



# TUM

TECHNISCHE UNIVERSITÄT MÜNCHEN  
INSTITUT FÜR INFORMATIK

## Determining the Most Harmful Roads in Search for System Optimal Routing

Jordan Ivanchev, Sorina Litescu, Daniel Zehe,  
Michael Lees, Heiko Aydt, and Alois Knoll

TUM-I1632

# Determining the most harmful roads in search for system optimal routing

Jordan Ivanchev<sup>1</sup>, Sorina Litescu<sup>1</sup>, Daniel Zehe<sup>1</sup>, Michael Lees<sup>2</sup>, Heiko Aydt<sup>3</sup>, and Alois Knoll<sup>4</sup>

<sup>1</sup>*TUM Create, Singapore*

<sup>2</sup>*University of Amsterdam, Netherlands*

<sup>3</sup>*Future Cities Laboratory, Singapore*

<sup>4</sup>*Technical University Munich, Germany*

## Abstract

Recent advances in Intelligent Transportation Systems (ITS), navigation tools and personal smart devices can be used to relieve congestion and thus improve traffic performance. Transportation networks, however, are seen as complex networks and therefore present a control challenge. In this paper we demonstrate that using an information dissemination technique and providing minimal but the right context to the population can steer the system into a more efficient operational state. As every commuter chooses the most optimal route from his/her own perspective, traffic distribution on the road network becomes heterogeneous, resulting in a small number of roads, which are largely overpopulated, while others remain underutilized [1]. In order to achieve a more homogeneous road utilization and thus reduce congestion levels, we propose a simple routing control strategy of informing drivers to avoid certain roads, which are chosen based on simulated outcomes of their closing. We demonstrate that the full removal of certain road segments from the network can redistribute traffic in a socially beneficial way, leading to an increase in transit performance on a city level. By considering a real road network and realistic traffic patterns we are able to validate our approach on a city scale. We identify the most harmful roads and quantify their negative effect on the system. Furthermore, since completely removing roads can be considered a rather extreme measure, we introduce the concept of *soft closing*. Instead of informing the whole population to avoid a certain road, we inform only a portion of the drivers, further improving the network utilization. We use the city of Singapore as a case study for our traffic assignment model which we calibrate and validate using both survey and GPS tracking devices data. By identifying and *soft closing* one road segment from the entire Singaporean road network (240,000 segments) we can reduce the average travel time of all 300,000 daily commuters by 6%, equating to 8,000 saved man hours.

## 1 Introduction and Existing Literature

In the context fast changing topologies of complex networks, this work examines the phenomenon of removing parts of a graph to improve the overall system's performance. In such cases the degree of improvement is several orders of magnitudes greater than the degree of change. This allows for effective steering mechanisms, which can exploit the complexity of the system. Using transportation systems as a case study, we confirm the existence of road segments that if closed can improve traffic conditions on a global level. Traffic demand changes dynamically throughout the day, however, the road infrastructure is static in the short term. We use the concept of *soft closing* of roads, to provide the capability of immediate adaptation to the road network. In this way information can be used as a steering tool in order to turn the previously static road infrastructure into a dynamically changing intelligent transportation system. Each participant, therefore, has an individual view of the network, depending on the information delivered. This new concept can be utilized by a central control system in order to efficiently make use of the network's resources and increase performance. We consider this a first step towards minimizing decentralization induced price of anarchy phenomena.

The dynamics of the most intriguing and significant examples of large complex social systems are governed by a human factor. The presence of free will, in particular, leads to highly stochastic behaviour, which can make modelling such systems challenging. The involvement of people in a system may further introduce a non-coordinated manner of operation. Those induced operational flaws, however, can be fixed by an efficient centralized control approach.

Control actions, on the other hand, may introduce reactions from the system, whose unpredictability increases with the size of the community and the magnitude of the intervention. Mechanisms for resolving inefficiencies, therefore, should be kept at minimal level of interaction in order to minimize the probability of introducing further problems. In this paper we aim at realizing a control strategy that reduces the negative effects of non-coordinated behaviour while keeping the degree of intervention at minimum, therefore maximizing efficiency and robustness.

In complex systems, inefficient states can be avoided by constructing adequate steering mechanisms [2]. Both the system's architecture, often represented by a network, and the time dependent dynamical interactions between the components make understanding and improving the performance of complex systems a challenging task. It has been shown that sparse inhomogeneous networks, which emerge in many real complex systems, are more difficult to control in comparison to dense and homogeneous systems [3].

As one example of complex systems, transportation systems are the subject of interest in variety of fields. There are many strategies of steering transportation systems that effectively increase traffic performance such as self-organizing traffic lights based on adaptation [4–6], or information dissemination techniques as [7–13], where commuters receive real time information about congestion in the network and adapt their routes accordingly.

The increasingly broader distribution of personal smart devices is a predisposition for the existence of Intelligent Transportation Systems (ITS). They progressively become more advanced [14] since the data availability provides a more complete view of the network, which leads to faster coordination [15]. Furthermore, since most drivers follow the advice provided by their navigation tool [16,17], the control of traffic can become more efficient and robust. The question stands, whether the system needs control at all. Drivers adapt to traffic conditions and tend to reach a Nash equilibrium state where a change of route for every participant would not be beneficial. In other words, the path choice of every driver is locally perceived as optimal and would not be voluntarily changed. This state of equilibrium is, however, not socially optimal when aiming at minimizing the overall population travel time. Our aim is to construct a centralized information dissemination system that counteracts this trend.

The phenomenon of non-coordinated social behaviour guided by individual optimal strategies is studied in [18]. While every actor in a scenario has an individually optimal strategy, the collection of those strategies results in a socially sub-optimal performance of the system. This discrepancy between local and global strategy outcomes is called *price of anarchy* (POA) and indicates inefficiency due to decentralization. In the case of transportation networks the standard uncoordinated behaviour is called *selfish routing* and is studied in [19]. It, however, also exists in other complex networks such as the Internet [20].

Measuring and reducing the POA has been object to numerous studies such as [19] and [21]. Furthermore, [22] develops a useful general theory for bounding the POA in games of incomplete information, where players are uncertain about each others' pay-offs. Finally a middle ground between centrally enforced solutions and completely unregulated behavior is sought in order to achieve stability in [23].

A trait of complex systems and their interaction is the emergence of the butterfly effect [24], where small changes in initial conditions can lead to performance alterations that are much bigger in magnitude [25]. We want to demonstrate that this holds true for traffic systems as well by removing a single road segment corresponding to one millionth of the size of a typical network. The effect of this modification on the systems performance in the sense of average travel time is then evaluated and compared in magnitude. Furthermore, we want to show that this change can, in fact, be beneficial for the system. In a way, we exploit the complexity of the system. First, we find the right road segment to remove using a brute-force search method (removing each link in turn and measuring the effect). Secondly, we try *soft closing* the road by informing a certain percentage of commuters that the road is closed.

Similarly to [13] and [8], in our study the control strategy is based on disseminating recommendations. Instead of providing information about traffic conditions, whose effects are highly unpredictable on a system level, we simply close a road for all users or just part of the population (*soft closing*). The choice of roads is not based on congestion levels but on the simulated outcome of those closures for the whole network. In this way the commuters are generically steered towards choosing more socially optimal routes and cases of local performance improvement that induce a negative effect on a global scale can be avoided.

Although unconventional, removing a road from the traffic infrastructure may lead to improved commuting conditions. The Braess paradox first mentioned in 1968 [26], states that adding extra capacity to a network where drivers act selfishly, can in some cases decrease performance. A generalization of this paradox [27] states that removing edges for large networks can produce an arbitrarily large improvement. It was further shown that the paradox can exist in all varieties of line-of-sight (LOS) networks as well [28]. Even the development of the human brain has a mechanism called synaptic pruning during which synapses (connections between neurons) are being removed in order to achieve more efficient learning [29]. It must be noted, however, that increasing the capacity of certain roads can lead to the avoidance of the paradox [30, 31].

There are numerous studies in real life cities that confirm the existence of the Braess paradox as in Stuttgart [32] and New York [33], where streets were closed for renovation or on purpose and better traffic conditions were observed. There are 70 more case studies from 11 countries that examine such conditions summarized in [34].

It is possible that, although well documented, in reality the Braess paradox may be stemming from secondary factors such as drivers taking less trips because of the reduced road network capacity. Also called disappearing traffic phenomenon [34], this translates in less overall usage of the road infrastructure. There seems to be no reasonable way to exclude the factor of willingness to travel when performing a real life experiment, which makes such empirical studies ambiguous. Furthermore, simulation based studies demonstrating the Braess paradox deal with artificial networks or just portions of real ones. In addition, a limited number of origin destination pairs are considered, thus making the results artificial. The chance that a road closure will be harmful to traffic conditions grows with increasing system size and generating authentic traffic that considers all participants and their diverse traffic demands, thus challenging the existence of the paradox in a realistic environment. We perform a complete city scale simulation, with systematic search of single road closure and provide a *soft closing* mechanism utilizing information dissemination tools to be able to dynamically control the system.

