

Analysing Routes Generated by Planning-based Centralised Traffic Routing

Andrii Nyporko^{1,2}, Lukáš Chrpa¹, Matyáš Švadlenka^{1,2}

¹Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague

²Faculty of Electrical Engineering, Czech Technical University in Prague

Abstract

Centralised Traffic Routing offers a global perspective of the entire region, enabling effective route finding for vehicles that navigate through the region. Recently, automated planning techniques were leveraged for centralised traffic routing. However, centralised routing techniques might generate different routes for vehicles even if they share the same origin and destination and arrive at a similar time. Also, centralised routing techniques are usually computationally demanding.

In this paper, we analyse the outputs of a planning-based centralised traffic routing method in terms of the number of different routes per journey, differences between traveled distance, and travel time among these (different) routes. Then, we propose a method that leverages the found routes and their relative distribution to vehicles for an online centralised routing.

Introduction

Urban populations have steadily grown over the past few decades, leading to higher traffic volumes, particularly during rush hours. This leads to ever-increasing economic problems and poses health risks (Chang et al. 2019).

Whereas modern systems, usually accommodating Decentralised Traffic routing approaches based on the Dynamic User Optimal principle (Friesz et al. 1989) (e.g., WAZE™), are quite effective, they can cause traffic congestion in network bottlenecks due to the lack of synchronization between vehicles. Centralized Traffic Routing, in contrast, offers efficient route generation from a global perspective of the entire region in which routing occurs, taking into account all vehicles at once. Recent works concerning leveraging automated planning techniques for Centralised Traffic Routing have shown promising results (Chrpa, Vallati, and Parkinson 2019; Vallati, Scala, and Chrpa 2021; Švadlenka, Chrpa, and Vallati 2023).

In this paper, we analyse the outputs of a planning-based Centralised Traffic Routing method (Švadlenka, Chrpa, and Vallati 2023) focusing on the differences of the generated routes for individual journeys (traveled distance and trip duration) as well as the distribution of vehicles among these

routes. The former analysis might help to identify situations in which some vehicles are routed on disproportionately “expensive” routes that drivers might tend to avoid (Chrpa 2024). The data about routes and vehicle distribution can be leveraged in a (naive) online method, we propose, that based on the distribution randomly allocate routes to vehicles.

Centralised Traffic Routing

In a nutshell, *Centralised Traffic Routing* is a problem in which vehicles with different origins and destinations are routed in a given road network such that specified global criteria are optimised. Technically speaking, we have a *road network* represented by a directed graph, in which vertices represent junctions and edges represent road segments (or links), a set of vehicles which we want to route, and for each of the vehicles, we have its locations of origin and destination (both are junctions in the given road network). Then, the *centralised traffic routing problem* deals with finding a route for each of the vehicles, represented by a path in the underlying graph representing the road network such that each route starts at the vehicle’s origin and finishes at the vehicle’s destination. The routes are optimised for given criteria such as minimising expected traffic intensity on particular road segments.

For the purpose of our analysis, we consider a recent technique that leverages automated planning for solving centralised traffic routing problems (Švadlenka, Chrpa, and Vallati 2023). The technique initially simplifies the road network such that it precomputes several “smart” routes for each origin-destination couple. These “smart” routes are diverse enough (i.e., they do not share many common road segments) and are within a specified suboptimality bound. In particular, a variant of the A* algorithm (Hart, Nilsson, and Raphael 1968) is used to identify routes that are within the specified suboptimality bound (Švadlenka and Chrpa 2023) and then the DBSCAN algorithm (Ester et al. 1996) clusters the routes based on their similarity, which is measured by Jaccard Index (that measures similarity of sets). For each cluster, we consider only the shortest route (as a “smart” route) since such a setting achieved the most promising results (Švadlenka, Chrpa, and Vallati 2023).

