Forecast-augmented Route Guidance in Urban Traffic Networks based on Infrastructure Observations

Matthias Sommer, Sven Tomforde and Jörg Hähner
University of Augsburg, Organic Computing Group, Augsburg, Germany

Keywords: Traffic Guidance, Proactive Vehicle Routing, Time Series Forecasting, Organic Traffic Control.

Abstract: Increasing mobility and raising traffic demands lead to serious congestion problems. Intelligent traffic management systems try to alleviate this problem with optimised signalisation of traffic lights and dynamic route guidance (DRG). One solution for the former aspect is Organic Traffic Control (OTC), offering a self-organised, decentralised traffic control system. Based on OTC, this paper presents two proactive routing protocols, resembling techniques known from the Internet domain, applied to the traffic routing problem: Distance Vector Routing and Link State Routing. These protocols were adapted to utilise forecasts of traffic flows to offer anticipatory and time-dependant DRG for road users. The efficiency of these protocols is demonstrated with simulations of two Manhattan-type road networks under disturbed and undisturbed conditions. The results indicate their benefit in terms of lower travel times and emissions, even under low compliance rates.

1 INTRODUCTION

Traffic congestion is a serious problem affecting all traffic participants. The resulting waste of time and fuel leads to billions of dollars lost annually[1]. Urban road networks come to their capacity limit due to increasing demands. The complexity of these traffic control systems, due to mutual influence of different traffic control strategies, is not longer feasible for a central instance. Thus, decentralised, self-organising and self-optimising approaches, that better utilise the existing infrastructure are needed. Organic Traffic Control (OTC) (Prothmann et al., 2011) represents such an approach. OTC selects the best known phase durations at intersections depending on the current traffic flow while learning the impact of its decision to improve the signalisation behaviour over time. OTC extends the existing traffic light controllers (TLC) at intersections via the Observer/Controller architecture (Tomforde et al., 2010). By communicating the local delays (i.e. occurring at the underlying intersection) and estimated travel times to nearby intersections, TLCs have the ability to determine the shortest paths to prominent destinations (such as the main hall or the main station) based on the current traffic flows within the network. The benefits of dynamic route guidance (DRG) systems are the alleviation of congestion, the enhancement of the performance of the road network, and the provision of navigational assistance for travellers which are unfamiliar with the network (Dong, 2011).

Routing protocols compute the fastest or shortest route from a starting point to a destination. This calculation is typically based on static information or recently monitored traffic data. The computed routes may then be visualised via Variable Message Signs (VMS) or on a navigational system and give drivers an indication how to traverse the network to reach their destination as fast as possible. This approach yields significant problems: First, the drivers need time to follow their proposed route. The traffic situation might have changed and so, the routes might be outdated. This leads to repeated re-routings, decreasing the acceptance of the route guidance mechanism. Second, especially in situations with developing congestion, a fast reaction is valuable to avoid further negative impacts on the traffic. Forecasts of the future traffic flow patterns help to detect capacity shortages in advance.

This paper presents two novel, anticipatory and time-dependent route guidance protocols for usage in the infrastructure: Temporal Link State Routing (TLSR) and Temporal Distance Vector Routing (TDVR). These protocols are used to distribute

[1]Urban mobility report, http://mobility.tamu.edu/ums/ (last access: 2016-02-02)

Sommer, M., Tomforde, S. and Hähner, J.
Forecast-augmented Route Guidance in Urban Traffic Networks based on Infrastructure Observations.
In *Proceedings of the International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS 2016)*, pages 177-186
ISBN: 978-989-758-185-4

knowledge about forecasts of the traffic flow for several future time steps in combination with current route states. Thereby, both concepts extend Distance Vector Routing (DVR) and Link State Routing (LSR) that were adapted for urban road networks (Prothmann et al., 2012) which just consider the current traffic conditions. Furthermore, the learning forecast module within the OTC system is presented which allows for proactive adaptation to upcoming changes in traffic demand. The major goals are the improvement of the network's robustness and the minimisation of the average travel time and emissions by preventing congestion. A better distribution of the traffic streams allows to use the capacity of the road network more efficiently. The performance of the novel protocols is investigated in a simulation-based evaluation with two artificial networks under free flow and disturbed conditions. The considered disturbances are due to high traffic demands and incidents modelled as abrupt blockages. The evaluation results indicate that proactive route guidance lowers the overall travel times and delays for road users. The algorithms are evaluated with several compliance rates (i.e. the degree to which drivers follow the route recommendations), showing the benefit of our approach even under low compliance rates.

The remainder of this paper is structured as follows: Section 2 presents a short overview of the state of the art in DRG in urban road networks. Section 3 briefly introduces the OTC system and its main components. The routing component and the new routing protocols are discussed in Section 4. Section 5 presents the evaluation results. The paper concludes with an outlook to future work in Section 6.

