modeling framework is either impractically data-intensive or merely speculative and so will not be considered here. Crash injury risk, however, is already part of transportation engineering practice and can be readily incorporated into a discrete planning framework provided the necessary data are available, namely risk and exposure by mode. Automobile crash safety has improved drastically in the past several decades [98]. Some of the improvement in crash risk per distance traveled is due to geometric

and other infrastructure modifications, however a great deal is the result of these vehicle design advances to cope with the collisions typical of high speed suburban crashes including a 17% reduction in crash injury risk per distance traveled from model year 2000 to 2008 vehicles alone [69]. This progress has only been helpful for people in cars, however.

For cyclists and pedestrians, the reduced visibility typical of newer cars only exacerbates the fundamental problems of mixing high-speed traffic with unprotected and relatively unexpected road users (given the low bicycle mode share typical of the U.S.) Motor vehicle speeds above 30kph are associated with more severe injuries when striking unprotected road users [102], as are collisions with sports utility vehicles, pickup trucks, and vans in comparison to passenger cars [49]. Since there is no practical bicycle or pedestrian equivalent to the armoring a modern vehicle represents, reducing crash risk for these road users will have to be achieved through some combination of improved awareness, better crash avoidance through electronic driver enhancements such as machine vision to detect bicycles [37], and infrastructure improvements. Crashes that damage motor vehicles or severely injure people are well documented, though often in uncoordinated databases, so development of safety performance functions (SPFs) appropriate for network-wide analysis of automobile crashes is quite feasible [65]. Travel demand models and DOT automobile count studies provide risk exposure for these studies in the form of vehicle volumes. Bicycle crashes, however, are routinely underreported, especially when the cyclist is not severely injured [96], so there is an empirical data gap from a risk perspective. The volume model in chapter IV can provide exposure values, however, which were also generally unknown, so significant progress could be made if proxy or surrogate measures for crash risk could be employed in lieu of unavailable data. This approach could be particularly fruitful when considering that cyclists make decisions based on their perceptions of risk, which may or may not be supported by empirical study of similar situations [104]. Cyclists

who must take evasive measures to avoid collisions will be dissuaded from cycling regardless of their empirical risk.

Bicycle facilities physically separated from motor vehicle traffic are known to encourage cycling, especially by new cyclists nervous about sharing street space with faster traffic [184]. These facilities can work well parallel to roadways with few or no driveways and infrequent intersections, however their treatment at intersections remains an unresolved problem [95] and can lead to situations that cyclists perceive as safe but are empirically more dangerous. In theory, dedicated signal phases for motorized vehicles and bicycles solve this problem, however the additional delay this introduces can reduce compliance among all road users. Assuming that bicycle facilities will not or should not be entirely grade separated at intersections, the task becomes a question of where and how to place potential conflict points. The vehicular cycling movement [62] favors having bicycles merge with the flow of motorized traffic upstream of intersections in order to avoid bicycle-specific conflict points that motorists are unused to. Advocates of separated facilities counter that this defeats the purpose of separated facilities and that alternative intersection designs are needed. This school of thought also points to the safety in numbers effect [93] whereby very low crash risks per distance cycled are only seen in areas with large numbers of cyclists [96] and so say that encouraging more cycling is the only means to real and lasting safety improvements since bicycle crashes will grow at a lesser rate than ridership. Because crashes are almost invariably the result of human error, intersection designs cannot be usefully evaluated outside their immediate context, i.e. the particular travel patterns and users of the facility. Observational studies that capture these behaviors are labor-intensive, but video sensing can be applied to collect detailed movements over a long time period and vastly reduce the cost of evaluating in-place infrastructure.

Video sensing has been applied to the conflict detection problem in the context of bicycles [110]. This work employed an initial automated filter to reduce the data-

storage requirements by only archiving video clips of potential interest, though that research still involved hundreds of hours of manual labor. That work only considered the limited context of a single street, but it identified one of the key challenges in this approach: the partial or full occlusion of one vehicle by other according to the vantage point of the camera. That work also failed to account for the distortion of a vehicle’s position in the ground plane inherent in orthorectifying the video frames into ground-based coordinate system.

In order to solve the occlusion problem and obtain greater spatial accuracy in vehicle movements, the work proposed here would employ multiple cameras with overlapping views from different vantage points to largely eliminate occlusion and allow a full 3d reconstruction of the scene [26, 86]. By computing vehicle height contours explicitly, accurate vehicle extents can be derived to measure precise distances between conflicting vehicles. This represents an improvement over prior work, e.g. [110]. In addition, separating moving features from fixed infrastructure would become a trivial exercise since vertical elevation would be known and the ground could be expected to maintain a constant z value. This would also make the results more robust with respect to changing lighting and environmental conditions than the 2-d background subtraction approaches necessary when using a single camera. Finally, by avoiding occlusion problems, more complex scenes such as intersections could be analyzed.

The previous work [110] identified potential conflict situations by considering the time gap between observed vehicle trajectories and a hypothetical collision. Safety surrogate measures such as this are often mentioned in the microsimulation literature in order to derive a crash risk from simulated roadways [67]. The goal of the work proposed here is to evaluate alternative intersection designs with explicit bicycle treatments in the North American context to learn about the likelihood of crashes as a result of the treatments, and whether these patterns change over time in response to greater user familiarity with them. Comparing these results with the same analyses of

unimproved intersections will provide valuable insight into the difficult and politically fraught decisions involved in reallocating urban space to non-motorized modes of transport. These studies will still be specific to individual locations, however. Their generalization depends on accurate traffic volumes, by mode, throughout the network so that SPFs can be developed that are transferable to other less studied locations. The volume model in chapter IV provides these data, but only to the extent that sufficient empirical observations are available to inform it.

5.2.2 Continuous counts

Many cities are transitioning from inductive loops to video-based vehicle detectors for actuating traffic signals. These detection cameras are installed above the traffic signal heads specifically to monitor the roadway in all conditions and must already operate continuously. Despite their being general purpose sensors, data from these cameras are effectively discarded after the signal detection process. Since the bicycle volume model in chapter IV depends on continuous counts for optimal fitting of temporal factors, and continuous counts, especially for bicycles, are generally unavailable in urban settings, these cameras represent a valuable and untapped resource for continuous collection of turning movement counts by vehicle type. Implementing this approach would only require a small industrial computer mounted in the signal control cabinet and the software proposed here.

With respect to turning movement counts, the primary disadvantage of signal detection cameras is that their fields of view will generally not overlap, and a significant area in the center of an intersection will not be visible in any view as is depicted in figure 5.2. Bicycle identification and tracking has been accomplished for a single camera view [188], and many actuated traffic signals can count vehicles already, however the equivalent task of tracking and counting individual vehicles between multiple views has not been published to the Author's knowledge. By fitting a



Figure 5.2: A rendering of the visible portions of a Charlottesville, VA intersection fitted with detection cameras covering all four approach legs.

probabilistic trajectory model to objects identified in each individual camera view, vehicles approaching in one view and departing in another can be matched based on the estimated trajectory through the intersection, physical size, and visual spectral parameters in order to record turning movement counts continuously by vehicle class. By doing this processing in real time, cumbersome storage and or data transmission requirements are avoided. Perhaps more importantly for municipalities, by retaining only aggregated count information, there are no privacy concerns beyond the existing use of the cameras for signal actuation. Deploying such a system in the field could be made relatively simple since camera calibration and orientation can be performed by identifying several known points in the image [92] and beyond those few points, precise information about site geometry, like lane markings, is not needed.

For a typical four-leg signalized intersection, meaning that there are three potential movements for each intersection approach, not counting u-turns, recording turning movements on a continuous basis represents twelve separate counts per unit time. Not only would six permanent bi-directional bicycle counters be expensive to install and

maintain, dedicated counters are also unable to capture turn information, which is particularly valuable in the context of the bicycle volume model presented in this work. By capturing turning movements, vehicle flows in the vicinity of a particular intersection can be much more readily estimated and correlations between travel on nearby network links verified. Expanding the sensor network beyond the intersection has value as well, however, since even continuous turning movement counts at a single intersection may miss significant directional correlations between nearby streets in areas with multiple alternative routes.

5.2.3 Dispersed sensing

For travel demand purposes, behavior at the level of intersection movements is below the relevant scale. Detailed behavior within the intersection is of little consequence to trip routing other than to note delay. Instead of focusing on movements within an intersection, capturing movements within the larger area can provide count information at certain locations as well as calibration for the directional-link correlation matrix central to the bicycle volume model. This problem is a logical extension of the view matching proposed above to an arbitrary number of sensors with much more distance between each view. In keeping with a low-cost approach, commodity smart phones could be used as cameras and video pre-processors. By using the computational power of the phone to identify individual vehicles and compressing this information into a visual signature for each vehicle, the data load would be small enough to utilize a cellular network to transmit these data to a cloud-based processing node that would perform the more computationally intensive probabilistic matching of visual signatures into observed routes and travel times of individual vehicles. There are examples of work in this area [155, 179], though these methods have not been practically implemented for the bicycle counting problem. Complete video footage would not be transmitted, preempting privacy concerns, and a profile view of traffic

taken mid-block would preclude license plate matching and make unskilled installation straightforward. The fundamentally distributed nature of this approach means that it would be extraordinarily scalable. Bicycle counts are often conducted by volunteers. In this scheme, rather than standing on a street corner for a few hours, volunteers could position a phone in front of their own house whenever it was convenient and the system would make use of a continually shifting suite of sensors.

This approach would not be necessary in every circumstance, however it has several notable advantages. First, by avoiding the confused environment of intersections, the robustness of the vehicle classification task is improved. Vehicles will not be turning with respect to the camera and vehicle profiles are generally mirror images while the front and back, which would be seen by signal detection cameras, look dramatically different. This system could also answer questions about drivers' and bicyclists' use of alternative routes that would be aggregated and effectively hidden in the data from the same number of traditional counters. Depending on number of sensors deployed at one time, full vehicle routes could be estimated over large areas which has implications for route selection and volume estimation.

5.3 Optimizing Investment

Transportation planning as it is practiced in the US is primarily concerned with planning investment, such that the final deliverable is a fiscally constrained list of projects. Reaching this goal can be broken into two broad tasks, each of which will be informed by the research outlined in this document. The first task is to identify weaknesses and or opportunities in the infrastructure. For automobile travel, the salient quantities are congestion and travel time. Bicycle networks in the developed world are rarely, if ever, more complete or of subjectively higher quality than their auto counterparts, however, and bicycle congestion plays a minor role. The principal metrics for non-motorized transportation are more fundamental: connectivity, safety,

and topography. After first identifying gaps and opportunities in the network and quantifying the performance of preliminary designs targeting these areas, the second task is a constrained integer programming problem to identify which proposals, implemented at particular times, will realize the greatest societal benefit from available funding streams. While the computational aspects of the second task are significant, it is the quantification of project costs and benefits to feed that optimization that is ultimately most challenging, especially when quantifying somewhat intangible concepts, such as safety.

Single use zoning and low-density land use mean that potential bicycling trips tend to be longer in the US than in countries with more compact development. Individual effort can overcome physical challenges, however, the risks associated with operating a bicycle on high-speed urban roadways designed for automobiles have to be addressed at the planning level. From this perspective, the question of safety also shifts from one of individual worry to one of many metrics that make up an overall tally of public health. While the automobile is well suited protecting its occupants from crash injury, it simultaneously deprives them of physical activity and social interaction, thereby degrading their overall health. Since sedentary lifestyles associated with high auto use lead directly to the chronic health problems driving ever-increasing health spending, a total-cost evaluation of transportation investments should include the economic value of injuries and induced fitness for all travel modes.

Even building on the work proposed above, evaluation of this hypothetical objective function will not yet be possible, especially with respect to the opportunity costs of mode shifts resulting from a particular investment. Any practical planning process will have to employ a hybrid of total-cost optimization with more traditional heuristic methods in order to make appropriate decisions in what will always be a data-poor environment due to the complexity of cities. Fortunately, high quality bicycle facilities can induce bicycle traffic just as their highway counterparts have done so effectively

for automobiles. Since bicycle trips are inherently shorter, however, there is no concern that this will encourage sprawl, which suggests that it is an appropriate means to increase the mode share of cycling even without full consequential modeling of individual projects. Proactive planning of this sort, as opposed to reactive responses to operational problems, does require a spatial analysis tool to identify key gaps in the bicycle network, however, because such efforts are necessarily more speculative and are too labor-intensive to perform manually without substantial bias.

