Centralised Urban Traffic Routing Using Mixed-Integer Programming

Andrii Nyporko[1,2][a], Matyáš Švadlenka[1,2][b], Nikolai Antonov[1,2][c], Mohammad Rohaninejad[1][d] and Lukáš Chrpa[1][e]

[1]*Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, Czechia*

[2]*Faculty of Electrical Engineering, Czech Technical University in Prague, Czechia*

{nyporand, svadlmat, antonni1}@fel.cvut.cz, {mohammad.rohani.nezhad, chrpaluk}@cvut.cz

Keywords: Centralised Traffic Routing, Mixed-Integer Programming, Urban Traffic Management.

Abstract: The increase in the urban population over the past decades led to an increase in the number of vehicles in urban road networks, especially in larger metropolitan areas. The problem is exacerbated during rush hours and when an unexpected or rare event occurs (e.g. accidents, concerts). Existing traffic routing methods, including those embedded in modern navigation systems, consider Dynamic User Optimal (DUO) traffic routing that generates routes in a decentralised fashion. Centralized traffic routing, which we consider in this paper, benefits from the global perspective of the situation that can utilise the road network more effectively. We propose a technique leveraging Mixed-Integer Programming (MIP) for distributing vehicles in the road network while minimizing traffic intensity on road segments. Our evaluation shows the potential of the proposed technique for centralized traffic routing.

1 INTRODUCTION

Over the past decades, the urban population has been steadily increasing. That contributed to an increase in traffic intensity in urban areas, especially during rush hours. Traffic congestion is one of the major economic problems as, for example, the cost of congestion in London exceeded £5 billion in 2020[1]. On top of that, heavy traffic in urban areas poses a major health threat (Chang et al., 2019). Occasional events, such as sports matches, rallies, or concerts, also have a major impact on urban traffic that might be more difficult to predict.

The concept of Smart Cities (Kirimtat et al., 2020) involves the need for effective traffic management, focusing on the proper distribution of traffic in road networks to minimize average travel time and distance traveled. In addition to initiatives such as car sharing to reduce the number of vehicles, effective traffic management utilizes road infrastructure through efficient traffic routing and traffic light control, which has been approached from a centralized perspective through scheduling (Xie et al., 2012), evolutionary methods (Pilát, 2018), and automated planning (Pozanco et al., 2021; McCluskey and Vallati, 2017; Antoniou et al., 2019). These strategies are incorporated into a framework that introduces pheromone-based traffic management (Cao et al., 2017). Modern navigation systems, such as those employing the Dynamic User Optimal (DUO) principle (Friesz et al., 1989), generate optimal routes in a decentralized manner by leveraging current traffic data. However, decentralized routing can cause issues, such as unsynchronised routing to network bottlenecks.

Centralized traffic routing aims to provide the optimal route for each vehicle from a global perspective of the controlled region, thereby utilizing the road network more effectively. In more detail, centralised traffic routing has to involve centralised infrastructure to vehicle (I2V, V2I) communication such that vehicles approaching the controlled region must broadcast their entry and exit points, allowing the infrastructure to generate routes, for all the vehicles approaching the region in a given time span, that are then broadcast back to the vehicles. Approaches in centralized traffic routing involve collecting data on vehicles' intended routes, predicting future traffic intensity, and broadcasting this prediction back to vehi-

[a] https://orcid.org/0000-0001-5535-4250
[b] https://orcid.org/0009-0004-9199-3024
[c] https://orcid.org/0009-0002-0156-3561
[d] https://orcid.org/0000-0002-0623-4890
[e] https://orcid.org/0000-0001-9713-7748

[1]https://www.london.gov.uk/press-releases/mayoral/cost-of-congestion-in-capital-revealed

Nyporko, A., Švadlenka, M., Antonov, N., Rohaninejad, M. and Chrpa, L.
Centralised Urban Traffic Routing Using Mixed-Integer Programming.
DOI: 10.5220/0013137100003890

In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025) - Volume 2*, pages 176-187
ISBN: 978-989-758-737-5; ISSN: 2184-433X

cles for route updates. Recent methods focus on using automated planning for centralized dynamic route allocation (Chrpa et al., 2019; Vallati et al., 2021; Svadlenka et al., 2023; Silva and Tang, 2024).

Current automated planning based techniques for centralised traffic routing reason on the individual vehicle level while computing routes (Chrpa et al., 2019; Svadlenka et al., 2023). Even though we need to compute the route for each individual vehicle, it is important to optimise traffic flows first (i.e., determine how many vehicles use which route) as this information is crucial to determine expected traffic intensity on particular road segments. Thus, the decision about traffic flows does not have to be made on the "microsimulation" level (the individual vehicle level) as it is currently done in the planning-based methods.

Hence, in this paper, we focus on the "microsimulation" aspect of planning-based centralized traffic routing such that we introduce a Mixed-Integer Programming (MIP) model (Wolsey, 1998), inspired by the multicomodity network flow problem (Ouorou et al., 2000), that distributes traffic across the road network, aiming to minimize traffic intensity on road segments. In contrast to planning-based techniques, we consider *macrosimulation* that, in other words, distribute traffic at the "flow" level. The main advantage of the proposed approach is that routes do not have to be recomputed multiple times and that many symmetries can be broken (e.g. the order in which vehicles are routed). To extract routes for individual vehicles, we propose an algorithm, based on Depth-First-Search, that searches through the allocated traffic flows (the solution of the MIP model).

Our approach was evaluated on scenarios from New York and Sydney metropolitan areas (Svadlenka et al., 2023) as well as on a central region of Dublin (Gueriau and Dusparic, 2020) by using the well-known SUMO simulator (Lopez et al., 2018). The results show the potential of our approach as it outperforms the decentralised (DUO) as well as the existing planning-based approaches for Centralised Traffic Routing in the New York and Sydney scenarios. We provide a thorough analysis of the results, discuss the limitations of the current approach as well as the lessons we learned, and provide some ideas for improvements that we plan to address in future work.