## 2 Data and Methods

### 2.1 Overview

In order to further study the examined phenomenon we perform a simulation based study, which allows us to control all factors in a systematic manner and state with certainty whether there are indeed harmful roads in a real world network scenario. By keeping the number of commuters and their origins and destinations constant within a single simulation run, we can isolate the phenomenon from all possible secondary influences and make sure that the measured changes in system performance are solely due to a change introduced in the network's topology.

The road network is modelled by means of a uni-direction graph. Nodes represent decision points at which a road may split or merge with another one. An intersection may be represented by a collection of nodes. Links represent road segments that connect two nodes. In order for a vehicle to traverse between its origin and destination, a route needs to be calculated based on the provided graph. The routes of all commuters are calculated using a stochastic routing approach. Every driver can have a preference for path choice, based on speed, distance or comfort, with assigned probabilities.

The travel time of every commuter is determined by the traverse times of all links included in its route. Those traverse times are calculated using a variation of the Bureau of Public Roads (BPR) function [35]. Realistic traffic is modelled by synthesizing a sufficiently large vehicle population based on Origin-Destination data available for the city of interest. Free flow velocities  $v_f$  are extracted from GPS tracking data. Parameters  $\alpha^s$  and  $\beta^s$  are calibrated for different types of roads depending on their speed limits using both GPS tracking data and a travel time distribution of the population for period of the day of interest.

Our case study examines the city of Singapore with population of 5.4 million people and around 1 million registered vehicles including taxis, delivery vans and public transportation vehicles [36]. It is an island city, which further simplifies our scenario since the examined system is relatively closed. We have used publicly available data to acquire a unidirectional graph of Singapore, that comprises of 240,000 links and 160,000 nodes representing the road system of the city. The number of lanes, speed limit and length of every link is available allowing us to extract information about its capacity.

For the purposes of our model we make use of two separate data sets. The first one is the Household Interview Travel Survey (HITS) conducted in 2012 in the city of Singapore, which studies the traffic habits of the population. Information about the origin destination pairs, their temporal nature, and

commuting time distribution during rush hour periods is extracted from it. The second data set consists of GPS trajectories of a 20,000 vehicle fleet for the duration of one month, providing information about recorded velocities on the road network during different times of the day.

## 2.2 Data Sets

### 2.2.1 HITS

The first data set that we have at our disposal is the Household Interview Travel Survey (HITS). It comprises of a large set of questions that aim at exploring the travelling habits of the population of Singapore. The survey covers slightly more than 0.67% of the population, which accounts to 35715 participants. Each person has answered 108 questions about demographics, commuting preferences and capabilities. We are, however, interested in the questions that deal with travel patterns. Every participant was asked to describe his/her trips for the whole day prior the day the survey was taken. This information is described in the following format:

Origin Code	Postal Code	Destination Code	Postal Code	Time of Start (hh:mm)	Duration (mins)	Means of Transportation
-------------	-------------	------------------	-------------	-----------------------	-----------------	-------------------------

The origin and destination locations are specified by a postal code. Singapore has a 6 digit postal code system, which allows for every building to have a unique postal code. Therefore the locations of interest can be pinned down with high precision. The column titled “means of transportation” can include various travel model such as private cars, taxi, public transportation, motorbike etc. Since our aim is to model the car population in Singapore and its behaviour, the entries that matter to us are the ones which create traffic . We are, therefore, looking for entries that put an extra vehicle on the road. All surveyed people that use public transportation are excluded from the data set since public transportation runs regardless of people that use it. Moreover, we also exclude the entries of passengers in private cars, in order not to count a vehicle multiple times. The trips that are left after the filtering process are used later for the agent generation process.

Furthermore, the information about the duration of the trips is used to create travel time distributions of the population, which is utilized in the calibration process. It is important to note that the survey was conducted on residents from different age groups, ethnicities, professions etc. It is, therefore, safe to assume that the data we can extract from the results is representative to an acceptable extent of the travel patterns in Singapore.

### 2.2.2 QI Data

The second data set that we use for this study is a GPS trajectory data from a commercial fleet tracking system. The size of the fleet is around 20,000. It comprises mainly of goods vehicles, truck and small lorries, however there is also data included from car leasing companies and personal trackers installed on private vehicles. The information about trip duration, origins and destinations, therefore, cannot be used to extract travel patterns reliably, since it is not representative for the commuting population. It can be used, however, to estimate average speeds on the roads with a good coverage of the whole network. The available data is for the duration of two months in 2014. Each entry has the following format:

Track id	Latitude	Longitude	Heading	Ground speed	Time stamp
----------	----------	-----------	---------	--------------	------------

The time difference between two consecutive signals (sampling period) from the same agent can vary between 1 second and 30 minutes. Typically there are more data points when a vehicle is turning and less data points (lower sampling rate) when the vehicle is going in a straight line. Vehicles usually have lower speeds at turns and move faster when going in a straight line. Therefore, there possibly is a slight bias towards lower velocities being recorded. In case of congestion, however, the vehicles typically go very slow in a straight line. As a result of this, it is possible that the number of samples from congested roads would be smaller than what would be expected if there was a fixed sampling frequency. Since the trackers send information 24 hours a day, there are time periods throughout which the vehicles are parked but still send out data. In order to exclude those samples we remove all data points, where there is no change in the position and velocity in the last 15 samples. After the samples from parked vehicles

have been removed, the size of the data set is around 120 million points. A map matching algorithm [37] is used on every trajectory in order to assign every sample point to an actual link in our road map of Singapore. All samples are then grouped according to links and time stamp in order to get a picture of the velocity profile of the city throughout the day.

## 2.3 Model and Simulation

### 2.3.1 Agent Generation

We are interested in simulating traffic conditions during rush hour. The chosen period of time is the morning rush hour since the traffic is more concentrated than during the evening and region of time during which it appears is more distinct. The simulated fragment of the day needs to be large enough in order to have enough samples from the HITS data in order to generate agents, however, we also want the traffic conditions during the time period to be as homogeneous as possible in order for the assumptions of our traffic model to hold. Therefore, we choose a one hour period from 7 to 8 that is centred around the time when most agents are starting their trips, which is around 7 : 30. Fig. 1 represents the distribution of trip starts throughout the day. The selected time period presents both the peak of trips starts and a homogeneous trip generation throughout.




Figure 1: Distribution of starting times of trips according to HITS data.

The number of agents that need to be generated on the road in order to have a quantitatively good representation of the traffic situation needs to be estimated as well. We use the HITS data in order to calculate this value. We assume that the data is representative for the portion of people who use cars or taxis from the whole population. Therefore we can say that,

$$\text{Car Users} = \text{Population} \times \frac{\text{Car Users in HITS}}{\text{Surveyed People in HITS}} \quad (1)$$

This states that we can find the number of people in Singapore that commute using a personal car or taxi (actively creating traffic) by knowing the percentage of people in HITS data that fulfil this requirement. The number of people who use cars and the total number of people in the HITS data are extracted from the itineraries by examining only trips with start time in the period 7 – 8 a.m. This gives us roughly 309,000 agents that should be generated during the examined time period.

The next step is to generate the origin, destination and starting time of every agent. As described in section 2.2.1 we have extracted a list of trips that actively create traffic on the roads. Using this list we create a distribution of postal codes being chosen as origins or destinations respectively, according to how many times they appear in the actual itineraries. We use a Bayesian estimator approach with a prior uniform distribution assumption. The process is as follows:

All existing postal codes in Singapore get an initial count of 1 and for every postal code that appears in our filtered trip data throughout the time period of interest we increment its counter by 1. Using those

counters we construct distributions for origin and destination postal codes such that the probability of a postal code being chosen as an origin or destination is proportional to its counter value.

This OD matrix construction approach is chosen in order to spread out the origins and destinations of population extracted from the HITS data set in order to represent reasonably well the traffic demands of the city. For every agent that we need to generate, we do the following steps:

1. Choose at random (uniformly sample) one of the trips that have been extracted from the HITS data.
2. Take the first two digits of the origin postal code and sample an existing postal code from the distribution of origins that have the same first two digits.
3. Take the first two digits of the destination postal code and sample an existing postal code from the distribution of destinations that have the same first two digits.

As a result of this the origins and destinations of all agents that need to be generated are determined. The next step is to calculate the routes that connect the start and end points of the agents' trips.

