A Methodology to Evaluate the Optimization Potential of Co-ordinated Vehicular Route Choices

Daniel Cagara Björn Scheuermann
Computer Engineering Group
Humboldt University of Berlin, Germany
{cagarada, scheuermann}@informatik.hu-berlin.de

Ana L. C. Bazzan
Instituto de Informática
Univ. Federal do Rio Grande do Sul, Brazil
bazzan@inf.ufrgs.br

Abstract— **A car navigation system's job is to plan a good route from an origin to a destination. There are many different options how this can be accomplished. Path choices can be calculated based on static road map representations, or they can take into account dynamic information like, e. g., known road blocks or the current traffic situation. More recently, the idea has gained ground that navigation systems could even cooperate in order to co-ordinate route choices so as to proactively avoid the formation of traffic jams. While several heuristics for algorithms to improve the vehicles' route choices have been proposed, little is known about the potential benefit of such optimizations. How much can we gain if dynamic information exchange and/or co-ordination between vehicles are used? Answering this question requires to obtain information on the travel times realized by "best possible", globally co-ordinated route choices—and therefore the solution of a highly complex optimization problem. Here, we propose a method to accomplish this. We use genetic algorithm optimization to jointly evolve the route choices of all cars in a street network iteratively towards an optimal solution, where the quality of each intermediate optimization step is assessed using a road traffic simulation.**

I. Introduction

Car navigation systems get increasing attention and popularity. The navigation system helps the user to find his best route from an origin to a destination. The selection of this optimal route is, in essence, the solution of a shortest path problem in directed graphs, and can thus be solved by algorithms like Dijkstra's algorithm, the A* algorithm [1] or the Bellman-Ford algorithm. Static map data forms the foundation for this functionality, but most modern devices use multiple sources of information and take various dynamic variables like the current traffic situation or historical traffic data into account when finding an optimal path. As of today, though, the decision in practice is generally a local one: the route choice is optimized per vehicle, minimizing the anticipated travel time of this one driver.

Such a local optimization from each driver's perspective, though, does not necessarily mean that an optimal point for the system as a whole is reached. In fact, it could mean that all drivers choose one route and a major traffic jam forms, while with proper distribution across multiple routes each single driver would be better off. Even with perfect information for all drivers, selfish route choices can lead to a highly suboptimal global situation—the *price of anarchy* is potentially high [2]. While several heuristic approaches to optimize route choices by cooperation between the individuals have been proposed [3], [4], there is no way to tell how far away from globally optimal routes (in whatever sense) those solutions actually are: if 5 % improvement is the best that can be reached, a heuristic which results in a 4 % improvement would be considered very good; if 20 % improvement are possible, the same heuristic would appear in a different light. The question how much improvement can be achieved in different kinds of road networks, if all the optimization potentials are fully used, has, so far, not been assessed in sufficient depth.

Here, we therefore do not aim for a route selection method that is directly applicable to the selection of routes in a real, live traffic scenario. But we discuss a methodology which is able to jointly optimize the route choices of all vehicles in a given scenario, in order to assess the optimization potential. To this end, we employ genetic algorithms [5]. Genetic algorithms have often be used very successfully to solve complex optimization problems. We argue that they are also well-suited to the structure of our problem. We employ a genetic algorithm in combination with microscopic traffic simulation in order to evaluate the quality ("fitness") of candidate solutions ("individuals"). Thereby, the set of routes for all vehicles is jointly evolved, so as to optimize the choices with regard to a global target function.

Besides their utility as a benchmark, we expect the obtained globally optimized routes to also provide valuable hints on how to *design* good heuristics: by analyzing the chosen optimal routes in different scenarios, we hope to identify patterns which can be taken up in the design of future cooperative routing algorithms.

This paper is structured as follows. After reviewing related work in Sec. II, we formally define the optimization problem in Sec. III. In Sec. IV, we then introduce the genetic algorithm which we use for optimization. In Sec. V, we present and discuss first results and experiences that we obtained from applying the proposed methodology to simple scenarios. Finally, in Sec. VI, we conclude this paper with a summary of our results and insights, along with a discussion of intended directions for future work.