2 STATE OF THE ART

Research in the field of vehicular route guidance can be categorised in a variety of ways (Schmitt and Jula, 2006). The routes can be computed in a static or in a dynamic fashion. First, static methods compute a fixed route before the start of the trip. Today, many cars come equipped with on-board GPS-based navigation systems (such as Garmin or TomTom) (Kaplan, 2005). The navigation relies on installed maps of the network which are mostly used by modified versions of the Dijkstra (Dijkstra, 1959) or the A* algorithm (Hart et al., 1968). They are applied to find the shortest route for a given starting point and a fixed destination (Prothmann et al., 2012; Nannicini et al., 2012). Dong (Dong, 2011) presents a brief summary of the published literature on in-vehicle route guidance systems until 2010. Second, reactive routing protocols also compute a route upfront. This can be subject to change during the trip, depending on the real-time informations received. The protocol reacts to changing traffic conditions and computes a new route based on these information. Third, predictive routing protocols go one step further. Apart from reacting to the current situation, but furthermore, they generate forecasts of the upcoming traffic streams and incorporate those forecasts into the computation of the proposed route. It was shown that reactive protocols and especially predictive approaches lower the average travel time in urban road networks (Fu, 2001), whereas reactive protocols are less complex than predictive systems, at the cost of decreased robustness against incidents and congestion (Schmitt and Jula, 2006).

Furthermore, DRG systems can be divided into centralised an decentralised approaches. A study of decentralised strategies for route guidance (F.S. Zuurbier and van Zuylen, 2006), comparing centralised an decentralised systems, points out that decentralised approaches have lower computational complexity, are easily scaled and extended, while being more robust against failures and measurement errors than centralised systems, but might only come up with a suboptimal system-wide solution. In a centralised system, all data is gathered in a traffic management centre where route proposals are computed for the entire network, taking into account system-wide objectives.

SACaNT-CNV (Simulation of Anticipatory Network Vehicle Traffic - Convergence) is a decentralised, proactive route guidance system (Wunderlich et al., 2000). Equipped vehicles are routed on a next-hop basis, each vehicle having a compliance rate of 100% (i.e. strictly following the proposed route). Non-equipped vehicles are assumed to statically follow the shortest free-flow routes. To determine the route proposals, time-dependent link travel times are used in a traffic simulation. It was shown that this approach leads to faster travel times for equipped vehicles.

(Dong et al., 2006) propose a user-equilibrium time-dependent traffic assignment algorithm, using the simulation assignment model DYNASMART. They estimate the current traffic conditions and derive short-term road utilisations to forecast travel times. Their results indicate that this route guidance strategy may reduce the total trip time.

Other decentralised approaches rely on car-2-car communication based on floating car data. A modified version of the Internet routing protocol *BeeHive* (Horst F. Wedde et al., 2007) resulting in the *BeeJamA* protocol routes traffic participants on a next-hop basis from intersection to intersection. Navigation servers store and manage regional routing tables with routes to other areas. These entries are updated based on data sent by vehicles that are assumed to repeatedly transmit their position, speed and destination. OTC does not demand additional hardware in vehicles and eliminates the single point of failure of a central server with a decentralised approach.

A delegate multi-agent system based on ant behaviour was presented in (Claes et al., 2011). It represents a decentralised approach for anticipatory vehicle routing. Every vehicle has to be equipped with a smart device, depicting the vehicle agent that gathers data from the vehicle, such as its state and location, and transmits these to infrastructure agents along the road and to nearby vehicle agents. The vehicle agents use this information to forecast road occupancies to determine the best route to their destination.

In contrast to some of the previous approaches, the DRG mechanism proposed in this paper does not rely on equipped vehicles. The flow and the signalisation data are available at the responsible TLC (e.g. via loop detectors or video cameras), representing a decentralised and proactive approach.

3 ORGANIC TRAFFIC CONTROL

Current traffic management systems usually rely on fixed-time signal plans. Thus, they are not able to adapt to the highly dynamic traffic patterns and to react to unforeseen situations, leading to longer travel times and higher emissions. OTC (Prothmann et al., 2011) extends parametrisable fixed-time controllers, offering a self-adapting, self-optimising system that transfers the traditional design-time decisions to runtime. OTC consists of four basic components: a) adaptive control of traffic lights, b) traffic-dependent establishment of Progressive Signal Systems, c) dynamic route guidance and d) forecasting of traffic situations.

3.1 Adaptive Control of Traffic Lights

OTC handles the adaptation of green times at traffic lights at intersections according to the present traffic conditions. The self-learning, self-optimising system follows a safety-oriented concept that allows OTC to adapt within certain controlled boundaries. Each individual instance of OTC is fully decentralised and controls one intersection only. Fig. 1 depicts the Observer/Controller architecture applied to traffic control. The *System under Observation and Control* (SuOC) situated at Layer 0 is a parametrisable fixed-time TLC. It offers interfaces for monitoring of detector data and adaptation of signal plans. The Observer

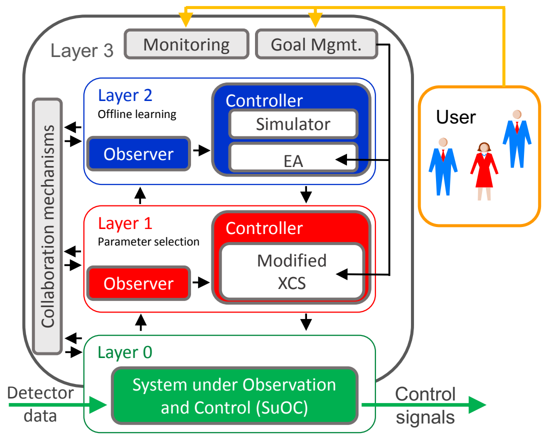

Figure 1: Architecture of an OTC-controlled TLC.

at Layer 1 retrieves raw data from the SuOC which is processed in the following (e.g. filter noise or generate forecasts). This component provides a situation description of the traffic flow of the intersection for the corresponding Controller on Layer 1. This Controller is represented by a learning classifier system with a database of rules (signal plans matched to situations). These are selected based on the current situation and actuated on Layer 0. Before a new signal plan is added to the rule base, it is simulated and evaluated based on an optimisation heuristic on Layer 2. The simulator is configured with the topology of the intersection and the current traffic situation. It evaluates several signalisation plans based on an evolutionary algorithm. The signal plan offering the lowest average delay is returned to Layer 1. As simulations tend to be time-consuming, Layer 2 acts in parallel to Layer 1. At last, Layer 3 provides an interface for monitoring and goal management. A more detailed description of the process is given in (Prothmann et al., 2012).