### 2.3.2 Routing

Since we aim to represent reality as much as possible the routing of the generated agents is stochastic. Some people prefer the shortest path, some the fastest and some prefer comfort rather than speed or time. We therefore have 3 distinct ways to calculate our routes. We are able to realize the various routing types by calculating the weights on our routing graph in different ways according to the respective preferences. After that, we use a shortest path algorithm that minimizes the sum of the weights for a path between origin and destination. The three types of weights are:

1.  $w_d$  = road length - minimizing distance
2.  $w_t = \frac{\text{road length}}{\text{road speed}}$  - minimizing time
3.  $w_c = \frac{\text{road length}}{\text{road speed} \times \text{number of lanes}}$  - maximizing comfort

After the generation of every agent one of the three preferences is chosen at random with probabilities  $p_d$ ,  $p_t$  and  $p_c$  respectively. The values of these three probabilities are calibrated since the preferences of routing choices vary from nation to nation. When the type of preference is chosen the corresponding route is calculated.

### 2.3.3 Traverse Time Calculation

We need to define some notation in order to proceed to describing the calculation of traverse times.

Variable	Description
$t_i$	time it takes to traverse road segment $i$ [s]
$l_i$	length of road segment $i$ [m]
$v_f^s$	free flow velocity on segment with speed limit $s$ [m/s]
$F_i$	flow on segment $i$
$w_i$	number of lanes on road $i$
$t$	simulation time from which the flow is calculated [hours]
$v_{min}$	minimum flow velocity at link $i$ at extreme congestion levels [m/s]
$I(i)$	function that checks if there is an intersection at the end of link $i$
$d^s$	intersection added delay for roads with speed limit $s$ [s]
$\alpha_s$	parameter from the BPR function for roads with speed limit $s$
$\beta_s$	parameter from the BPR function for roads with speed limit $s$
$S(i)$	number of successors of road segment $i$
$P(i)$	number of predecessors of road segment $i$

After calculating the routes of all agents, the number of vehicles that must pass through every link in the network can be extracted. The time needed to traverse a link  $t_i$  for the link  $i$  is calculated by using an extended version of the Bureau of Public Roads (BPR) function:

$$t_i = \min \left( \frac{l_i}{v_f^s} \left( 1 + \alpha^s \left( \frac{F_i}{2000w_it} \right)^{\beta^s} \right), \frac{l_i}{v_{min}} \right) + I(i)d^s \quad (2)$$

and

$$I(i) = \begin{cases} 1 & \text{if } S(i) + P(i) > 2 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

### 2.3.4 Assumptions

There are four assumptions that we make regarding our traffic model:

1. Agents are rational in the sense that they would choose the best possible route with respect to their preferences.
2. There is no re-routing in our model. Once the route of an agent is chosen, there will be no changes to it throughout the trip.
3. In order to model the traverse time on every link using the capacity and the estimated flow we need to assume that traffic is homogeneous within the simulation period. This may lead to a reduction of congestion levels since homogeneous traffic flows are a best case scenario.
4. The minimum possible velocity at a link in extreme congestion is set to  $5\text{km/h}$ . The BPR function that we use to estimate traverse times is known not to represent realistically extremely congested situations as the traverse time exponentially goes to infinity when the flow gets bigger. This is why we have set a minimum possible speed for all links that we believe is realistic, which means that in all cases agents keep moving forward with an average velocity of at least  $5\text{km/h}$ .

### 2.3.5 Extraction of free flow speeds

We split the links according to their speed limits into 3 categories with speed limit  $s = [50, 70, 90]$  km/h. The free flow velocities are extracted from the QI data set, where we have calculated the time variation of average velocities on all roads with the respective speed limits as shown on Fig.2. We take the maximum average velocity for each group of roads throughout the day and set it to be the free flow velocity  $v_f^s$ .




Figure 2: Extraction of free flow velocities for different road categories. Fig. 2a) shows average velocities throughout the day for roads with speed limit 50 km/h. Fig. 2b) shows average velocities throughout the day for roads with speed limit 70 km/h. Fig. 2c) shows average velocities throughout the day for roads with speed limit 90 km/h. The red dotted lines are used to mark the maximum velocity, which we consider to be also the free flow velocity for the respective type of roads.



Parameter	Value in Simulation	Value from Data
$\mu_t$	2.99	2.90
$\sigma_t$	0.851	0.881
$v^{50}$ [km/h]	22.4	22.8
$v^{70}$ [km/h]	39.1	35.5
$v^{90}$ [km/h]	64.3	59.3

Table 1: Comparison between simulation and real world data

Parameter	Calibrated Value
$\alpha^{50}$	0.8
$\beta^{50}$	2
$\alpha^{70}$	1
$\beta^{70}$	3
$\alpha^{90}$	1.2
$\beta^{90}$	5
$p_d$	0.31
$p_t$	0.33
$p_c$	0.36
$d^{50}$ [s]	1
$d^{70}$ [s]	4
$d^{90}$ [s]	1

Table 2: Calibrated Parameters and their values

### 2.3.6 Calibration

In order to calibrate the parameters of our simulation we use real world data from the HITS and QI data sets. First, the travel time distribution of people who start their trips within the time period of interest is constructed. The aim of the calibration process is to minimize the difference between this distribution and the one acquired from our simulation. Next, the average velocities on the roads grouped by their speed limits is extracted from the QI data. Once again, the difference between those velocities and the ones calculated in our simulation should be minimized. This multi-objective optimization problem is solved using grid search.

The first set of calibrated parameters are the  $\alpha$  and  $\beta$  parameters of the BPR function. Those can vary widely depending on the road conditions and drivers' behaviour, that is why they have to be calibrated for a specific population and infrastructure profile. The values that we have acquired after the calibration step fall well into the range of accepted values for the parameters [38].

The next set of parameters that are calibrated are the probabilities that agents have to choose a certain preferred type of route. We use as a starting point the values mentioned in [39]. The final calibrated values show that roughly one third of the agents prefer the shortest routes and the rest have a preference for fastest paths and wide roads, which may in some cases coincide.

The last set of parameters are the delays due to intersections for the three groups of roads. The calibrated values show that the most time on average is lost at major road intersections (usually due to traffic lights), while small roads and highways do not exhibit such large delays.

The calibrated parameters and their final values are noted in Table 2. On Fig. 3 the comparison between real and simulation data is presented. The specific values are also shown in Table 1. It can be observed that there is a slight tendency for lower velocities in results from the QI data set. This, as already mentioned in the data set description, may be due to the sampling algorithm employed in order to collect the data and its tendency to sample more when vehicles are turning and therefore have lower velocities.

### 2.3.7 Validation

In order to validate our results we have chosen the three most congested road segments according to our simulation. The velocity on those segments calculated using the traverse function either reaches the




Figure 3: Outcome of the calibration process. on Fig.3a) we can see a log normal distribution with the parameters extracted from the HITS data compared to a log normal distribution with the parameters extracted from the simulation. On Fig. 3b) we can compare the average speeds of the three groups of roads extracted from QI data and from the simulation.

critical preset minimum of 5 km/h or is very close. All three examined roads have a speed limit of 90 km/h. All samples of vehicles that have passed on those roads between 7 and 8 am on weekdays are extracted from the QI data. Their velocity profile is shown on Fig. 4. It can be observed that in reality as well as in our simulation those road segments are experiencing heavy congestion and low average speeds.

It should be noted that in cases of severe congestion the vehicles are mostly still, which results in a decrease in the number of samples for such periods. This means that the velocity profiles of the examined road segments in reality might show even higher degree of congestion. It can be observed from the graph that all three road segments that are severely congested in our simulation seem to be congested in reality as well according to the GPS tracking data, which means that our traffic assignment strategy and traverse time function have given appropriate approximations of reality for the desired level of detail.

Furthermore, in Fig. 5 we present a congestion map produced by our simulation and compare it with typical traffic pictures from Google Maps for the desired period of time. Since this service presents traffic averaged over 10 minutes intervals, we show three pictures from the beginning, middle and the end of the examined period. Our congestion map represents closely what is provided by the real data estimations of the Google Maps service.