II. Related Work

In recent years, there have been numerous papers that focused on improving route choices in road traffic networks. For example, [4] introduces an approach to optimize traffic flow using a genetic algorithm. This approach is targeted to practical use in a real world traffic information system by periodically repeating short-time forecasts of the traffic situation. We, in contrast, do not want to optimize routes online based on limited knowledge. We rather aim to quantify how much can be gained by searching for the best solution given perfect knowledge.

In [6], the authors' goal is again finding the optimal distribution of traffic in road networks. Their approach is to use an evolutionary game, Minority Game. They show that near-optimal traffic distribution can be achieved even when drivers choose their routes independently and without communication. However, they are working in a simple scenario with several highly simplified abstractions, whereas we use microscopic traffic simulations and therefore operate on a much more realistic view of road traffic.

In [7], a mathematical approach to optimizing traffic from the systems perspective is pursued. Again, do the complexity of the problem, several abstractions are made—for instance, static traffic flows are assumed. By using our genetic algorithm based optimization strategy, we are able to avoid such simplifications.

The existing body of work on traffic flow improvement in a more general sense includes, for example, approaches like [3], where an ant-hierarchical fuzzy system is applied. Yet, it is generally unknown how far those approaches are from an optimal solution—and this is just the question we are targeting here.

III. Problem Formulation

Assume the road map being represented by a graph $G = (V,E)$, consisting of a set of vertices V which represent the intersections of the road network, and a set of directed arcs $E \subset V \times V$ which describes the existing roads as directed connections between pairs of vertices. Further assume that there are z cars $c_1,\ldots,c_z$ in the road network, with car c starting its journey at time t_0^c. The route of car c is a sequence of k_c links, i. e., $R^c = (l_1^c,\ldots,l_{k_c}^c)$, where $\forall i \in \{1,\ldots,k_c\} : l_i^c \in E$ and for $i \in \{2,\ldots,k_c\}$ it holds that the i-th link starts at the vertex where the $i-1$-st link ended.

Each link $l \in E$ is associated with a function f_l; $f_l(t)$ is the travel time that it will take a vehicle entering link l at time t to traverse the link. This function depends on the traffic density and the characteristics of the respective road segment. It relates to the fundamental diagram used in traffic engineering. The functions f_l are therefore, in turn, influenced by the route choices of the cars in a non-trivial way. This is what makes the optimization problem so complex.

Consequently, the time at which car c has traversed the i-th link along its route is

$$t_i^c = t_{i-1}^c + f_{l_i^c}(t_{i-1}^c), \tag{1}$$

and the car's total travel time is given by

$$T_c = \sum_{i=1}^{k_c} f_{l_i^c}(t_{i-1}^c) = t_{k_c}^c - t_0^c. \tag{2}$$

In order to find the "best" routes for all cars, it is first necessary to define by which measure this decision is to be made. There are many possible choices for the target function. It is for example conceivable to minimize the mean of the absolute travel times, or the average (or maximal) level of congestion on the road networks links. Here, we propose to optimize the route choices in such a way that the mean relative improvement in the cars' individual travel times is maximized. The relative improvement is measured compared to the situation in which each car is individually and egoistically driving the shortest path. That is, for each individual car we determine the travel time if no optimization is in place; for car c, we denote this travel time by $\widehat{T_c}$. We then consider the ratio $T_c/\widehat{T_c}$, which describes the relative improvement (or deterioration) for car c: a ratio smaller than one means that the car's travel time has improved (i. e., it has reduced), a ratio larger than one corresponds to a deterioration for the individual car. Our target function is the mean relative gain of the cars, i. e., we minimize the expression

$$\sqrt[z]{\prod_{i=1}^{z} \frac{T_{c_i}}{\widehat{T}_{c_i}}}. \tag{3}$$

By using the relative changes instead of the absolute ones, we prevent longer routes from being favored during the optimization at the cost of ignoring shorter routes. Other target functions are equally well usable in our framework, so that this choice is not critical for the applicability of the method in general.