In the second stage, the Centralised Traffic Routing problem is compiled into a planning task (modelled in the PDDL language (Fox and Long 2003)), where the vehicles have

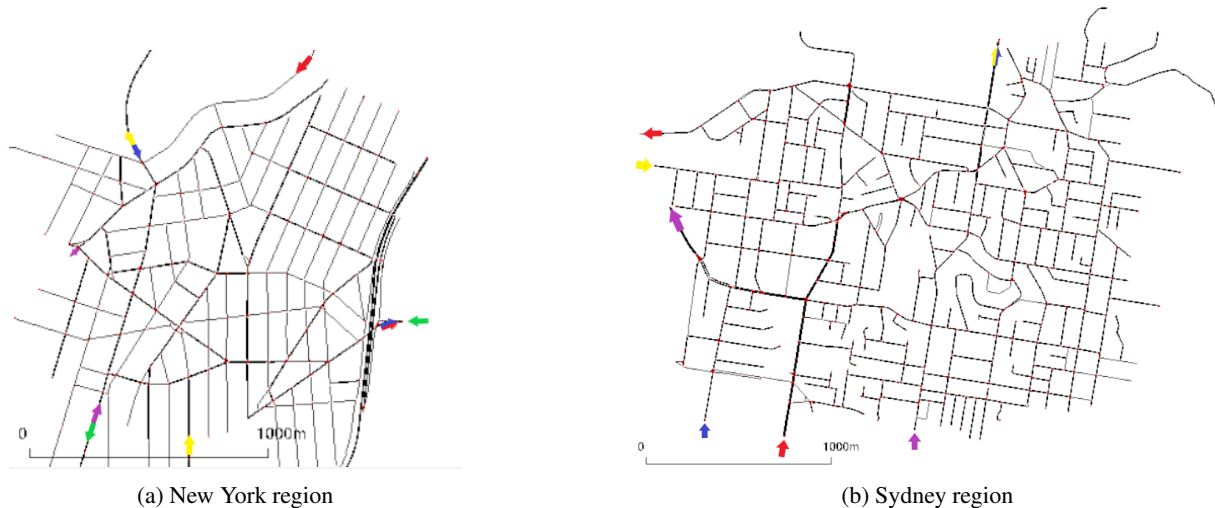


Figure 1: The New York (left) and Sydney (right) scenarios. The colored arrows illustrate sample traffic flows (their entry and exit points to the controlled area). The scenarios are taken from (Svadlenka, Chrpa, and Vallati 2023).

to be routed through the road network such that the traffic intensity on particular road segments is minimized. In particular, vehicles are navigated by “drive” actions that represent elementary moves (via a single road segment) within the network. The compilation is adopted from the work of Chrpa, Vallati, and Parkinson (2019). Note that vehicles of each traffic flow can be routed only on pre-identified “smart” routes. To solve the compiled planning task (to generate vehicle routes) we use the Mercury planner (Domshlak, Hoffmann, and Katz 2015).

In our case, each Centralised Traffic Routing problem represents a 30-second time window of traffic. To solve the problem, we imposed a time limit of 27 seconds (for “smart” route pre-processing and planning together) and, if no solution is produced, the shortest routes are applied. Such a setting aligns with related work (Svadlenka, Chrpa, and Vallati 2023; Chrpa, Vallati, and Parkinson 2019)

Offline vs. Online Routing

For the online use, the above-described approach requires collecting data from vehicles and estimating when the vehicles enter the controlled region before constructing a Centralised Traffic Routing problem and generating routes. Although modern navigation systems can provide a good estimate of when a vehicle might enter the controlled region, it still might be non-trivial overhead to collect data on all incoming vehicles. On top of that, solving a Centralised Traffic Routing problem is computationally demanding and requires non-trivial time.

On the other hand, the approach can be used offline on historic traffic patterns. In particular, we can extract information about vehicles passing the controlled region from the data, and compute routes (by Centralised Traffic Routing) offline, and possibly evaluate the impact of the routing on the traffic (Svadlenka and Chrpa 2024).

From the results of the offline (centralised) routing, we

can extract information about extracted routes for each considered origin-destination couple (or a journey) as well as the percentages of vehicles assigned to each of the routes. That information can be (naively) leveraged for online routing, i.e., after the vehicle broadcasts its intentions to a centralised routing infrastructure, the infrastructure randomly assigns one of the precomputed routes to the vehicle according to the distribution. For example, for a given origin-destination couple l_o, l_d , the offline routing approach found three routes - R1, R2, R3 – with the vehicle distribution of 20%, 50%, 30%, respectively. Then, for the vehicle approaching the region such that it enters the region in the location l_o and leaves the region in the location l_d , our (naive) online approach will assign R1 with probability 20%, R2 with 50%, and R3 with 30%.