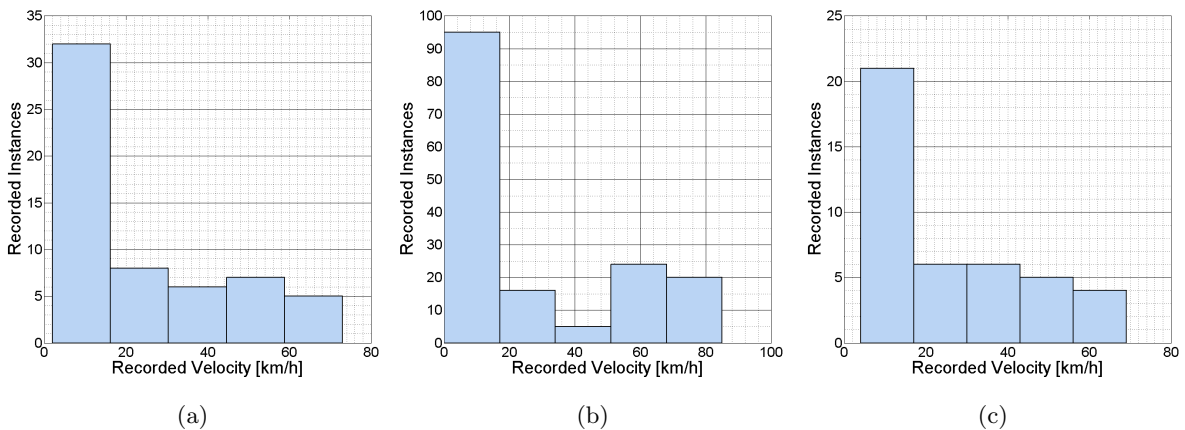


Figure 4: Velocity profile according to QI data of the most congested links from simulation results. The velocity samples are taken for the period between 7 and 8 a.m. for weekdays.



Figure 5: Comparison of congestion maps from our simulation 5d and Google Maps typical traffic service for 7 a.m. 5a, 7 : 40 a.m. 5b, and 8 : 15 a.m. 5c

## 3 Experiments

### 3.1 Study 1: Systematic Road Removal

The initial experiment aims at identifying links(road segments) whose closure would result in better traffic performance. It consists of examining all 240,000 links one by one and removing them from the routing graph. For every link removal a separate simulation run is performed and the routes and travel times of the population are recalculated according to the new road network. The results are compared to the initially simulated scenario, while the origins and destinations of all drivers are kept the same for all simulation runs.

The procedural sequences of actions for the systematic road closures are described in a step-by-step manner below in Algorithm 1. The algorithm describes the process of closing every link in the road network individually and re-simulating the traffic conditions, in the sense of re-calculating the traverse times for every link based on the flows extracted from the calculated routes of the agent population. In the case of soft closing, instead of closing every link, we inform only half of the agents that are using it, that the link is closed, which forces them to find an alternative route. Due to the stochastic element in choosing the group of informed agents, we perform the experiment with different randomly chosen groups 10 times in order to evaluate the variation in the results. The total number of links in the road network is 240,000, which amount to 2,640,000 simulations runs that were needed in order to perform the desired experiments.

**Data:**

$G$	Road network graph consisting of nodes and links
$A$	Set of all agents in the population
$L$	Set of all links in $G$
$ComputeRoutes$	Set of agents $\times$ Graph $\rightarrow$ Set of routes
$ComputeTravelTimes$	Graph $\times$ Routes $\rightarrow$ Set of travel times
$RemoveLink$	Link $\times$ Graph $\rightarrow$ Graph
$RandomSample$	Set of agents $\times$ Ratio $\rightarrow$ Set of agents

**Result:** Set of average population travel times for respective link closures at 100% -  $t^{100}$  and at 50% -  $t^{50}$

```

// Compute routes  $R^A$  and travel times  $T^A$ 
 $R^A \leftarrow ComputeRoutes(A, G)$ 
 $T^A \leftarrow ComputeTravelTimes(G, R^A)$ 
 $t^0 \leftarrow mean(T^A)$ 
 $t^{100} \leftarrow t^{100} \cup t^0$ 
 $t^{50} \leftarrow t^{50} \cup t^0$ 
foreach  $l \in L$  do
     $G^l \leftarrow RemoveLink(l, G)$  // Remove link  $l$  from the road network
     $A^l \leftarrow \forall a \in A : l \subseteq R^a$  // Identify agents that pass through link  $l$ 
    // Re-caluculate routes of affected agents and population travel times
     $R^{A^l} \leftarrow ComputeRoutes(A^l, G^l)$ 
     $R_l \leftarrow R^{A^l} \cup R^{A \setminus A^l}$ 
     $T^A \leftarrow ComputeTravelTimes(G, R_l)$ 
     $t^l \leftarrow mean(T^A)$ 
     $t^{100} \leftarrow t^{100} \cup t^l$  // Store the computed population average travel time
     $i \leftarrow 1$  while  $i < 10$  do
         $\tilde{A}^l \leftarrow RandomSample(A^l, 0.5)$  // Randomly sample half of the agents that pass through the
        link
        // Re-caluculate routes of affected agents and population travel times
         $R^{\tilde{A}^l} \leftarrow ComputeRoutes(\tilde{A}^l, G^l)$ 
         $\tilde{R}_l \leftarrow R^{\tilde{A}^l} \cup R^{A \setminus A^l}$ 
         $T^A \leftarrow ComputeTravelTimes(G, \tilde{R}_l)$ 
         $\tilde{t}^l \leftarrow mean(T^A)$ 
         $t^{50} \leftarrow t^{50} \cup \tilde{t}^l$  // Store the computed population average travel time
         $i \leftarrow i + 1$ 
    end
end

```

**end**

**Algorithm 1:** Quantifying population travel time change for partial and full closure of links

For each experiment a high performance cluster node was used. Each experiment ran on 32 threads on two Intel Xeon E5 (@ 2.60GHz) CPUs. The entire system has 192 GB of memory. The simulation used a bi-directional dijkstra implementation from the SEMSim traffic nano-scopic traffic simulation [40]. Since routing requests can be paralleled for each trip, the performance of the simulation benefited from the large number of threads. In total each all threads handled around 21,000,000 routing requests for each simulation experiment. Since we used 3 different metrics for weight calculation (distance, travel time and comfort), each thread had to load all three networks in order to ensure maximum performance. Therefore, these experiments had to be run on high performance hardware to ensure a quick turnaround time on the experiments.

As a result of this procedure, we can evaluate the effects of every link's closure on the average travel time of the population. We have found that the closing of a road segment can indeed lead to a reduction in the average travel time. In 21 cases the closure of a link in the network leads to a decrease of 1 minute or more in the average travel time, which corresponds to 3.73% overall system performance increase. The most harmful link gives a 74.25 seconds decrease of overall trip duration translating to 6400 saved hours for the driver population on a daily basis, solely from the morning rush hour period. Although part of the backbone of the network in a topological sense, the removal of certain major road segments, would decrease overall travel time. In other cases, however, removing such important roads significantly increases commuting time for the population. Those links are identified as crucial for the traffic system.

A closer look at the implications of closing one of the beneficial links can be seen in Fig. 6. The closure of the examined link reduces the amount of traffic on the roads in its proximity. The collective length of roads receiving traffic is far greater than that of the roads that experience reduced flow volumes. The vehicles that are taken from the closed road and the ones in its vicinity are spread over a larger portion of the city thus increasing the homogeneity of traffic flows. Since average travel time is reduced we can conclude that the streets that receive traffic produce less additional time than what is gained by the streets that give away traffic. Due to the non-linear nature of the *traverse time vs. flow* relationship taking an agent from a very congested street and putting it on a less congested one, indeed makes a difference and reduces overall trip duration. Therefore, as a result of the road closure traffic from a group of highly congested roads is distributed along less populated parts of the network thus relieving traffic conditions.

In most cases, however, traffic that has to be reassigned has to go through a smaller road than the initial one and congests the system even further. This does not happen in only 639 out of the 240,000 examined cases, corresponding to a 99.73% probability of worse traffic conditions arising from a road closure. Intuitively, after considering the complexity of a large city road network and the hundreds of thousands of vehicles on the roads, decreasing connectivity of the road infrastructure will rather increase congestion than relieve it. As noted in our results, the likelihood of observing commuters being forced into more socially beneficial paths given a road closure is indeed small, but this event is not impossible. It seems that in such cases the additional bottlenecks, if any, introduced by the road network alteration are less harmful than the traffic stress that is relieved from the initial state of the system.

One can consider two main reasons for the occurrence of the observed phenomenon. The first one is selfish routing whose effects can be diminished if a certain road is closed leaving the drivers no alternatives other than choosing a more socially optimal path. The second reason is the nature of transportation networks evolution, which adapt in an incremental manner in time to the changes in traffic demand. As a natural consequence of that as new roads are added, rarely old ones are removed. And although, evolutionary processes undeniably offer working solutions, their optimality is not guaranteed. It is, therefore, not improbable that some of the old roads become obsolete or even harmful.

