Katholieke Universiteit Leuven

Departement Elektrotechniek ESAT-SCD (SISTA) / TR 06-030

Data Quality, Travel Time Estimation, and Reliability

Sven Maerivoet[†] and Bart De Moor[†]

March 2006
Technical report

This report is available by anonymous ftp from *ftp.esat.kuleuven.be* in the directory *pub/sista/smaerivo/reports/paper-06-030.pdf*

[†]Katholieke Universiteit Leuven
Department of Electrical Engineering ESAT-SCD (SISTA)
Kasteelpark Arenberg 10, 3001 Leuven, Belgium
Phone: (+32) (0) 16 32 86 64 Fax: (+32) (0) 16 32 19 70
E-mail: {*sven.maerivoet,bart.demoor*}*@esat.kuleuven.be*
WWW: *http://www.esat.kuleuven.be/scd*
Our research is supported by: **Research Council KUL**: GOA AMBioRICS, several PhD/postdoc & fellow grants,
Flemish Government: **FWO**: PhD/postdoc grants, projects, G.0407.02 (support vector machines), G.0197.02 (power islands), G.0141.03 (identification and cryptography), G.0491.03 (control for intensive care glycemia), G.0120.03 (QIT), G.0452.04 (new quantum algorithms), G.0499.04 (statistics), G.0211.05 (Nonlinear), research communities (ICCoS, ANMMM, MLDM),
IWT: PhD Grants, GBOU (McKnow),
Belgian Federal Science Policy Office: IUAP P5/22 ('Dynamical Systems and Control: Computation, Identification and Modelling', 2002-2006), PODO-II (CP/40: TMS and Sustainability),
EU: FP5-Quprodis, ERNSI,
Contract Research/agreements: ISMC/IPCOS, Data4s,TML, Elia, LMS, Mastercard.

Please use the following BibTEX entry when referring to this document:

```
@techreport{MAERIVOET:06,

    author      =  "Sven Maerivoet and Bart De Moor",
    title       =  "Data Quality, Travel Time Estimation, and Reliability",
    institution =  "Katholieke Universiteit Leuven"
    number      =  "06-030"
    year        =  "2006",
    month       =  mar

}
```

Data Quality, Travel Time Estimation, and Reliability

Sven Maerivoet* and Bart De Moor
Department of Electrical Engineering ESAT-SCD (SISTA)†,
Katholieke Universiteit Leuven
Kasteelpark Arenberg 10, 3001 Leuven, Belgium

(Dated: March 1, 2006)

In this paper, we turn our attention towards what is called *exploratory data analysis* (EDA) of all traffic flow measurements gathered on Flanders' motorways. In a first section, we describe how all these measurements are obtained by detectors either embedded in the concrete, or by cameras positioned alongside the road, how they are stored in a central database, and how we can query this database, e.g., to give a visualisation of weekly patterns. We then discuss the quality of the measurements, from a statistical point of view. To this end, we give a technique that tracks outliers and give some pointers for dealing with missing values. Subsequently, we provide a methodology for quickly assessing structural and incidental detector malfunctioning; this is done by creating maps that give a clear visual indication of when and where the problems occurred. The final section of this paper provides a methodology for the off-line estimation of travel times, based on flow measurements (as opposed to the much used technique based on speed measurements). To conclude, we provide some reliability and robustness properties related to travel times and traffic flow dynamics, which gives us an extra instrument for the analysis of recurrent congestion.

PACS numbers: 02.50.-r,89.40.-a

Keywords: traffic flow theory, exploratory data analysis, travel time estimation

Contents

I. Acquisition of traffic flow measurements 2
A. Aggregation procedures 2
1. Single inductive loop detectors 2
2. Cameras 4
3. Operational characteristics of single inductive loop detectors 4
4. Some remarks on speed estimation techniques 7
B. Storage of the measurements in a central database 7
C. Visualising weekly patterns 8

II. Quality assessment of the measurements 11
A. Comparing estimations of mean speeds 12
B. Measurement errors and outlier detection 14
1. Outliers in a statistical sense 14
2. Explanation of the methodology 14
3. Discussion of the results 15
4. Dealing with missing values 16
C. Assessing detector malfunctioning 19
1. Explanation of the methodology 20
2. Illustrative detector maps 20