The proceeding discussion of total cost investment optimization presumes a list of projects with corresponding costs and benefits. As a practical matter, staff time to develop and evaluate alternative projects is limited, and an idea that never makes it on the short list will not be funded, even if it would have been the most optimal use of resources. The final component proposed in this work is a spatial analysis tool to compare bicycle impedance by TAZ pair to idealized infrastructure connecting those origins and destinations in order to identify highly leveraged points for investment and therefore the most worthwhile areas to invest limited staff time in exploring more detailed concepts. Meagre infrastructure funding doesn't just affect project funds. It will limit the planning process as well, just as the very same scarcity is making project development and selection more complex. Besides making more societally optimal choices when selecting projects to fund, ensuring that highly beneficial projects are on the list by focusing limited personnel resources may have a similarly large impact.

APPENDIX A

Supporting Information for Pavement Management

A.1 Paving LCA Details

Several other normative assumptions were made to develop the LCA model. The energy used to hold mixed asphalt at process temperature awaiting transport from plant to job-site was not included because this energy is small compared to other flows and because it is difficult to estimate representative average values for such a circumstantial variable. In addition, it was assumed that the milled paving material was trucked from the job site to the mix plant as recycled asphalt pavement (RAP), however this material was not incorporated in any of the treatment mixes. While many DOTs are using increasing amounts of RAP, typically up to 20% of the mix, this practice is not yet standardized.

The model includes all stages of the life cycle from material extraction to completion of the treatment operation. Table A.1 provides the key variables used in the model development. Fine and coarse aggregates are quarried, transported to the mix plant, and heated. Coarse and fine fractions are considered separately because they are generally sourced from different locations and contain different moisture fractions.

Fine aggregates retain more moisture in stockpiles due to their increased surface area, so the energy required to vaporize the entrained water in the mixing drum varies by aggregate size [10]. A specific heat value for both the coarse and fine aggregate material is also provided in Table A.1. Aggregate moisture content was estimated based on communication with VDOT pavement engineers and assuming uncovered stockpiles. Aggregate heating at the asphalt plant includes the energy required to heat both the mineral and its entrained moisture from ambient temperature to 100°C. Additional heat is needed to overcome the latent heat of vaporization of the moisture and to heat the aggregate from 100°C to process temperature. Asphalt binder is heated directly assuming a constant specific heat. Total fuel requirements are estimated based on the efficiency of the plant burner assembly. Additional impacts result from operating the mixing equipment, delivering the material at the job site, and the pave-train placing it. Average haul distances are used in this LCA for all process steps. Milling is treated as a separate operation with GHG contributions occurring from the diesel fuel burned by the milling machine and by trucks used to haul the milled material back to the asphalt plant.

Table A.1: Model inputs for the LCA developed to estimate the GWP associated with pavement maintenance activities.

Parameter	Value	Units	Reference
Specific heat of aggregate (granite)	0.79	kJ / (kg °C)	[181]
Aggregate at quarry	4.44E-3	kg CO ₂ / kg Aggregate	[48]
Specific heat of bitumen	1.96	kJ / (kg °C)	[97]
Bitumen at refinery	0.574	kg CO ₂ / kg Bitumen	[48]
Mix temperature	180	°C	
Ambient temperature	35	°C	
density of asphalt	29.87	tonne / mi ft in	
diesel used in paving	1.41	gal / (in ft mi)	[84]
CO ₂ of burning diesel	11.090	kg CO ₂ / gal	[51]
CO ₂ of milling	0.365	tonne CO ₂ / mi ft in	[84]
Haul RAP to plant	1.424E-3	tonne CO ₂ / mi ft in / km	
Fuel oil	9E-5	kg CO ₂ / kJ heat	[48]
Capacity of a truck	16	tonnes	
Truck emissions	0.763	kg CO ₂ / km operation	[48]
Quarry to plant haul	10	km	
Refinery to plant haul	300	km	
Plant to job-site haul	80	km	
Plant burner efficiency	85	%	
Coarse aggregate moisture	8	%	VDOT
Fine aggregate moisture	12	%	VDOT
Plant electricity	2.00	kWh / tonne mixed	[10]
USA electrical power plant	4.93E-4	tonne CO ₂ / KWh	[48]

Table A.2: Summary totals for each plan for the reference single site.

id	CCI_{final}	\overline{CCI}	$Econ_{annual}$ (\$)	GHG_{total} (tonne CO_2 eqv.)
0	31.18	55.83	0	0.00
1	25.16	58.49	2226	15.78
2	62.73	79.51	4788	52.32
3	71.44	84.08	7014	68.10
4	32.15	67.15	4453	31.56
5	47.15	76.15	6679	47.34
6	59.51	83.03	8906	63.12
7	80.88	89.43	9576	104.64
8	75.00	89.43	9576	104.64
9	80.88	89.90	11467	99.66
10	71.90	87.77	11132	78.90
11	83.93	91.85	11802	120.42
12	80.88	91.85	11802	120.42
13	71.44	87.35	9241	83.88
14	80.12	89.96	11467	99.66
15	90.00	91.48	13693	115.44
16	84.33	90.37	13358	94.68
17	88.77	91.72	13693	115.44
18	91.55	93.28	14028	136.20

APPENDIX B

Justification for Network-wide Pavement Management

B.1 Introduction

Conventional life cycle assessment (LCA) and life cycle costing (LCC) techniques have a number of important limitations when applied to the management of large infrastructure networks like roadways or water systems [190]. The operation, maintenance, and upgrade of networks of public assets have temporal and spatial characteristics that do not fit well within the ISO 14040 framework in several ways [90]. First, infrastructure networks have problematic temporal boundaries because they deliver services for which there is no finite or planned end. Second, spatial boundaries present challenges when considering networks of infrastructure since similar components (e.g., road segments) can have different characteristics based on geography. Third, when LCA is combined with LCC, there are additional complications because infrastructure is typically maintained by public agencies that have different financial considerations relative to private firms. Fourth, conventional functional units often fail to consider the connections between life cycle phases e.g., a paved road (production phase) will result

in greater traffic flow (use phase) than if the road was never built. These limitations have come into focus as LCA/LCC is increasingly applied in policy settings.

To overcome some of these limitations, past work has considered individual network components over a finite time horizon [191, 164]. The service that infrastructure provides has an indefinite lifespan creating practical challenges around whether to apply discounting. Intergenerational discounting is at odds with many definitions of sustainability [79]. Since it is infeasible to plan specific actions for very long time horizons [153] and discounting is standard economic practice, there is a tension between future and present. Compounding these challenges, agency budgets vary over time in unpredictable ways [105], yet agencies lack the ability to restructure assets in order to maximize delivered value. This precludes innovative solutions to achieve performance goals, for example abandoning low-volume roadways.

Another consequence of studying individual components of infrastructure networks and then multiplying by some n number of components to yield an estimate for the entire system is that the impacts may not scale linearly. Network metrics are computed over all assets, including constraints such as annual cash flow, which means that a network-optimal plan for a given level of investment will be comprised of varying per-asset actions. These cannot be known below the network-level. Also, variations between assets must be considered in a network optimization study because even when assets are nominally similar, differential use or loading [35] will produce variations in optimal management plans for each asset.

In contrast to private sector LCA, public infrastructure generally exists to maximize positive externalities for a given resource input. Transportation infrastructure supports economic activity, and this raises challenges in terms of selecting system boundaries that reliably estimate these externalities [47]. While these desired externalities are often measured in economic units, they constitute separate metrics from, and are not readily combined with, agency resource inputs. This has policy ramifications since the

objectives of public agencies can conflict, e.g. considering air quality together with economic development. Infrastructure scale actions also shift the background context in which to apply LCA. Transportation investments shape travel demand [47, 135] and such traffic growth is more significant in terms of overall life-cycle roadway impact than potential improvements in fuel economy [192].

The choice of functional unit also has important implications for teasing apart the conflation of public policies, actions, and consequential impacts in a LCA framework. In the case of roadways, for example, some number of lane-km is a typical functional unit [164, 88, 192] since the user actions are unknown when considering generalized LCA. However, when considering agency emissions in maintenance planning, emissions from queued vehicles upstream of the work zone can be on par with the embodied emissions of the maintenance work itself [88] so the use phase cannot be over-generalized in unit selection. In light of these limitations of conventional LCA/LCC, new methods are needed to more properly understand the environmental consequences of large infrastructure networks.

B.1.1 Proposed Framework

A multi-objective general model for long-term management of a network of public infrastructure assets considering economic externalities (user cost), public investments (agency cost), and environmental impacts can address shortcomings of conventional practice in this area. Figure B.1 shows the data flows between the LCA, LCC, optimization, and policy domains. Actions in this general model represent maintenance, upgrades, or retirement of assets. A subset of this general model may be appropriate given budgetary or scope constraints on the part of the decision maker, however one should always consider the long-term driving relationships and feedback loops present in large systems. In the case of transportation, agency funding formulas, and thus current investment, are typically based on existing assets and travel demand, both of

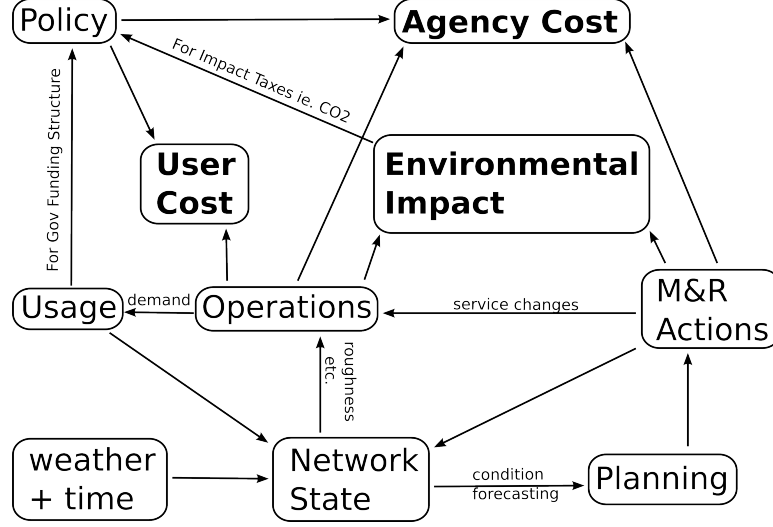


Figure B.1: A general framework for LCA in asset management of networks of public infrastructure. M&R = maintenance and rehabilitation.

which have been shaped by past agency investments.

Transportation infrastructure is a large category of public resource expenditure and it results in significant environmental impact. Since the focus in many countries has shifted from new construction to maintenance, a pavement management system (PMS) is a useful case study to apply the principles shown in Figure 1 [173]. Infrastructure in the United States suffers from deferred maintenance while agency budgets are shrinking in real dollars, so an optimization approach is particularly needed. This work is based on an environmentally optimized PMS developed by the authors where user cost is comprised of vehicle operating costs (VOC) and the value of time lost in transit. Agency costs result from maintenance activities, and greenhouse gas emissions (GHG) resulting from maintenance actions and motor fuel consumption serve as a proxy for environmental impact.

B.2 Method

In Gosse et al. (2012) [71], the authors presented a pavement management framework to program discrete segment-linked maintenance actions for a network

of pavement segments over a finite planning horizon. This effort is differentiated primarily by its use of genetic algorithms (GA) to avoid the computational limits of mathematical programming, its consideration of both preventative and more intensive maintenance treatments, and its resulting multi-objective Pareto surface. This work applies that model with modifications to include fuel consumption in the use phase and relating pavement deterioration to heavy truck volumes. The solution space of the discrete asset management problem is large, so a Pareto front (PF) of non-dominated solutions giving actions for the duration of the plan is first computed for each asset individually. Network-wide solutions consist of one asset-specific plan for each asset in the network. If interactions between assets are neglected, this two-step optimization will not exclude solutions that could dominate those in the larger pool of solutions considering all assets and times in a single solution space. The model was initially validated using a large network of 430 asphalt highway segments maintained by the Virginia Department of Transportation (VDOT). A smaller, synthetic network of only 8 segments representing the conditions in the prior study but enabling the study of asset-specific behaviors is considered here. All segments have two 3.44m lanes of asphalt pavement and, for simplicity, no shoulder.

VDOT measures pavement condition using the critical condition index (CCI), which is a composite of several pavement distress ratings and ranges in scale from 0 (complete failure) to 100 (perfect new pavement). The study area average in 2007 of 90, used here as the initial condition of all segments, is considered very good. VDOT has calibrated an exponential deterioration model for bituminous interstate pavements that is a function of time only [162]. Since load-related distress is primarily a function of heavy truck traffic [5] it is reasonable to model deterioration using truck traffic as well [108]. To capture truck impacts in a simple and representative manner, pavement age in the VDOT model is inflated in this work by the percentage of heavy trucks.