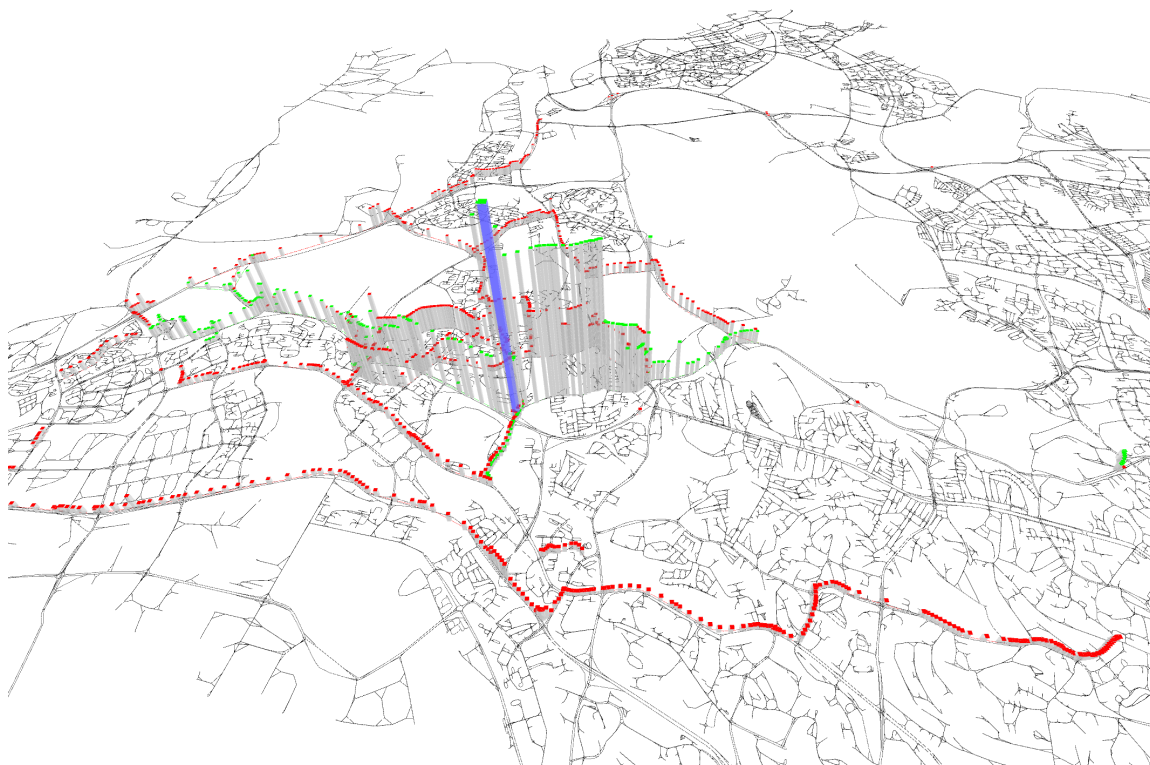


Figure 6: A map representing the changes that occur in traffic due to the closing of the road indicated in blue. Green and red bars represent reduced and increased traffic respectively. The height of the bars illustrates the magnitude of the change.

Since our aim is to maximize utilization of the road network, it may be contradictory and inefficient

Improvement threshold [s]	Links in Study 1	Links in Study 2	Links in common	Percentage with respect to Study 2 [%]
10	639	487	427	87.68
20	289	276	221	80.07
30	221	167	126	75.45
40	126	111	48	43.24
50	50	78	27	34.62
60	21	64	2	3.13

Table 3: Percentage of common links between the sets found in Studies 1 and 2 depending on the improvement threshold.

to completely remove parts of it. Although improving traffic conditions by closing roads is a remarkable phenomenon, greater social benefits can be achieved with milder measures. It is possible that a driver can have a only one viable path between his origin and destination. In such a case, if a removal of a link on this path is globally beneficial, the specific driver would be forced into an unnecessary d-tour. Therefore, it might be more appropriate to not fully dispose of already existing infrastructure, which can help avoid such extreme cases.

### 3.2 Study 2: Soft Closing

We suggest the concept of *soft closing* of links. Rather than removing a link from the road network completely, we remove it only for a fraction of the agents that initially pass through it. In this way traffic demand in the city is re-distributed more homogeneously. An extension of our initial experiment is performed where instead of informing all agents that a certain road is unavailable we do so only for half of the drivers passing through it. The experiment is done in order to examine whether partial closing of certain links can further decrease overall travel time. The figure of 50% closure was chosen since it is the middle ground between completely closed and opened. We refer to drivers that cannot pass through the link anymore as *informed* and the to the ones that can stay as *uninformed*. Since the two groups are chosen at random, we perform the experiment for every link 10 times in order to evaluate the effects of informing different subgroups.

We have found 64 links whose partial closure leads to more than 1 minute decrease of average trip duration. The link that shows best performance if partially closed gives 100 seconds decrease of the population average travel time. This amounts to 6.25% increase in system-wise performance. It must be noted that the links, which lead to biggest improvement if completely closed and those who are only half closed do not coincide. Only 2 out of the 64 links with best performance from the *soft closing* experiment are in the list of the most harmful links from the systematic road removal experiment, which indicates an underlying categorization of roads according to their optimal flows.

This categorization can be observed by examining the sets of best worst links from the two studies. 87% percent of the links that lead to more than 10 seconds improvement of average travel time are the same for both experiments. If we examine the common links that lead to more than 1 minute improvement, however, the percentage is just 3% (more detailed data in Table 3). This shows that although the general sets of best worst links are very similar, the ones that obtain the best results for a given percentage are indeed the links that “specialize” in this particular percentage.

In order to evaluate the induced system effects arising from the variation of percentage of informed agents, we examine in detail the top two links representing different roads from each of the sets of the most harmful road segments from our two previous experiments. We perform a sweep of the percentage of *informed* agents and for every step evaluate the average trip duration for the population described in Algorithm 2. The percentage of informed agents ( percentage of closure) is varied from 0 to 100 in 10% steps. For every percentage, similarly to the *soft closing* approach, we have performed 10 separate runs in order to evaluate the variation of the results .

Fig. 7 provides an overview of the effects of changing the accessibility of a link to the whole system. The curves of the links that come from *Study 1* (Road 1 and 2) reach their minimum in proximity to 100% and resemble a linear function. The links coming from *Study 2*(Road 3 and 4) have convex curves with optimal percentage of redirected agents between 40% and 50%. In the latter cases by further reducing

**Data:**

$G$	Road network graph consisting of nodes and links
$A$	Set of all agents in the population
$S$	Set of all links chosen from previous studies
$ComputeRoutes$	Set of agents $\times$ Graph $\rightarrow$ Set of routes
$ComputeTravelTimes$	Graph $\times$ Routes $\rightarrow$ Set of travel times
$RemoveLink$	Link $\times$ Graph $\rightarrow$ Graph
$RandomSample$	Set of agents $\times$ Ratio $\rightarrow$ Set of agents

**Result:** Set of average population travel times for varying percentages of link closures  $t_l^0, t_l^{0.1} \dots t_l^1$  for all chosen links  $l$

```

// Compute routes  $R^A$  and travel times  $T^A$ 
 $R^A \leftarrow ComputeRoutes(A, G)$ 
 $T^A \leftarrow ComputeTravelTimes(G, R^A)$ 
foreach  $l \in S$  do
   $t_l^0 \leftarrow mean(T^A)$ 
   $p \leftarrow 0.1$ 
   $A^l \leftarrow \forall a \in A : l \subseteq R^a$  // Identify agents that pass through link  $l$ 
  while  $p \leq 1$  do
     $c \leftarrow 0$  while  $c < 10$  do
       $\tilde{A}^l \leftarrow RandomSample(A^l, p)$  // Randomly sample the indicated by p percentage of agents
      // Re-calculcate routes of affected agents and population travel times
       $R^{\tilde{A}^l} \leftarrow ComputeRoutes(\tilde{A}^l, G^l)$ 
       $\tilde{R}_l \leftarrow R^{\tilde{A}^l} \cup R^{A \setminus A^l}$ 
       $T^A \leftarrow ComputeTravelTimes(G, \tilde{R}_l)$ 
       $\tilde{t}^l \leftarrow mean(T^A)$ 
       $t_l^p \leftarrow t_l^p \cup \tilde{t}^l$  // Store the computed population average travel time
       $c \leftarrow c + 1$ 
    end
     $p \leftarrow p + 0.1$ 
  end
end

```

**Algorithm 2:** Computation of information variation profile of chosen best worst links

the traffic on the selected links the system's performance starts to deteriorate due to other congestion spots created as a result of the traffic re-distribution. Fig. 7 also provides a visual evidence for the categorization of links according to optimal flows.