3.2 Progressive Signal Systems

TLCs that are located at nearby intersections which are directly connected via streets may communicate with each other to form a distributed coordination of intersection controllers. Through this collaboration, Decentralised Progressive Signal Systems (DPSS, also called "green waves") may be established. The mechanism is a three-step-process: 1) identification of possible partners, 2) negotiation of parameters and timing restrictions, and 3) establishment of the DPSS for the most prominent streams within the network (streams with the highest traffic flows) leading to lower travel times while increasing the throughput.

3.3 Dynamic Route Guidance

To turn OTC in an even more robust traffic control system, a self-organised route guidance mechanism has been integrated that computes the fastest routes through the network to prominent places based on the current traffic conditions. Techniques from the Internet domain, such as the Distance Vector Routing (DVR) and the Link State Routing (LSR) protocol (Tanenbaum, 2002) were adapted to road traffic guidance (Prothmann et al., 2012). Since these protocols work well for a complex network with a huge number of nodes, such as the Internet, we see it as an appropriate approach for urban road networks. TLCs use the existing communication infrastructure to send their locally monitored traffic situations to nearby TLCs. This situation description contains the turning delays (approximated with a formula from (Webster, 1958), Equation 1) and the estimated travel times for outgoing sections (based on a formula from the U.S. Dept. of Transportation (USDOT)[2], Equation 2). Assuming that M corresponds to the turning's current traffic flow in vehicles per hour, S denoting the saturation flow (the maximal flow assuming permanent green), t_c representing the cycle time of the intersection and t_g denoting the turning's effective green time, the turning's delay is calculated as:

$$t_d = 0.9 * [\frac{t_C * (1 - t_g/t_C)^2}{2 * (1 - M/S)} + \frac{1800 * g^2}{M * (1 - g)}] \quad (1)$$

Finally, $g = \frac{M}{t_g/t_c * S}$ corresponds to the degree of saturation of the turning for the current green time and traffic flow. The estimated travel time for a section is computed in dependence of the monitored traffic flow M as:

$$t_d = t_F * (1 + (M/C)^2) \quad (2)$$

where $t_F = \frac{s}{v}$ denotes the travel time during free flow conditions based on the length s and the speed limit v of the section. C is the estimated maximal capacity of the section calculated according to formulas given by the USDOT. The result is an up-to-date description of the networks traffic situation, from which routes to arbitrary destinations can be derived. Each TLC locally determines the routes with the lowest travel time which are then visualised through VMSs at each intersection. Momentarily, only the route with the lowest travel time is displayed, but it can be easily extended to output alternative route recommendations. This approach showed to be especially profitable during disturbed conditions (Prothmann et al., 2012), lowering the network-wide travel times and the number of stops.

OTC provides new route recommendations at each intersection on a next-hop basis. It is assumed that not all road users follow these suggestions as each individual driver optimises his route without respect to the network-wide optimum. Thus, decentralised route guidance may result in an user-optimised equilibrium but not an optimum for the entire system. Previous research reports a widely varying acceptance of VMS-based route recommendations, ranging between 20% (Erke et al., 2007) to 70% (Emmerink et al., 1996).

Momentarily, only the current traffic flows are considered for the route proposals. So, drivers can be confronted with several route changes during high-load and quickly changing traffic conditions. This is due to the continuous change of traffic conditions while the vehicles traverse the network. Therefore, the initial route recommendation might already be outdated at the next intersection, possibly leading to a reduction of the system's acceptance. Our proactive protocols take forecasts of traffic flows into account. They consider current and upcoming traffic flows, turning the existing reactive DRG system into a more robust, proactive and anticipatory one. It is assumed that this approach reduces the re-routing demands, leading to a broader acceptance of the system, and making it more reliable. Forecasts of the future traffic flow enable the detection of capacity shortages in advance.

3.4 Forecast Module

Recent work (Sommer et al., 2015) focused on the development of a self-optimising forecast module for time series. This module is situated in the Observer at Layer 1. It offers a dynamic weighting of forecasts from several forecasting techniques based on historic knowledge. Furthermore, it classifies the time series based on their characteristics, such as trend, seasonality or non-stationarity. If necessary, it automatically processes time series to normalise them or to make them stationary. Several forecast methods independently compute forecasts based on their individual model and data. These forecasts are then combined by a combination strategy. The applied strategies range from a simple average to sophisticated machine learning algorithms. Finally, the combined forecast is returned. The module itself learns the best configuration during runtime. Therefore, no system expert is needed to determine the best combination of the applied forecasting techniques during design time. Only the active forecast methods and the combination strategy have to be specified. In this context, the forecast component is used to estimate the future traffic flows on sections and turnings.