III. Off-line travel time estimation and reliability indicators 27
A. Common approaches towards travel time estimation 27
B. Estimating travel times based on flow measurements 28
1. Constructing the cumulative curves 28
2. Dealing with synchronisation issues and systematic errors 28
3. Estimating the distribution of the travel time 31
C. Indicators of reliability 31
1. Overview of the case study area 32
2. Travel time reliability 32

† Phone: +32 (0) 16 32 17 09 Fax: +32 (0) 16 32 19 70
URL: http://www.esat.kuleuven.be/scd
*Electronic address: sven.maerivoet@esat.kuleuven.be

3. Constructing reliability maps 33

IV. Conclusions 41

Acknowledgements 42

References 43

In this paper, we turn our attention towards what is called *exploratory data analysis* (EDA) of all traffic flow measurements gathered on Flanders' motorways. In a first section, we describe how all these measurements are obtained by detectors either embedded in the concrete, or by cameras positioned alongside the road, how they are stored in a central database, and how we can query this database, e.g., to give a visualisation of weekly patterns.

We then discuss the quality of the measurements, from a statistical point of view. To this end, we give a technique that tracks outliers and give some pointers for dealing with missing values. Subsequently, we provide a methodology for quickly assessing structural and incidental detector malfunctioning; this is done by creating maps that give a clear visual indication of when and where the problems occurred.

The final section of this paper provides a methodology for the off-line estimation of travel times, based on flow measurements (as opposed to the much used technique based on speed measurements). To conclude, we provide some reliability and robustness properties related to travel times and traffic flow dynamics, which gives us an extra instrument for the analysis of recurrent congestion.

I. ACQUISITION OF TRAFFIC FLOW MEASUREMENTS

Since the last decade, a tremendous amount of traffic data is being gathered by detectors in Flanders' motorway road network (see Table I for a summary of the total length of all roads in Flanders); this data stems from over 1600 sensors in total (see Fig. 1) [66]. Until now, only data collection has been performed, but recently the Flemish government expressed interest in analysing this data. More specifically, due to the presumed high level of data corruption, it becomes worthwhile to perform quality assessments of the available data and provide corrections if possible. This will allow the Flemish motorway operating agency to use the detector data for fine tuning certain control measures pertaining to optimal flows and incident avoidance, as well as on- and off-line travel time prediction and the assessment of network reliability.

Province	*Motorways*	*On-/Off-ramps*	*Normal roads*	*Total*
Antwerpen	230 km	98 km	970 km	1298 km
Oost-Vlaanderen	203 km	80 km	1030 km	1313 km
West-Vlaanderen	187 km	94 km	1286 km	1567 km
Vlaams-Brabant	194 km	104 km	604 km	902 km
Limburg	102 km	44 km	1062 km	1208 km
Total	916 km	420 km	4952 km	6288 km

TABLE I: The total length of all roads in Flanders, the Dutch-speaking northern part of Belgium (information cited from [2]).

A. Aggregation procedures

In Belgium, there are mainly two types of detectors employed: single inductive loop detectors (SLD) embedded in the concrete and cameras positioned above or alongside the road. In the following two sections, we briefly discuss each of these devices. The third section and fourth section deal with the operational characteristics of the single inductive loop detectors employed, and some remarks on the traditional way of estimating speeds.

1. Single inductive loop detectors

These are inductive loops of copper wire embedded in the concrete, typically in a rectangular setup (see the left part of Fig. 3); they create an induced electromagnetic field that changes whenever a vehicle passes over the loop. Comparing the change in a loop's total inductance against a calibrated treshold, allows the associated controller logic (which energises the physical loop with a periodic signal) to count vehicles each time the current settles again to its stationary state. Counting the number of successive pulses corresponds to a vehicle's on-time (see the left part of Fig. 2) [58]. In Belgium, the SLDs are provided by the company *Macq électronique*[71].