A question that arises when utilizing a *soft closing* strategy is whether the choice of agents that are not allowed on the link affects the results. For every percentage of closure of every link that we examine, the group of *informed* agents is sampled from all drivers that initially need to use the examined road segment. Therefore, every time a different set of agents is informed and redirected. It is intuitive to expect variations in the results since every driver has a distinct origin and destination and would affect the system differently after re-routing. The computed coefficient of variation  $\sigma / \mu$  on those experiments is  $4 \times 10^{-4}$ . In Study 3, we have picked 2 of the best worst links identified during the previous two studies and varied the percentage of agents that are informed of their closure. The measured coefficients of variation  $\frac{\sigma}{\mu}$  is recorded in Table 4.

It can be noted the deviation is surprisingly small, which means that the choice of agents, which need to find other routes is not a decisive factor. This simplifies significantly the analysis of our *soft closing* strategy. This unexpected discovery may be explained with the fact that by considering a real world scenario we also get a great variety of origin destination pairs. The apparent homogeneity of agents on this level of abstraction thus allows us to consider them as groups rather than individuals.

### 3.3 Spatial distribution of effects of partial and full road closures

Fig. 8a depicts the spatial distribution of road closure effects on the average travel time of the population. Most of the segments that lead to a significant change in population travel time, if closed, would have a

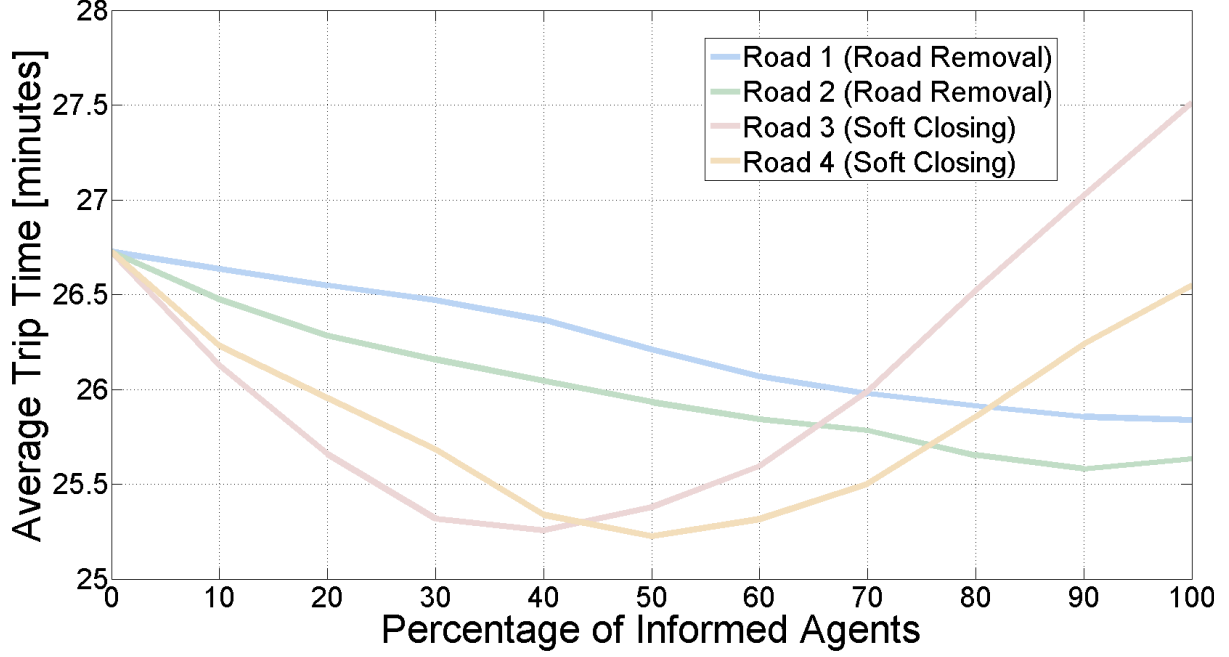


Figure 7: Variation of the average population travel time for different percentages of “soft” closing of examined links.

Road ID	10%	20%	30%	40%	50%	60%	70%	80%	90%
1	$4.3 \times 10^{-4}$	$4.5 \times 10^{-4}$	$3.1 \times 10^{-4}$	$3.3 \times 10^{-4}$	0.0017	$6.2 \times 10^{-4}$	$3.4 \times 10^{-4}$	$3.4 \times 10^{-4}$	$4.2 \times 10^{-4}$
2	$1.9 \times 10^{-6}$	$3 \times 10^{-6}$	$5.3 \times 10^{-6}$	$5.8 \times 10^{-6}$	$4.6 \times 10^{-6}$	$3 \times 10^{-6}$	$3 \times 10^{-6}$	$2.6 \times 10^{-6}$	$3 \times 10^{-6}$
3	$1.5 \times 10^{-4}$	$2.5 \times 10^{-4}$	$3.4 \times 10^{-4}$	$2.5 \times 10^{-4}$	$2.4 \times 10^{-4}$	$4.3 \times 10^{-4}$	$3.6 \times 10^{-4}$	$2.9 \times 10^{-4}$	$1.9 \times 10^{-4}$
4	$5.8 \times 10^{-4}$	$3.9 \times 10^{-4}$	$3 \times 10^{-4}$	$6.8 \times 10^{-4}$	$2.7 \times 10^{-4}$	$3.4 \times 10^{-4}$	$3.7 \times 10^{-4}$	0.0016	$1.4 \times 10^{-4}$

Table 4: Variation coefficients for chosen links from Study 3 for different percentages of closure

negative impact on the system. In this way using the results of the study, we can also identify the crucial portions of the network, whose removal is highly undesirable. Although, smaller in numbers, the links who are harmful to system seem to cover some of the backbone roads of the city. Those roads would be considered important, by just looking at the network topology, however, according to our results their removal can reduce congestion levels. There are regions of the city where we can even observe an alternation on a single road of beneficial and harmful segments. One of those sensitive regions is shown in more detail on Fig. 9a. In contrast to this, Fig. 9b shows the effects of half closures in the same region, where it can be observed that using a soft closing strategy for the whole portion of the road is beneficial in the sense of average population travel time.

The spatial profile from Study 2 illustrated on Fig. 8b, shows a similar set of road segments that are significant (cause a change in average travel time of more than 10 seconds), however some of the links that were crucial in Study 1 (should not be closed), are coloured in blue, which means that their half closure would benefit the population. Apparently the significant links can be either crucial for traffic or harmful depending on the percentage of their closure. There probably is a underlying categorization of those significant links according to their optimal throughput.

Fig. 10a shows that the full or partial removal of the majority of road segments in the network would have almost no effect on the average travel time of the population. The removal of a single road segment does not present a drastic change to the network, therefore it is expected that the effect of such



a removal is minuscule. There are, however, counterexamples of this, which can be observed in more detail by removing the unaffected links from the set of results so that only the road segments that have a more significant effect on traffic if closed are left (illustrated on Fig. 10b). The distribution from Study 2 seems to be slightly shifted towards the negative side, which means that soft closing strategies can be more beneficial and less harmful than complete road removal ones.

### 3.4 Study 3: Local Effects and Equilibria analysis

As a final step, the implications for affected agents are studied for the four previously examined links. Fig. 11 displays how the *informed* and *uninformed* agents perform for different degrees of closure of links. Examining travel times at the point of social optimum, the groups of *informed* agents save between 23% and 41% travel time, while the *uninformed* agents benefit the reduced congestion on the initial path and get between 23% and 50% improvement. This shows that none of the primary affected groups of agents experiences negative effects. On the contrary, the improvements in their average travel times are between 4 and 8 times higher than the overall population performance increase.

Given the invariance of our results to the *informed* agents selection, it is expected that the curve of *uninformed* agents has a negative slope, since congestion levels along their paths decrease. It can be noticed that the *informed* and *uninformed* agents' curves cross on Fig. 11. At the point of crossing it can be assumed that an agent who has perfect information about the traffic situation would not make the choice to change from *informed* to *uninformed* or vice versa. This point can be perceived as a Nash equilibrium for this link for a single commuter.

Furthermore, the point of equilibrium for the collective group of *informed* and *uninformed* agents, which can be considered as a local social optimum, is also identifiable by locating the minimum of the affected agents group average travel time function. The three points of interest (single agent equilibrium, affected agents group equilibrium and social optimum) do not coincide in the 3 out of the 4 studied cases as seen on Table 5. The desired percentage of closure that should be chosen in general is the percentage for social optimum to occur since it saves the biggest amount of total time. If, however, agents choose their routes selfishly or even in local groups a different equilibrium point will be reached, thus leading to a sub optimal traffic distribution resulting in society paying the POA due to lack of centralization.