[2]http://www.fhwa.dot.gov/ohim/hpmsmanl/appn7.cfm

4 DYNAMIC ROUTE GUIDANCE IN URBAN ROAD NETWORKS

Communication links between neighbouring intersections allow TLCs to exchange estimated travel times and delays to calculate alternative routes to prominent destinations under the current traffic demand and signalisation. A route is defined by an origin, a destination, and the connecting roads in-between. Each TLC is extended by a routing component (RC) that allows for a self-organised, fully decentralised route guidance. With the help of a routing protocol, travel times are distributed in the network and routing tables containing the proposed routes are managed. Each intersection manages its own routing tables. Its entries are of the form: ”to destination X, turn right, estimated arrival time: y seconds”. This information is then displayed by VMS. In the following, these protocols and their modifications towards proactive route guidance based on traffic flow forecasts are explained.

4.1 Requirements for a Real-world Deployment

To deploy OTC in a real-world scenario, some preconditions have to be met. First, sensors are needed that monitor the current traffic conditions. Inductive loop detectors resemble a cheap and well established monitoring technique in traffic management applications (Parkany and Xie, 2005). Second, the OTC logic works as an extension of the existing TLCs at signalised intersection. We assume that these controllers are parametrisable. For a fast and reliable computation, more computational power might be needed. With the help of the communication infrastructure, travel costs can be distributed in the network. A shortest path algorithm has to utilise this data to determine the routes with the lowest travel time. Furthermore, routing protocols are needed to derive and distribute the route recommendations. At last, we need devices providing the route recommendations for the drivers. This can be realised with VMS (collective systems), or in case car-to-infrastructure communication is available, via direct communication to smartphones or the cars’ navigational system (individual systems). Besides these characteristics, OTC does not need further changes – especially no sophisticated detection and analysis devices.

4.2 Link State Routing (LSR)

A modified version of the Internet protocol LSR serves as route guidance heuristic. It broadcasts estimated travel times for each section and each turning movement whereby the best routes are derived. The protocol works as a three step process: 1) Each RC estimates the local delays for each turning of the intersection and communicates those to all other RCs in the network using broadcast messages (so-called *advertisements*). These delays are calculated based on the current waiting times during red light phases and the estimated travel time to a next intersection based on the current traffic flow. These advertisements contain *link states* describing a path from a starting intersection to a destination and its estimated travel time. 2) After receiving all advertisements from the other RCs, each RC builds a graph by connecting the subgraphs obtained from the link states. It represents the topology, the current traffic flows and the approximated travel times within the network. 3) Finally, every RC locally computes the best routes. The Dijkstra algorithm is used to calculate the paths with the lowest travel time from each TLC to all reachable destinations based on the previously generated network graph. In a final step, the interior routing tables for the approaching roads are updated with the best route to all reachable destinations. Each table entry now contains entries for each incoming section, the destination, the recommended next turning and the estimated travel time to this destination. Further details on the existing LSR mechanism are given in (Prothmann et al., 2012).

4.3 Temporal Link State Routing (TLSR)

The novel TLSR protocol resembles an extension of the LSR protocol, utilising current traffic demands and forecasts of future traffic flows. By broadcasting graph-series that encode the current and the forecasted traffic flows, TLSR is able to consider the time-dependant changes of traffic.

Time-dependent Representation. 1) The local delays are estimated for the current traffic flow and for a number of traffic flow forecasts for future points in time. Based on previous forecasts and the actual values, the forecast accuracy and its standard deviation can be determined. These forecasts and forecast errors are added to the advertisements and also broadcasted to other RCs. 2) The edges (representing roads) of the network graph represent the current traffic flow and the forecasts for different points of time in the future, converting the network graph into a *time-dependant representation*. 3) The benefit of the incorporation of the received forecasts is highly dependant of their accuracy and of the degree to which they are taken into account for the calculation of route recommendations. Thus, the Dijkstra algorithm computes the fastest routes with respect to the forecast accuracies.

Reliability-considering Dijkstra. The derivation of qualitative route recommendations depends heavily on the degree to which forecasts are taken into account. This ranges from only relying on current traffic conditions (which is similar to the previous protocol) to considering long periods of predicted traffic demands. To estimate the reliability of the forecast F_t, the sending RC calculates the Mean Absolute Scaled Error (MASE) (Hyndman and Koehler, 2006) where the scaled error q_t is defined as

$$q_t = \frac{Y_t - F_t}{\frac{1}{n-1}\sum_{i=2}^{n} |Y_i - Y_{i-1}|} \quad (3)$$

with Y_t is the current traffic flow and Y_{t-1} is the traffic flow of the previous measurement. The absolute error is scaled based on the in-sample mean absolute error from a benchmark forecast method. The MASE is then calculated as

$$MASE = mean(|q_t|) \quad (4)$$

with $mean(|q_t|)$ denoting the sample mean of q_t over a certain period. A scaled error of less then one arises if the forecast is better than the average naive one-step forecast computed in-sample. This error measure is sent to other controllers which determine the degree to which they consider the current flow, respectively the forecasts. The lower the MASE, the higher the trust in its accuracy. Likewise, the influence of the forecast on the estimated travel time calculation raises. Finally, the current flow Y_t and the forecast F_{t+1} are combined based on a smoothing function computed as

$$x_t^* = \alpha Y_t + (1-\alpha)F_{t+1} \quad \text{where} \quad 0 \leq \alpha \quad (5)$$

where α is the MASE. An optimal MASE of 0 results in only considering the forecast and a MASE of 1 or higher in only considering the currently monitored flow. We limit the maximum MASE to 1. The result x_t^* then serves as the new estimated delay for the according turning or section. Previously, Dijkstra considered only one value per edge of the network. Now, the edges contain costs for the current time step and several forecasts for different points in time. Our modified Dijkstra chooses the closest entry for the point in time where a value is needed. This means, at the first intersection (start of the route), the current costs Y_t will be considered. At the next intersection, the forecast for the estimated arrival time $t + \Delta t$ will be used, and so on (Fig. 2). Finally, an approximated travel time for the whole route is computed.