Based on Sathaye (2011), this work considers the steady state of periodic pavement

rehabilitation [153]. This is consistent with the VDOT deterioration model [162] in which corrective maintenance (CM) restores the pavement to perfect condition, provided it was applied before permanent damage occurs to the underlying structure. VDOT considers more intensive treatments but their cost vs. performance makes CM the preferred alternative [38]. Unlike Sayathe (2011), the model presented here also considers preventative maintenance (PM) in addition to CM. VDOT models PM as improving pavement condition by 15 points CCI without altering the effective age, and thus deterioration rate, of the pavement [153]. LCA data for these treatments are taken from Gosse et al. (2012) [71]. “Green” paving technologies, such as warm-mix asphalt and others [190, 192, 191], were not considered here since the focus is on scheduling optimization and so only current VDOT treatments were included.

User cost is modeled as vehicle operating cost (VOC), with fuel use as a proxy metric, and the user value of time, since there is concern with using incremental VOC due to pavement condition alone as the objective of PMS [6]. In this work, vehicles are considered either as medium automobiles or heavy trucks, with trucks including an additional delay cost due to a cargo value of \$200,000 after AASHTO [7]. Delay cost is assumed to be 50% of the prevailing wage for autos and 100% for trucks. The wage used is \$18.16, which is the U.S. Census 2010 private industry average adjusted to 2007 dollars.

Incremental VOC due to pavement condition is not well characterized in the literature. OPUS found that small changes in roughness in pavements in good condition had negligible effect on VOC [134], while Santero and Horvath found pavement condition to be one of the primary contributors to use-phase pavement GHG emissions [151]. In this work we adopt a simple model of increasing VOC (fuel usage) from a base case [7] according to a function of roughness [15]. To estimate roughness from CCI, we first applied a linear transformation from CCI to the pavement serviceability rating (PSR) which was in turn converted to the international roughness

index (IRI) using an exponential relationship [107].

It is assumed that automobiles use regular unleaded gasoline and trucks use diesel with CO2 emissions of these fuels taken from US EPA [174]. Fuel prices were taken from October 2011 US averages from the American Automobile Association and deflated to 2007 dollars: \$0.83/liter gasoline and \$0.93/liter diesel. The additional fuel use per unit delay time was taken from AASHTO (2010) for 105km/hr: 0.25 liters / min auto delay and 2.19 litres / min truck delay [7]. Delay values were computed using a VDOT spreadsheet tool that used 2007 data from the study area and highway capacity manual methods for rolling terrain to compute total delay by vehicle class for nighttime (7pm to 7am) single-lane closures [168]. For this work, the relationship was reduced to equation 1, which gives total delay in minutes per closure night. Delay thus scales linearly with segment length, but is nonlinear with respect to average annual daily traffic (AADT) and truck ratio as shown in Eq. B.1.

$$delay = 1.93E5e^{3.007E-4*truckRatio*(AADT-3.566E4)} \quad (B.1)$$

B.2.1 Modeling Horizon

In Gosse et al. (2012), the authors demonstrated that short modeling horizons affected the final PF, which is inconsistent with a goal of long-term optimal management [71]. For pavements, the planning horizon should be long enough to include the first major rehabilitation [5]. Planning for an indefinite service life means that the metrics computed over the modeled period should be representative of the long-term steady state. The computational burden of extending this period is large but the VDOT deterioration model used here suggests that 20 years is adequate, particularly given the addition of the truck-related effective age introduced in this work [162]. All costs are calculated in constant 2007 dollars and no discounting is applied, consistent with the steady state assumption.

The residual pavement condition at the end of the planning period is not addressed in conventional finite-horizon models, such that optimal solutions will allow the pavement to deteriorate near the end of the plan, which violates the requirement that the entire period be representative of the steady state. To account for this artifact of optimization, we continue to compute use and deterioration for an additional 5 years beyond the initial 20, however no maintenance actions are programmed during this salvage period. In this way, segments that are deteriorating rapidly at year 20 will be penalized when average metrics are computed over all 25 years of plan plus salvage period. While this approach has limitations, agencies re-run these models annually and thus only act on the year 1 actions, so small variations in the later years of the plan have little practical impact. As a first-order approach, salvage value computed this way represented the most expedient means of approximating a steady state condition with finite computational resources and avoids the need to discount economic flows in contrast with the treatment of environmental measures.

B.3 Results

The modeling framework developed here generates a PF of solutions for the network that are non-dominated with respect to all three metrics of interest: user cost, agency cost, and GHG emissions. Figure B.2 shows the trajectory of individual segment condition over time for a single selected reference plan also identified in subsequent figures. Improvements in pavement condition over time correspond to maintenance activity. Degradation is predicted using the modified VDOT model. Each non-dominated plan in the network PF (Figure B.3 is comprised of a set of segment-specific actions such as those in Figure B.2.

Variation among the segments of the network considered here was minor in comparison to the more general case managing a heterogeneous group of assets lacking even common descriptors. Nevertheless, the non-linear response of the system to




Figure B.2: Condition trajectories for each segment for a reference plan. Large improvements represent CM, smaller improvements PM.




Figure B.3: Non-dominated network solutions comprising the PF. The circled plan is noted in Figures B.2, B.4, and B.5.

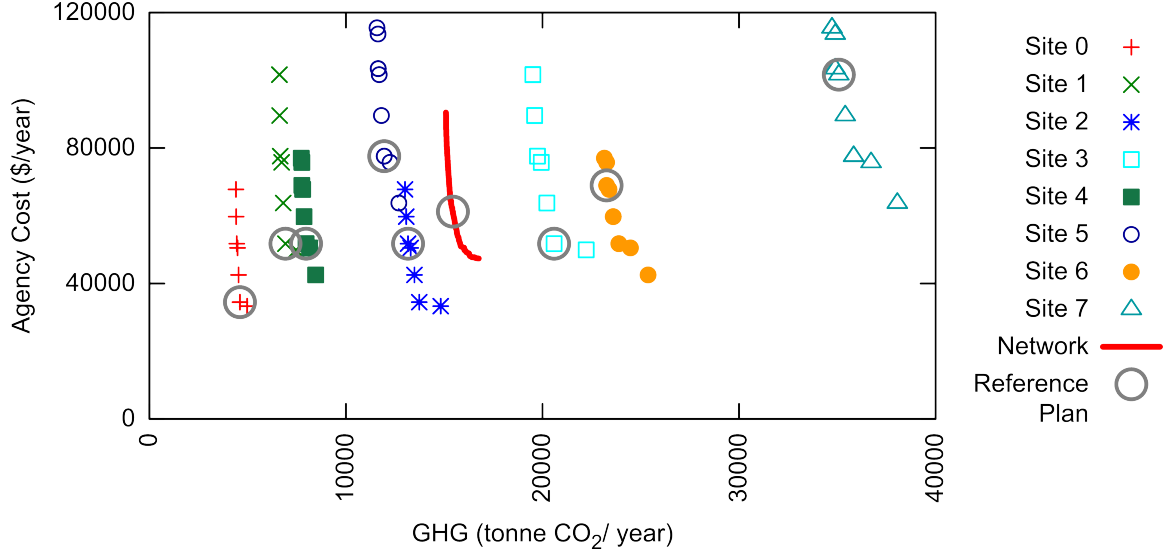


Figure B.4: Individual PFs for each segment with the network PF normalized by the number of segments. The components of the reference plan from the previous figures are circled.

AADT, truck ratio, VOC increase due to roughness, and budget (since budget influences roughness), means that net work-optimal solutions differ from those selected by examining each segment's individual PF. Figure B.4 shows the individual PFs of each segment, with the components of the reference plan circled. The shape of the network PF is reflected in the individual segment PFs, however, it can be seen from the variation in the relative PF position of reference plan components shown in Figure B.4 that non-dominated network-wide solutions do not necessarily follow from considering each segment independently, even though the network-wide values are simple sums of the per-segment values.

Figure B.5 shows two plans selected by inspection of the individual PFs in Figure B.4 in contrast to the network PF. The “knuckle” plan selects the per-segment plan at the turn of each segment PF illustrated in Figure B.4. The “middle” plan selects per-segment plans in the center of their respective PFs. We see in Figure B.5 that these plans, selected according to the properties of the individual PFs, are dominated by plans in the network PF. The “knuckle” and “middle” plans incurred 3.7% and




Figure B.5: The network-optimal PF and reference plan from Figures B.2 - B.4 in contrast to two solutions derived from examining the per-segment PFs in Figure B.4.

7.2% greater agency costs or 1.0% and 0.6% greater GHG emissions respectively than solutions on the PF. Agency cost shows more response to optimization than GHG emissions because baseline fuel consumption dominates emissions, but is not entirely sensitive to policy, as agency cost is. Because these plans are selected from per-segment PFs, they are not drastically sub-optimal in the network context, however they do illustrate that to achieve the network-optimal, one must optimize over the network as a whole. Examination of the components of the reference plan in Figure B.4 also reveals a tendency to select progressively more intensive maintenance plans for higher-GHG segments. This result is intuitive for the simple network here, but would be impractical in real-world application without a numerical framework such as is presented here.

PFs in Figures B.2, B.4, and B.5 were presented without user cost for legibility. Figure B.6 shows the strong linear relationship between user costs and GHG emissions that allows this simplification for this particular problem formulation. While agency maintenance actions incur user costs through delay, use-phase fuel consumption,




Figure B.6: Per-segment and network-wide relationship between GHG emissions and user cost. Per-segment symbols from Figure B.4.

and corresponding emissions, dominate and are negatively correlated with agency investment through pavement condition.

B.4 Conclusions

- A general model was presented to address the challenges of incorporating conventional LCA into network asset management for the public sector.
- A subset of this model was used to present a roadway asset management case study.
- Computation of asset-specific PFs allows a greatly reduced computation load when considering the network as whole.
- Considering each asset individually, even with respect to its PF of non-dominated possible actions, produces sub-optimal solutions relative to a network-wide optimization.

- Management of varied networks of assets requires additional tools beyond multiplying a single asset LCA by the number of assets in the network.

B.5 Acknowledgements

The authors thank Drs. Brian Diefenderfer and Michael Fontaine (Virginia Center for Transportation Innovation and Research) and Kevin Chisnell (VDOT District 8) for technical expertise and network data.

APPENDIX C

Supporting Information for Mobility

C.1 Traffic Microsimulation

Microsimulation was carried out using the commercial VISSIM v5.4 microsimulation software because it has the capability to simulate lateral behavior and thus vehicle passing movements within a single lane. VISSIM has a COM interface that allows interactions with running simulations, however this does not allow modification of the network which was necessary to run the many cases explored in this analysis. Instead, a Perl script was used to generate VISSIM input files directly for all 576 simulation runs. These simulations were run in parallel on a desktop PC with each instance logging data directly into a transactional database.

Vehicle speeds were assumed to be uniformly distributed and constant with grade, with the exception of bicycles. Motor vehicles all had an expected velocity of 16.7 m/s with buses deviating by ± 0.28 m/s and automobiles deviating by ± 1.39 m/s based on the higher speed variations generally observed for private vehicles. For bicycles, a cyclist with power ranging from 50-80W riding a commuter type bicycle was assumed to estimate lower and upper bounds on the bicycle speed distribution used in VISSIM.

Resulting speeds by grade were computed using a first principles prediction of bicycle speed [194]. Since 4% is a considerable grade, speeds varied considerably between the uphill case with an average of 1.6 m/s, level ground with an average of 4.5 m/s, and downhill with an average of 9.2 m/s.

C.2 Vehicle Emissions

Because vehicle emissions are known to vary with roughness, and roughness is directly modeled in the framework through the pavement management module, a means of estimating emissions at various rolling resistances was required. The Comprehensive Modal Emission Model (CMEM) [16] is well regarded, based on the US vehicle fleet, and most importantly power based, which allows varying rolling resistance independently of other parameters. The model is divided into a light duty code, which was used for automobiles, and a heavy duty code which was used to model a 10.67m transit bus.

CMEM is unfortunately not available as source code, nor is the distributed code able to handle very large volumes of input data. As a result, a detailed lookup table was generated for each vehicle type at each roughness considered. Compiled C++ code was then able to efficiently search this table in memory to compute second by second fuel use and emissions from detailed vehicle movements output by the VISSIM microsimulation.

The CMEM emissions model is additive which results in a linear relationship between increased roughness and resulting emissions, so only a base and rough cases were computed. Final values were interpolated linearly using those data points with the base case providing a minimum value since increased smoothness has no additional impact beyond a certain point. In terms of table lookup metrics, emissions were most sensitive to acceleration, so nearest-neighbor search was used for all parameters except acceleration. Final emissions values were then computed using interpolation between

the two bounding acceleration values.