It must be noted that this discrepancy does not result from lack of information. The alternatively calculated personal Nash equilibrium and group equilibrium are based on full knowledge of the system. We can thus conclude that simply choosing the fastest route even in the presence of perfect information does not lead to optimal traffic distribution. It is, therefore, vital that the system is always considered as a whole because the collection of local optimal solutions may not produce the expected result due to the high complexity.

Road	Social Optimum [%]	Nash Equilibrium [%]	Affected Group Optimum [%]
Road 1	100	70	90
Road 2	90	90	90
Road 3	40	27	50
Road 4	50	44	60

Table 5: Points of Equilibrium



(a)



(b)

Figure 8: Comparison of road closure effects from Study 1 Fig. 8a and Study 2 Fig. 8b. Road segments coloured in the red gamma represent increase in average population travel time. The thicker and redder road segments are represented the higher the increase of average population travel time. Road segments coloured in the blue gamma analogically represent a decrease of average population travel time. The closure of roads that are not coloured would result in a change of average travel time of less than 10 seconds.

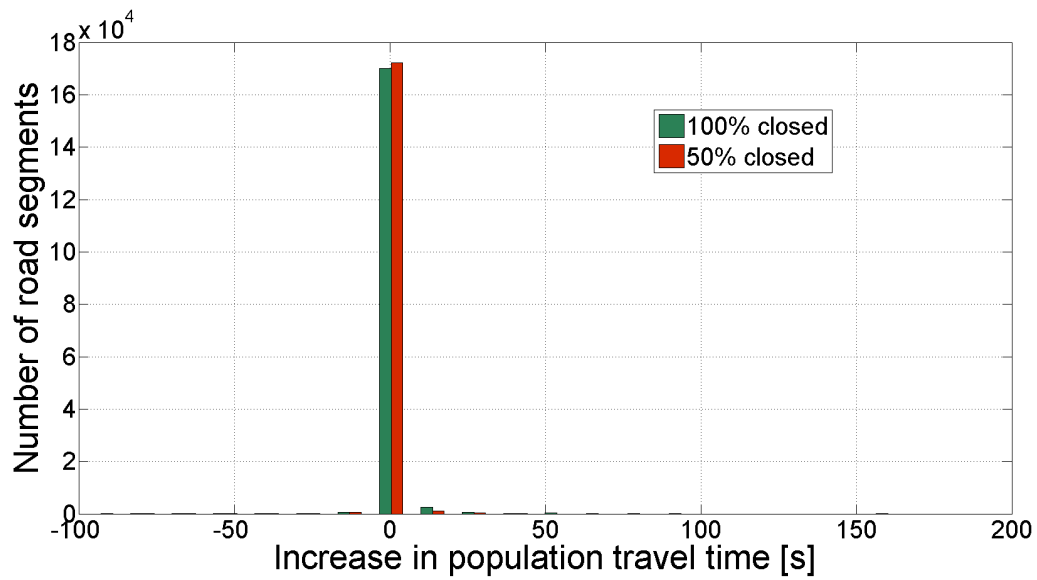


(a)

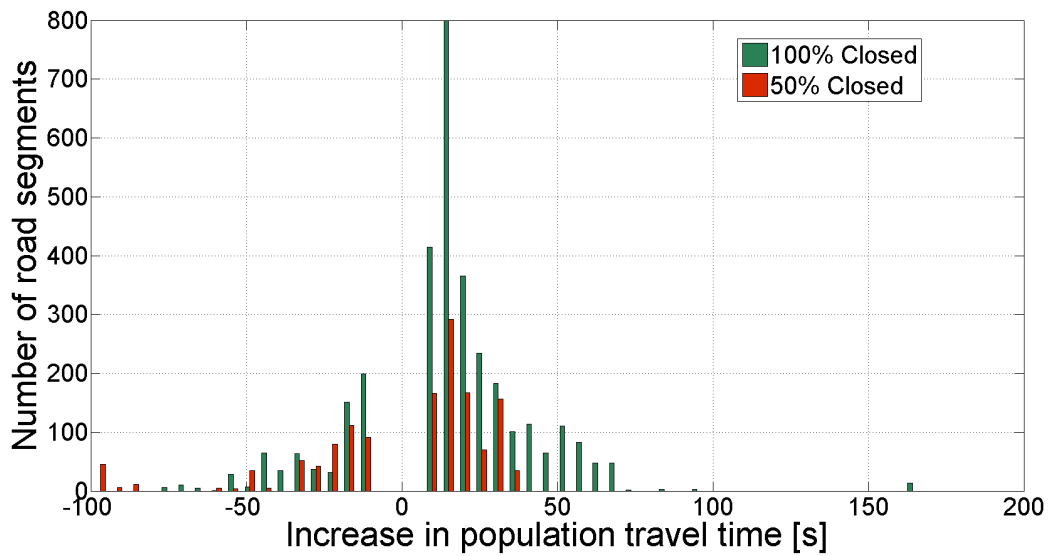


(b)

Figure 9: Zoomed version of Fig. 8 in order to observe in more detail a sensitive area from Study 1 Fig. 9a and Study 2 Fig. 9b



(a)



(b)

Figure 10: Results of Study 1 and 2 summarized in histograms. Fig. 10a shows the distribution of links according to their effect on the average population travel time. Fig. 10b shows the same distribution, however all links that have an effect smaller than 10 seconds in magnitude are excluded from the distribution.

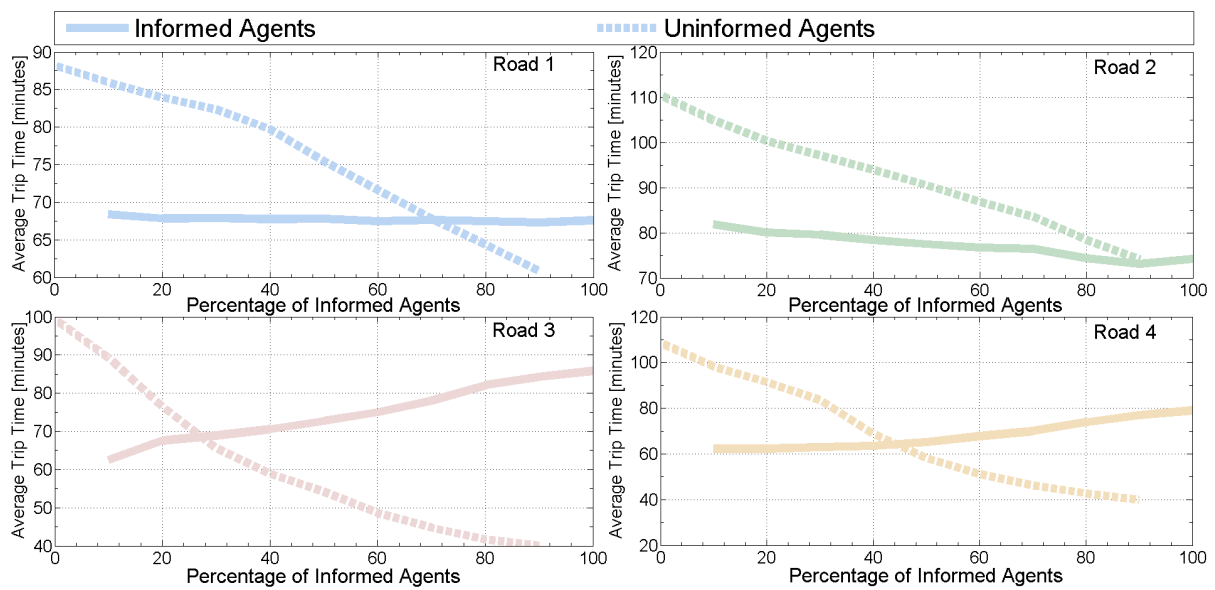


Figure 11: Comparison of the average travel time of informed and uninformed agents for different percentages of *soft closing*. The solid line represents the progression of average travel time of the agents that are re-routed as the percentage of information increases. The dashed line represents the performance of the group of agents that can stay on the link. Every sub graph is specific to one of the examined links from Fig. 7 mapped with the respective colors.

## 4 Conclusion

In conclusion, we have confirmed that by disseminating information about a road that should be avoided to all traffic participants on a large city scale the system's performance can be improved. The closure of just a single road segment might reduce overall travel time for the whole commuting population by as much as 4% which corresponds to thousands of saved hours on a daily basis. This rather extreme outcome of a small change in initial conditions is a manifestation of the complexity of the system. By disseminating the information only partially to the population we perform *soft closing* of roads in order to further improve traffic conditions (up to 6%). This strategy provides the ability of the road network to behave dynamically, at zero infrastructure construction cost, via information dissemination.