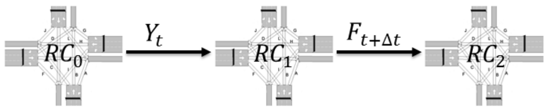

Figure 2: Time dependant use of actual values and forecasts.

4.4 Distance Vector Routing (DVR)

As an alternative to the standard LSR protocol, DVR has been considered. The DVR protocol maintains routing tables for each of the intersection's approaches, updates them based on messages received from neighbouring RCs and communicates changes to its own neighbours. The protocol works as follows: 1) Initially, each intersection checks if it is directly connected to prominent destinations (e.g. the main hall). In case such a prominent destination is detected, the RC creates new routing table entries for all approaching roads leading to that destination. Each entry contains the destination, the approaching road, the proposed turning and the estimated travel time (based on the delay caused by red lights plus the estimated travel time to the destination). The travel time can be estimated in a static way, derived from the length and the speed limit of the connecting road, or dynamically, based on the current traffic flow. 2) New or updated entries are then sent upstream to neighbouring RCs where matching routing table entries are updated iteratively. At last, each recipient has the estimated travel time from itself to prominent destinations. If a destination is yet unknown, a new entry is created, otherwise the existing entry is updated (the costs and the proposed next turning are updated if the costs for the received route are lower than the previous ones). 3) Finally, each RC knows the estimated fastest route to each destination that is reachable from itself.

4.5 Temporal Distance Vector Routing (TDVR)

TDVR tries to cover the time-dependant traffic conditions for future time steps considering traffic flow forecasts. DVR processes the network upstream, starting at a destination. This process is not applicable for TDVR. Each RC has to know the travel time from an initial RC_0 to itself. This is necessary to determine the point in time for which a traffic flow forecast is computed (Fig. 2). To determine the arrival time of a vehicle at RC_1, the estimated travel time (see Section 3) from RC_0 to RC_1 has to be calculated. RC_1 receives the request and estimates the turning's delays and travel times for outgoing sections to neighbouring destinations for the point in time the vehicle is estimated to arrive at RC_1. Similar to DVR, RC_1 forwards the updated request to further RCs in its proximity. Furthermore, it returns the discovered route (from RC_0 to the found destination) and its estimated travel time. RC_0 receives this information and updates its routing tables. Its entries are therefore based on the current traffic flow conditions and the respective traffic flow forecasts.

4.6 Adaptation for Regional Routing

Broadcasting local traffic informations to other RCs leads to high communicational overhead which raises quadratic with the number of RCs. In the following, m denotes the number of prominent destinations and n represents the number of intersections in the network. In the worst case, DVR has to send n messages per destination, resulting in a communication complexity of $O(n*m)$. LSR broadcasts the link states of an intersection with a single message. As each controller has to forward the data of each other RC at most once, the communication complexity is $O(n^2)$.

To minimise this overhead, the protocols were extended based on the concept of the Border-Gateway-Protocol (Tanenbaum, 2002), also known from the Internet domain. The concept of the distinction between intra- and inter-network routing was transferred to regions of cities. RCs near to each other form a region. RCs can be sorted into regions based on their relative distance to each other, or manually via a configuration file. Each RC can only belong to exactly one region. If all its neighbours are in the same region, it is an *interior*, otherwise an *exterior* node. Only exterior nodes are allowed to communicate with exterior nodes from other regions. Consequently, RCs have to propagate less messages and the complexity for the calculation of the shortest paths decreases. Fig. 3(b) shows three regions: A, B and C with each 9 nodes and 10 centroids (a centroid models a source/sink of traffic). Dark dots are exterior, light dots represent interior nodes. Lines between two nodes show the possibility to communicate with each other (i.e. the corresponding road segments are normally both-way).

5 EVALUATION

In the following, the evaluation scenarios and their results are presented. The evaluation was done with AIMSUN 8.0.9 (Barcelo and Casas, 2002), a professional traffic modelling and simulation software widely used by traffic experts. The simulations were executed on an Intel Core i7 quad-core CPU with 2.6 GHz and 8 GB RAM.

5.1 Experimental Setup

The DRG system has been evaluated by comparing OTC-controlled intersections with and without DRG against a reference run with fixed-time signalisation. The simulation was done with a Manhattan-style network and a regional variant (Fig. 3), simulating moderate traffic demands.

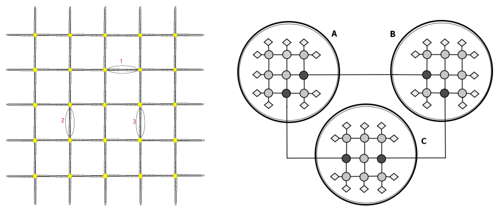

(a) 25 junctions and 20 destinations. (b) 3 connected regions.