C.3 Probabilistic Delay

A reasonable validation of the VISSIM simulation results can be obtained by computing a probabilistic expected delay due to buses encountering a bicycle they cannot pass, cars and buses getting stuck behind that bus, and both vehicle types encountering a bicycle passing another bicycle when they would otherwise have been able to pass a single bicycle. The goal of this probabilistic work is not to replace microsimulation, but to demonstrate that the large differences seen in the simulation results between configurations where buses get stuck and where they do not are indeed plausible. This approach may also be beneficial where microsimulation is impractical. Because we assume that vehicles may not cross the centerline of the roadway, simulations were carried out for each characteristic length, lane configuration, grade, and traffic volume for a single direction. While vehicles are able to cross the centerline in real traffic, we have chosen not to model this behavior based on the assumption that the study area consists of a congested and narrow corridor where such behavior would be impractical at best and a safety risk at worst. These results were later aggregated to produce daily values, which were then inflated to annual values for purposes of combining them with pavement management costs and emissions.

The following sections include code from the R statistical programming language that was used to perform the analyses presented in this work. Much of this should be self explanatory, however the reader should be aware that `<-` is an operator equivalent to an equals sign and `$` is the subsetting operator. The first section outlines the variable names and constants that will be used throughout.

```
widthCar <- 1.75
widthBus <- 2.6
widthBike <- 0.5
widthPassing <- 1
```

```

widthParking <- 2.5
# following distance comes from the microsimulation car following model
lFd <- 3.5
# vehicle lengths also affect queue length
lBike <- 1.5
lCar <- 4.2
lBus <- 10.677
# bicycle velocity is in the data itself, auto is set at 60 kph
vMotorized <- 60/3.6
vBus <- vMotorized
vCar <- vMotorized
# the queue seen by the stuck bus
lQBike <- lFd + lBike

```

The first step is to compute the delay a bus stuck behind a bicycle would likely experience. This is a function of where along the segment the bus encounters the bicycle, with a later encounter causing less delay. Finally, this probability is conditioned on the likelihood of it's occurrence. Vehicle speeds are uniformly distributed within a range, which can have a small affect on the results but is considered negligible over reasonable characteristic lengths for this analysis. Distances are measured from the front of the bicycle behind which the queue has formed and input vehicle volumes are hourly.

```

# tp is the time gap between bike and car it's maximum value is when the bus catches
# the bike at the end of the segment
SRs$tPMax <- (SRs$length * (vBus - SRs$vBike) + lQBike * SRs$vBike)/vBus/SRs$vBike
# segment length affects how likely it is for there to be a vehicle at all the
# probability is the lower tail of the exponential distribution
SRs$pBike <- pexp(SRs$tPMax, SRs$volumeBike/3600)
SRs$pBus <- pexp(SRs$tPMax, SRs$volumeBus/3600)
SRs <- ddply(SRs, .(simRun), transform, tPBus = integrate(function(t) t * dexp(t, volumeBus/3600),
  0, tPMax)$value/pBus)
# position of the nose of the bike when the bus encounters it
SRs$xe <- (SRs$tPBus * vBus * SRs$vBike - lQBike * SRs$vBike)/(vBus - SRs$vBike)
# expected delays /vehicle/km when it is has delay at all
SRs$delayBusStuck <- (vBus - SRs$vBike)/vBus/SRs$vBike * (SRs$length + lQBike - SRs$xe) *
  1000/SRs$length

```

Once a bus is stuck behind a bicycle, following vehicles will build a queue behind the bus until the bicycle exits the segment and the queue is allowed to clear. Bicycles are not significantly affected since their desired speed is approximately that of the queue. Cars and buses in the following queue follow essentially the same analysis with their respective volumes and lengths varying. An iterative approach using a root-finding function is necessary to determine the following queue length. We do not assume that vehicles will progressively begin to accelerate in a clearing shockwave after the stuck bus departs the link, both for simplicity and because the presence of additional bicycles in the queue would hinder this process. We also neglect the affects of vehicle length on starting position on the segment.

```
# starting queue length for the first vehicle behind the stuck bus
lqf0 <- lQBike + lFd + lBus
SRs$lqfCar <- lqf0
SRs$lqfBus <- lqf0

# the expected value of the queue length based on a given following time.
Elqf <- function(lqf0, tpf, Vdbus, Vdcar, Vdbike) {
  nBus <- Vdbus * tpf
  nCar <- Vdcar * tpf
  nBike <- Vdbike * tpf
  return(lqf0 + nBus * (lBus + lFd) + nCar * (lCar + lFd) + nBike * (lBike + lFd))
}

SRs <- ddply(SRs, .(simRun), transform, tPFollowingMax = uniroot(function(tpf) length -
  xe - vBike/(vBus - vBike) * (vBus * tpf + lQBike - Elqf(lqf0, tpf, volumeBus/3600,
    volumeCar/3600, volumeBike/3600)) - Elqf(lqf0, tpf, volumeBus/3600, volumeCar/3600,
    volumeBike/3600), lower = 0, upper = length/vBike)$root)

# there is a minimum tP value for folowing vehicles which is the bus tP value plus a
# correction for the length of the bus to cross the starting line. This would be a
# bigger offset in a longer queue tPFollowingMin <- tPBus +
# (lFd+lBus)/velocityMotorized but we assume that vehicles will stack instead and
# let them start at x=0, therefore
tPFollowingMin <- 0

# neglect the auto headway 'in front of' the stuck bus in computing expected
# following headways
SRs$pCarFollowing <- pexp(SRs$tPFollowingMax - tPFollowingMin, SRs$volumeCar/3600)
SRs <- ddply(SRs, .(simRun), transform, tPFollowingCar = integrate(function(t) t * dexp(t,
  volumeCar/3600), tPFollowingMin, tPFollowingMax)$value/pCarFollowing)
```

```

# bus headway is definitely this since we know we're measuring from the last bus
# entry
SRs$pBusFollowing <- pexp(SRs$tPFollowingMax - tPFollowingMin, SRs$volumeBus/3600)
SRs <- ddply(SRs, .(simRun), transform, tPFollowingBus = integrate(function(t) t * dexp(t,
  volumeBus/3600), tPFollowingMin, tPFollowingMax)$value/pBusFollowing)
# given expected arrival times, we can compute an expected queue length seen by each
# vehicle type
SRs$lqfBus <- Elqf(lqf0, SRs$tPFollowingBus, SRs$volumeBus/3600, SRs$volumeCar/3600, SRs$volumeBike/3600)
SRs$lqfCar <- Elqf(lqf0, SRs$tPFollowingCar, SRs$volumeBus/3600, SRs$volumeCar/3600, SRs$volumeBike/3600)
SRs$nBusFollowing <- SRs$volumeBus/3600 * SRs$tPFollowingBus
# a car or a bus both have their encounter at different expected times
SRs$xeFollowingCar <- SRs$xe + SRs$vBike * (vBus * SRs$tPFollowingCar + lQBike - SRs$lqfCar)/(vBus -
  SRs$vBike)
SRs$delayFollowingCar <- (vBus - SRs$vBike)/vBus/SRs$vBike * (SRs$length + SRs$lqfCar -
  SRs$xeFollowingCar) * 1000/SRs$length
SRs$xeFollowingBus <- SRs$xe + SRs$vBike * (vBus * SRs$tPFollowingBus + lQBike - SRs$lqfBus)/(vBus -
  SRs$vBike)
SRs$delayFollowingBus <- (vBus - SRs$vBike)/vBus/SRs$vBike * (SRs$length + SRs$lqfBus -
  SRs$xeFollowingBus) * 1000/SRs$length

```

The final delay component considered is that of faster bicycles passing slower bicycles and thus blocking the travel lane when a single bicycle would not have done so. The likelihood of this is based on the uniform distribution of bicycle speeds while the impact is dependant on whether the lane configuration allows two bicycles to pass abreast without blocking buses and or cars. For the lane configurations considered here, a car can only pass a single bicycle within either the narrow or wide lane, so it could only pass two bicycles abreast in the presence of a bicycle lane. A bus can never pass two bicycles abreast. The likelihood of an encounter is approximated here using the motorized occupancy.

```

# expect the faster bike to be 1/4 of the spread above the mean given a uniform
# distribution
SRs$vBike2 <- SRs$vBike + SRs$vBikeSpread/4
SRs$vBike1 <- SRs$vBike - SRs$vBikeSpread/4
# length of the moving passing zone
lBPass <- lFd + lBike
SRs$delayBikePass <- (vBus - SRs$vBike2)/vBus/SRs$vBike2 * lBPass * SRs$vBike1/(SRs$vBike2 -

```

```

SRs$vBike1)
SRs$tFbikeMax <- SRs$length * SRs$vBikeSpread/(SRs$vBike + SRs$vBikeSpread/2)/(SRs$vBike -
  SRs$vBikeSpread/2)
SRs$pSlowBike <- 0.5 * pexp(SRs$tFbikeMax, SRs$volumeBike/3600)
# we'll use the occupancy as a probability of being on a given spot
SRs$spaceOcc <- (SRs$volumeBus/3600 * lBus + SRs$volumeCar/3600 * lCar)/vMotorized
# probability of overlap between the passing even and the vehicles
SRs$delayBikePass <- (SRs$spaceOcc * lBPass/SRs$vBike2) * SRs$pSlowBike * 0.5 * (SRs$tFbikeMax *
  SRs$volumeBike/3600 * SRs$delayBikePass * 1000/SRs$length)

```

Having computed the various delay components, we can test how much of the variance in the microsimulation results between various lane configurations is attributable to these values. A high correlation indicates that these are significant mechanisms in the overall outcomes and so justifies the focus on the stuck condition in this framework.

```

# establish boolean switches to weight lane conditions
SRs$pass2BikesCar <- SRs$laneWidth > 3.5 & SRs$bikeLane
# can a bus pass a bike in the same lane?
SRs$passBikeBus <- SRs$laneWidth > 2.6 + 1 + 0.5
# are buses getting stuck in the first place?
SRs$busStuck <- as(!SRs$passBikeBus & !SRs$bikeLane, "integer")
# large travel time magnitudes will skew results without weighting
SRs <- ddply(SRs, .(grade, hourlyVolumeId, length), transform, lmWeight = 100/max(travelTimeBus))
# car following delay factor
SRs$cfdF <- SRs$pBike * SRs$pBus * SRs$pCarFollowing * SRs$delayFollowingCar * SRs$busStuck
# bicycles passing each other delay factor for cars
SRs$cpbdf <- SRs$pBike * SRs$delayBikePass * as(!SRs$pass2BikesCar, "integer")
# fit a linear model
lmTtCar <- lm(travelTimeCar ~ cfdF + cpbdf + volumeRequestedMotorized:length, data = SRs,
  weights = lmWeight)

```

From the model fit, we see a high contribution from the delay factors, which is a significant improvement over fitting a more basic model using scenario parameters alone, such as length, grade, etc.

```

Call:
lm(formula = travelTimeCar ~ cfdF + cpbdf + volumeRequestedMotorized:length,
  data = SRs, weights = lmWeight)

```

```

Weighted Residuals:
      Min       1Q   Median       3Q      Max
-9.642 -0.338 -0.038  0.307  7.426

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)    5.97e+01   9.94e-02  600.46 < 2e-16 ***
cfd           3.59e-01   1.23e-02   29.12 < 2e-16 ***
cpb           1.48e+00   1.66e-01    8.95 < 2e-16 ***
volumeRequestedMotorized:length 8.08e-06  1.10e-06    7.35 2.2e-12 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 1.54 on 284 degrees of freedom
Multiple R-squared:  0.846, Adjusted R-squared:  0.844
F-statistic:  518 on 3 and 284 DF,  p-value: <2e-16

```

Bus delay is similar to auto delay, except that while all autos that are delayed follow a bus, bus delay must be apportioned between the likelihood that a bus is following compared to the likelihood that a bus is stuck itself, to avoid double counting.

```

# delay of the stuck bus
SRs$bsdf <- SRs$busStuck * SRs$pBike * SRs$delayBusStuck * (1/(1 + SRs$nBusFollowing))

# delay of following buses
SRs$bfdf <- SRs$busStuck * SRs$pBike * SRs$delayFollowingBus * SRs$pBusFollowing * (SRs$nBusFollowing/(1 +
  SRs$nBusFollowing))

SRs$busStuckPassingBikes <- as((!SRs$passBikeBus & SRs$bikeLane) | (SRs$passBikeBus &
  !SRs$bikeLane), "integer")

# delay due to encountering a bicycle passing another
SRs$bpbdf <- SRs$delayBikePass * SRs$busStuckPassingBikes

# fit a linear model to the delay factors
lmTtBus <- lm(travelTimeBus ~ bsdf + bfdf + bpbdf, data = SRs, weights = lmWeight)

```

As with car delay, we see a significant contribution from the probabilistic delay factors indicating their consistency with the delay predicted through simulation.

```

Call:
lm(formula = travelTimeBus ~ bsdf + bdf + bpdf, data = SRs,
    weights = lmWeight)

Weighted Residuals:
    Min       1Q   Median       3Q      Max
-70.92   0.24    0.97    1.63   84.45

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  58.9875     0.7205   81.87  <2e-16 ***
bsdf          0.4826     0.0317   15.22  <2e-16 ***
bdf          125.5681    11.6955   10.74  <2e-16 ***
bpdf          4.5173     1.5131    2.99   0.0031 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 14 on 284 degrees of freedom
Multiple R-squared:  0.736, Adjusted R-squared:  0.734
F-statistic:  264 on 3 and 284 DF,  p-value: <2e-16

```

Aggregating the auto and bus predicted delay, we see good agreement with simulated motorized delay. Microsimulation captures much more detail with regards to vehicle interactions, such as delays caused by vehicle passing maneuvers and periods of locally high congestion. As a result, the probabilistic analysis was employed simply to provide additional verification of the simulated results in the absence of ground observations of these phenomena. Figure C.1 summarizes this relationship and supports the notion that the stuck bus phenomena is a significant contributor to overall delay in the cases considered in this work.