We aim at simulating a realistic scenario by taking a complete road network of a large city and populating it with agents according to collected real world data that is also used for the calibration and validation of our model. By using a simulation study we gain control over the environment of the scenarios, thus allowing us to remove all secondary factors and influences so that the examined phenomenon can be isolated and studied in a profound way. The completeness of the scenario ensures that the effects of the paradox are not only local but global. Changes in the modelling of traverse times will affect the whole system in the same way, which may slightly change the results quantitatively but not qualitatively. It must be noted that if the disappearing traffic due to reduced capacity of the network is modelled, the congestion will decrease. Exploiting the computational power available in present days will be a key tool for the planning, control and support of future smart cities. Simulation based methods such as *soft closing* can be used to ensure efficient utilization of resources and fast instantaneous adaptability to demand changes. It is important to study further such information dissemination techniques since they provide flexibility and dynamic properties to the physically static road infrastructure, a trait that is highly desired in the future of transportation.

The method of simulating outcomes of network changes can also be used for future infrastructure planning to avoid building roads that produce congestion and for personalized road pricing that can be used to balance traffic flows and achieve a social equilibrium state. We expect that future smart cities will rely heavily on simulation approaches enabled by the increase of computational power availability. A futuristic ITS making use of our "soft closing" approach will be able to dynamically change the road network and thus continuously steering the system dynamics into optimal states.

Finally, if excessive traffic congestion is viewed as the marginal costs to society of congestion exceeding the marginal costs of efforts to reduce congestion (such as adding to road or other transport infrastructure), our approach presents a strategy that can both ease traffic conditions and is cost free in the sense that no construction of infrastructure is necessary.

## Acknowledgments

This work was financially supported by the Singapore National Research Foundation under its Campus for Research Excellence And Technological Enterprise (CREATE) programme.

## References

- [1] Stefan Lämmer, Björn Gehlsen, and Dirk Helbing. Scaling laws in the spatial structure of urban road networks. *Physica A: Statistical Mechanics and its Applications*, 363(1):89–95, 2006.
- [2] Wen-Xu Wang, Xuan Ni, Ying-Cheng Lai, and Celso Grebogi. Optimizing controllability of complex networks by minimum structural perturbations. *Physical Review E*, 85(2):026115, 2012.
- [3] Yang-Yu Liu, Jean-Jacques Slotine, and Albert-László Barabási. Controllability of complex networks. *Nature*, 473(7346):167–173, 2011.
- [4] Stefan Lämmer and Dirk Helbing. Self-control of traffic lights and vehicle flows in urban road networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(04):P04019, 2008.
- [5] Carlos Gershenson. *Design and control of self-organizing systems*. CopIt Arxivs, 2007.
- [6] Seung-Bae Cools, Carlos Gershenson, and Bart DHooghe. Self-organizing traffic lights: A realistic simulation. In *Advances in applied self-organizing systems*, pages 41–50. Springer, 2008.

- [7] Giovanni Petri, H Jeldtoft Jensen, and John W Polak. Global and local information in traffic congestion. *EPL (Europhysics Letters)*, 88(2):20010, 2009.
- [8] Javed Aslam, Sejoon Lim, and Daniela Rus. Congestion-aware traffic routing system using sensor data. In *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on*, pages 1006–1013. IEEE, 2012.
- [9] Sejoon Lim, Hari Balakrishnan, David Gifford, Samuel Madden, and Daniela Rus. Stochastic motion planning and applications to traffic. *The International Journal of Robotics Research*, page 0278364910386259, 2010.
- [10] Sejoon Lim and Daniela Rus. Stochastic motion planning with path constraints and application to optimal agent, resource, and route planning. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4814–4821. IEEE, 2012.
- [11] Sejoon Lim and Daniela Rus. Congestion-aware multi-agent path planning: distributed algorithm and applications. *The Computer Journal*, 57:825–839, 2013.
- [12] Salvatore Scellato, Luigi Fortuna, Mattia Frasca, Jesús Gómez-Gardeñes, and Vito Latora. Traffic optimization in transport networks based on local routing. *The European Physical Journal B*, 73(2):303–308, 2010.
- [13] Sorina Litescu, Vaisagh Viswanathan, Michael Lees, Alois Knoll, and Heiko Aydt. Information impact on transportation systems. *Journal of Computational Science*, 9:88–93, 2015.
- [14] Woodrow Barfield and Thomas A Dingus. *Human factors in intelligent transportation systems*. Psychology Press, 2014.
- [15] Daniel Enemark, Mathew D McCubbins, and Nicholas Weller. Knowledge and networks: An experimental test of how network knowledge affects coordination. *Social Networks*, 36:122–133, 2014.
- [16] Verena Franken and Transportation Research Board. Use of navigation systems and consequences for travel behaviour . In *ECTRI-FEHL-FERSI Young Researcher Seminar, Brno, Czech Republic*, 2007.
- [17] Verena Franken and Barbara Lenz. Influence of mobility information services on travel behavior. In *Societies and Cities in the Age of Instant Access*, pages 167–178. Springer, 2007.
- [18] Hyejin Youn, Michael T Gastner, and Hawoong Jeong. Price of anarchy in transportation networks: efficiency and optimality control. *Physical review letters*, 101(12):128701, 2008.
- [19] Tim Roughgarden. How unfair is optimal routing? In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 203–204. Society for Industrial and Applied Mathematics, 2002.
- [20] E Friedman. A generic analysis of selfish routing. Technical report, Working paper, 2001.
- [21] Tim Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [22] Tim Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 862–879. ACM, 2012.
- [23] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [24] Edward N Lorenz. *The essence of chaos*. University of Washington Press, 1995.
- [25] Steven D Gribble. Robustness in complex systems. In *Hot Topics in Operating Systems, 2001. Proceedings of the Eighth Workshop on*, pages 21–26. IEEE, 2001.
- [26] NG Bean and PG Taylor. Can braess’s paradox occur in loss networks. *University of Adelaide (November 1994), can be found online at <http://citeseerx.ist.psu.edu/viewdoc/summary>*, 1994.

- [27] Tim Roughgarden. On the severity of braess’s paradox: designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72(5):922–953, 2006.
- [28] NG Bean, FP Kelly, and PG Taylor. Braess’s paradox in a loss network. *Journal of Applied Probability*, pages 155–159, 1997.
- [29] Gal Chechik, Isaac Meilijson, and Eytan Ruppin. Synaptic pruning in development: a computational account. *Neural computation*, 10(7):1759–1777, 1998.
- [30] Eitan Altman, Rachid El Azouzi, and Odile Pourtallier. Avoiding paradoxes in routing games. *Teletraffic Science and Engineering*, 4:643–654, 2001.
- [31] Jeff Alstott, Sinisa Pajevic, Ed Bullmore, and Dietmar Plenz. Opening bottlenecks on weighted networks by local adaptation to cascade failures. *Journal of Complex Networks*, page cvv002, 2015.
- [32] Gina Kolata. The new york times. [online], 2015. Available at [://www.nytimes.com/1990/12/25/health/what-if-they-closed-42d-street-and-nobody-noticed.html](http://www.nytimes.com/1990/12/25/health/what-if-they-closed-42d-street-and-nobody-noticed.html).
- [33] Walter Knödel. *Graphentheoretische Methoden und ihre Anwendungen*, volume 13. Springer-Verlag, 2013.
- [34] Sally Cairns, Stephen Atkins, and Phil Goodwin. Disappearing traffic? the story so far. *Proceedings of the ICE-Municipal Engineer*, 151(1):13–22, 2002.
- [35] Stella C Dafermos and Frederick T Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, 73(2):91–118, 1969.
- [36] Land Transport Authority. Annual vehicle statistics 2014, 2014.
- [37] Abhinav Sunderrajan, Heiko Ayt, Wentong Cai, and Aaron Knoll. Map stream: initializing what-if analyses for real-time symbiotic traffic simulations. In *Simulation Conference (WSC), 2014 Winter*, pages 2896–2906. IEEE, 2014.
- [38] Alan J Horowitz. Delay/volume relations for travel forecasting based upon the 1985 highway capacity manual. *FHWA Report*, 1991.
- [39] Shlomo Bekhor, Moshe E Ben-Akiva, and M Scott Ramming. Evaluation of choice set generation algorithms for route choice models. *Annals of Operations Research*, 144(1):235–247, 2006.
- [40] Daniel Zehe, Alois Knoll, Wentong Cai, and Heiko Ayt. {SEMSim} cloud service: Large-scale urban systems simulation in the cloud. *Simulation Modelling Practice and Theory*, pages –, 2015.