Figure 3: Simulated networks (Ellipses show incident locations).

The first network consists of a 5-on-5 grid of 25 intersections and 20 prominent destinations at the border of the network (Fig. 3(a)). Every hour, 15 vehicles travel from every origin to every destination, resulting in 5700 vehicles per hour traversing the network. To simulate disturbed traffic conditions, we also investigated road blockages. The congested scenario simulates disturbed traffic flows through temporary blockages of roads, resulting in traffic congestion. The locations of these incidents are marked with circles. Three streets were each blocked for 40 minutes, forcing vehicles to take alternative routes. Incident 1 starts at simulation minute 15, incident 2 at 45 minutes and incident 3 at 75 minutes. Ten vehicles per hour travel from every origin to every destination. In total, 3800 vehicles per hour traverse the network. The simulation duration was 2 hours and 15 minutes.

The second network (Fig. 3(b)) allows for an evaluation of the regional routing protocols. Three equally shaped 3-on-3 manhattan-style regions with 9 intersections are each connected by one or two streets. The simulation time spanned 2 hours. For every origin-destination pair, 8 vehicles per hour are generated, resulting in 6048 cars per hour traversing the network.

The routing protocols are executed every 2 minutes. The routing compliance is set to 10% (low), 40% (medium) and 70% (high). Cars not following the DRG proposals, use the static shortest path to their destination. The traffic flow forecasts were created by the following methods: Exponential Smoothing, Double Exponential Smoothing, Double Smoothing Average, Moving Average and Kalman Filter. Their forecasts were combined with the simple average. The protocols have been evaluated with respect to the vehicles' mean delay and the travel time averaged over all trips. The fuel consumption and pollution emissions of the simulated vehicles have been investigated to estimate the environmental impact of DRG. The emissions have been determined with the help of AIMSUN's environmental model, configured according to (Panis et al., 2006).

5.2 Evaluation Results

The following section presents the results of the evaluation for the regular and the congested scenario in both networks executed as described before. As the data is gathered when the simulated vehicles have completed their trip, the effects of the incidents show up in the figure with an approximate delay of 10 minutes after the incidents occurrence.

5.2.1 Manhattan Scenario

Fig. 4 indicates that OTC with and without routing is able to drastically reduce the overall delay compared to fixed-time signalisation. The figure shows the average travel time in seconds per kilometre for a simulated scenario of 2 hours. The compliance rate was set to 70%. The mean delay is given in brackets behind the protocol's name.

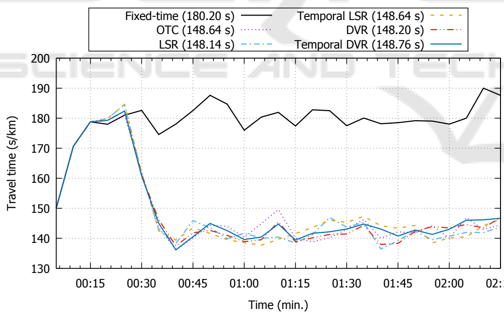

Figure 4: Travel times for the regular Manhattan scenario.

The first 15 minutes of the simulation represent the warm-up time, where OTC gathers data to calibrate the forecast methods. Each observer/controller needs to populate its initial empty database of mappings between optimised signal plans and monitored traffic demands. Therefore, the overall performance is identical. Afterwards, it can be seen that the static fixed-time signal plans are not able to reduce the negative impacts of the traffic demands. Table 1 presents the average travel times over all trips, the average fuel consumption and the average CO_2 emissions per vehicle for the reference run and the proactive routing protocols with different compliance rates. Reductions compared to the reference run are given in brackets. Not only is the travel time reduced significantly (17%) in comparison to the reference run, but so are the fuel consumption (4% to 6%) and the pollution emissions (2% to 3.3%). These results must be interpreted with caution. During undisturbed conditions, an improved signalisation alone is enough to guarantee a reduction of queues. OTC without routing already reduces the travel time to 227.6 seconds.

In contrast to the undisturbed scenario, the congested scenario clearly shows the benefit of DRG during incidents. Table 2 summarises the evaluation results. The routing mechanism gives drivers indications how to avoid congested areas. It reduces the average travel time by 10% (TLSR) to 11% (TDVR) (Fig. 5) for a compliance rate of 70%.

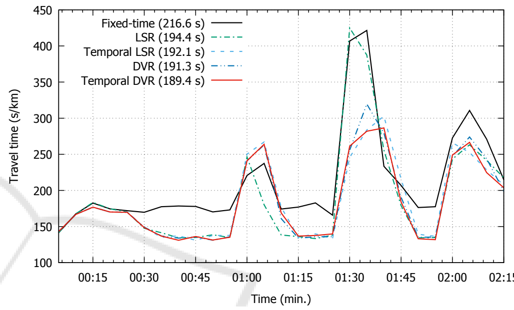

Figure 5: Travel times for the congested scenario.

The best performance was delivered by TDVR achieving the highest travel time reduction. At 1:45, during a severe incident, TDVR reduces the average travel time to 282 seconds per kilometre (33% improvement over the reference run with 421 seconds and 30% improvement over OTC without routing with 408 seconds). This indicates that TDVR correctly forecasted the upcoming congestion due to the incident, preventing more severe disturbances. This resembles an improvement of 4.0% compared to OTC without routing with an average trip travel time of 320 seconds. The decrease in travel time is achieved by re-routing drivers over alternative routes, which can be longer than the planned one. Therefore, the use of routing protocols sometimes leads to slightly higher fuel consumption and CO_2 emissions.