Experimental Evaluation

The experiments aim at i) analyzing the output of the planning-based Centralised Traffic Routing method in terms of the number of alternative routes per origin-destination couple (journey), and the difference of the distance and travel time among these alternatives, and ii) evaluating the (naive) online routing method that leverages route distribution data (acquired by the offline approach).

We have considered two different regions of New York and Sydney depicted in Figure 1 (Svadlenka, Chrpa, and Vallati 2023). We generated multiple scenarios per region differing in origin-destination couples, or journeys (5 per scenario in New York and 4 in Sydney), and traffic intensity (ranging between 760 and 1208 per hour). In total, we considered 16 scenarios for New York and 40 for Sydney. The amount of vehicles varies in each scenario - from 2446 to 5590. Each scenario concerned 1 hour of traffic. The sub-optimality bound for the route length was set to 1.3 and the parameters of the DBSCAN algorithm were the same as (Svadlenka, Chrpa, and Vallati 2023).

Traffic was simulated in the well-known SUMO simulator (Lopez et al. 2018) such that the simulation was run until all routed vehicles reached their destinations. We have also enabled “teleports” in the simulations that prevent deadlocks (e.g. merging from side roads, which might be impossible in heavier traffic if traffic rules are followed rigidly).

The experiments were run on a computer equipped with AMD Ryzen 5000 7, with a memory limit of 32GB¹.

Traffic Distribution Analysis

By analysing routes generated by the planning-based Centralised Traffic Routing method, we identified several trends. Interestingly, for each origin-destination couple (or a journey) in each scenario (for both considered regions), at most 2 different routes were generated. In particular, in the Sydney region, **80.6%** journeys had 2 routes while in the New York region, it was **97.3%**. The distribution among routes in different scenarios ranged from **99.9% - 0.01%** (only 1 out of 1072 vehicles in a journey was navigated to one of the routes) to **49% - 51%**, on average (488 to 490 vehicles).

Another interesting trend we observed was that the shortest route (for a journey) is always among the generated routes (for that journey). Number-wise, on average, in the Sydney region, the vehicles were routed on the shortest routes in **75%** of cases, while in the New York region, only in **30.67%**. Part of the reason is that about 25% of Centralised Traffic Routing problem instances were unsolved in New York and about 53% in Sydney. In such cases, the shortest routes were considered for all journeys in unsolved instances. Yet, in some cases, the planning-based routing method also generated the shortest routes.

The average difference between the traveled distance on the two routes (per journey) was **11.1%** and **12.1%** for New York and Sydney, respectively. In other words, taking the alternative route (to the shortest one) adds more than 10% to the traveled distance, yet it is unlikely to be critical for the risk of drivers not following those (alternative) routes.

Average trip duration is another metric, usually more important than traveled distance. We calculate the average percentage increase of trip duration of the shortest route with respect to the alternative route per journey as follows:

$$R = \frac{1}{N \cdot M} \sum_{j=1}^N \sum_{i=1}^M r^{i,j}$$

$$r^{i,j} = \left(\frac{\text{dur}_s^{i,j} - \text{dur}_a^{i,j}}{\text{dur}_a^{i,j}} \right) \times 100, \quad \text{where:}$$

- R is the overall average ratio across all journeys,
- N is the total number of journeys,
- M is the number of scenarios,
- $r^{i,j}$ is the ratio for a scenario i and a journey j ,
- $\text{dur}_s^{i,j}$ is the trip duration for the shortest route,
- $\text{dur}_a^{i,j}$ is the trip duration for the alternative route.

¹Benchmark data are provided in https://gitlab.com/xankrig/flairs38_benchmark/

	Increase (%)	STD
New York (Full)	47.6	48.65
New York (Intermediate)	44.25	48.27
Sydney (Full)	2.76	10.98
Sydney (Intermediate)	3.46	11.14

Table 1: The average percentage increase of trip duration between the shortest and the alternative route. The STD stands for the standard deviation.