2 RELATED WORKS

The concept of Smart Cities (Kirimtat et al., 2020) incorporates the need for effective traffic management involving a proper distribution of traffic in road networks, minimising average travel time and average driven distance. Besides initiatives that aim at reducing the number of vehicles (e.g. by car-sharing), effective traffic management has to utilize road infrastructure through effective traffic routing and effective traffic light control. Both traffic optimisation strategies are considered in a framework that introduced pheromone-based traffic management (Cao et al., 2017). Traffic light control has been tackled from a centralised perspective by means of scheduling (Xie et al., 2012), evolutionary approaches (Pilát, 2018), or automated planning (Pozanco et al., 2021; McCluskey and Vallati, 2017; Antoniou et al., 2019). Traffic routing methods that are embedded in modern navigation systems (e.g. WAZE™) usually follow the Dynamic User Optimal (DUO) principle (Friesz et al., 1989) that, in a nutshell, generates (optimal) routes in a decentralised fashion while leveraging current traffic data (Du et al., 2014; Claes et al., 2011). A possible issue of decentralised (DUO) routing might involve unsynchronised routing to "network bottlenecks" that might not (yet) be busy when routing takes place.

The centralised traffic routing approaches, on the other hand, aim to provide the optimal route for each vehicle from the global perspective of the (controlled) region and hence can utilise the controlled road network more effectively. Each vehicle approaching the controlled region has to broadcast its intentions, i.e., where it enters the network and where it plans to leave the network. The infrastructure that collects the information from approaching vehicles has to provide routes for the vehicles across the controlled region. Yamashita et al. (Yamashita et al., 2005) proposed an approach that collects data about vehicles' intended routes, based on the collected data it makes a prediction of future traffic intensity in the area, and the prediction is broadcasted back to vehicles, so the vehicles might update their routes according to the prediction. Recent Centralised Traffic Routing approaches are mostly based on Automated Planning (Chrpa et al., 2019; Vallati et al., 2021; Svadlenka et al., 2023) and, most recently, on centralised dynamic route allocation (Silva and Tang, 2024).

Although planning-based approaches achieved promising results (Chrpa et al., 2019; Svadlenka et al., 2023), they tend to struggle with scalability and might be suitable only for smaller regions (Chrpa and Vallati, 2023). The size of the road network in which centralised traffic routing techniques have to reason in can be reduced by precomputing suitable routes for each traffic flow (Svadlenka et al., 2023; Silva and Tang, 2024). Another aspect contributing to the poor scalability of centralised traffic routing techniques is the need to provide a route for each individual vehicle, i.e., considering *microsimulation*.

There are approaches addressing the vehicle routing problem, which optimizes routes to serve customers, usually by minimising the total travel distance, or mitigating possible delays. This draws a parallel with our work, although the goals are different – optimising traffic flows in traffic routing versus optimising delivery services. One of the most recent approaches (Polimeni and Vitetta, 2024) proposed an approach tackling the problem of integrating road network design and vehicle routing (public and freight transport). Also, we can pinpoint (Wenning et al., 2006) and (Krishnan et al., 2017), where this problem is tackled in decentralised fashion that draws some parallel between decentralised and centralised traffic routing.

3 CENTRALISED TRAFFIC ROUTING

In a nutshell, the problem of *Centralised Traffic Routing* can be understood as finding routes for a set of vehicles in a road network, where each vehicle has its locations of origin and destination while optimising for specified criteria such as minimising average travel time or minimising traffic intensity on the road segments. In this paper, we use the concept of Centralised Traffic Routing described by Chrpa et al. (Chrpa et al., 2019).

3.1 Problem Specification

Formally, the *Centralised Traffic Routing* problem is a tuple $\chi = (\mathcal{N}, V, O, D, C)$, where $\mathcal{N}$ represent a *road network* in form of a labeled directed graph $\mathcal{N} = (N, E, R, \rho)$, where vertices N represent *junctions* and edges E connect the adjacent junctions by *road segments* from R by a mapping $\rho : E \to R$. Note that we admit that more edges can be mapped to a single road segment (because of the representation of road networks in SUMO and simplifications we applied such as "merging" roundabouts into a single junction). Let $V = \{v_1, \ldots, v_k\}$ be a set of *vehicles* that approach the network such that each vehicle has its location (junction) of *origin*, specified by a function $O : V \to N$, and its *destination* location (junction), specified by a function $D : V \to N$. Note that a pair (origin,destination) refers to a *traffic flow*. A function $C : R \times \mathbb{N}_0 \to \mathbb{R}_0^+$ represents the *cost* for using a road segment by a given number of vehicles.

We say that $\Sigma = \{p_1, \ldots, p_k\}$ is a *solution* of χ if and only if for each $i \in \{1, \ldots, k\}$ it is the case that p_i is a sequence of edges from E forming a path in $\mathcal{N}$ starting at $O(v_i)$ and finishing at $D(v_i)$. The *cost* of Σ is determined as $\sum_{x \in R} C(x, |\{i \mid e \in p_i,\ x = \rho(e)\}|)$. We would like to note that the cost of navigating through junctions (e.g. traffic lights) is relaxed out.

A Centralised Traffic Routing problem is generated according to the current traffic situation periodically, i.e., every n seconds, as proposed in (Chrpa et al., 2019; Svadlenka et al., 2023). In particular, vehicles' intentions, i.e., where they enter the controlled region and where they leave the region, are collected for vehicles that are approaching the region. Generated routes are then assigned to the vehicles (before they enter the region). Automated planning-based approaches, in a nutshell, tackle the problem by considering "drive" actions, specifying the "elementary" moves of vehicles between adjacent junctions using a corresponding road segment (Chrpa et al., 2019; Svadlenka et al., 2023). In practice, vehicles collect information with their sensors about traffic status and transmit this information and their intentions to the centralised system. This information is used to generate data for solving the problem of distributing the traffic in the system - creating/updating the paths, and finally assigning them to each vehicle.