Fig. 6 and Fig. 7 present the average travel times evaluated for several compliance rates. A higher compliance rate means that drivers are more likely to follow the given routing proposals. The figure's horizontal axis shows the simulation time and the vertical axis shows the average travel time in seconds per kilometre. Our results suggest that the benefit of all routing protocols increases for higher compliance rates.

Table 1: Results for the regular scenario with TDVR and TLSR.

	Fixed-time	TDVR			TLSR		
		0.7	0.4	0.1	0.7	0.4	0.1
Travel time [s/veh]	275	228 (17.1%)	228 (17.1%)	228 (17.1%)	227 (17.5%)	230 (16.4%)	227 (17.5%)
Fuel [l/100km]	15.1	14.5 (4.0%)	14.4 (4.6%)	14.2 (5.8%)	14.5 (4.0%)	14.6 (3.4%)	14.2 (5.8%)
CO_2 [g/veh]	518.5	506.0 (2.4%)	504.7 (2.6%)	501.4 (3.3%)	504.7 (2.7%)	508.3 (2.0%)	501.6 (3.3%)

Table 2: Results for the congested scenario with TDVR and TLSR.

	Fixed-time	TDVR			TLSR		
		0.7	0.4	0.1	0.7	0.4	0.1
Travel time [s/veh]	345	306 (11.3%)	310 (10.1%)	320 (7.2%)	310 (10.1%)	313 (9.2%)	335 (2.9%)
Fuel [l/100km]	16.6	16.4 (1.2%)	16.3 (1.8%)	16.6 (0.0%)	16.6 (0.0%)	16.4 (1.2%)	17.0 (-2.4%)
CO_2 [g/veh]	549.6	543.6 (1.1%)	545.7 (0.7%)	551.8 (-0.4%)	547.2 (0.4%)	547.0 (0.5%)	562.3 (-2.3%)

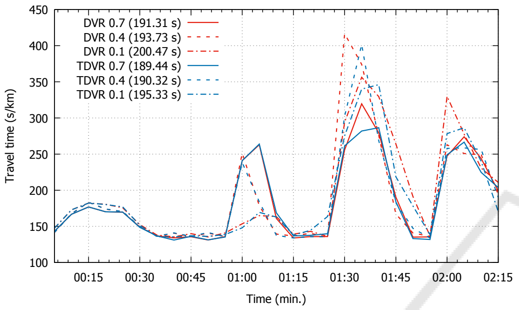

Figure 6: Compliance rates for the congested Manhattan scenario with DVR.

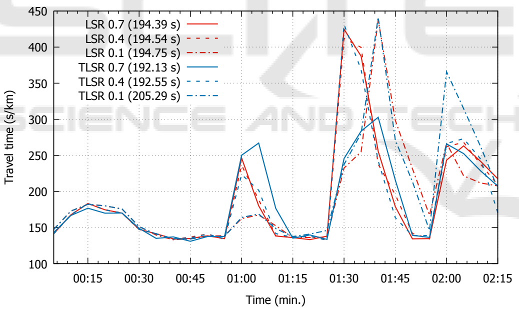

Figure 7: Compliance rates for the congested Manhattan scenario with LSR.

5.2.2 Regional Scenario

Table 3 depicts the comparison of communicational and computational effort between regional protocols and the basic variants. The table shows the number of messages each TLC has to send during one iteration of the executed routing protocol and the average runtime in seconds of a complete protocol run. The results clearly depict that the regional protocols decrease the number of messages and the computational overhead. The reference run has no communication between TLCs and therefore no sent messages. The anticipatory protocols need more computational power to compute the forecasts for all sections and turnings of the network, leading to longer runtime.

As Table 4 indicates, the regional protocols do not or only to a slight extend increase the average travel time. The simplification of the communication due to the regional aggregation of RCs offers equally good route recommendations while reducing the communication and computational effort.

6 CONCLUSION

This paper presented two novel time-aware, anticipatory routing protocols (TLSR and TDVR) for dynamic, proactive traffic guidance in urban road networks. The well-known Internet protocols Distance Vector Routing and Link State Routing have been extended, utilising traffic flow forecasts to compute the best routes through a road network. The routes are determined by a self-organised approach, extending parametrisable traffic light controllers. The route recommendations are visualised by Variable Message Signs at each intersection, guiding drivers from intersection to intersection on a next-hop basis. The protocols were evaluated with compliance rates of 10%, 40% and 70%. A simulation study investigated the benefits of these protocols under disturbed and undisturbed conditions in two different networks. Our findings strongly support the view that the consideration of traffic flow forecasts leads to a decrease in system-wide travel times for urban vehicular traffic. Consequently, this leads to a reduction in pollution emissions and fuel consumption. In general, this counts especially for disturbed and congested conditions, but to a limited extend also for medium and low traffic saturations. The benefit increases for higher compliance rates. During undisturbed conditions, an improved signalisation alone is enough to guarantee a reduction of queues. The dynamic route guidance improves the network's robustness by guiding road users on alternative routes around blocked areas. The Temporal Distance Vector Routing protocol showed to be the most beneficial approach, not only for congested

Table 3: Comparison of regional and basic protocols in terms of sent messages and computational time.