IV. The Genetic Algorithm

Due to the complex codependencies, the optimization problem is very difficult. Here, we tackle it by using a genetic algorithm. Genetic algorithms are a way to probabilistically search for an optimum regarding a predefined optimization criterion. The underlying model is borrowed from the evolution of living organisms. They operate by iteratively refining a fixed-size *population* of solution candidates, which are called *individuals*. In our case, an individual corresponds to a set of route choices: a specific driving route for each single car. These route choices are encoded in a bit field, termed the individual's *chromosome*.

The first generation of individuals is generated randomly. For all the individuals, the target function is evaluated; in the context of genetic algorithms, the value of the target function is also called *fitness*. A number of operations (selection, crossover, and random mutation) are then executed on the population, taking into account the determined fitness values. These operations generate new individuals forming a new generation. The intention is to preserve the beneficial properties of good individuals, while adding sufficient randomness in order to find yet better points in the optimization space.

Consequently, properties of individuals with high fitness values are more likely to end up in the newly generated individuals. The fitness of all individuals in the new generation is assessed again, and the process is repeated.

In order to encode the cars' route choices for an individual into a chromosome without generating a too complex solution space, we limit the number of possible route choices for each car to a fixed set of k alternative routes. This seems appropriate, since most routes that are theoretically conceivable (huge detours, routes with cycles,...) are not viable options anyway. We obtain the k path alternatives for a given car by an iterative penalty method [8]: Dijkstra's algorithm is applied k times. Initially, a road segment's cost is set to its length divided by the maximum allowed speed on the segment. Whenever a segment is used on a path, its weight is increased by multiplying it by a constant factor y. This makes the segment less likely to be chosen again on alternative routes. y has to be chosen in such a way that a segment is re-used only in cases where no suitable alternative exists. In the road network used in our evaluations, $y = 4$ has shown to give good results.

In a practice, one should choose k as a power of two (we use $k = 8$ here), so that the route choice for a car can be encoded in $\log_2 k$ bits, and each combination of bit values is a valid entry. Consequently, a chromosome is a bit string of length $z \cdot \log_2(k)$ bits. The bits $[(i-1)\cdot\log_2(k),\ldots,i\cdot\log_2(k)[$ represent the route choice for car c_i.

In order to obtain the resulting travel times of all cars in different configurations, we use the microscopic traffic simulator SUMO [9]. Before starting the optimization process, we once run a simulation where cars drive along their individually and independently chosen shortest paths. From this simulation run, we obtain the baseline travel time values $\widehat{T}_c$ used in the target function. During the optimization, in order to evaluate the fitness of a given individual, we run simulations in which cars choose the routes as determined by the chromosome under evaluation. Thereby we obtain the values T_c, so that we can determine the fitness according to (3). In order to reduce the simulation time per generation, our implementation distributes the simulations for assessing the individuals in a populations across the machines in a cluster.

The crossover operation generates new individuals for the next generation. To do so, it picks two candidates using the Roulette Wheel Selection method [10], which prefers individuals with good fitness. The chromosomes of these individuals are cut at a random point. By concatenating the left part of the first individual's chromosome and the right part of the second individual's, a new one is formed. After an individual has been sampled in this way, a mutation operation flips each bit in the chromosome with some probability (we use $p = 0.001$ here), causing random route changes for a small fraction of the cars.

V. Results

For our results presented here, we use a street map of the German city of Eichstädt grabbed from the OpenStreetMap [11] project. We imported it into SUMO using netconvert, SUMO's map conversion tool.