Moving from simulations in a single direction to bi-directional scenario evaluation, Figure C.2 provides an additional view of the distribution of the scenarios considered and relative proportion of stuck conditions with respect to grade and dominated status.

C.4 Time and Discount Rate

Transportation investment decisions have traditionally been made using a fixed positive discount rate typically exceeding that of inflation. This implies both a recognition of the time value of capital, as well as a pure time preference that values future monetary flows less than present flows even when considered in present dollars. If we assume that inflation reflects the real purchasing power of money, then an analysis inflating future costs according to projected inflation and then discounting them to present dollars is equivalent to using a value of zero for both inflation and discount rate. As is typical in LCA practice, we feel that a pure time preference inappropriately burdens future generations where the infrastructure has no foreseeable end of life and can be expected to serve a similar function in the future as the present. Thus all flows computed in the present study are annualized using a 0% discount rate.

This analysis was formulated on such an annualized basis to describe a particular moment in time. As such, we did not attempt to forecast trends in vehicle fleet technology, and relied on scenarios of bicycle mode share that reflect the present day and a large increase from the present case. The case of pavement management is particular, however, in that investments are discrete in time and space, and so do not readily lend themselves to annualization. Unlike other long-term capital investments, such as buildings and bridges, that can be expected to experience a complete life-cycle including demolition and disposal, a city street is rarely if ever completely reconstructed to the extent a bridge may be replaced, nor does it have a foreseeable end of life. As described below, we computed pavement maintenance plans for a time period that was both long enough to capture the full range of pavement deterioration and rehabilitation while still being computationally feasible in a discrete network-wide formulation. This value is somewhat arbitrary and we deemed 30 years to be a good compromise. A two year salvage period was also used where no treatments were programmed, though use and deterioration was still computed, so that the optimization would not preference

the pavement's falling into disrepair at the conclusion of the 30 year planning period. This approach is somewhat laborious, however it is necessary in order to capture the non-linear effects of dynamic vehicle loading, since pavement deterioration due to heavy vehicles accelerates as roughness increases, as well as to develop a framework that could capture the transition period from existing conditions to potentially different long-term optimal ones. This framework could readily incorporate changing vehicle fleets, mode share, overall passenger volumes, and costs, however we sought to present the analysis in as general a manner as possible and so omitted site-specific forecasts from the analysis.

C.5 Pavement Management

The pavement management optimization was taken largely from previous work by the authors [71], but with the addition of dynamic pavement loads as a function of the result of the international roughness index (IRI) [56] given by equation C.1:

$$ESAL_{dynamic} = ESAL_{static}(1 + 0.003153 v IRI) \quad (C.1)$$

where $ESAL_{static}$ is the AASHTO load equivalency, $ESAL_{dynamic}$ is the impact the pavement sees due to dynamic interactions between the vertical profile of the pavement (IRI in m/km) and the vehicle at speed v . At freeway speeds, $ESAL_{dynamic}$ actually begins to decrease as the excitation frequency increases beyond typical resonance points, however the urban segments considered in this work do not operate in the freeway regime. As a result, interstate data was excluded from the linearization presented in equation C.1.

Pavement deterioration is predicted using a combination of pure aging, and load

related distresses [137]:

$$IRI_y = 1.04e^{C_{env}t}(IRI_{y-1} + 263(1 + S_n)^{-5}ESAL_{cum}) \quad (C.2)$$

where y is the year at which IRI is computed, t is the age since rehabilitation y , S_n is the structural number of the pavement, C_{env} is the environmental coefficient = 0.1, and $ESAL_{cum}$ is total $ESAL_{dynamic}$ since rehabilitation in millions. C_{env} varies according to precipitation and freeze-thaw cycles and while there is considerable variation in this parameter, the value chosen here is the near the middle of the range. For the present analysis, the total contribution of pavement management is small, so the sensitivity to this value is correspondingly low. Lanes of the same type were managed identically for both directions of the segment. Feasible treatment plans were evaluated over 30 years with 2 year salvage period where loads and deterioration was computed but no treatments were programmed in order to avoid an undesirable decline in pavement condition at the end of the maintenance forecasting period. Final values for pavement costs and emissions were annualize totals including the entire 30 + 2 years. Constraints on feasible plans included a minimum of four years between resurfacing treatments, or in the absence of a resurfacing within the four year window, a maximum of 2 preventative treatments. The maximum allowable IRI was defined as 3m/km for bicycle lanes, 6m/km for auto lanes, and 8m/km for parking lanes. While an IRI value for a parking lot isn't an entirely appropriate measure, the deterioration model computed age-related distresses only for parking and so the constraint triggers reasonable maintenance levels over time.

The small network size of this problem, i.e. a maximum of 3 unique treatment plans, makes an exhaustive search for Pareto-optimal network plans covering all lane types feasible. For larger problems, the code employs a genetic algorithm computed using a massively parallel graphics processing unit to estimate a network Pareto front

of annual maintenance actions over many segments. In either case, non-dominance was determined based on the GHG emissions from fuel and pavement management separately with annual mean IRI used to break ties between plans that contained the same total number of each pavement treatment and thus had equivalent agency GHG emissions.




Figure C.1: Predicted travel time aggregated by person trips compared with simulation results.

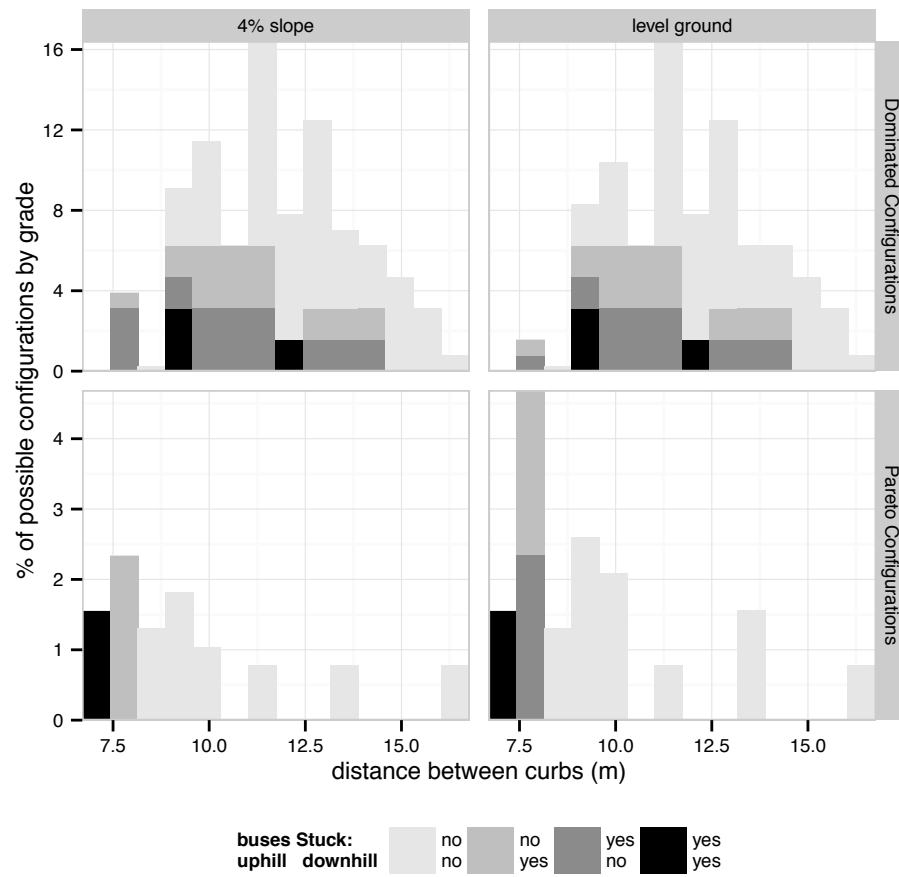


Figure C.2: Distribution of stuck and unstuck lane configurations

APPENDIX D

Publications resulting from this work

Gosse, C. A., & Clarens, A. F. (in review). Estimating Spatially and Temporally Continuous Bicycle Volumes Using Sparse Data. *Transportation Research Record, Journal of the Transportation Research Board*.

Gosse, C. A., & Clarens, A. F. (2013). Quantifying the total cost of infrastructure to enable environmentally preferable decisions: the case of urban roadway design. *Environmental Research Letters*, 8(1). doi:10.1088/1748-9326/8/1/015028

Gosse, C. A., & Clarens, A. F. (2012). Network Aspects of Pavement Management Optimization Using Life Cycle Assessment. In *LCA and Construction 2012, Nantes, France* pp. 108-115. RILEM, Bagneux France and IFSTTAR, Paris France. ISBN: 978-2-35158-127-8

Gosse, C., Smith, B., & Clarens, A. (2012). Environmentally preferable pavement management systems. *Journal of Infrastructure Systems*, 19(3), pp. 315-325. doi:10.1061/(ASCE)IS.1943-555X.0000118

Garber, N., Haas, P. R., & Gosse, C. (2010). *Development of Safety Performance Functions for Two-Lane Roads Maintained by the Virginia Department of Transportation*. Report: VTRC 10-R25, Virginia Center for Innovation in Transportation Research, Charlottesville, VA

BIBLIOGRAPHY

- [1] (2005), District of Columbia Bicycle Master Plan, *Tech. rep.*, District Department of Transportation.
- [2] (2006), Bicyclist Fatalities and Serious Injuries in New York City 1996-2005, *Tech. rep.*, New York City Departments of Health and Mental Hygiene, Parks and Recreation, Transportation, and the New York City Police Department, New York, NY.
- [3] (2011), NYC Commuter Cycling Indicator, *Tech. rep.*, New York City Department of Transportation, New York, NY.
- [4] AASHTO (1990), *Guidelines for Pavement Management Systems*, Association of American State Highway and Transportation Officials, Washington, DC.
- [5] AASHTO (1993), *AASHTO Guide for Design of Pavement Structures*, American Association of State Highway and Transportation Officials, Washington, DC.
- [6] AASHTO (2001), *Pavement Management Guide*, 254 pp., Washington DC.
- [7] AASHTO (2010), *User and Non-User Benefit Analysis for Highways*, American Association of State Highway and Transportation Officials.
- [8] Alta Planning and Design (2009), Count Adjustment Factors, *Tech. Rep. March*, National Bicycle and Pedestrian Documentation Project, Portland, OR.
- [9] and ACEC, I. (2009), Illinois - Livable and Sustainable Transportation Rating System and Guide, *Tech. rep.*, Illinois Department of Transportation with Illinois Joint Sustainability Group.
- [10] Ang, B. W., T. F. Fwa, and T. T. Ng (1993), Analysis of Process Energy Use of Asphalt Mixing Plants, *Energy*, 18(7), 769–777.
- [11] Arpad Horvath (2003), Life-Cycle Environmental and Economic Assessment of Using Recycled Materials for Asphalt Pavements, *Tech. Rep. 510*, University of California Berkely.
- [12] ASCE (2009), Infrastructure Fact Sheet, *Tech. Rep. 3*, doi:10.1126/science.1.3.84.
- [13] Association, A. A. (2007), Fuel Gauge Report.

