

Traffic prediction based on a local exchange of information

Bruno Barbieri¹, Guillaume Sartoretti^{1,2}, Jean-Luc Falcone¹, Bastien Chopard¹, and Martin J. Gander²

¹ Computer Science Department

² Department of Mathematics,
University of Geneva, Switzerland

Abstract. We propose a decentralized method for traffic monitoring, fully distributed over the vehicles. An algorithm is provided, specifying which information should be tracked to reconstruct an instantaneous map of traffic flow. We test the accuracy of our method in a simple cellular automata traffic simulation model, for which the traffic condition can be controlled and analyzed theoretically. We show how local communication parameters affect the method accuracy.

1 Introduction

Nowadays in big cities, the need for an overview of the traffic situation is critical in order to avoid queues, accidents, and establish an efficient routing strategy. Current solutions often involve a centralized node which collects information from the drivers within the traffic, and broadcast it to let everyone make reasonable decisions. While this strategy works well, it raises the question of privacy: the central node knows everything about everyone in the network, and this centralization of the information may hit back if the knowledge is in the wrong hands.

In this paper, we extend the fully decentralized communication method described in [1, 2]. This method allows drivers to gather information and share it anonymously with the neighboring drivers throughout their journey, letting everyone have a good view of the status of the traffic. Due to this type of information distribution, the disadvantages of centralization and the privacy violation can be avoided. However, the problem is not trivial because the system must converge quickly to an accurate and stable representation of the current traffic conditions, and it must be able to detect rapid changes due to accidents or rush hours.

Early experiments have tested the capability of our method to build an accurate traffic map, using a simple traffic model in which the cars accumulate information as they move, through mutual exchanges. Using a simple circuit (fully described below), our method was able to predict the traffic condition, expressed as the rate of cars able to cross an intersection. Moreover, our method was able to adjust the prediction when the actual value was suddenly modified

(for instance due to an accident). For instance, results for traffic increase and decrease are shown in Fig. 1 taken from [2].

In all preceding simulations, we did assume that cars communicate only with immediate neighbors at a pace corresponding to the car movement itself. We can estimate what those restrictions mean in terms of real units, by deriving the model time step and the cell size. Assuming that street cells are 10 meters long, and that cars have a speed allowing them to fully stop in a street cell, we can compute the simulation time step (Δt) given the car braking deceleration. Using a deceleration value of 8 ms^{-2} (in the range of current car models), we find that $\Delta t \approx 1.12$ and a maximum car speed of $8.9 \text{ ms}^{-1} \approx 32 \text{ kmh}^{-1}$. The communication radius is then equal to 10 m and the bandwidth needed for a map of 200 cells and 128 temporal iterations is 22.3 kb/s. This is however very conservative compared to the currently available technology. For instance, Bluetooth 3 devices can communicate with an effective radius of 100 meters at a theoretical speed of 24 Mbits/s.

In this paper, we investigate how our method can be improved by increasing the communication radius and the communication frequency. The paper is organized as follows. We first propose a cellular automata model, from which the traffic flow can be predicted with the knowledge of only a few quantities. Then we define how the cars communicate and how they combine the information to progressively build a global traffic map. Finally we provide some results about the effect of the communication parameter on the prediction accuracy.

2 Traffic Model and Theoretical Framework

In order to produce simple traffic conditions suitable to test our communication method, we consider a Cellular Automata (CA) model. This approach has been widely used in traffic simulations, with excellent results (see for instance [3–7]).


2.1 Basic model

Here we consider several road segments, interconnected through simple junctions. A road segment is a 1D vector of length L , where each component (cell) represents a possible location for a car.