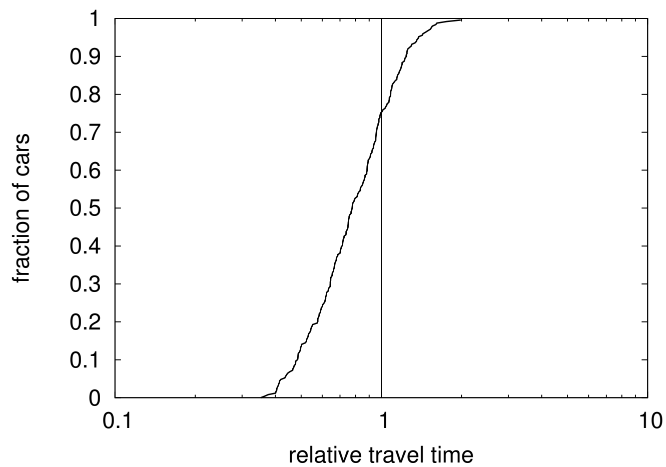

Fig. 1. CDF of the relative travel times in the congested scenario, after ten generations.

In order to test our methodology, we concentrated on two extreme cases. In both scenarios, the simulations include 250 cars, and we simulate a duration of 250 minutes. The scenarios differ in the distributions of the cars' starting times and their origin-destination pairs, though. In the congested case, scenario A, there is highly concentrated traffic, where all cars drive along the same route. Moreover, all cars enter the simulation close to its beginning, with 1 s fixed spacing. In the other case, scenario B, sources and destinations are chosen uniformly over the whole road network, and vehicles enter at uniformly distributed points in time. In this scenario, no congestion builds up and traffic flows smoothly even if each cars individually chooses the best route.

For the initial results presented here, we use a population size of 100 individuals for the genetic algorithm. On our machines, each simulation run takes ca. 5 min; distributed over six machines, an evolution over ten generations takes approximately 14 hours. The time for the operations of the genetic algorithm itself is negligible.

In Figure 1, we plot the per-car relative travel time gains in scenario A after ten generations. The plot shows the cumulative distribution of these values. As discussed above, values below one mean that the car has improved its travel time. As can be seen, this is the case for the majority of cars; improvements of up to 65 % are reached for individual cars. For some cars, the travel time deteriorates. This may indicate that a globally good solution indeed requires some cars to actually drive longer than if they entered the traffic jam. If this picture is confirmed once we obtained more results with our simulation methodology, an interesting research question arises: how could drivers be incentivized to voluntarily accept a deterioration, in order to improve the global outcome?

During the simulation we have found that the genetic algorithm converges very quickly to a virtually optimal point. The reason for this behaviour can be seen in the nature of this scenario: due to very high congestion, the possible travel speed

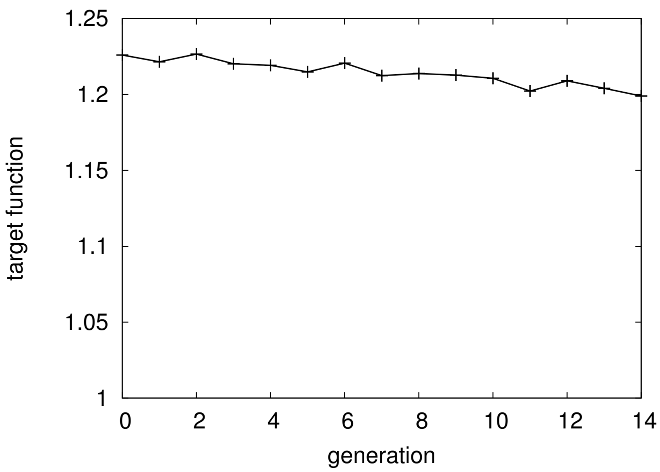

Fig. 2. Progress of the target function in the non-congested scenario.

on the shortest path gets so low that picking almost any other alternative route leads to an improvement in the cars' travel times. In fact, since alternative routes are virtually free (recall that all vehicles concentrate on the same origin-destination pair), a randomized choice between one of the k alternative paths distributes vehicles well and is thus not unlikely to yield a good result. This is just what happens when individuals for the first generation of the genetic algorithm are sampled—so, the existence of a very good individual from the very beginning is a likely event.