3.2 Determining Cost Through Traffic Intensity

Each road segment has its *capacity*, i.e., the maximum theoretical number of vehicles that can fit into the segment while also considering a minimum space between vehicles. If the number of vehicles routed to the given road segment exceeds its capacity, then we say that the road segment is *congested*. Then we specify two thresholds that divide the traffic intensity into (additional) three levels – *light, medium* and *heavy*. To draw a parallel between the categories and the well-known Level of Service, the light intensity level corresponds to grades A and B, the medium intensity level to C and D, and the heavy intensity level to E and F. Such a categorization has also been proposed in (Chrpa et al., 2019; Svadlenka et al., 2023).

The cost of the road segment is then determined by its length, its number of vehicles, and its traffic intensity category. Since we distinguish four categories of traffic intensity, i.e., *light, medium, heavy, congested*, we define for each road segment $r \in R$ thresholds u_r^l, u_r^m, u_r^h determining the maximum number of vehicles for the light, medium, and heavy traffic intensity level, respectively, and constants l_r, m_r, h_r, C_r representing the cost of a single vehicle for respective traffic intensity level.

Then, the cost function is a piecewise linear function as follows.

$$C(r,n) = n \cdot l_r \ (n \in [0, u_r^l]),\ n \cdot m_r \ (n \in (u_r^l, u_r^m]),\\ n \cdot h_r \ (n \in (u_r^m, u_r^h]),\ n \cdot C_r \ (n \in (u_r^h, \infty)) \tag{1}$$

3.3 Road Network Simplification

Automated-planning-based centralised traffic routing techniques reason on the level of individual vehicles, i.e., on the micro-simulation level. Therefore, the size of the road network in which these techniques reason is one of the considerable factors that affect their performance as the route has to be computed for each individual vehicle (Chrpa and Vallati, 2023).

A straightforward way how to simplify the road network is to precompute routes for each traffic flow that have bounded suboptimality (Svadlenka et al., 2023; Silva and Tang, 2024). Arguably, long routes might not be very efficient even in light traffic. Such routes with bounded suboptimality can be found by an algorithm combining Floyd-Warshall algorithm with Branch and Bound algorithm (Silva and Tang, 2024) or by a variant of A* with the Euclidean-distance heuristic (Svadlenka et al., 2023). The latter approach is considered in our experiments.

Further reduction of the size of the network to reason with can be done by precomputing "smart" routes that besides being bounded suboptimal are diverse enough so they might not share "common bottlenecks" (Svadlenka et al., 2023). In particular, bounded suboptimal routes are clustered according to the Jaccard Index by which the diversity of routes (based on comparing sets of their road segments) is determined and then, from each cluster, one route is selected (e.g. the shortest one) (Svadlenka et al., 2023).

Note that even though our MIP model works on the level of traffic flows, i.e., macro-simulation, we can still leverage the above simplification methods.

4 CENTRALISED TRAFFIC ROUTING MIP MODEL

The proposed MIP model draws inspiration from multicommodity flow problems (Ouorou et al., 2000), which are the type of network flow problem (Ahuja et al., 1993) with multiple commodities. Multicommodity flow problem is a fundamental class of optimization problems, often forming the backbone of other more intricate applications in telecommunications, logistics, and transportation (Fortz et al., 2017). Our MIP model goes beyond traditional multicommodity flow formulations for urban traffic systems, as the routing cost functions are continuous, convex, and piecewise linear. These cost functions are well-suited for capturing the different levels of urban traffic flow. This incorporation is particularly time-sensitive in centralised traffic routing, requiring real-time planning due to the dynamic environment. Therefore, we restricted our approach to a tight optimization time, which is well-designed for dynamic environments. To the best of our knowledge, this work is the first to integrate network flows with multiple commodities for centralized (urban) traffic routing while offering a novel perspective on real-time traffic flow routing in metropolitan environments.

4.1 Multiflow MIP Model

We aim at optimizing multicommodity flows to minimize traffic intensity on road segments while ensuring that vehicles entering a junction also leave it, maintaining flow integrity. A single traffic flow, defined by a specific origin and destination, is restricted to a subgraph of the road network created through preprocessing methods. We identify the number of vehicles that have to be routed in the considered traffic flow. For the junction of origin/destination, the sum of outgoing/incoming traffic flow equals to the number of routed vehicles. For any other junction, the flow conservation law holds, i.e. the sum of incoming traffic flow equals the sum of outgoing one. Decision variables for traffic intensity in each flow are modeled separately to maintain consistency and ensure all vehicles reach their destinations. While different flows can share road segments, the traffic intensity and cost on each segment depend on the total traffic from all flows using that segment.

Given the Centralised Traffic Routing problem $\chi = (\mathcal{N}, V, O, D, C)$, the initial step is to identify traffic flows. Each traffic flow is represented by a pair of (origin,destination) junctions in which vehicles (of that traffic flow) enter and leave the region, respectively. A traffic flow i is represented by a triple (o^i, d^i, n^i), where $o^i, d^i \in N$ are origin and destination junctions of i, respectively, and $n^i = |\{v \mid v \in V, O(v) = o^i, D(v) = d^i\}|$ is the number of vehicles in the flow i. We consider only traffic flows that are not empty, i.e., there is at least one vehicle in them. The number of non-empty traffic flows (of χ) is denoted as N_F. A subset of edges that might be generated by any preprocessing technique that simplifies the road network (as described in Section 3.3) for the flow i is denoted as E^i ($E^i \subseteq E$).