- [14] Association, B. M. (2012), Healthy transport = Healthy lives, *Tech. Rep. July*.
- [15] Barnes, G., and P. Langworthy (2003), The Per-mile Costs of Operating Automobiles and Trucks, *Tech. rep.*, Minnesota Department of Transportation.
- [16] Barth, M., F. An, T. Younglove, G. Scora, C. Levine, M. Ross, and T. Wenzel (2000), NCHRP PROJECT 25-11 DEVELOPMENT OF A COMPREHENSIVE MODAL EMISSIONS MODEL, *Tech. rep.*, National Cooperative Highway Research Program.
- [17] Behrisch, M., L. Bieker, J. Erdmann, and D. Krajzewicz (2011), SUMO - Simulation of Urban MObility: An Overview, in *SIMUL 2011, The Third International Conference on Advances in System Simulation*, pp. 63–68, Barcelona, Spain.
- [18] Bertsimas, D., and R. Demir (2002), An approximate dynamic programming approach to multidimensional knapsack problems, *Management Science*, 48(4), 550–565.
- [19] Bevan, Tim and Reid, Lisa and Anderson, Jaralee and Muench, Stephen and Se-skin, S. (2011), Compendium of Project Delivery and Operations & Maintenance Criteria, Pilot Test Version, *Tech. Rep. July*, Federal Highway Administration.
- [20] Boilé, M., P. Narayanan, and K. Ozbay (2007), Impact of Buses on Highway Infrastructure Case Study for New Jersey State, *Transportation Research Record, Journal of the Transportation Research Board*, 1841, 32–40.
- [21] Bolstad, W. M. (2010), *Understanding Computational Bayesian Statistics*, 1 ed., 315 pp., John Wiley & Sons, Hoboken (New Jersey).
- [22] Boyles, S. D., Z. Zhang, and S. T. Waller (2010), Optimal Maintenance and Repair Policies under Nonlinear Preferences, *Journal of Infrastructure Systems*, 16, 11–20.
- [23] Brandenburg, C., A. Matzarakis, and A. Arnberger (2007), Weather and cycling - a first approach to the effects of weather conditions on cycling, *Meteorological Applications*, 14, 61–67, doi:10.1002/met.
- [24] Broach, J., J. Gliebe, and J. Dill (2009), Development of a Multi-class Bicyclist Route Choice Model Using Revealed Preference Data, in *12th International Conference on Travel Behavior Research*, Jaipur, India.
- [25] Broach, J., J. Gliebe, and J. Dill (2011), Bicycle route choice model developed using revealed preference GPS data, in *TRB 2011 Annual Meeting*, Washington, DC.
- [26] Brown, M., D. Burschka, and G. Hager (2003), Advances in computational stereo, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25, doi:10.1109/TPAMI.2003.1217603.

- [27] Buehler, R., and J. Pucher (2012), Big City Cycling in Europe, North America, and Australia, in *City Cycling*, chap. 13, p. 295, The MIT Press, Cambridge, MA.
- [28] Buehler, R., and J. Pucher (2012), International Overview: Cycling Trends in Western Europe, North America, and Australia, in *City Cycling*, chap. 2, pp. 9–30, MIT Press, Cambridge, MA.
- [29] Buehler, R., J. Pucher, D. Merom, and A. Bauman (2011), Active travel in Germany and the U.S. Contributions of daily walking and cycling to physical activity., *American journal of preventive medicine*, *41*(3), 241–50, doi:10.1016/j.amepre.2011.04.012.
- [30] Bureau of Labor Statistics (2012), May 2012 National Occupational Employment and Wage Estimates United States.
- [31] Campbell Scientific (2001), Application Note 2T-M: Wind Chill, *Tech. Rep. 435*, Campbell Scientific, Logan, UT.
- [32] Campbell Scientific (2002), Application Note 5-S: Heat Index, *Tech. Rep. 435*, Campbell Scientific, Logan, UT.
- [33] Carnegie Mellon University Green Design Institute (2008), Economic Input-Output Life Cycle Assessment (EIO-LCA), US 1997 Industry Benchmark model.
- [34] Cavill, N., S. Kahlmeier, H. Rutter, F. Racioppi, and P. Oja (2008), Economic analyses of transport infrastructure and policies including health effects related to cycling and walking: A systematic review, *Transport Policy*, *15*(5), 291–304, doi:10.1016/j.tranpol.2008.11.001.
- [35] Chan, W. (2003), Optimal fund-allocation analysis for multidistrict highway agencies, *Journal of infrastructure systems*, (December), 167–175.
- [36] Chantrill, C. (2013), Government Spending: the Details.
- [37] Cho, H., P. E. Rybski, and W. Zhang (2011), Vision-based 3D bicycle tracking using deformable part model and Interacting Multiple Model filter, *2011 IEEE International Conference on Robotics and Automation*, pp. 4391–4398, doi: 10.1109/ICRA.2011.5980482.
- [38] Chowdhury, T. (2008), Supporting Document for the Development and Enhancement of the Pavement Maintenance Decision Matricies Used in the Needs-Based Analysis, *Tech. rep.*, Virginia Department of Transportation, Maintenance Division.
- [39] Clawson, M. (1962), Urban Sprawl and Speculation in Suburban Land, *Land Economics*, *28*(2), 99–111.

- [40] Council, U. G. B. (2009), *LEED Reference Guide for Green Neighborhood Development*, leed 2009 ed., 502 pp., US Green Building Council, Washington, DC.
- [41] Department, W., and E. Services (2003), City of Toronto Bicycle/Motor-Vehicle Collision Study, *Tech. rep.*, Transportation Services Division Transportation Infrastructure Management Section, Toronto, ON.
- [42] Diefenderfer, B. K., A. K. Moruza, M. C. Brown, D. S. Roosevelt, E. C. Andersen, A. A. O’Leary, and J. P. Gomez (2010), Development of a weight-distance permit fee methodology for overweight trucks in Virginia, *International Journal of Pavement Research and Technology*, 2, 236–241.
- [43] Dill, J. (2009), Bicycling for transportation and health: the role of infrastructure., *Journal of public health policy*, 30 Suppl 1(1), S95–110, doi:10.1057/jphp.2008.56.
- [44] Dill, J., and J. Gliebe (2008), Understanding and Measuring Bicycle Behavior: a Focus on Travel Time and Route Choice, *Tech. Rep. December*, Oregon Transportation Research and Education Consortium, Portland, OR.
- [45] Dor, A., C. Ferguson, C. Langwith, and E. Tan (2010), A Heavy Burden : The Individual Costs of Being Overweight and Obese in the United States, *Tech. rep.*, George Washington University School of Public Health Services, Washington, DC.
- [46] Dowling, R., A. Flannery, P. Ryus, T. Petrisch, and N. Rouphail (2010), Nield Test Results of the Multimodal Level of Service Analysis for Urban Streets, *Tech. Rep. January 2010*, National Cooperative Highway Research Project.
- [47] Echenique, M. (2011), Land use/transport models and economic assessment, *Research in Transportation Economics*, 31(1), 45–54, doi:10.1016/j.retrec.2010.11.007.
- [48] Ecoinvent Centre (2007), ecoinvent data v2.0. ecoinvent reports No. 1-25, *Tech. rep.*, Swiss Centre for Life Cycle Inventories, Dübendorf.
- [49] Eluru, N., C. R. Bhat, and D. A. Hensher (2008), A mixed generalized ordered response model for examining pedestrian and bicyclist injury severity level in traffic crashes, *Accident Analysis & Prevention*, 40(3), 1033–1054.
- [50] Elvik, R. (2011), Publication bias and time-trend bias in meta-analysis of bicycle helmet efficacy: A re-analysis of Attewell, Glase and McFadden, 2001, *Accident Analysis & Prevention*, 43(3), 1245–1251.
- [51] EPA (2005), Average Carbon Dioxide Emissions Resulting from Gasoline and Diesel Fuel, *Tech. Rep. EPA420-F-05-001*, U.S. Environmental Protection Agency.

- [52] Ewing, R. (1993), Transportation Service Standards – As If People Matter, *Transportation Research Record, Journal of the Transportation Research Board*, (1400), 10–17.
- [53] Ewing, R. (1997), Is Los Angeles-Style Sprawl Desirable?, *Journal of the American Planning Association*, 63(1), 107–126, doi:10.1080/01944369708975728.
- [54] Faghri, A., and E. Egyhaziova (1999), Development of a Computer Simulation Model of Mixed Motor Vehicle and Bicycle Traffic on an Urban Road Network, *Transportation Research Record, Journal of the Transportation Research Board*, (1674), 86–93.
- [55] Federal Highway Administration (2009), National Household Transportation Survey, *Tech. rep.*, U.S. Department of Transportation, Washington, DC.
- [56] Fekpe, E. (2006), Pavment Damage from Transit Buses and Motor Coaches, *Tech. rep.*, Batelle Memorial Instititue, Columbus, OH.
- [57] Fellendorf, M. (1994), VISSIM: A microscopic simulation tool to evaluate actuated signal control including bus priority, in *64th Institute of Transportation Engineers Annual Meeting*.
- [58] Ferreira, A., A. Antunes, and L. Picado-Santos (2002), Probabilistic Segment-Linked Pavement Management Optimization Model, *Journal of Transportation Engineering, ASCE*, pp. 568–577.
- [59] FHWA (2008), Warm-Mix Asphalt: European Practice, *Tech. rep.*, US Department of Transportation.
- [60] Finkelstein, E. A., M. daCosta DiBonaventura, S. M. Burgess, and B. C. Hale (2010), The costs of obesity in the workplace., *Journal of occupational and environmental medicine / American College of Occupational and Environmental Medicine*, 52(10), 971–6, doi:10.1097/JOM.0b013e3181f274d2.
- [61] Flynn, B. S., G. S. Dana, J. Sears, and L. Aultman-Hall (2012), Weather factor impacts on commuting to work by bicycle., *Preventive medicine*, 54(2), 122–124, doi:10.1016/j.ypmed.2011.11.002.
- [62] Forester, J. (2012), *Effective cycling*, MIT Press, Cambridge MA.
- [63] Fwa, T. F., W. T. Chan, and K. Z. Hoque (2000), Multiobjective optimization for pavement maintenance programming, *Journal of Transportation Engineering*, 3(October), 367–374.
- [64] Gao, L., and Z. Zhang (2008), Robust Optimization for Managing Pavement Maintenance and Rehabilitation, *Transportation Research Record, Journal of the Transportation Research Board*, 2084, 55–61.

- [65] Garber, N., P. R. Haas, and C. Gosse (2010), Development of Safety Performance Functions for Two-Lane Roads Maintained by the Virginia Department of Transportation, *Tech. Rep. 434*, Virginia Department of Transportation, Transportation Research Council, Richmond, VA.
- [66] Gesch, D., M. Oimoen, S. Greenlee, C. Nelson, M. Steuck, and D. Tyler (2002), The National Elevation Data Set, *Photogrammetric Engineering and Remote Sensing*, 68(1), 5–11.
- [67] Gettman, D., and L. Head (2003), Surrogate Safety Measures from Traffic Simulation Models, *Transportation Research Record, Journal of the Transportation Research Board*, 1840, 104–115.
- [68] Gharaibeh, N. G., Y. Zou, and S. Saliminejad (2010), Assessing the Agreement among Pavement Condition Indexes, *Journal of Transportation Engineering*, 136(8), 765–772, doi:10.1061/(ASCE)TE.1943-5436.0000141.
- [69] Glassbrenner, D. (2012), An Analysis of Recent Improvements to Vehicle Safety, *Tech. Rep. June*, National Highway Traffic Safety Administration, Washington, DC.
- [70] Golob, T. F., W. W. Recker, and V. M. Alvarez (2004), Freeway safety as a function of traffic flow., *Accident; analysis and prevention*, 36(6), 933–46, doi:10.1016/j.aap.2003.09.006.
- [71] Gosse, C., B. Smith, and A. Clarens (2012), Environmentally preferable pavement management systems, *Journal of Infrastructure Systems*, 19(3), 315–325, doi:10.1061/(ASCE)IS.1943-555X.0000118.
- [72] Gosse, C. A., and A. F. Clarens (2013), Quantifying the total cost of infrastructure to enable environmentally preferable decisions: the case of urban roadway design, *Environmental Research Letters*, 8(1), doi:10.1088/1748-9326/8/1/015028.
- [73] Gotschi, T. (2011), Costs and benefits of bicycling investments in Portland, Oregon., *Journal of physical activity & health*, 8(Suppl 1), S49–58.
- [74] Gouldson, A., A. Morton, and S. J. T. Pollard (2009), Better environmental regulation – contributions from risk-based decision-making, *Science of The Total Environment*, 407(19), 5283–5288, doi:DOI:10.1016/j.scitotenv.2009.06.013.
- [75] Griswold, J. B., A. Medury, and R. J. Schneider (2011), Pilot Models for Estimating Bicycle Intersection Volumes, *Transportation Research Record: Journal of the Transportation Research Board*, 2247, 1–7, doi:10.3141/2247-01.
- [76] Gu, W., Y. Ouyang, and S. Madanat (2012), Joint optimization of pavement maintenance and resurfacing planning, *Transportation Research Part B*, 46(4), 511–519, doi:10.1016/j.trb.2011.12.002.