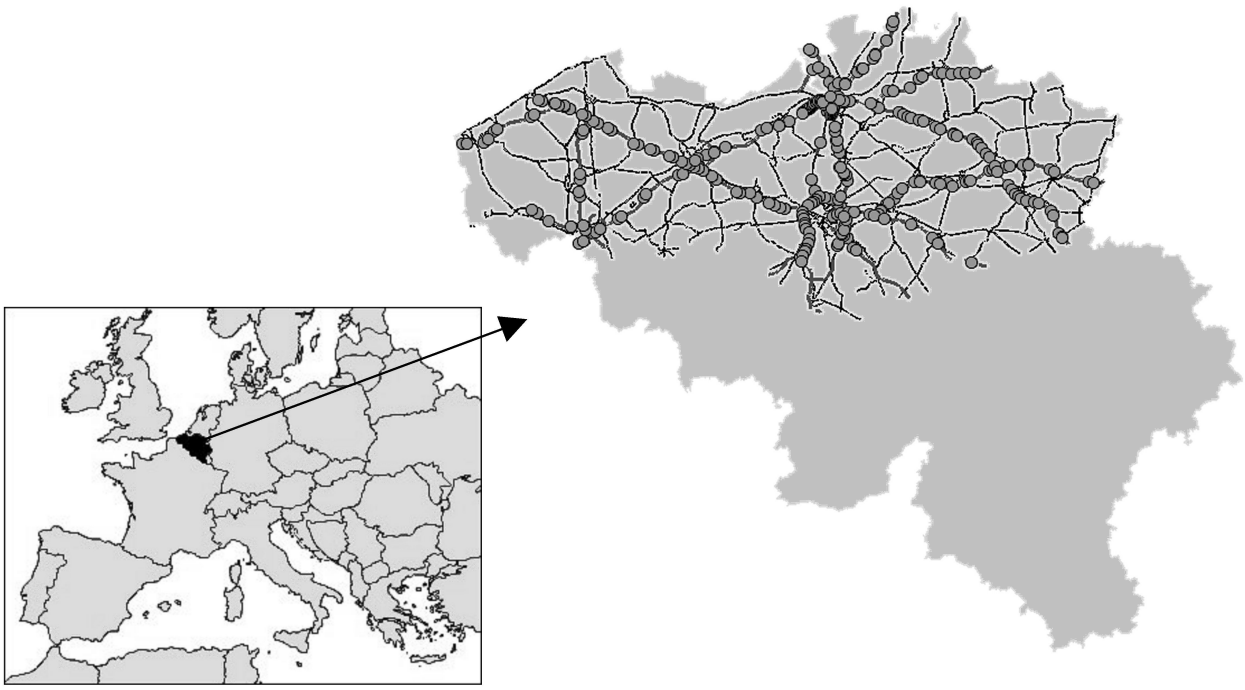

FIG. 1: Flanders' motorway road network and its underlying secondary network of national roads, located in the northern part of Belgium. All motorways are equipped with more than 1600 sensors in total, as indicated by the locations of the gray circles (mostly single inductive loop detectors and some cameras), each minute measuring local flows, occupancies and time-mean speeds (for all lanes separately).

An SLD is sometimes called a *presence-type detector*, and is correspondingly only able to measure flows and occupancies. In order to get a reliable estimation of a vehicle's speed, a double inductive loop detector (DLD), consisting of two closely spaced single inductive loop detectors, can be used. The vehicle's speed is computed based on the distance between both loops and the time needed for the vehicle to travel this distance. As such, these DLDs are also called *speed traps*. Typical dimensions for an SLD are a width of 1.8 metres (i.e., half the width of a typical lane in Belgium and The Netherlands), with a length of 1.5 metres. The width assures that a typical vehicle can not avoid a detector when changing lanes. The length is taken large enough such that and that a small truck is considered as a single vehicle, and at the same time it is assumed to be small enough such that the individual vehicles are still counted under congested conditions. Double inductive loop detectors are spaced 1 metre apart [8]. Each loop detector is connected to a circuit board that contains the controller logic which processes the changes in the coils' inductances as vehicles drive by (see the right part of Fig. 2).