The above idea of leveraging the concept of network flows for routing traffic (in each flow separately) is modelled as follows. We introduce x variables representing the number of vehicles of a given flow routed on a given edge, i.e., $x^i_{u,v}$ represents the number of vehicles of flow i routed on edge $(u,v) \in E^i$. Then for each non-empty traffic flow $i \in \{1,...,N_F\}$, we define the following equations.

$$\forall (u,v) \in E : x^i_{u,v} \in \mathbb{N}_0,$$

$$\sum_{w \in V, (o^i,w) \in E^i} x^i_{o^i,w} = \sum_{(w,d^i) \in E^i} x^i_{w,d^i} \tag{2}$$

$$\sum_{(o^i,w) \in E^i} x^i_{O,w} = n_i \tag{3}$$

$$\sum_{(v,o^i) \in E^i} x^i_{v,o^i} = 0 \tag{4}$$

$$\sum_{(d^i,w) \in E^i} x^i_{d^i,w} = 0 \tag{5}$$

$$\sum_{(u,v) \in E^i, u \neq o^i} x^i_{u,v} = \sum_{(v,w) \in E^i, w \neq d^i} x^i_{v,w} \tag{6}$$

$$\forall (u,v) \in E \setminus E^i : x^i_{u,v} = 0 \tag{7}$$

Equation (2) represents that the sum of all outgoing subflows from the origin junction equals the sum of all incoming subflows into the destination junction. Equation (3) determines the number of vehicles in the flow. Equations (4) and (5) ensure that none of the vehicles in the flow can return to the origin junction as well as no vehicle can reenter the network from the destination junction. Equation (6) represents that for all junctions, other than origin and destination, the sum of all incoming traffic must be equal to the sum of outgoing traffic. The x^i variables defined on edges that are not part of E^i are set to 0 (Equation (7)).

The objective function that we want to optimise in our model minimises the cost of the allocation of traffic flows in the road network. The original optimization equation, presented in Section 3.2 makes our model quadratic, so, in order to tackle that as solvers handle linear models more effectively, we will represent that equation in a linearised form as our model objective function, i.e.:

$$\min \sum_{r \in R} y_r$$

Note that the cost of traffic allocation on an edge depends on the number of vehicles in it and that multiple edges might share a single corresponding road segment.

One common approach involves linearising piecewise functions, where distinct regions defined by binary variables cause nonlinearity. Concerning that, we will formulate the equation (1) so that the cost for each edge and each traffic level is calculated as follows:

$$y_r \geq l_r \sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} - M(1 - z^l_r) \tag{8}$$

$$y_r \geq m_r \sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} - M(1 - z^m_r) \tag{9}$$

$$y_r \geq h_r \sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} - M(1 - z^h_r) \tag{10}$$

$$y_r \geq C_r \sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} - M(1 - z^C_r) \tag{11}$$

In the presented model the cost of two junctions is calculated by a piecewise function (1) where the binary variables z^l_r, z^m_r, z^h_r, and z^C_r are used (light, medium, heavy, congested, respectively), each representing a region of the piecewise function and determining the level of traffic intensity on the road segment r according to the number of vehicles allocated on it. By imposing constraints that link these binary variables to the objective function through linear inequalities, the nonlinear objective function can be obtained with a series of linear segments (constraints (8) to (11)). Note that M is a (very) big constant, the usage of which combined with the constraint (12) ensures that only one segment (traffic intensity level) is active at a time, effectively capturing the behavior of the original nonlinear objective function. The model showed the best performance with M equal to 10 million after testing it with different values.

$$\forall r \in R$$

$$z^l_r \in \{0;1\}, z^m_r \in \{0;1\}, z^h_r \in \{0;1\}, z^C_r \in \{0;1\}$$

$$z^l_r + z^m_r + z^h_r + z^C_r = 1 \tag{12}$$

$$\sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} \leq z^l_r u^l_r + z^m_r u^m_r + z^h_r u^h_r + z^C_r M \tag{13}$$

$$\sum_{i=1}^{N_F} \sum_{\substack{(u,v) \\ r=\rho((u,v))}} x^i_{u,v} \geq z^m_r u^l_r + z^h_r u^m_r + z^C_r u^h_r \tag{14}$$

Equation (12) represents that only one traffic intensity level can be chosen for a given road segment. Equations (13) and (14) represent the upper and lower bounds for the number of vehicles (from all flows) for the respective traffic intensity level. The thresholds u^l_r, u^m_r, u^h_r represent the maximum number of vehicles on the road segment r that can be considered as a light, medium, or heavy level of traffic intensity, respectively. The model successfully works without the equation (14), but it helps with the continuous relaxation of the model, enhancing the performance and computational speed.

4.2 Extracting Routes for Vehicles

The solution of the MIP model provides information on how particular traffic flows are distributed in the road network. Specifically, the decision variables x contain information about the number of vehicles from a given traffic flow allocated to a given road segment. Such information is however insufficient as we need to know the exact routes for all the vehicles.

Algorithm 1 describes the procedure of extracting routes alongside the number of vehicles allocated to each of the routes for a single traffic flow. Therefore, the algorithm is called for each traffic flow separately. As an input, the algorithm takes a description of the road network, the solution of the MIP model, the identifier of the given traffic flow (i), and its starting junction (the location of origin).

The algorithm performs a Depth First Search while looking for individual routes, which is implemented by recursively calling the SearchRoute function. The current (partial) route is stored in p and the current junction in j.

We initially look for outgoing edges from j (in the road network graph) that have some vehicles allocated to them, i.e., the value of x^i for a given edge is greater than zero (Line 2). If none such outgoing edge exists, then we reached the destination (j is the destination location since there was an incoming traffic to j but there is no outgoing traffic). We extract the route p and the number of vehicles (veh_count) for p equals the minimum of allocated vehicles per each edge on the route p (additional traffic on some edges belongs to other routes). Then, we decrement the values of x on the route by veh_count. If some outgoing edge from j is already in p, then we have a loop (Line 10). We might encounter loops if the MIP solver returns a suboptimal solution that might happen, for example, if we impose time limits on the solver. If a loop is identified, we remove it by decrementing the x values on the loop by the minimum of these x values (Lines 12–13). In other situations, we iterate through the outgoing edges (with some traffic on them) such that we append the edge to p and call the SearchRoute with a subsequent junction (Lines 15–18).