Table 1 shows how much, on average, the routes differ in trip duration for the journeys in the given urban regions (New York Full, Sydney Full). Also, to address a possible bias concerning an empty road network for the first and last routed vehicles, we calculate the ratios only for vehicles arriving after 10 minutes from the start of the simulation and up to 10 minutes before the last vehicles enter the simulation (New York Intermediate, Sydney Intermediate).

For New York, it was always the case that the trip duration was shorter for the alternative routes. On average, using the shortest route (for a journey) means more than 47% (or 44% for the “Intermediate” setting) trip duration increase with respect to the alternative route. In particular, the trip duration increase ranges from **7%** to **121%** (respectively, from **4%** to **67%** in the “Intermediate” setting) per journey.

In the Sydney case, the trip duration results were more diverse. Using the shortest route led to a shorter average trip duration for **52%** of journeys (respectively, **37%** in the “Intermediate” setting). On the other hand, the maximum trip duration increased while using the shortest route was **46%** (respectively, **43%** in the “Intermediate” setting). On average, using the alternative route was only marginally better (around 3%) as shown in Table 1. Notably, the distribution of traffic among the two routes (for each journey) affects the (average) trip duration of the vehicles. Specifically, very underrepresented routes (i.e., less than 1% of the routed vehicles) might affect the average results more considerably.

Concerning the possibility of drivers’ disobedience, i.e., that a driver follows a different route than assigned, there is a higher chance for such behavior in the Sydney scenarios in which taking the shortest route leads to a shorter trip duration than taking the alternative route. That might motivate drivers to take the shortest route rather than the assigned alternative route (Chrpa 2024).

Online Routing Results

As described earlier, the results of (planning-based) Centralised Traffic Routing that can be generated offline can be leveraged for online routing such that routes to vehicles are allocated randomly according to the distribution.

Table 2 shows the results of the comparison of Default (the shortest routes only), Offline (the planning-based centralised routing (Svadlenka, Chrpa, and Vallati 2023)), and the (naive) Online method we introduced. To account for randomness, each online scenario was run 10 times.

Traveled distance of the offline routing approach is higher than the Default (only shortest routes) by **9.8%** in New York and by **3.8%** for Sydney. Trip duration, on average,

	New York			Sydney		
	Default	Offline	Online	Default	Offline	Online
Distance	2233.89	2452.41	2451.96± 0.09	2629.91	2730.91	2731.01±0.22
Duration	3918.63	2953.24	2884.51±14.53	2083.47	1708.27	1618.49±38.46
Speed	1.06	2.12	2.20±0.01	2.78	3.20	3.5±0.06
Teleports	79.88	85.44	136.79±11.69	473.6	287.05	280.12±24.47

Table 2: The average results for the default, offline routing, and online routing. Due to randomness, each online routing scenario was run 10 times. The online routing results also include the standard deviation.

is higher by **32.6%** for New York and **21.9%** for Sydney for the Default setting (than for the offline one). Interestingly, the average speed for the offline routing doubled in New York scenarios with respect to the Default setting. The number of teleports (referring to the number of times the simulation “broke” a deadlock) was slightly higher in New York but considerably smaller in Sydney (by almost **40%**). Such results indicate the viability of centralised traffic routing approaches, as also shown in the recent literature (Vallati, Scala, and Chrpa 2021; Svadlenka and Chrpa 2023; Svadlenka, Chrpa, and Vallati 2023).

The online method yields almost identical average traveled distance compared to the offline method. This is, of course, expectable as the online method follows the vehicle distribution of the offline method. Interestingly, for the online method, the average trip duration decreased by **2.32%** for New York and by **5.25%** for Sydney (in comparison to the offline method), and the average speed increased by **3.77%** for New York and by **9.37%** for Sydney. The number of teleports increased on average by **60%** in New York while remaining almost the same in Sydney. This indicates that the New York scenarios are more prone to deadlocks for the online routing method.

These results indicate that even naive online routing that leverages the data generated by offline methods can maintain (or even improve) the traffic situation that would occur if centralised traffic routing methods were applied online.