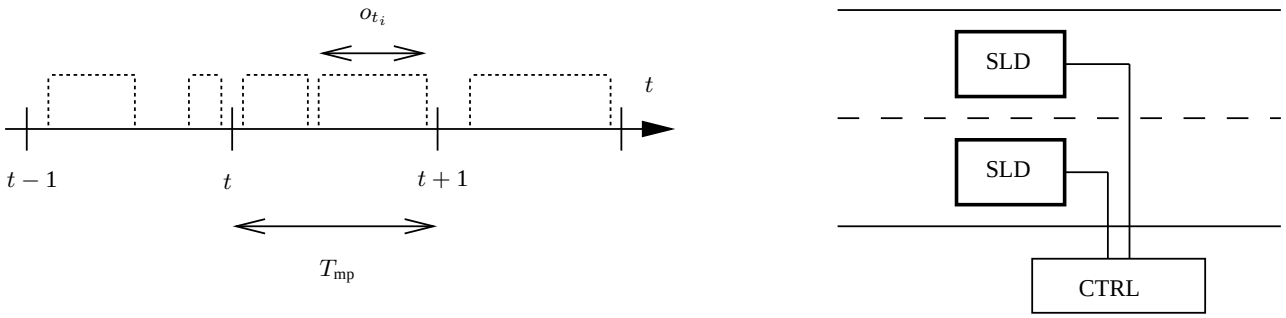

FIG. 2: *Left:* each time a vehicle i passes over the area of a single inductive loop detector, the controller logic records the vehicle's on-time o_{t_i} as a consecutive number of pulses. The detector aggregates these on-times during measurement periods of length T_{mp}. *Right:* two single inductive loop detectors (marked as SLD) embedded in a road; the coils are connected to the controller logic (marked as CTRL) which processes the changes in the coils' inductances as vehicles drive by.

2. Cameras

These are mounted above or alongside the road and record all traffic that drives over a certain section of the road (see the middle part of Fig. 3). As vehicles pass by, the image processing algorithms embedded in the camera's software detect and count these vehicles in real-time. Cameras are able to easily outperform inductive loop detectors in terms of quality of the measurements (which is of course dependent on the capability of the software to deal with varying road and weather conditions). In Belgium, there are some 200 cameras in use and all of them (as well as their accompanying software) are provided by *Traficon*[72]. Traficon essentially provides a detector board that contains a *video image processor* (VIP); this processor detects vehicles that cross lines that are superposed on the camera's picture (see the right part of Fig. 3).

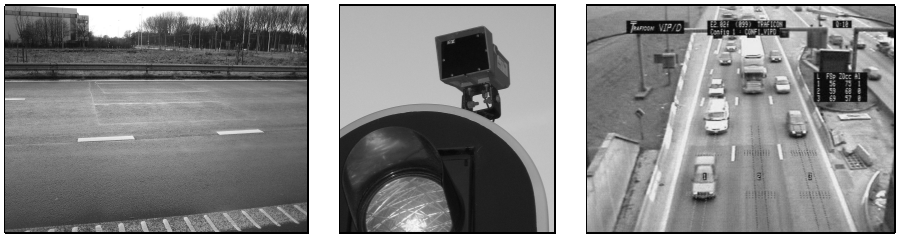

FIG. 3: Some images of traffic detectors typically encountered in the Belgian road network. *Left:* two single inductive loop detectors embedded in the concrete. *Middle:* a traffic camera mounted on top of a traffic light. *Right:* an image sequence of a camera that is processed by a Traficon video image processor to extract local traffic data.

Considering the spatial and temporal measurement regions, we note that an SLD corresponds to region R_t, a DLD corresponds to two such successive regions, whereas a camera corresponds to region $R_{t,s}$ (put more correctly, it resembles sequences of R_s regions that correspond to the video's individual frames) [39].

Other possible detectors are pneumatic tubes which detect changes in pressure, detectors based on infrared beams, radar devices using the Doppler effect, ... Different detection schemes require different installation and maintenance costs. Nowadays, the Belgian government chooses to replace faulty single inductive loop detectors with cameras, as these latter can quickly be installed without having to completely block one or more lanes of the road.

3. Operational characteristics of single inductive loop detectors

Most of the detectors are located right before and after each complex of on- and off-ramps at motorways. This clearly gives a very sparse spatial distribution because there are many kilometres of road compared to the kilometres spanned by these complexes. In total, the number of detectors present in Flanders' motorway network amounts to 1654 for the year 2001 and 1800 for the year 2003: each detector is responsible for a single lane that can be located on the main road of the motorway, or on an on- or off-ramp. Measurements for each detector are aggregated every minute, i.e., $T_{\text{mp}} = 60$ s. Noting that traffic flows vary in time and space, the operations of the detectors can be seen as the stochastic sampling of these traffic flows. So we should always keep in mind that the obtained measurements are not absolute values, but samples from a stochastic distribution.

There are four *macroscopic* variables[73] that each detector i in the motorway network outputs after the elapse of each measurement period t:

- $q_{c_i}(t)$, the number of cars driving by,
- $q_{t_i}(t)$, the number of trucks driving by,
- $\rho_i(t)$, the occupancy of the detector,
- and $\overline{v}_{t_i}(t)$, the time-mean speed of all vehicles driving by.