We can see that the algorithm performs an exhaustive Depth First Search such that it only follows edges (road segments) that have some traffic allocated to them. Since the sum of outgoing traffic for a junction (other than the origin or destination) equals the

Function `SearchRoute`(p, j, i, *all_routes*)**:**
  $outgoing\_edges = \{(j, j') \mid (j, j') \in E, x^i_{j,j'} > 0\}$;
  if $outgoing\_edges = \emptyset$ **then**
    ;   `/* We extracted a route */`
    veh_count $\leftarrow \min_{(j_x,j_y)\in p} x^i_{j_x,j_y}$;
    if *veh_count*>0 **then**
      all_routes $\leftarrow$ all_routes$\cup\{(p, \text{veh_count})\}$;
      $\forall(j_x, j_y) \in p : x^i_{j_x,j_y} \leftarrow x^i_{j_x,j_y} -$ veh_count;
    end
  else
    if $\exists(j, j') \in p$ **then**
      ;   `/* We found a loop */`
      $p^j \leftarrow p$.subpath_from(j);
      veh_count $\leftarrow \min_{(j_x,j_y)\in p^j} x^i_{j_x,j_y}$;
      $\forall(j_x, j_y) \in p^j : x^i_{j_x,j_y} \leftarrow x^i_{j_x,j_y} -$ veh_count;
    else
      for $(j, j') \in outgoing\_edges$ **do**
        p.append((j, j'));
        SearchRoute(p, j', i, all_routes);
      end
    end
  end
SearchRoute($\langle\rangle$, start_junction, i, $\{\}$);

Algorithm 1: Algorithm for extracting routes for a given traffic flow.

sum of incoming traffic, there always exists an outgoing edge that has some traffic allocated to it (unless we are in the destination junction). We have two situations in which we stop searching in a given branch (and then backtrack and continue searching another branch until we explore all the branches). Firstly, we reach the destination junction. In such a way we can extract the route (from the origin) and the amount of traffic equals the minimum of the current allocation on the edges of the route. Updating the traffic allocation then reflects the extracted route and the traffic on it. The second "stopping" situation concerns loop detection (that might happen if the solution of the MIP model is suboptimal). We then identify the amount of traffic in the loop and update the traffic allocation by removing the traffic that is on that loop. It should be noted that loop detection (and elimination) improves the quality of the MIP solution since the resulting routes do not contain loops (albeit the MIP solution does). Since we perform an exhaustive search, we eventually extract routes that cover all routed vehicles. Note that decisions made during the search, i.e., the order in which the successors are explored might affect resulting routes (and traffic allocation on them). In other words, the solution of the MIP model might yield multiple valid routing solutions.

The final step, after Algorithm 1 finishes for every traffic flow, is to allocate extracted routes to the vehicles. For each traffic flow, we select vehicles that belong to this traffic flow and then we iterate over all_routes (from Algorithm 1) and, in an iteration step, we assign the current route to the respective number of (yet unallocated) vehicles. We would like to note that there might not be a single unique route allocation for the vehicles if the number of routes is greater than one. In our case, we allocate routes in the order they are stored in all_routes to vehicles ordered by their ids.

5 EXPERIMENTAL EVALUATION

The aim of the experimental evaluation is to show the potential of our MIP-based technique for Centralised Traffic Routing. In particular, we compare against a planning-based technique that uses "smart route" preprocessing (Svadlenka et al., 2023) to show that our macrosimulation-based MIP approach is more scalable and capable of more effective traffic routing than microsimulation-based planning approaches. On top of that, we also compare against the decentralised (DUO) approach for dynamic traffic routing that is implemented within the SUMO simulator (Lopez et al., 2018).

The quality of routing is measured by the average travel time that is extracted from the simulation (by SUMO). To provide a bigger picture, we also measured the average traveled distance and the average speed. Note that SUMO simulations consider aspects that we relaxed out (e.g. traffic lights). Simulations concerning Centralised Traffic Routing are done offline, i.e., vehicles routes are computed in advance.

5.1 Scenarios and Settings

For the experiments, we used two scenarios - New York (located between Grand Concourse and Sheridan Boulevard) and Sydney (southeast of Centennial Park) - that were introduced by Svadlenka et al. (Svadlenka et al., 2023) (depicted in Figure 1). For New York and Sydney, we have considered 16 and 40 scenarios, respectively, that differ by traffic flows (origin and destination) and how flows evolve in time (static or increasing). For all scenarios, we considered a 1-hour time window that was divided into 30-second "episodes" where each corresponded to a single instance of a Centralised Traffic Routing problem. Hence, for each scenario, we considered 120 episodes. For New York, we considered 5 traffic flows per scenario while for Syndey 4 traffic flows per scenario. The traffic intensity ranged between 760 and 1208 vehicles per hour per traffic flow. The simulation of (routed) traffic ran for 2 hours in SUMO. The road network, in both cases, is initially empty and the first vehicles arrive into the network at time zero (of the simulation time).

Additionally, we considered a scenario from central Dublin, where we used a real historical dataset (Gueriau and Dusparic, 2020). The central region was selected according to the most intense traffic (depicted in orange in Figure 2). We considered an almost 2-hour time window spanning from 7am to 9am (the morning rush hour). Again, we divided the time window into 30-second episodes, yielding 238 episodes in total. The total number of routed vehicles was 10 778. The simulation of (routed) traffic also ran for 2 hours in SUMO. The road network, again, is initially empty and the first vehicles arrive into the network at time zero (of the simulation time).

As described in Section 3.3, we preprocessed the original road network in two ways. One involves generating bounded suboptimal routes for each traffic flow (the bound was set to 1.3 as in (Svadlenka et al., 2023)), which we later denote as BR. The other "smart routes" preprocessing method involves clustering of potential routes based on their diversity (i.e., having fewer road segments in common) and selecting the shortest route for each cluster (Svadlenka et al., 2023). This preprocessing method is later denoted as SR. Both BR and SR preprocessing methods are used with MIP method as well as with the planning method we compare against (for which the preprocessing methods were originally designed) (Svadlenka et al., 2023).