As one might expect, the results in the non-congested scenario B are quite different. Clearly, if there is no noteworthy congestion, there is not much optimization potential: if every car follows its individually calculated shortest path route, we will be very close to the optimal point. Therefore, it is to be expected that the target function converges to one over the generations. Indeed, we notice that the genetic algorithm does by far not as quickly converge to a stable level as in the congested scenario. In contrast to the congested situation, initial random choices may be expected to be quite far off the ideal—which is again just what we observe.

Figure 2 shows the progress of the target function over 14 generations. The y axis shows the target function of the best individual in the respective generation, which, slowly, moves towards the anticipated long-term limit of 1—which, in essence, means that each car will again choose its individually shortest path. Over the generations evaluated here, there is continuous improvement, but the results are still far off the ones obtained with individually optimized routes. Many cars are still worse off. Due to limited available computation time, we were not yet able to verify the long-term evolution, and leave this for future work.

VI. Discussion and Conclusion

In summary, the first results provided here clearly underline that the potential for optimization heavily depends on the traffic situation: if there is no substantial congestion, then very obviously the possible gain of coordinated route choices is close to zero. If, on the other hand, there is substantial congestion and if viable alternative paths are available, much can be gained. In this case, though, our observations corroborate the suspicion that a randomized choice between some reasonable route alternatives might indeed not be far off the optimal strategy. Moreover, our results already raise the problem how to deal with those drivers who's travel times increase—and a plethora of follow-up questions related to the acceptance of such route choices, or appropriate incentive or compensation mechanisms. Using our methodology, we now have the tools to assess these issues in more depth.

Our future work, consequently, will first focus on obtaining more results with larger simulation scenarios, more generations, different target functions, and a greater variety of different traffic situations. We plan to use computing clusters in order to simulate more complex city scenarios, like for instance the TAPASCologne scenario with realistic origin-destination matrices [12].

We also hope to learn about the structure of "good" route choices from our results. Even if our approach is not directly usable for route choices in real navigation systems—because it requires global knowledge, even about cars entering in the future—we intend to extract hints about how optimal routing can be done from a global systems perspective.

References

[1] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[2] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[3] H. M. Kammoun, I. Kallel, A. M. Alimi, and J. Casillas, "Improvement of the road traffic management by an ant-hierarchical fuzzy system," in *CIVTS '11: Proceedings of the IEEE Symposium on Computational Intelligence in Vehicles and Transportation Systems*, Apr. 2011, pp. 38–45.

[4] Y. Shigehiro, T. Miyakawa, and T. Masuda, "Road traffic control based on genetic algorithm for reducing traffic congestion," *Electronics and Communications in Japan*, vol. 95, no. 4, pp. 11–19, 2012.

[5] M. Srinivas and L. Patnaik, "Genetic algorithms: a survey," *IEEE Computer*, vol. 27, no. 6, pp. 17–26, June 1994.

[6] S. M. Galib and I. Moser, "Road traffic optimisation using an evolutionary game," in *GECCO '11: 13th Annual Genetic and Evolutionary Computation Conference*, July 2011, pp. 519–526.

[7] O. Jahn, R. Möhring, A. Schulz, and N. Stier-Moses, "System-optimal routing of traffic flows with user constraints in networks with congestion," Massachusetts Institute of Technology (MIT), Sloan School of Management, Working papers 4394-02, 2004.

[8] V. Akgun, E. Erkut, and R. Batta, "On finding dissimilar paths," *European Journal of Operational Research*, vol. 121, no. 2, pp. 232–246, 2000.

[9] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "SUMO – simulation of urban mobility: An overview," in *SIMUL '11: Proceedings of the Third International Conference on Advances in System Simulation*, Oct. 2011.

[10] J. Zhong, X. Hu, M. Gu, and J. Zhang, "Comparison of performance between different selection strategies on simple genetic algorithms," in *CIMCA-IWATIC '05: Proceedings of the International Conference on Computational Intelligence for Modelling Control and Automation / International Conference on Intelligent Agents, Web Technologies and Internet Commerce*, Nov. 2005, pp. 1115–1121.

[11] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, Oct. 2008.

[12] "Tapascologne project," http://sf.net/apps/mediawiki/sumo/index.php?title=Data/Scenarios/TAPASCologne.