It is important to realise that an SLD is not capable of measuring the speed of a single vehicle. This stems from the fundamental fact that the measurements are taken at a single point in space (i.e., measurement region R_t). And without knowing a vehicle's length, its speed can not be derived. So either the length or the speed can be calculated (provided one of the two is known), but not both. As such, only $q_{c_i}(t)$, $q_{t_i}(t)$, and $\rho_i(t)$ are measured directly; an estimate of the time-mean speed is $\overline{v}_{t_i}(t)$ is derived from these values. The detectors operate with a resolution of 50 Hz, so each $1 \div 50 = 0.02$ s the detectors record pulses due to the changing current in the loop. All detectors within one complex are connected to a counting station that contains a microprocessor to handle the signals of at most 20 SLDs (one such station

controls 4 groups of at most 5 SLDs). We now explain the operation of a single inductive loop detector i that is installed in the Belgian motorway network since 1980 [57, 58].

1. All vehicles are assumed to have the same mean speed $\overline{v}_{t_i}(t-1)$, calculated during the previous measurement period.

2. When the jth vehicle passes over the loop, its on-time o_{t_j} is recorded as a number of pulses (see the left part of Fig. 2 for a schematic overview). As such, this on-time corresponds to an integer multiple of the sampling period, i.e., 2 ms.

3. The jth vehicle is then classified as being either a car or a truck, based on its recorded on-time o_{t_j} and a *treshold* $\tau_i(t-1)$ that was calculated during the *previous* measurement period:

$$\begin{aligned} \text{if} \;\; o_{t_j} \leq \tau_i(t-1) \;\; &\Longrightarrow \;\; \text{car} \\ \text{else} \;\; o_{t_j} > \tau_i(t-1) \;\; &\Longrightarrow \;\; \text{truck} \end{aligned}$$

   This treshold essentially is the 'trick' that the counting stations for SLDs use to discriminate between cars and trucks. Because no vehicle lengths and speeds are known (they can not be measured by an SLD), the vehicles' on-times are used in comparison with a dynamic treshold, in order to obtain a vehicle type classification.

4. At the end of the measurement period $t + T_{\text{mp}}$, the detector's controller logic has determined the number of cars $q_{c_i}(t)$ and trucks $q_{t_i}(t)$ counted, and the occupancy $\rho_i(t) = \frac{1}{T_{\text{mp}}} \sum_{i=1}^{N} o_{t_i}$ [39]. Furthermore, it has also calculated the *average on-time for a car*:

$$\overline{o}_{t_{c_i}}(t) = \frac{1}{q_{c_i}(t)} \sum_{j=1}^{q_{c_i}(t)} o_{t_j}. \tag{1}$$

5. Based on the average on-time for a car, the controller logic then determines the dynamic treshold, using a so-called *control curve*, as shown in Fig. 4. This new treshold is then to be used during the *next* measurement period. As such, there is a lag of 1 minute before the measurements can adapt to changing traffic conditions (which change the treshold).

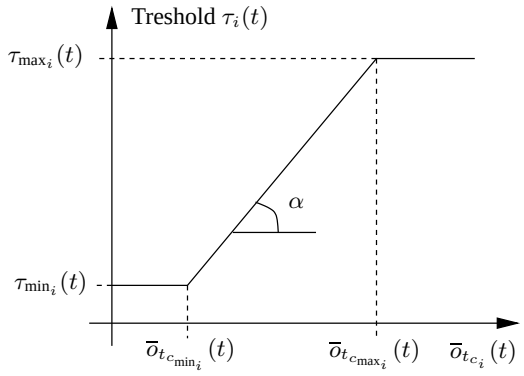

FIG. 4: An illustration of the control curve used to calculate the dynamic treshold $\tau_i(t)$ for SLD i at measurement period t. The curve is assumed to have lower and upper boundaries, as well as a linear part that relates the treshold to the average on-time $\overline{o}_{t_{c_i}}(t)$ for a car.