The thresholds determining the level of traffic intensity for road segments are derived from the physical capacity of the road segment that is measured by the number of "standard" vehicles that can physically fit into the road segment while considering a minimum distance between them. Light traffic intensity of a road segment corresponds to the number of routed vehicles up to 40% of its capacity. Medium traffic intensity of a road segment corresponds to the number of routed vehicles from 40% to 60% of its capacity. Heavy traffic intensity of a road segment corresponds to the number of routed vehicles higher than 60% of its capacity but not more than the capacity. Congested traffic intensity of a road segment corresponds to the number of routed vehicles higher than its ca-

(a) New York region

(b) Sydney region

Figure 1: The New York (left) and Sydney (right) scenarios. The colored arrows illustrate sample traffic flows (their entry and exit points to the controlled area). The scenarios are taken from (Svadlenka et al., 2023).

pacity. These thresholds are determined in analogy with (Vallati et al., 2021). The cost of a road segment (for a single vehicle) with light traffic intensity equals the length of the road segment. The cost of a road segment (for a single vehicle) with medium traffic intensity equals 10 times the length of the road segment. The cost of a road segment (for a single vehicle) with heavy traffic intensity equals 100 times the length of the road segment. The cost of a road segment (for a single vehicle) with congested traffic intensity equals 100000. Note that the cost is determined in analogy to (Chrpa et al., 2019; Svadlenka et al., 2023).

For solving the MIP model, the global nonlinear solver Gurobi was also chosen for his ability to work with quadratic models[2]. Since we need to obtain the solution within a certain time limit (one episode considers 30 seconds of traffic), process it, and send it back, the runtime is limited to 25 seconds. For the planning-based approach, we use the Mercury planner (Domshlak et al., 2015) with a time limit of 25 seconds as well. In case we fail to find any solution for a given episode within the time limit, we assign the shortest routes (for New York and Sydney), or the routes specified in the historical dataset (for Dublin) for all the vehicles considered in the failed planning episode.

In a nutshell, we have extracted the following metrics from simulating the different routing methods in SUMO (Lopez et al., 2018). Average trip distance, i.e., how many meters the vehicles had to travel on average, average trip duration, i.e., how many seconds the vehicles had to take to reach their destinations, and average speed (in meters per second). The experiments were run on a computer equipped with AMD Ryzen 5000 7, with a memory limit of 32GB[3].

5.2 Results

The results of the experiments are summarised in Table 1. Noteworthy, for New York and Sydney, the results are averaged for all considered scenarios with different traffic flows (16 for New York and 40 for Sydney). It can be seen that in all scenarios and both variants of pre-processing optimisations of the road network (BR and SR), MIP approach solves (not necessarily optimally) all episodes (instances of the Centralised Traffic Routing problem). Compared to the planning approach, MIP model has shown that it is capable of scaling (much) better, which is especially apparent for the BR cases (where we consider bounded sub-optimal routes for each traffic flow) where the planning approaches solved (not necessarily optimally) only a few episodes.

Focusing on the average trip duration, which is the metric we optimise for, in contrast to planning, MIP achieved better results for the BR preprocessing. Whereas in planning, the SR preprocessing, which pre-computes several promising routes for each traffic flow, is essential for improving the coverage (i.e., the percentages of solved episodes), for MIP model SR is too restrictive and introduces suboptimality (as

[2] https://www.gurobi.com/documentation/

[3] Benchmark data are provided here: https://github.com/xankr/utc-mip-icaart2025

Figure 2: Image of the whole Dublin metropolitan area (left) and the considered city center scenario (right)

Table 1: Simulation results for all city benchmarks. “Solved” denotes the percentage of solved episodes (not necessarily optimally) by the MIP and Planning approaches. “Distance”, “Speed” and “Duration” denote the average trip distance, average speed, and average trip duration, respectively.

Parameters	Baseline		MIP		Planning	
	Naive	DUO	BR	SR	BR	SR
			New York			
Solved [%]	-	-	**100**	**100**	9.1	74.5
Distance [m]	**1915**	2355	2385	2136	2042	2361
Speed [m/s]	1.95	2.76	**3.18**	2.69	2.26	3.07
Duration [s]	1670	1356	**1193**	1388	1605	1235
			Sydney			
Solved [%]	-	-	**100**	**100**	1.77	46.9
Distance [m]	2608	2980	3002	2901	**2454**	2770
Speed [m/s]	3.38	5.6	**6.74**	5.75	3.35	3.86
Duration [s]	1324	746	**641**	761	1245	1160
			Dublin			
Solved [%]	-	-	**100**	**100**	0	96.6
Distance [m]	828	904	**812**	822	828	835
Speed [m/s]	5.13	**6.15**	5.92	5.52	5.13	5.81
Duration [s]	360	229	276	283	360	**226**

can be seen from the results for MIP). Comparing the best results (of average trip duration) between MIP and planning, we can see that in Sydney the MIP-based approach is better by about 45%, in New York slightly better (about 4%), while in Dublin, it is worse by about 20%. In comparison to the Naive approach that, in the New York and Sydney scenarios, considers shortest routes, the MIP approach improves the average trip duration by about 30% and more than 50%, respectively. In Dublin, where the Naive approach consists of historical traffic data (hence the routes are not necessarily the shortest), the improvement of the MIP-based approach was roughly 25%. Concerning DUO, the MIP-based approaches outperformed it in New York and Sydney by roughly 12% and 14%, respectively, while in Dublin, they were worse by about 17%. Interestingly, in the Dublin scenario, the planning approach (with the SR preprocessing) was slightly better than DUO. To provide a better perspective we compared the best approach in each category in Figure 3.