Discussion

The results showed that even by considering a single additional route (per journey), the traffic situation can considerably improve. This observation is in line with recent results concerning centralised selecting routes from several alternatives (Silva and Tang 2024). However, the generation of at most two different routes per journey demonstrates the current limitations of planning-based approaches. The “smart route” preprocessing can considerably limit the number of possible routes per journey, yet it is important to reduce the size of the road network the planning-based techniques have to reason with (Svadlenka, Chrpa, and Vallati 2023). Also, we observed that for a single instance of the Centralised Traffic Routing problem, the planning-based method always generated only one route per journey (as it is simpler for planners).

Another important aspect is that larger differences among routes for a single journey might discourage drivers from following the assigned route. That might, however, mitigate the benefits of centralised traffic routing (Chrpa 2024). In

our analysis, we have identified some cases in the Sydney region in which some routes are both longer and take more time than their alternatives. That might be problematic as drivers might be less likely to follow such routes.

Leveraging the results of centralised routing methods acquired offline on historical data for online routing has shown its promise. In practice, it would mean that it might not be necessary to use computationally expensive centralised routing techniques online but just leverage lessons learnt from them on historical data. However, such a claim has to be taken with a grain of salt as the online approach we described might be (very) prone to changes of traffic patterns and especially to unexpected events (e.g. traffic accidents).

Conclusion

In this work, we analysed the outputs of a planning-based centralised traffic routing method (Svadlenka, Chrpa, and Vallati 2023) in terms of the number of generated routes per journey, vehicle distribution among these routes, and the differences between average traveled distance and average trip duration among these routes. The latter two metrics are important to assess the risk that drivers might not follow the routes that were generated by centralised traffic routing techniques, since some of the routes might be considerably more “expensive” to follow. Our analysis has shown that in some Sydney scenarios, there might be such a risk.

Acquired data about the generated routes and vehicle distribution among these routes (for all journeys) can be leveraged in a (naive) online routing technique, which we proposed in this paper, that allocates routes to vehicles according to that distribution. Our experimental results indicate that such an approach is viable and might alleviate the need to use computationally expensive centralised traffic routing techniques online. However, as we also pointed out, such an approach might have limitations as it might be vulnerable to traffic pattern changes or when an unexpected event occurs.

In the future, we plan to investigate a very recent centralised routing technique based on Mixed-Integer Programming that even outperformed a decentralised one on the same scenarios (Nyporko et al. 2025). Furthermore, we plan to investigate how to classify traffic patterns in order to make the online technique more robust, and how to determine when it might be necessary to resort to route generation (by centralised or decentralised routing techniques).

Acknowledgements This research is supported by Czech Science Foundation (project no. 23-05575S) and by the Grant Agency of the Czech Technical University (project no. SGS24/115/OHK3/2T/37).

References

- 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.; 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*.
- Chrpa, L. 2024. *Centralized Intelligent Traffic Routing in the Light of Disobedience of Drivers*. Springer International Publishing. 87–97.
- Domshlak, C.; Hoffmann, J.; and Katz, M. 2015. Red-black planning: A new systematic approach to partial delete relaxation. *Artif. Intell.* 221:73–114.
- Ester, M.; Kriegel, H.-P.; Sander, J.; and Xu, X. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *the Second International Conference on Knowledge Discovery and Data Mining*, 226–231.
- Fox, M., and Long, D. 2003. Pddl2.1: An extension to pddl for expressing temporal planning domains. *Journal of Artificial Intelligence Research* 20:61–124.
- 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.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Sci. Cybern.* 4(2):100–107.
- 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*.
- Nyporko, A.; Švadlenka, M.; Antonov, N.; Rohaninejad, M.; and Chrpa, L. 2025. Centralised urban traffic routing using mixed-integer programming. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*, 176–187. INSTICC.
- 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., and Chrpa, L. 2023. Towards a framework for intelligent urban traffic routing. In Franklin, M., and Chun, S. A., eds., *Proceedings of the Thirty-Sixth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2023, Clearwater Beach, FL, USA, May 14-17, 2023*. AAAI Press.
- Svadlenka, M., and Chrpa, L. 2024. A framework for centralized traffic routing in urban areas. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 8810–8814. ijcai.org.
- 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*, 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*, 3821–3826. IEEE.