The parameters $\overline{o}_{t_{c_{\min_i}}}(t)$, $\overline{o}_{t_{c_{\max_i}}}(t)$, $\tau_{\min_i}(t)$, $\tau_{\max_i}(t)$, and the slope α of the linear function, are received from a central computer that sends these values every measurement period T_{mp} to *all* the single loop detectors in the motorway network. In Belgium, their values are set by an operator at respectively 9, 72, 18 and 144 (expressed in pulses of 2 ms, i.e., multiples of the sampling period). For a typical car length of 4.5 m, the values 9 and 72 correspond to speeds of approximately 120 km/h and 15 km/h, respectively. An important consequence of this, is that the detector's logic can not detect speeds below 15 km/h. The slope α is taken equal to 2, giving the following equation for the linear function:

$$\tau_i(t) = 2\,\overline{o}_{t_{c_i}}(t). \tag{2}$$

As indicated, the calculation of the new treshold is purely based on the number of cars, as it is assumed that the majority of the vehicles are cars, and that their individual lengths are more or less constant with an average of $\overline{l}_c$ = 4.5 m. In case the number of counted cars $q_{c_i}(t)$ is strictly less than a predefined lower bound (in Belgium, this bound is set at 6 cars), then the calculation of a new treshold value (as shown by equation (2) and in Fig. 4) is omitted and the previous value is maintained, i.e., $\tau_i(t) = \tau_i(t-1)$.

6. Once the car and truck counts and the detector's occupancy are known, only one variable is missing in order to calculate the time-mean speed $\overline{v}_{t_i}(t)$: the *average vehicle length* $\overline{l}(t)$ needs to be known. The effective vehicle length, *as seen by the detector*, is actually the sum of the vehicle's length l_j, and the length K_{ld} of the loop detector. This corresponds to the following equation:

$$l_j + K_{\text{ld}} = o_{t_j}\ v_j, \tag{3}$$

with as previously stated, $K_{\text{ld}} \approx 1.5$ m for SLDs in Belgium.

Just as the treshold $\tau_i(t)$ for the on-time is used for the classification of cars and trucks, we can look at the equivalent treshold $\lambda_i(t)$ related to the vehicle length:

$$\lambda_i(t) + K_{\text{ld}} = \tau_i(t)\ \overline{v}_{t_i}(t). \tag{4}$$

Multiplying both sides of equation (2) with the mean speed $\overline{v}_{t_i}(t)$ and applying equation (4), yields:

$$\lambda_i(t) + K_{\text{ld}} = \alpha\ \overline{o}_{t_{c_i}}(t)\ \overline{v}_{t_i}(t), \tag{5}$$

which is by equation (3) equivalent to:

$$\lambda_i(t) + K_{\text{ld}} = \alpha\ (\overline{l}_c + K_{\text{ld}}), \tag{6}$$

or:

$$\lambda_i(t) = 2\ \overline{l}_c + 1.5\text{ m}. \tag{7}$$

In other words: the calculated treshold for the classification based on the perceived vehicle length corresponds to trucks which have a minimum length $\overline{l}_t$ of 2 × 4.5 m + 1.5 m = 10.5 m. These values can now be used to determine the average vehicle length $\overline{l}(t)$, which is a mixture of the proportions of the counted cars and trucks, and is expressed as the following weighted average (with the flows now expressed as hourly counts):

$$\overline{l}(t) = \frac{(q_{c_i}(t)\ \overline{l}_c) + (q_{t_i}(t)\ \overline{l}_t)}{q_{c_i}(t) + q_{t_i}(t)}. \tag{8}$$

7. The final step now estimates the time-mean speed of the vehicles. It is assumed that individual vehicle lengths and speeds are uncorrelated, and that all vehicles passing the SLD during one minute have the same speed. Applying the relation between occupancy, flow, and density as expressed by $\rho = \overline{l}\ k \Longrightarrow k = \frac{\rho}{\overline{l}}$, to the fundamental relation of traffic flow theory as expressed by $q = k\ \overline{v}_s$, results in the following *estimation* for the time-mean speed[74]:

$$\overline{v}_{t_i}(t) = \frac{\cancel{(q_{c_i}(t) + q_{t_i}(t))}\ \dfrac{(q_{c_i}(t)\ \overline{l}_c) + (q_{t_i}(t)\ \overline{l}_t)}{\cancel{(q_{c_i}(t) + q_{t_i}(t))}}}{\rho_i(t)} \tag{9}$$

4. Some remarks on speed estimation techniques

As mentioned in the previous section, the estimation of the mean speed is based on an assumed average vehicle length. The inverse of this length is called the *g-factor*, which converts occupancy to density [27, 49]:

$$\text{space-mean speed} = \frac{\text{flow}}{\text{occupancy} \times g}. \tag{10}$$