In terms of average trip distance, it is no surprise that all routing approaches tend to generate longer routes in order to mitigate traffic intensity on exposed

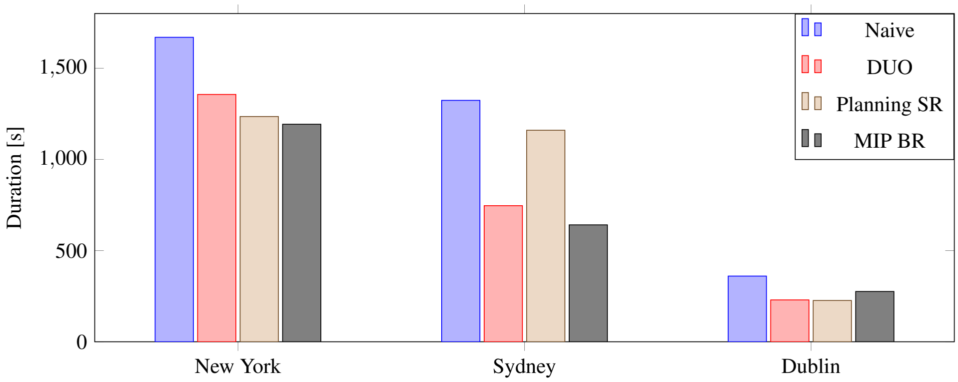

Figure 3: Average Trip Duration Comparison for Different Methods per Scenario.

Table 2: Average total runtime (PT [s]), average number of optimally solved scenarios (OPT [%]) and average optimality gap (GAP [%]) of MIP.

	New York		Sydney		Dublin	
	BR	SR	BR	SR	BR	SR
PT	328	62	247	45	1005	452
OPT	100	100	99.8	100	97.1	100
GAP	0	0	0.95	0	2.8	0

road segments. In the New York scenario, the routes get more than 20% longer in comparison to the Naive approach. On the other hand, the average speed increased by more than 50% which not only compensated for longer driven distance but also saved some time. In the Dublin scenario, the average trip distance for all centralised routing approaches is roughly the same as in the Naive approach, yet it leads to better trip duration results (note that the planning results for BR are the same as Naive as none of the episodes was solved). DUO, on the other hand, had longer routes by about 10%, which led to better utilisation of the road network (measured by the average speed).

5.3 Discussion

The results have shown that leveraging MIP-based techniques is a viable option for dealing with Centralised Traffic Routing problems as these techniques provide better scalability than techniques based on automated planning. That said, the macrosimulation type of reasoning on the level traffic (sub)flows alleviates some symmetries that are associated with the microsimulation type of reasoning since there is no difference in terms of the objective function which individual vehicle takes which route. For example, if v_1 takes route r_1 and v_2 takes r_2 (assuming that v_1 and v_2 belong to the same traffic flow), the value of our objective function will be the same as if v_2 takes r_1 and v_1 takes r_2.

Better performance of the MIP approach led to improvement of the overall results, measured by the average trip duration, in New York and Sydney, outperforming all the other methods. In Dublin, however, MIP underperformed both DUO and the planning approaches. The main difference between the New York and Sydney scenarios and the Dublin scenario is that in the former we deal with several traffic flows (5 and 4, respectively) each with several vehicles (at most 20) per episode, while in Dublin, the number of traffic flows is larger often containing a single vehicle (in a single episode). This aspect mitigates the benefits of macrosimulation as for "single-vehicle" flows there is no difference to microsimulation. The reason for "scattered" traffic flows in Dublin is that most vehicles start or finish their trip within the region.

Another aspect that affects the results is the accuracy of the objective function, i.e., how it reflects the actual traffic situation. As Table 2 summarises, the MIP approach generated optimal solutions in most cases and almost optimal in the rest. The objective function we use in this paper has been specified in the literature (Chrpa et al., 2019; Vallati et al., 2021; Svadlenka et al., 2023) with the rationale to reduce traffic intensity for the road segments in the controlled (urban) region. The objective function, however, might not accurately capture some nuances such as the shape of road segments and, more importantly, how the segments are connected. For example, if traffic from a side road is merging with the traffic on a main road on an uncontrolled junction, it might introduce additional bottlenecks as the traffic from the side road might not (easily) merge if the traffic on the main road is (slightly) more intense. We have observed such situations in the simulations that had a detrimental impact on the results. Also, in contrast to DUO, our objective function does not consider the current traffic situation outside the given planning episode (e.g., while routing we do not get information about heavy traffic that is currently on some road segments).

The lessons we have learned indicate that Centralised Traffic Routing (via MIP) is a viable way to effectively route traffic in urban regions suffering from heavy traffic (especially in rush hours). Despite the above drawbacks, the results have shown that our MIP method can outperform the decentralised approaches (DUO) in scenarios in which we route several more intense traffic flows (such as in the New York and Sydney scenarios). In other words, Centralised Traffic Routing seems to work effectively in scenarios in which we route transit traffic from multiple traffic flows that might interfere with each other. We believe that centralised traffic routing can complement the decentralised one as we might identify common traffic flows (that interfere with each other) and route only vehicles in these flows by centralised routing techniques while the other vehicles by decentralised routing techniques.

6 CONCLUSION

In this paper, we have addressed the Centralised Traffic Routing problem by means of Mixed-Integer Programming by modelling the problem as a combination of multiple network flows. We designed a MIP model that naturally captures the cost function (as specified in the literature (Chrpa et al., 2019; Svadlenka et al., 2023)). We have shown that the macrosimulation level reasoning that MIP allows improves scalability over the microsimulation-based approaches such as those based on automated planning (Chrpa et al., 2019; Svadlenka et al., 2023). In terms of Centralised Traffic Routing in general, our experiments (especially those on the New York and Sydney scenarios) showed that it has a good potential to outperform distributed routing methods that are nowadays routinely exploited in navigation systems. The lessons learned from the experiments indicate that Centralised Traffic Routing has more potential in routing several more intense traffic flows rather than a large number of "scattered" traffic flows (as happened in the Dublin scenario).