- [77] HDR (2009), Assessing the Full Costs of Congestion on Surface Transportation Systems and Reducing Them through Pricing, *Tech. Rep. February*, U.S. Department of Transportation.
- [78] Heinen, E., K. Maat, and B. van Wee (2011), Day-to-Day Choice to Commute or Not by Bicycle, *Transportation Research Record: Journal of the Transportation Research Board*, 2230, 9–18, doi:10.3141/2230-02.
- [79] Hellweg, S., B. Hofstetter, and K. Hungerbühler (2003), Case Study Discounting and the Environment Should Current Impacts be Weighted Differently than Impacts Harming Future Generations ?, *International Journal of Life Cycle Assessment*, 8(1), 8 – 18.
- [80] Hendrickson, C., A. Horvath, S. Joshi, L. Lave, and G. Design (1998), Economic input-output models for environmental life-cycle assessment, *Environmental Science & Technology*, 32(7), 184.
- [81] Hendrickson, C. T., L. B. Lave, and H. S. Matthews (2006), *Environmental life cycle assessment of goods and services: an input-output approach*, Resources for the Future.
- [82] Hiroyasu, T., M. Miki, J. Kamiura, S. Watanabe, and H. Hiroyasu (2002), *Evolutionary Multiobjective Optimisation*, chap. 9, pp. 202–227, Springer.
- [83] Hood, J., E. Sall, and B. Charlton (2011), A GPS-based bicycle route choice model for San Francisco, California, *Transportation Letters: The International Journal of Transportation Research*, 3(1), 63–75, doi:10.3328/TL.2011.03.01.63-75.
- [84] Horvath, A. (2003), Pavement Life-cycle Assessment Tool for Environmental and Economic Effects (PaLATE), Microsoft Excel Document.
- [85] Horvath, A., and C. Hendrickson (1998), Comparison of Environmental Implications of Asphalt and Steel-Reinforced Concrete Pavements, *Transportation Research Record, Journal of the Transportation Research Board*, 1626(1), 105–113.
- [86] Hossain, M., M. Akbar, and J. Starkey (2007), Inexpensive construction of a 3D face model from stereo images, *2007 10th international conference on computer and information technology*, doi:10.1109/ICCITECHN.2007.4579387.
- [87] Huang, Y., R. Bird, and M. Bell (2009), A comparative study of the emission by road maintenance works and the disrupted traffic using life cycle assessment and micro-simulation, *Transportation Research Part D*, 14, 197–204.
- [88] Huang, Y., R. Bird, and O. Hendrich (2009), Development of a life cycle assessment tool for construction and maintenance of asphalt pavements, *Journal of Cleaner Production*, 17, 283–296.

- [89] International Energy Agency (2009), *Transport Energy and CO₂*, IEA Publications, Paris, France.
- [90] International Standards Organization (2006), ISO 14040, Life Cycle Assessment - Principles and Framework, *Tech. rep.*
- [91] IPCC (2013), IPCC Fifth Assessment Report: Summary for Policymakers, *Tech. Rep. September 2013*, Intergovernmental Panel on Climate Change, Geneva, Switzerland.
- [92] Ismail, K., T. Sayed, and N. Saunier (2010), Camera Calibration for Urban Traffic Scenes: Practical Issues and a Robust Approach, in *Transportation Research Board 89th Annual Meeting*.
- [93] Jacobsen, P. L. (2003), Safety in numbers: more walkers and bicyclists, safer walking and bicycling, *Injury Prevention*, 9(3), 205–209, doi:10.1136/ip.9.3.205.
- [94] Jacobsen, P. L., F. Racioppi, and H. Rutter (2009), Who owns the roads? How motorised traffic discourages walking and bicycling., *Injury prevention : Journal of the International Society for Child and Adolescent Injury Prevention*, 15(6), 369–73, doi:10.1136/ip.2009.022566.
- [95] Jensen, S. U. (2008), Bicycle Tracks and Lanes: A Before-and-After Study, in *Transportation Research Board 87th Annual Meeting*.
- [96] Johan de Hartog, J., H. Boogaard, H. Nijland, and G. Hoek (2010), Do the health benefits of cycling outweigh the risks?, *Environmental health perspectives*, 118(8), 1109–16, doi:10.1289/ehp.0901747.
- [97] Junhong, L., L. Zhizhang, G. Jianming, and L. Zhiwei (2003), Truck waste heat recovery for heating bitumen used in road maintenance, *Applied Thermal Engineering*, 23(4), 409–416, doi:DOI:10.1016/S1359-4311(02)00213-2.
- [98] Kahane, C. J. (2004), Lives Saved by the Federal Motor Vehicle Safety Standards and Other Vehicle Safety Technologies , 1960-2002, *Tech. Rep. October*, National Highway Traffic Safety Administration, Washington, DC.
- [99] Karlsson, R., U. Hammarström, H. Sörensen, and O. Eriksson (2011), Road surface influence on rolling resistance Coastdown measurements for a car and an HGV, *Tech. rep.*
- [100] Karp, R. M. (1972), *Reducibility among combinatorial problems*, chap. 8, pp. 219–241, Springer.
- [101] Keoleian, G. A., K. Kar, M. M. Manion, and J. W. Bulkley (1997), Industrial Ecology of the Automobile: A Life Cycle Perspective, *Tech. rep.*, Society of Automotive Engineers (SAE), Warrendale, PA.

- [102] Kim, J.-K., S. Kim, G. F. Ulfarsson, and L. A. Porrello (2007), Bicyclist injury severities in bicycle–motor vehicle accidents, *Accident Analysis & Prevention*, 39(2), 238–251.
- [103] Krizec, K. (2007), Estimating the Economic Benefits of Bicycling and Bicycle Facilities: an Interpretive Review and Proposed Methods, in *Essays on Transport Economics*, edited by P. Coto-Millán and V. Inglada, Contributions to Economics, pp. 219–248, Physica-Verlag HD.
- [104] Krizek, K. J., et al. (2006), Guidelines for analysis of investments in bicycle facilities, *Tech. rep.*, Transportation Research Board, Washington, DC.
- [105] Kuhn, K. D. (2010), Network-Level Infrastructure Management Using Approximate Dynamic Programming, *Journal of Infrastructure Systems*, 16, 103–111.
- [106] Kuhnimhof, T., J. Armoogum, R. Buehler, J. Dargay, J. M. Denstadli, and T. Yamamoto (2012), Men Shape a Downward Trend in Car Use among Young Adults - Evidence from Six Industrialized Countries, *Transport Reviews*, 32(6), 761–779, doi:10.1080/01441647.2012.736426.
- [107] Labi, S., and K. C. Sinha (2003), Measures of Short-Term Effectiveness of Highway Pavement Maintenance, *Journal of Transportation Engineering, ASCE*, 129(6), 673–683.
- [108] Lamprey, G., S. Labi, and Z. Li (2008), Decision support for optimal scheduling of highway pavement preventive maintenance within resurfacing cycle, *Decision Support Systems*, 46, 376–387.
- [109] Lang, R. E., and R. R. Sohmer (2000), Legacy of the Housing Act of 1949 : The Past , Present , and Future of Federal Housing and Urban Policy, *Housing Policy Debate*, 11(2), 291–298.
- [110] Laureshyn, A., H. Ardo, A. Svensson, and T. Jonsson (2009), Application of automated video analysis for behavioural studies: concept and experience, *IET Intelligent Transport Systems*, 3(3), 345, doi:10.1049/iet-its.2008.0077.
- [111] League of American Bicyclists (2013), 70 Largest Cities for Bike Commuting.
- [112] Lemer, A. C. (1996), Infrastructure Obsolescence and Design Service Life, *Journal of Infrastructure Systems*, 2(4), 153–161.
- [113] Li, Y., and S. Madanat (2002), A steady-state solution for the optimal pavement resurfacing problem, *Transportation Research Part A: Policy and Practice*, 36(6), 525–535, doi:DOI:10.1016/S0965-8564(01)00020-9.
- [114] Lidicker, J., N. Sathaye, S. Madanat, and A. Horvath (2012), Pavement Resurfacing Policy for Minimization of Life-Cycle Costs and Greenhouse Gas Emissions, *Journal of Infrastructure Systems*, doi:10.1061/(ASCE)IS.1943-555X.0000114.

- [115] Litman, T. (2010), Evaluating Public Transportation Health Benefits, *Tech. Rep. June*, Victoria Transport Policy Institute.
- [116] Menghini, G., N. Carrasco, N. Schüssler, and K. W. Axhausen (2010), Route choice of cyclists in Zurich, *Transportation Research Part A: Policy and Practice*, 44(9), 754–765.
- [117] Mesinger, F., et al. (2006), North American Regional Reanalysis, *Bulletin of the American Meteorological Society*, 87(3), 343–360, doi:10.1175/BAMS-87-3-343.
- [118] Meyer, M., and E. Miller (2000), *Urban Transportation Planning*, 656 pp., McGraw-Hill Science/Engineering/Math, New York, NY.
- [119] Millard-Ball, A., and L. Schipper (2011), Are We Reaching Peak Travel? Trends in Passenger Transport in Eight Industrialized Countries, *Transport Reviews*, 31(3), 357–378, doi:10.1080/01441647.2010.518291.
- [120] Miller, J. S., and L. D. Evans (2011), Divergence of potential state-level performance measures to assess transportation and land use coordination, *Journal of Transport and Land Use*, 4(3), 81–103, doi:10.5198/jtlu.v4i3.96.
- [121] Mohring, H., J. Schroeter, and P. Woonchutikula (1987), The values of waiting time, travel time, and a seat on a bus, *Rand Journal of Economics*, 18(1), 40–56.
- [122] Morcous, G., and Z. Lounis (2005), Maintenance optimization of infrastructure networks using genetic algorithms, *Automation in Construction*, 14, 129–142.
- [123] Mroueh, U. M., P. Eskola, and J. Laine-Ylijoki (2001), Life-cycle impacts of the use of industrial by-products in road and earth construction, *Waste management*, 21(3), 271–277.
- [124] NCHRP (1979), National Cooperative Highway Research Program Report 215: Pavment Managment System Development.
- [125] Noland, R. B. (2003), Traffic fatalities and injuries: the effect of changes in infrastructure and other trends., *Accident; analysis and prevention*, 35(4), 599–611.
- [126] Noland, R. B., and L. L. Lem (2002), A review of the evidence for induced travel and changes in transportation and environmental policy in the US and the UK, *Transportation Research Part D: Transport and Environment*, 7(1), 1–26, doi:10.1016/S1361-9209(01)00009-8.
- [127] Nordback, K., and B. Janson (2010), Automated Bicycle Counts, *Transportation Research Record: Journal of the Transportation Research Board*, 2190, 11–18, doi:10.3141/2190-02.

- [128] Nordback, K., W. Marshall, B. Janson, and E. Stolz (2013), Estimating Annual Average Daily Bicyclists : Error and Accuracy, in *TRB 2013 Annual Meeting*, Transportation Research Board, Washington, DC.
- [129] Norman, J., H. L. MacLean, and C. a. Kennedy (2006), Comparing High and Low Residential Density: Life-Cycle Analysis of Energy Use and Greenhouse Gas Emissions, *Journal of Urban Planning and Development*, 132(1), 10, doi: 10.1061/(ASCE)0733-9488(2006)132:1(10).
- [130] Norton, P. D. (2011), *Fighting Traffic: The Dawn of the Motor Age in the American City*, 408 pp., The MIT Press.
- [131] Office of Transportation and Air Quality (2013), Truck Carrier Partner 2.0.12 Tool: Quick Start Guide 2012 Data Year - United States Version, *Tech. rep.*, U.S. Environmental Protection Agency, Washington, DC.
- [132] OpenPlans (2013), Open Trip Planner.
- [133] OpenStreetMap Contributors (2013), Open Street Map.
- [134] Opus Central Laboratories, and Transportation Research Laboratory (1999), Review of VOC-pavement roughness relationships contained in Transfund's project evaluation, *Tech. rep.*, Opus International Consultants Limited, Central Laboratories, Lower Hutt, New Zealand.
- [135] Ouyang, Y. (2007), Pavement Resurfacing Planning for Highway Networks: Parametric Policy Iteration Approach, *Journal of Infrastructure Systems*, 1, 65–71.
- [136] Ouyang, Y., and S. Madanat (2004), Optimal scheduling of rehabilitation activities for multiple pavement facilities: exact and approximate solutions, *Transportation Research Part A: Policy and Practice*, 38(5), 347–365, doi: DOI:10.1016/j.tra.2003.10.007.
- [137] Paterson, W. D. O., and B. Attoh-Okine (1992), Summary Models of Paved Road Deterioration Based on HDM-III, *Transportation Research Record, Journal of the Transportation Research Board*, (1344), 99–105.
- [138] Pawlovich, M. D., W. Li, A. Carriquiry, and T. Welch (2006), Iowa's Experience with Road Diet Measures: Use of Bayesian Approach to Assess Impacts on Crash Frequencies and Crash Rates, *Transportation Research Record, Journal of the Transportation Research Board*, (1953), 163–171.
- [139] Persaud, B., and C. Lyon (2010), Evaluation of Lane Reduction "Road Diet" Measures on Crashes and Injuries, *Tech. Rep. 1*, Federal Highway Administration, Washington DC, doi:10.3141/1784-11.