In order to describe the motion of the cars on a road, we modified Wolfram’s rule 184 [8] such as to be able to track individual cars. At initialization, an identifier ID_j is assigned to each car j . Identifiers are positive integers ranging from 1 to N where N is the total number of cars. Let $s_i(t)$ be the state of road cell i at time t , and $n_i(t)$ is the occupation of cell i at time t . Both values are defined as:

$$s_i(t) = \begin{cases} 0 & \text{if cell } i \text{ is empty,} \\ ID_j & \text{if car } j \text{ is in cell } i \text{ at time } t. \end{cases}$$

$$n_i(t) = \begin{cases} 0 & \text{if } s_i(t) = 0, \\ 1 & \text{otherwise.} \end{cases}$$



(a) Case with $\pi = 0.9$ when $t \leq 1320$ and $p_i = 0.1$ for $t > 1320$



(b) Case with $\pi = 0.1$ when $t \leq 1320$ and $\pi = 0.9$ for $t > 1320$

Fig. 1. Estimated value of traffic flow (π) in a two-road traffic circuit. The local estimation is shown with a continuous curve and the real value with a dashed curve. We can see that the estimation matches the real value 200 iterations after the event causing the traffic flow modification.

The car traffic can then be described by the update function:

$$s_i(t+1) = n_{i-1}(t)(1 - n_i(t))s_{i-1}(t) + n_i(t)n_{i+1}(t)s_i(t) \quad \text{for } 1 < i < L. \quad (1)$$

This means that a car moves to the next cell if and only if this next cell is free; otherwise it stays still. Roads are considered as one-way lanes. The two special locations are the beginning and the end of a road. These cells obey other rules that implement the chosen boundary conditions, for instance the junction that interconnects two segments.

Circuits A simple example of interconnection between two roads is the circuit shown in Fig. 2. It is composed of two side by side roads of length $L_1 = L_2 = L$. A car reaching the end of one of the roads may cross to the other road with probability π (as if there was a stop sign, or a traffic light), and given that the beginning of the other road is free.

Let s_i and s'_i be the states of each cell in road-one and road-two respectively. Accordingly, the rule for cells 1 and L in each road can be expressed as:

$$\begin{aligned} s_1(t+1) &= n'_L(t)(1 - n_1(t))s'_L(t) \cdot X(t) + n_1(t)n_2(t)s_1(t), \\ s'_L(t+1) &= n'_{L-1}(t)(1 - n'_L(t))s'_{L-1}(t) + n'_L(t)(1 - n_1(t))s'_L(t) \cdot X'(t). \end{aligned}$$

The quantities $X(t)$ and $X'(t)$ are Boolean random variables which are 1 with probability π . The equations for s'_1 and s_L are obtained by symmetry.

Occupation density of a cell On a road segment, the car density $\rho_i(t)$ on cell i is defined as $\langle n_i(t) \rangle$, the average of n_i (over time or space). We can write an evolution equation for $\rho_i(t)$ by taking the average of eq. (1):

$$\rho_i(t+1) = \rho_{i-1}(t) - \langle n_{i-1}(t) n_i(t) \rangle + \langle n_i(t) n_{i+1}(t) \rangle. \quad (2)$$

Note that $\langle n_i(t) n_{i+1}(t) \rangle \neq \rho_i(t) \rho_{i+1}(t)$ in general, because $n_i(t)$ and $n_{i+1}(t)$ are highly correlated.




Fig. 2. Diagram of a two-road circuit, with an initial number M_1 and M_2 of cars on each road.

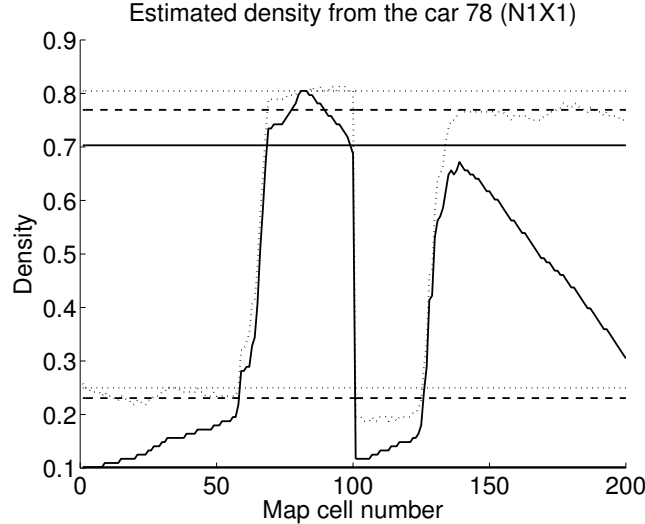


Fig. 3. Traffic density, computed over the last 128 iterations. The solid curve is the density estimated from the local map of a single car drawn at random at the end of a simulation with a car density of 0.5 and $\pi = 0.3$. The dotted curve is the real density obtained from the global map. Horizontal lines represent the values of p (bottom) and q (top). Solid, dotted and dashed lines corresponds respectively to the local estimates, the real estimates and the imposed values.