	Fixed-time	DVR	Reg. DVR	TDVR	Reg. TDVR	LSR	Reg. LSR	TLSR	Reg. TLSR
Messages (#/RC/iter.)	0	1437	133	1463	133	35	25	35	24
Runtime (s)	10	360	105	1400	400	280	270	330	370

Table 4: Travel time of regional and basic protocols for the regional scenario.

	DVR	Reg. DVR	TDVR	Reg. TDVR	LSR	Reg. LSR	TLSR	Reg. TLSR
Undisturbed: Travel time [s/veh]	119.83	119.27	119.49	119.30	119.73	119.76	119.39	121.68
Disturbed: Travel time [s/veh]	127.83	128.44	128.41	127.15	128.69	129.14	129.00	130.18

but also during free flow conditions. The communicational overhead and the computational costs can be reduced by partitioning a larger network into smaller sub-networks via the Border-Gateway-Protocol.

REFERENCES

Barcelo, J. and Casas, J. (2002). Dynamic network simulation with AIMSUN. In *Proc. of the Int. Symp. on Transp. Simul.*, pages 1–25. Kluwer.

Claes, R., Holvoet, T., and Weyns, D. (2011). A decentralized approach for anticipatory vehicle routing using delegate multiagent systems. *IEEE Trans. on Int. Transp. Systems*, 12:364–373.

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271.

Dong, J., Mahmassani, H., and Lu, C.-C. (2006). How reliable is this route: Predictive travel time and reliability for anticipatory traveler information systems. *Journal of the Transp. Research Board*, (1980):117–125.

Dong, W. (2011). An overview of in-vehicle route guidance system. *34th Australasian Transp. Research Forum (ATRF 2011), Adelaide, Australia, 28-30 Sept. 2011*.

Emmerink, R. H. M., Nijkamp, P., and Rietveld, P. (1996). Variable message signs and radio traffic information, an integrated empirical analysis of drivers route choice behaviour. In *Transp. Research Part A: Policy and Practice*, volume 30, pages 135–153.

Erke, A., Sagberg, F., and Hagman, R. (2007). Effects of route guidance variable message signs (vms) on driver behaviour. *Transp. Research Part F: Traffic Psychology and Behaviour*, 10(6):447 – 457.

F.S. Zuurbier, J. v. L. and van Zuylen, H. (2006). Comparison study of decentralized feedback strategies for route guidance purposes. In *Proc. of the 9th TRAIL Congress*, pages 1–15.

Fu, L. (2001). An adaptive routing algorithm for in-vehicle route guidance systems with real-time information. *Transp. Research Part B*, 35(8):749–765.

Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. on Systems Science and Cybernetics*, 4(2):100–107.

Horst F. Wedde et al. (2007). Highly Dynamic and Adaptive Traffic Congestion Avoidance in Real-Time Inspired by Honey Bee Behavior. In *Mobilität und Echtzeit – Fachtagung der GI-Fachgruppe Echtzeitsysteme*, pages 21–31. Springer.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, pages 679–688.

Kaplan, E. (2005). *Understanding GPS - Principles and applications*. Artech House, 2nd edition edition.

Nannicini, G., Delling, D., Schultes, D., and Liberti, L. (2012). Bidirectional a* search on time-dependent road networks. *Netw.*, 59(2):240–251.

Panis, L. I., Broekx, S., and Liu, R. (2006). Modelling instantaneous traffic emission and the influence of traffic speed limits. *Science of The Total Environment*, 371(13):270 – 285.

Parkany, E. and Xie, P. C. (2005). A complete review of incident detection algorithms & their deployment: What works and what doesn't. Technical report, Fall River, MA: New England Transp. Consortium.

Prothmann, H., Tomforde, S., Branke, J., Hähner, J., Müller-Schloer, C., and Schmeck, H. (2011). Organic Traffic Control. In *Organic Computing – A Paradigm Shift for Complex Systems*, chapter 5.1, pages 431–446. Birkhäuser.

Prothmann, H., Tomforde, S., Lyda, J., Branke, J., Hähner, J., Müller-Schloer, C., and Schmeck, H. (2012). Self-organised routing for road networks. In *Proc. of the 6. Int. Workshop on Self-Organizing Systems*, LNCS 7166, pages 48–59. Springer.

Schmitt, E. J. and Jula, H. (2006). Vehicle route guidance systems: classification and comparison. *2006 IEEE Intelligent Transp. Systems Conference*, 1:242–247.

Sommer, M., Tomforde, S., Hähner, J., and Auer, D. (2015). Learning a dynamic re-combination strategy of forecast techniques at runtime. In *Proc. of IEEE Int. Conf. on Autonomic Computing (ICAC)*, pages 261–266.

Tanenbaum, A. S. (2002). *Computer Networks*. Pearson, 4th edition.

Tomforde, S., Prothmann, H., Branke, J., Hähner, J., Müller-Schloer, C., and Schmeck, H. (2010). Possibilities and limitations of decentralised traffic control systems. In *IEEE World Congress on Comp. Int.*, pages 3298–3306. IEEE.

Webster, F. V. (1958). *Traffic signal settings*. Road Research Technical Paper No. 39, Road Research Laboratory, England, HMSO.

Wunderlich, K. E., Kaufman, D. E., and Smith, R. L. (2000). Link travel time prediction for decentralized route guidance architectures. *IEEE Transactions on Intelligent Transp. Systems*, 1(1):4–14.