- [140] Persaud, B., C. Lyon, and T. Nguyen (1999), Empirical Bayes Procedure for Ranking Sites for Safety Investigation by Potential for Safety Improvement, *Transportation Research Record: Journal of the Transportation Research Board*, 1665, 7–12.
- [141] Polzin, S. E. (2006), The Case for Moderate Growth in Vehicle Miles of Travel : A Critical Juncture in U.S. Travel Behavior Trends, *Tech. Rep. April*, U.S. Department of Transportation, Washington, DC.
- [142] Price Waterhouse Coopers, and European Business School Supply Chain Management Institute (2009), Transportation & Logistics 2030 Volume 2: Transport Infrastructure - Engine or hand brake for global supply chains?, *Tech. rep.*
- [143] Priedhorsky, R., D. Pitchford, S. Sen, and L. Terveen (2012), Recommending Routes in the Context of Bicycling : Algorithms , Evaluation , and the Value of Personalization, in *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pp. 979–988, New York, NY.
- [144] Putnam, R. D. (2001), *Bowling Alone: The Collapse and Revival of American Community*, 544 pp., Touchstone Books by Simon & Schuster.
- [145] Rahul, T., and A. Verma (2013), Economic impact of non-motorized transportation in Indian cities, *Research in Transportation Economics*, 38(1), 22–34, doi:10.1016/j.retrec.2012.05.005.
- [146] Raksuntorn, W., and S. I. Khan (2007), Saturation Flow Rate , Start-Up Lost Time , and Capacity for Bicycles at Signalized Intersections, *Transportation Research Record, Journal of the Transportation Research Board*, 1852, 105–113.
- [147] Robinson, D. (2007), Bicycle helmet legislation: Can we reach a consensus?, *Accident Analysis & Prevention*, 39(1), 86–93.
- [148] Rojas-Rueda, D., a. de Nazelle, M. Tainio, and M. J. Nieuwenhuijsen (2011), The health risks and benefits of cycling in urban environments compared with car use: health impact assessment study, *Bmj*, 343(aug04 2), d4521–d4521, doi:10.1136/bmj.d4521.
- [149] Rosey, F., J.-M. Auberlet, O. Moisan, and G. Dupre (2009), Impact of Narrower Lane Width: Comparison Between Fixed-Base Simulator and Real Data, *Transportation Research Record: Journal of the Transportation Research Board*, (2138), 112–119.
- [150] Sæ lensminde, K. (2004), Cost-benefit analyses of walking and cycling track networks taking into account insecurity, health effects and external costs of motorized traffic, *Transportation Research Part A: Policy and Practice*, 38(8), 593–606, doi:10.1016/j.tra.2004.04.003.

- [151] Santero, N. J., and A. Horvath (2009), Global warming potential of pavements, *Environmental Research Letters*, 4(3), 034,011, doi:10.1088/1748-9326/4/3/034011.
- [152] Santero, N. J., E. Masanet, and A. Horvath (2011), Life-cycle assessment of pavements Part II: Filling the research gaps, *Resources, Conservation and Recycling*, 55(9-10), 810–818, doi:10.1016/j.resconrec.2011.03.009.
- [153] Sathaye, N., and S. Madanat (2011), A bottom-up solution for the multi-facility optimal pavement resurfacing problem, *Transportation Research Part B: Methodological*, 45(7), 1004–1017, doi:10.1016/j.trb.2011.03.002.
- [154] Sathaye, N., and S. Madanat (2012), A bottom-up optimal pavement resurfacing solution approach for large-scale networks, *Transportation Research Part B*, 46(4), 520–528, doi:10.1016/j.trb.2011.12.001.
- [155] Sawhney, H., and R. Kumar (2005), Vehicle identification between non-overlapping cameras without direct feature matching, *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, pp. 378–385 Vol. 1, doi:10.1109/ICCV.2005.247.
- [156] Schrank, D., B. Eisele, and T. Lomax (2012), Urban Mobility Report, *Tech. Rep. December*, Texas A&M Transportation Institute, College Station, TX.
- [157] Sener, I., N. Eluru, and C. Bhat (2009), An analysis of bicycle route choice preferences in Texas, US, *Transportation*, 36(5), 511–539, doi:10.1007/s11116-009-9201-4.
- [158] Shoup, D. (2011), *The High Cost of Free Parking*, 2nd ed., 765 pp., American Planning Association.
- [159] Silva, C., T. Farias, H. Frey, and N. Rouphail (2006), Evaluation of numerical models for simulation of real-world hot-stabilized fuel consumption and emissions of gasoline light-duty vehicles, *Transportation Research Part D: Transport and Environment*, 11(5), 377–385, doi:10.1016/j.trd.2006.07.004.
- [160] Smart Growth America (2011), Virginia, *Tech. rep.*, Smart Growth America, Washington, DC.
- [161] Squires, G. D. (2002), Urban Sprawl and the Uneven Development of Metropolitan America, in *Urban Sprawl: Causes, Consequences, and Policy Responses*, chap. 1, pp. 1–22.
- [162] Stantec Consulting Services and H.W. Lochner (2007), Development of Performance Prediction Models for Virginia Department of Transportation Pavement Management System, *Tech. rep.*, Virginia Department of Transportation, Richmond, VA.

- [163] Stephens, T. (1988), Physical activity and mental health in the United States and Canada: Evidence from four population surveys, *Preventive Medicine*, 17(1), 35–47.
- [164] Stripple, H. H. (2001), Life Cycle Assessment of Road A Pilot Study for Inventory Analysis, second edition, *Tech. Rep. March*, IVL Swedish Environmental Research Institute Ltd.
- [165] SULLIVAN, J. M., and M. SIVAK (2012), CARBON CAPTURE IN VEHICLES: A REVIEW OF GENERAL SUPPORT, AVAILABLE MECHANISMS, AND CONSUMER-ACCEPTANCE ISSUES, *Tech. Rep. May*, University of Michigan Transportation Research Institute, Ann Arbor, MI.
- [166] Thenoux, G., G. Lvaro, and R. Dowling (2007), Energy consumption comparison for different asphalt pavements rehabilitation techniques used in Chile, *Resources, Conservation and Recycling*, 49, 325–339.
- [167] Thurston, D. L., and S. Srinivansan (2003), Constrained Optimization for Green Engineering Decision-Making, *Environmental Science and Technology*, 37(23), 5389–5397.
- [168] Transportation Research Board (2010), *Highway Capacity Manual*, Transportation Research Board of the National Academies, Washington DC.
- [169] Truffer, C. J., S. Keehan, S. Smith, J. Cylus, A. Sisko, J. A. Poisal, J. Lizonitz, and M. K. Clemens (2010), Health spending projections through 2019: the recession’s impact continues., *Health affairs (Project Hope)*, 29(3), 522–9, doi: 10.1377/hlthaff.2009.1074.
- [170] University of Washington, and CH2M Hill (2011), Greenroadsâ€Ž Version 1.5 Manual, *Tech. rep.*, University of Washington.
- [171] U.S. Census Bureau (2009), Census of Government Finance, Electronic.
- [172] U.S. Census Bureau (2012), American Communities Survey, *Tech. rep.*, U.S. Census Bureau, Washington, DC.
- [173] U.S. Environmental Protection Agency (1996), *Indicators of the Environmental Impacts of Transportation*, U.S. Environmental Protection Agency.
- [174] U.S. Environmental Protection Agency (2005), Emission Facts: Average Carbon Dioxide Emissions Resulting from Gasoline and Diesel Fuel, *Tech. Rep. February*, U.S. Environmental Protection Agency, Washington DC.
- [175] U.S. Environmental Protection Agency (2009), Mandatory Reporting of Greenhouse Gases; Final Rule, *Federal Register*, 74(209), 56,259–56,519.
- [176] U.S. Environmental Protection Agency (2013), Frequently Asked Questions on Mortality Risk Valuation.

- [177] US EPA, C. C. D. (), Greenhouse Gas Emissions: Transportation Sector Emissions.
- [178] U.S. Supreme Court (2007), Massachusetts v. U.S. Environmental Protection Agency, *United States Reports*, 549, 497–560.
- [179] Wang, X., K. Tieu, and W. E. L. Grimson (2010), Correspondence-free activity analysis and scene modeling in multiple camera views., *IEEE transactions on pattern analysis and machine intelligence*, 32(1), 56–71, doi:10.1109/TPAMI.2008.241.
- [180] Wang, Y. C., K. McPherson, T. Marsh, S. L. Gortmaker, and M. Brown (2011), Health and economic burden of the projected obesity trends in the USA and the UK, *The Lancet*, 378(9793), 815–825.
- [181] Waples, D., and J. Waples (2004), A Review and Evaluation of Specific Heat Capacities of Rocks, Minerals, and Subsurface Fluids. Part 1: Minerals and Nonporous Rocks, *Natural Resources Research*, 13(2), 97–122.
- [182] Wells, J. (2009), Paying Our Way: A New Framework for Transportation Finance, *Tech. rep.*, National Surface Transportation Infrastructure Commission, Washington, DC.
- [183] Wilson, J. H., L. Z. Williams, J. J. Schreiber, M. A. Mullen, T. D. Peterson, R. Strait, and J. James H. Wilson (2009), State approaches to reducing transportation sector greenhouse gas emissions, in *Transportation, Land Use, Planning, and Air Quality, Transportation, Land Use, Planning, and Air Quality: Selected Papers of the Transportation, Land Use, Planning, and Air Quality Conference 2009*, vol. 347, p. 9, ASCE, doi:doi:10.1061/41059(347)9.
- [184] Winters, M., G. Davidson, D. Kao, and K. Teschke (2010), Motivators and deterrents of bicycling: comparing influences on decisions to ride, *Transportation*, 38(1), 153–168, doi:10.1007/s11116-010-9284-y.
- [185] Wu, Z., and G. W. Flintsch (2009), Pavement Preservation Optimization Considering Multiple Objectives and Budget Variability, *Journal of Transportation Engineering, ASCE*, 135(5), 305–315, doi:10.1061/(ASCE)TE.1943-5436.0000006.
- [186] Youn, H., M. Gastner, and H. Jeong (2008), Price of Anarchy in Transportation Networks: Efficiency and Optimality Control, *Physical Review Letters*, 101(12), 128,701, doi:10.1103/PhysRevLett.101.128701.
- [187] Zaabar, I., and K. Chatti (2010), Calibration of HDM-4 Models for Estimating the Effect of Pavement Roughness on Fuel Consumption for U.S. Conditions, *Transportation Research Record: Journal of the Transportation Research Board*, 3(2155), 105–116.

- [188] Zaki, M. H., T. Sayed, and A. Cheung (2013), Automated Collection of Cyclist Data Using Computer Vision Techniques, in *TRB 2013 Annual Meeting*, Transportation Research Board, Washington, DC.
- [189] Zapata, P., and J. A. Gambatese (2005), Energy Consumption of Asphalt and Reinforced Concrete Pavement Materials and Construction, *Journal of Infrastructure Systems*, 11(1), 9–20.
- [190] Zhang, H., G. A. Keoleian, and M. D. Lepech (2008), An integrated life cycle assessment and life cycle analysis model for pavement overlay systems, in *Life-Cycle Engineering*, edited by F. Biondini and D. M. Frangopol, pp. 907–912.
- [191] Zhang, H., G. A. Keoleian, M. D. Lepech, and A. Kendall (2010), Life-Cycle Optimization of Pavement Overlay Systems, *Journal of Infrastructure Systems*, 16(4), 310–322, doi:10.1061/(ASCE)IS.1943-555X.0000042.
- [192] Zhang, H., M. D. Lepech, G. A. Keoleian, S. Qian, and V. C. Li (2010), Dynamic Life-Cycle Modeling of Pavement Overlay Systems: Capturing the Impacts of Users, Construction, and Roadway Deterioration, *Journal of Infrastructure Systems*, 16(4), 299–309, doi:10.1061/(ASCE)IS.1943-555X.0000017.
- [193] Zhang, H., G. A. Keoleian, and M. D. Lepech (2012), Network-Level Pavement Asset Management System Integrated with Life Cycle Analysis and Life Cycle Optimization, *Journal of Infrastructure Systems*, doi:10.1061/(ASCE)IS.1943-555X.0000093.
- [194] Zorn, W. (2008), Speed and Power Calculator.