2.2 Theoretical framework

When running a two-road circuit simulation with the update rules defined above, the traffic in each road segment will start with a free traffic region followed by a congested traffic region when cars form a queue. A typical traffic situation is illustrated in Fig. 3. This structure will always happen at steady state but the actual length of the queues will fluctuate during the simulation.

Let p be the occupation density at the beginning of the road and q the occupation density at the end of the road, we have shown that both quantities form the following relation with π [2]:

$$J = q\pi = p = 1 - q, \quad (3)$$

from which we concluded that $\pi = p/q$ and $p = \frac{\pi}{1+\pi} = 1 - q$.

3 Decentralized flow estimation

Using eq. (3) we can reconstruct the value of π knowing the value of p and q . Thus, we can estimate the flow of each road segment, if we can estimate the values of p and q . In this section, we present a way to estimate those quantities by a decentralized exchange of information between neighboring cars.

3.1 Global map

We call *global map*, the real traffic map that each car will try to estimate. According to our discussion, we can restrict this map to a few cells in the network, those which are enough to predict the traffic flow in each road segment. Here, the map is composed of the measured densities at each road end, $\hat{\rho}_2$ and $\hat{\rho}_{L-1}$. Those densities are computed as the average occupation of the cell during the last 128 iterations:

$$\hat{\rho}_i(t) = \frac{1}{128} \sum_{k=t-127}^t s_i(k). \quad (4)$$

The use of a “sliding time window” (of size 128 in our case) to compute the density of a cell, allows us to forget the remote past and to get a dynamic view at traffic evolutions. The size of 128 was chosen for practical implementation reasons, and provides good results as shown in the next section. If the window size is too small, fluctuations will dominate time-stable structures. If it is too large, temporary (but relevant) structures will become undetectable.

3.2 Local Map

Every car in the model maintains an estimated map of traffic density, termed *local map*. This map is updated at each iteration by direct observations and communications with other cars. Cars can exchange information with neighboring cars. Here, during a single iteration, each car can communicate with at most 5 other cars, the leading and trailing ones, and the three cars in the opposite lane.

For a street map with $2L$ -cell road segments, the local map m_c of a car c at iteration t is a matrix with 128 rows and $2L$ columns. The elements of the map defined as

$$m_c(i, j) = \begin{cases} 1 & \text{if a car was recorded by } c \text{ at location } j \text{ at time } t - i + 1, \\ 0 & \text{else.} \end{cases} \quad (5)$$

An estimate of the density of each location can be computed by averaging the corresponding column. We chose to track position along each road in order to study the dynamic of the whole local maps. However, the method presented here is able to reconstruct the traffic condition with just two density values at beginning ($i = 2$) and at the end of ($i = L - 1$) of each road segment. The local map could then be restricted to two columns per road segment.

Local map update All cars start with a local map filled with zeroes. At each iteration, they update the map according to the following steps: **(1)** The local map matrix rows are shifted down, eliminating the oldest row. The first row is filled with zeroes. **(2)** Each car adds 1 to the element of the first row corresponding to its current location. **(3)** Each car sends its map to each neighboring car and receives maps from them, merging the information into its own local map with an element-wise logical **OR**.

It is easy to see that the logical **OR** of the first 128 lines of the local maps of all the cars in our circuit will produce the global map. Indeed, as all cars only put their own position at each step of the simulation in their local maps, the logical **OR** of all the local maps will give the occupation of each of the cars in the last 128 iterations. The question is whether our car to car communications will be fast enough to provide an accurate **OR** of all the local maps, and thus the global map.