It is now possible to tune the SLD's processor by estimating this g-factor, so it can be used for the calculation of the mean speed. The algorithm elaborated upon in section I A 3 assumes constant average vehicle lengths, which implies the use of a fixed g-factor. However, it is considered bad practice to use a constant for this critical parameter, e.g., setting the mean speed fixed during free-flow conditions and estimating the g-factor, while using this fixed g-factor during congested conditions (for these latter, the fleetmix, e.g., the percentage of long vehicles, plays an important role) [14, 16, 36]. The use of a constant g-factor can lead to flawed results, as examined by Mikhalkin et al. [41], Hall and Persaud [27], and Pushkar et al. [1, 49]; Dailey addresses this problem by explicitly taking into account the statistical nature of the measurements, thereby providing criteria that help to evaluate their reliability [20, 21]. Other possible approaches are those elaborated by Coifman et al., who provide better estimations for the average vehicle length and the speed, e.g., by tuning it with estimations coming from double loop detectors [14], or try to estimate the median speed instead of the mean speed at SLDs [16]. It is important to realise that flawed measurements due to a wrong g-factor can and will lead to faulty predictions of travel times based on the measured speeds [32]. In light of the algorithm elaborated in the previous section, Kwon et al. describe a methodology for the real-time estimation of the portion of truck traffic on motorways, based on data from single inductive loop detectors; they assume the existence of a truck-free lane and a high lane-to-lane speed correlation [36].

Note that the operation of the single inductive loop detectors as described above, has a negative side effect: when the average length $\overline{l}(t)$ of a vehicle is calculated, it is done using information collected during the *previous* minute. This leads to the fact that the implemented algorithm shows incorrect behaviour when the mean speed fluctuates abruptly, as the newly calculated treshold has a lag of one minute. This leads to an overestimation of the number of trucks counted when the speed suddenly drops. This can be seen in Fig. 5, where the upper part shows the mean speed and the lower part shows the percentage of detected trucks in the traffic stream: we observe a strong correlation between these two at times when the speed drops to very low values (e.g., in congestion periods). Currently, the only way to resolve this problem is to post-process the data (excluding the application of real-time corrections), as for example elaborated in the work of De Ceuster and Immers, who recalculate the average vehicle length using an exponentially weighted moving average, taken over a longer averaging period that one minute [11]. Regardless of the problem here indicated, it can be assumed that the total vehicle count, i.e., $q_{c_i}(t) + q_{t_i}(t)$, can be considered as the most reliable measurement from a single inductive loop detector.

B. Storage of the measurements in a central database

As already mentioned in the introduction of this section, there are over 1600 sensors located in Flanders' motorway road network (see Fig. 1); they are mostly single inductive loop detectors, with some 200 Traficon cameras (each sensor accounts for one lane). All these sensors are grouped into *measurement posts*; these posts group sensors at a single location over all lanes in the same driving direction. They are typically located right before and right after motorway on-ramps, off-ramps, merges, and diverges. A collection of measurements posts is called a *measurement complex* [66].

Several front-end computers query these sensor complexes each minute, after which the measurements are relayed to a central computer in Flanders' traffic centre (located in Antwerp). This central processor creates exchange files that get sent to the traffic centres in the Flemish and Walloon[75] regions in Belgium. These files have a lifetime of one hour, after which they get overwritten; at regular intervals they are stored into a central database that is kept at the traffic centre. As such, this database (called *MINDAT*) contains raw, unprocessed, and unvalidated data (note that no distinction is made between measurements coming from single inductive loop detectors and cameras) [66]. With respect to the ranges of all stored measurements, we note that $q_{c_i}(t), q_{t_i}(t), \rho_i(t) \in \{0, \ldots, 127\}$, and $\overline{v}_{t_i}(t) \in \{0 \ldots, 255\}$. As an exception, a value of 127 for either the car flow, truck flow, or occupancy measurement is used as a sentinel value that is automatically placed in the database in case of a transmission failure. When such a failure occurs at the level of a measurement complex, all results stemming from the sensors corresponding to the measurement posts will have this value.