In the future, we plan to investigate how effectively we can identify bottlenecks (e.g. merging from the side road on an uncontrolled junction) and how these bottlenecks can be effectively represented in the objective function. Also, we plan to investigate how we can effectively identify "common traffic flows" in larger urban areas and how to integrate Centralised Traffic Routing techniques on these flows into other (decentralised) routing approaches.

ACKNOWLEDGMENTS

This research is supported by Czech Science Foundation (project no. 23-05575S), by the European Union under the OP JAK project ROBOPROX (reg. no. CZ.02.01.01/00/22_008/0004590), and by the Grant Agency of the Czech Technical University (project no. SGS24/115/OHK3/2T/37).

REFERENCES

Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. (1993). *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., USA.

Antoniou, G., Batsakis, S., Davies, J., Duke, A., McCluskey, T. L., Peytchev, E., Tachmazidis, I., and Vallati, M. (2019). Enabling the use of a planning agent for urban traffic management via enriched and integrated urban data. *Transportation Research Part C: Emerging Technologies*, 98:284 – 297.

Cao, Z., Jiang, S., Zhang, J., and Guo, H. (2017). A unified framework for vehicle rerouting and traffic light control to reduce traffic congestion. *IEEE Transactions on Intelligent Transportation Systems*, 18(7):1958–1973.

Chang, K.-H., Hsu, P.-Y., Lin, C.-J., Lin, C.-L., Juo, S.-H. H., and Liang, C.-L. (2019). Traffic-related air pollutants increase the risk for age-related macular degeneration. *Journal of Investigative Medicine*, 67(7):1076–1081.

Chrpa, L. and Vallati, M. (2023). Centralised vehicle routing for optimising urban traffic: A scalability perspective. In *IEEE Intelligent Vehicles Symposium, IV 2023, Anchorage, AK, USA, June 4-7, 2023*, pages 1–6. IEEE.

Chrpa, L., Vallati, M., and Parkinson, S. (2019). Exploiting automated planning for efficient centralized vehicle routing and mitigating congestion in urban road networks. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC 2019*.

Claes, R., Holvoet, T., and Weyns, D. (2011). A decentralized approach for anticipatory vehicle routing using delegate multiagent systems. *IEEE Transactions on Intelligent Transportation Systems*, 12(2):364–373.

Domshlak, C., Hoffmann, J., and Katz, M. (2015). Red-black planning: A new systematic approach to partial delete relaxation. *Artif. Intell.*, 221:73–114.

Du, L., Han, L., and Li, X.-Y. (2014). Distributed coordinated in-vehicle online routing using mixed-strategy congestion game. *Transportation Research Part B: Methodological*, 67:1–17.

Fortz, B., Gouveia, L., and Joyce-Moniz, M. (2017). Models for the piecewise linear unsplittable multicommodity flow problems. *European Journal of Operational Research*, 261(1):30–42.

Friesz, T. L., Luque, J., Tobin, R. L., and Wie, B.-W. (1989). Dynamic network traffic assignment considered as a continuous time optimal control problem. *Operations Research*, 37(6):893–901.

Gueriau, M. and Dusparic, I. (2020). Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic. In *23rd IEEE International Conference on Intelligent Transportation Systems*.

Kirimtat, A., Krejcar, O., Kertesz, A., and Tasgetiren, M. F. (2020). Future trends and current state of smart city concepts: A survey. *IEEE Access*, 8:86448–86467.

Krishnan, A., Markov, M., and Bonakdarpour, B. (2017). Distributed vehicle routing approximation. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 503–512.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., and Wießner, E. (2018). Microscopic traffic simulation using sumo. In *Proceedings of ITSC*.

McCluskey, T. L. and Vallati, M. (2017). Embedding automated planning within urban traffic management operations. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017*, pages 391–399.

Ouorou, A., Mahey, P., and Vial, J.-P. (2000). A survey of algorithms for convex multicommodity flow problems. *Management science*, 46(1):126–147.

Pilát, M. (2018). Evolving ensembles of traffic lights controllers. In *IEEE 30th International Conference on Tools with Artificial Intelligence, ICTAI 2018*, pages 958–962. IEEE.

Polimeni, A. and Vitetta, A. (2024). Network design and vehicle routing problems in road transport systems: Integrating models and algorithms. *Transportation Engineering*, 16:100247.

Pozanco, A., Fernández, S., and Borrajo, D. (2021). Online modelling and planning for urban traffic control. *Expert Systems*.

Silva, M. D. R. L. and Tang, M. (2024). A new framework for centralized coordinated multi-vehicle dynamic routing. *IEEE Access*, 12:24243–24253.

Svadlenka, M., Chrpa, L., and Vallati, M. (2023). Improving the scalability of automated planning-based vehicle routing via smart routes identification. In *8th International Conference on Models and Technologies for Intelligent Transportation Systems, MT-ITS 2023*, pages 1–6. IEEE.

Vallati, M., Scala, E., and Chrpa, L. (2021). A hybrid automated planning approach for urban real-time routing of connected vehicles. In *24th IEEE International Intelligent Transportation Systems Conference, ITSC*, pages 3821–3826. IEEE.

Wenning, B.-L., Timm-Giel, A., and Pesch, D. (2006). A distributed routing approach for vehicle routing in logistic networks. In *IEEE Vehicular Technology Conference*, pages 1 – 5.

Wolsey, L. A. (1998 - 1998). *Integer programming / Laurence A. Wolsey.* Wiley-Interscience series in discrete mathematics and optimization. J. Wiley, New York.

Xie, X.-F., Smith, S., and Barlow, G. (2012). Schedule-driven coordination for real-time traffic network control. In *Proceedings of ICAPS*.

Yamashita, T., Izumi, K., Kurumatani, K., and Nakashima, H. (2005). Smooth traffic flow with a cooperative car navigation system. In *4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*, pages 478–485. ACM.