4 Simulation Results

To investigate how a better communication scheme could improve our method, we have performed several simulations by varying the communication radius and the communication frequency. Each trial parameter is represented by the notation $NnXx$ where n is the neighborhood radius (in cell units) and x is the communication frequency in number of communications per traffic iteration. For instance, N3X2 means that we use a three-cell wide neighborhood (corresponding to a communication radius of 30 meters) and we perform two communication iterations per traffic update (corresponding to a frequency of 1.7 updates per second).

In this study, both parameters were varied from 1 to 5, using a two-road system of 200 cells ($L = 100$) with a car density of 0.5 and $\pi = 0.3$. Maps were limited to the last 128 iterations. To quantify the impact of the communication radius and frequency, we measured the estimation error average among all cars. The estimation density error was computed as:

$$err_\rho = \frac{1}{2nL} \sum_{i=1}^n \sum_{j=i}^{2L} |\hat{\rho}_i(j) - \rho_G(j)|, \quad (6)$$

where n is the number of cars, $\hat{\rho}_i(j)$ the density estimation of car i at location j and $\rho_G(j)$ the value computed from the global map for the same location (exact value). Results are presented in Fig. 4 (left). We can see that both communication parameters reduce the error when increased, although the communication radius has a more important effect. An increase of the latter will also strongly reduce the error standard deviation, meaning that local maps reach a strong consensus.

We repeated the same experiment by measuring the π estimation error:

$$err_\pi = \frac{1}{n\pi_G} \sum_{i=1}^n |\hat{\pi}_i - \pi_G|, \quad (7)$$

where n is the number of cars, $\hat{\pi}_i$ the estimation of car i and π_G the value computed from the global map (exact value). Results shown in Fig. 4 (right) are fully consistent with the plots of density estimation error discussed above.

To provide a closer view on the effects of both communication parameters, we did focus on a car picked at random. Using a record of all cars time/position during a single simulation, we replayed exactly the same traffic evolution, changing

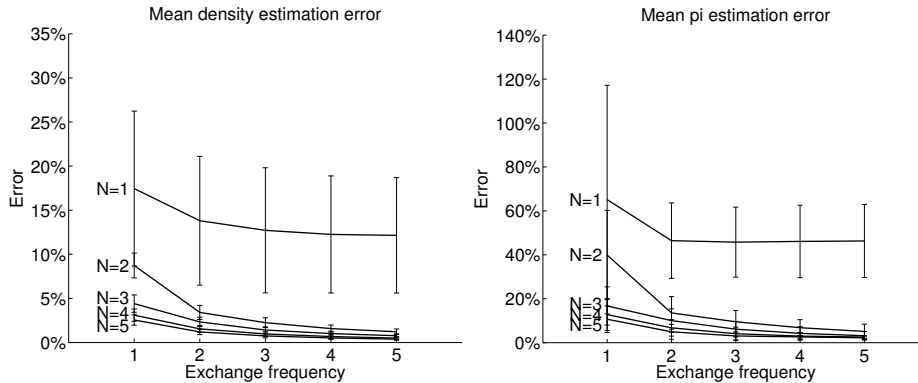


Fig. 4. Relative error averaged over all the cars for 25 traffic simulations in a two-road circuit with $L = 100$. Left: Relative density error; Right: Relative π error. The error bars represent the standard deviation.

only the communication scheme (this is possible because in our toy circuit cars behave independently of the map content). The π estimation by the reference car is shown in Fig. 5 for the N1X1, N1X5, N5X1, N5X5 and N2X2 communication settings. Again, increasing the communication parameters clearly improves the π estimations. Moreover, we can note that when the communication radius is low, the reference car is prone to two kinds of errors: (i) the low density area (p) is often underestimated leading to an infinite value of π when the estimations reach zero. This effect could be easily reduced by adding a pseudocount to each map column. (ii) Temporarily good estimates of π tend to “drift down” until they suddenly improve again. This means that a car tends to forget good estimations because of a lack of interactions with other cars. The sudden rises observed can be explained by encounters with a car on the opposite lane which carries the missing information.

Finally, we present how errors are distributed in the local map of the reference car at the end of each simulation. Since a map time/space location is considered empty until a car records its presence, a local map can only underestimate street density but never overestimate it. We show how the underestimates are distributed by taking the logical XOR operation between the global map and a local map. Every resulting dot will be a missed car occupancy in a given position at a given time. Local map error diagrams of the reference car are displayed in Fig. 6 for several communication settings. Since the information travels at a constant maximum speed depending directly on the communications parameters, “light-cones” appear on the error diagrams. They encompass a space-time zone where the local information had time to be exchanged with neighbors and thus corresponds perfectly to the global map. In each diagram the cones appear to originate from the location of the reference car at the end of the simulation. Further analysis of the diagram is required, but we can clearly see that increasing the communication parameters increases the cones aperture. In the case of

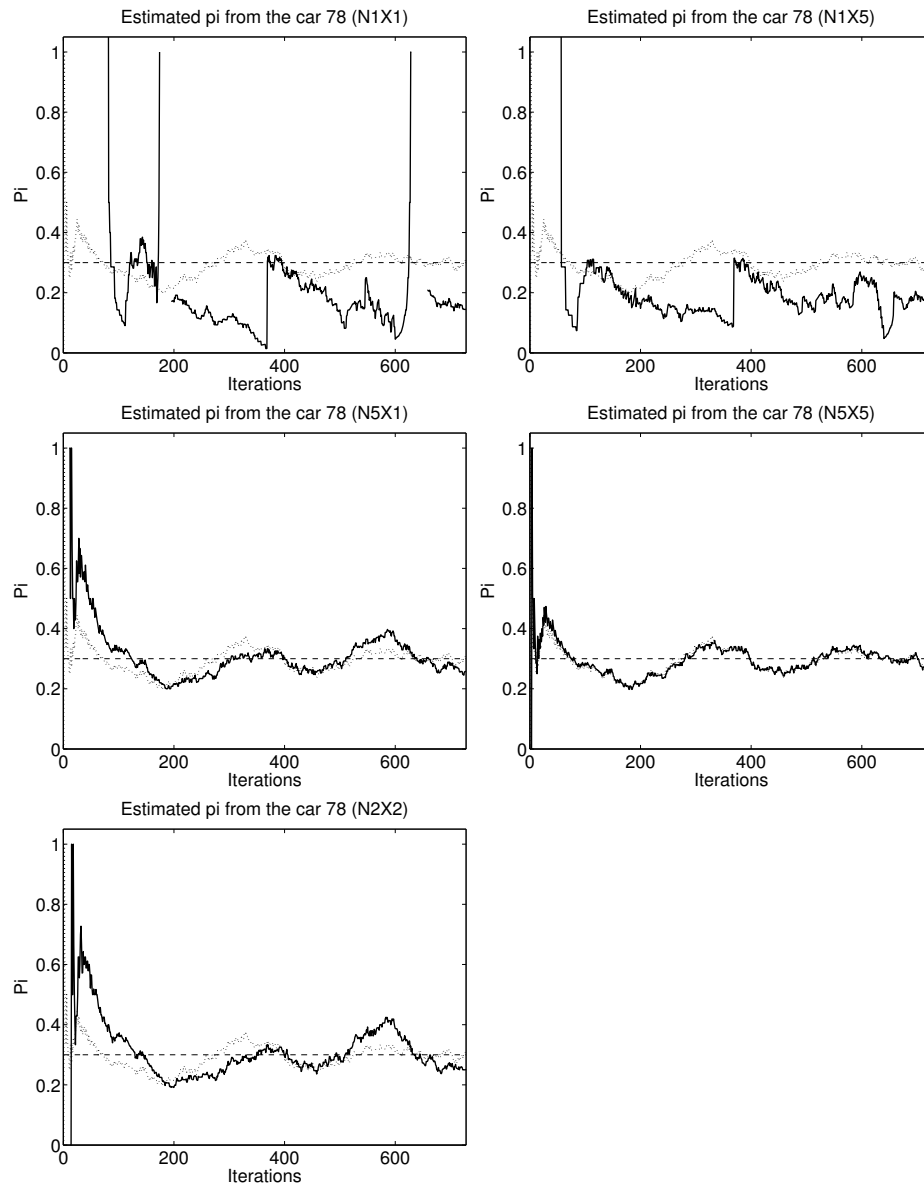


Fig. 5. Estimation of π according to a reference car picked at random. A single traffic simulation took place and was used to build five different simulations of map updates with different settings. The solid black curves show the estimation of π from the reference car local map, the dotted curves are the estimated π using the global map (correct empirical value), and the horizontal dashed lines are the imposed value of π .

N5X5, the cone is wide enough to reach the system boundaries after a couple of iterations. Even with N2X2, the local map is perfect after less than 28 iterations.

5 Discussion

We have proposed a decentralized method allowing the cars traversing a road network to build an estimate of the traffic situation, by exchanging local and anonymous information. We shown the traffic conditions in a simple circuit can be described with only the knowledge of the car densities at the extremities of the road segments. Within this framework, we have tested the capability of the cars to predict the global traffic condition.¹ The method is robust enough to detect flow variations caused by traffic accidents² [2].

We have also shown that the traffic estimation errors can be strongly reduced by allowing a larger communication radius and a higher frequency of exchange. In particular increasing the radius allows cars in free flow traffic to exchange information. Indeed, in this traffic regime found in low density areas, cars are always separated by at least a blank cell. In roads with a density of $\rho = 0.5$ and evenly spaced cars, no car will be able to exchange information with a N1X x communication scheme. The estimated density will then be very small (equal to 1 over the map temporal size). By taking at least N2X x , cars will be able to exchange information and merge their maps to get a good estimate. We should note that increasing only the communication frequency will not help in this particular situation because although the map exchanges will happen faster, they will not be able to communicate across gaps between cars.

The method is very economic in terms of used memory, since only 2×128 bits per road segment are stored. A local map for the whole Borough of Manhattan in New-York city could then be stored in less than 150 kB (11 avenues, 220 streets, and two lanes per road segment give us $11 \times 220 \times 2 \times 2 \times 128$ bits). The computational cost is also very low: (i) most of the operations are element-wise and thus may run in parallel; (ii) bitwise **OR** operations are executed as a single instruction in all current processors. If we use a communication scheme of N2X2, we need to exchange information at a rate of 0.45 updates per second which corresponds for the whole Manhattan map to a bandwidth of 334 kB/s.

We have also tested the method using a more realistic street map and promising results are described in [1].

References

1. Sartoretti, G.: Communication method for decentralized estimation of traffic density. Technical report, University of Geneva (2012) Master Thesis.
2. Sartoretti G., Falcone J.-L., B. Chopard, Gander M. J: Decentralized Method for Traffic Monitoring. In: G. Sirakoulis S. Bandini (Eds.), Cellular Automata - 10th International Conference on Automata for Research and Industry, Springer Verlag (2012), 464-473.

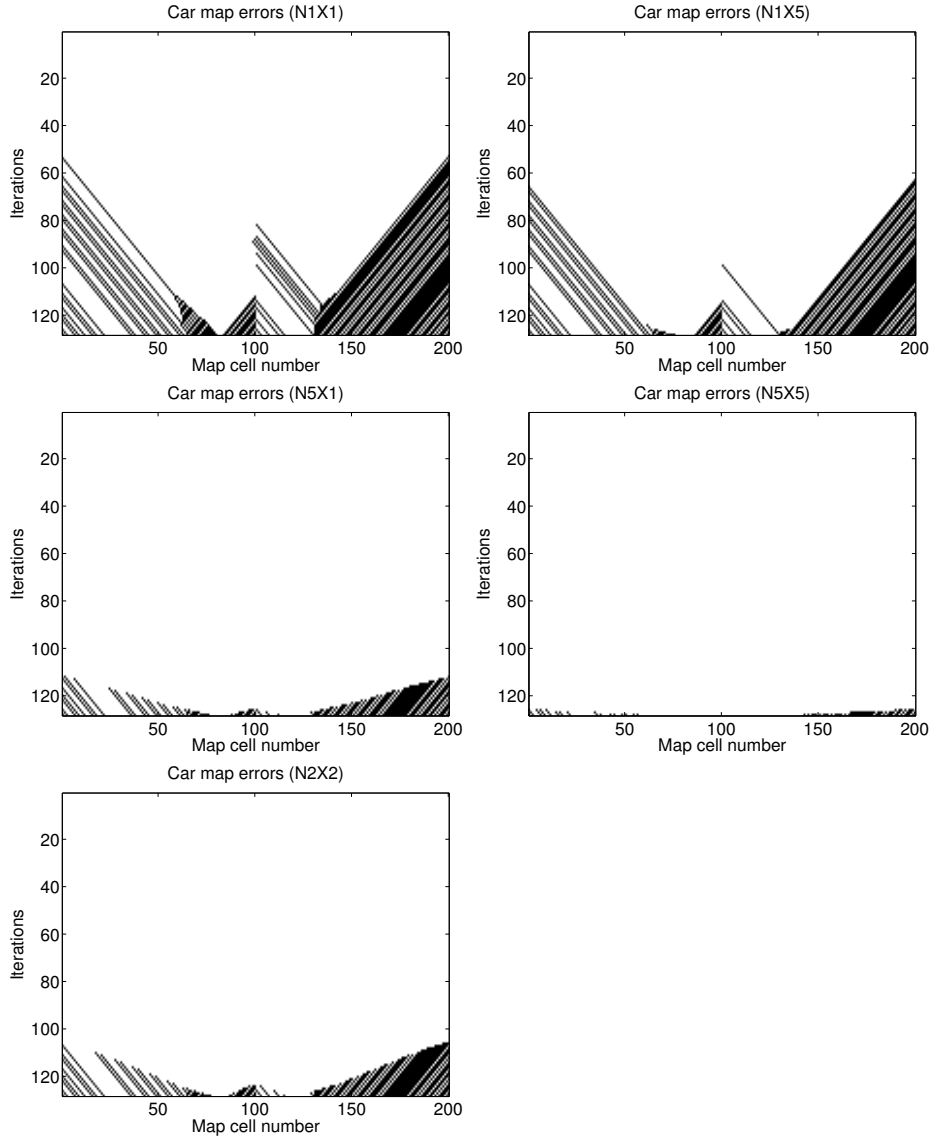


Fig. 6. Errors in the local map of a single car at the end of a simulation. Each diagram corresponds to the five settings presented in Fig. 5. For each case, the local map of the reference car at the end of the simulation was compared to the global map using a logical XOR operation. Every black pixel corresponds to an error of the map for a given location at a given time. Since only underestimate errors are possible, they represent a missing car observation.

3. Chopard, B., Luthi, P.O., Quelo, P.A.: Cellular automata model of car traffic in a two-dimensional street network. *Journal of Physics A: Mathematical and General* **29**(10) (1996) 2325
4. Dijkstra, J., Timmermans, H.J.P., Jessurun, A.J.: A multi-agent cellular automata system for visualising simulated pedestrian activity. In: in S. Bandini and T. Worsch (Eds.), *Theoretical and Practical Issues on Cellular Automata - Proceedings on the 4th International Conference on Cellular Automata for research and Industry*, Springer Verlag (2000) 29–36
5. Dupuis, A., Chopard, B.: Cellular automata simulations of traffic: A model for the city of geneva. *Networks and Spatial Economics* **3** (2003) 9–21 [10.1023/A:1022044932736](https://doi.org/10.1023/A:1022044932736).
6. Nagel, K., Schreckenberg, M.: A cellular automaton model for freeway traffic. *J. Phys. I France* **2**(12) (1992) 221–2229 <http://dx.doi.org/10.1051/jp1:1992277>.
7. Nishinari, K., Fukui, M., Schadschneider, A.: A stochastic cellular automaton model for traffic flow with multiple metastable states. *Journal of Physics A: Mathematical and General* **37**(9) (2004) 3101
8. Wolfram, S.: *Theory and Application of Cellular Automata*. World Scientific, Singapore (August 1986)