From the above description, we know that each detector collects a measurement quadruple each minute. For one year, this corresponds to 60 minutes/hour $\times$ 24 minutes/day $\times$ 365 days/year, i.e., 525,600 measurements. Given the fact that each

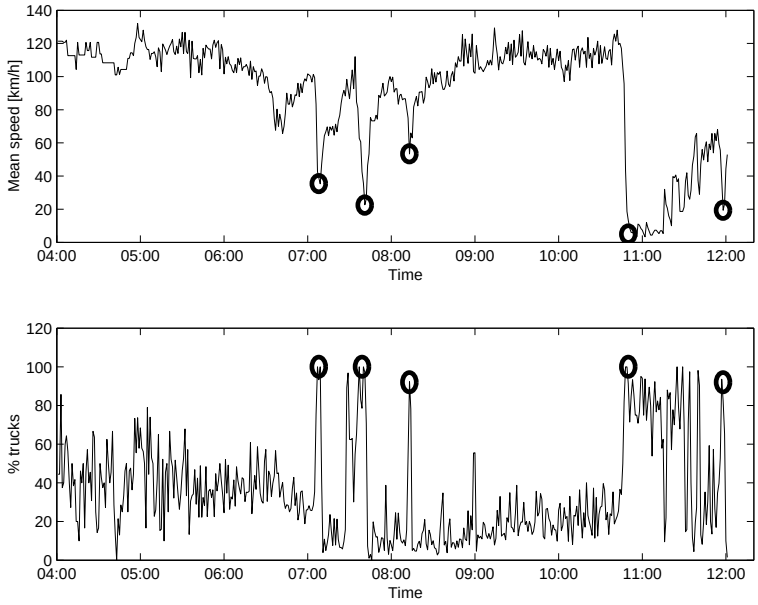

FIG. 5: The algorithm implemented in Flanders' single inductive loop detectors shows incorrect behaviour: each time the mean speed suddenly drops, the percentage of detected trucks in the traffic stream increases (data taken from loop detectors 810 – 812 at Ternat near Brussels, during 12 september 2001 between 04:00 and 12:00).

quadruple comprises 4 bytes, this corresponds to 2,102,400 bytes, or 2.01 MB. In the database, there are 1654 (for the year 2001), respectively 1800 (for the year 2003) sensors stored, resulting in a grand total of 869,342,400, respectively 946,080,000 measurements, corresponding to 3,477,369,600, respectively 3,784,320,000 bytes, or 3.24, respectively 3.52 GB. Compare this to another system, e.g., the famous *California Freeway Performance Measurement System* (PeMS), which has some 26,000 loop detectors, aggregating data at 30 second intervals into a database of 2 GB per day [64].

For our study, we were able to obtain a database containing all the measurements of the SLDs and camera's in Flanders' motorway network (as depicted in Fig. 1) for the years 2001 and 2003. It is interesting to note that it took nearly six months of extensive lobbying before the bureaucratic administration was able to handle and grant our request.

C. Visualising weekly patterns

Depending on the day of the week, the transportation demand will vary from location to location and time period to time period. In order to have a quick look at these phenomena, we have provided several three-dimensional charts in Fig. 6 and Fig. 7, plotting the total flows (cars plus trucks) for all days in 2001 and 2003, grouped together by day of the week. As such, we pooled together all *similar weekdays* for both years, each time resulting in seven different data sets that contain all Mondays, Tuesdays, Wednesdays, Thursdays, Fridays, Saturdays, and Sundays, with holidays filtered from them [13]. A blind-box approach would be to cluster these weekdays automatically based on the data itself, but in this case we know each day of the year *a priori*, so the clustering could be done manually[76].

In both figures, the flows during each day of a year are represented for all weekdays separately. For each of these days, the shown flows are an average taken over all detectors in Flanders' motorway road network; as such, they represent an average travel behaviour over the complete network. Looking at Fig. 6, we can see that Mondays, Tuesdays, Wednesdays, and Thursdays have very similar patterns. Fridays however are a bit different, in that the evening rush hour starts a earlier and has a broad peak. Considering the weekends, we note that Saturdays and Sundays are characterised by the absence of a morning rush hour. Furthermore, Saturdays have an evenly distributed afternoon peak, starting at approximately 10:00, lasting until 20:00; the peak for Sundays starts a bit later and is more intense towards the end. Comparing the weekly patterns of 2001 in Fig. 6 and those of 2003 in Fig. 7, we note that weekdays exhibit a slightly lower traffic demand, as opposed to the weekends where the traffic demand is more or less the same.

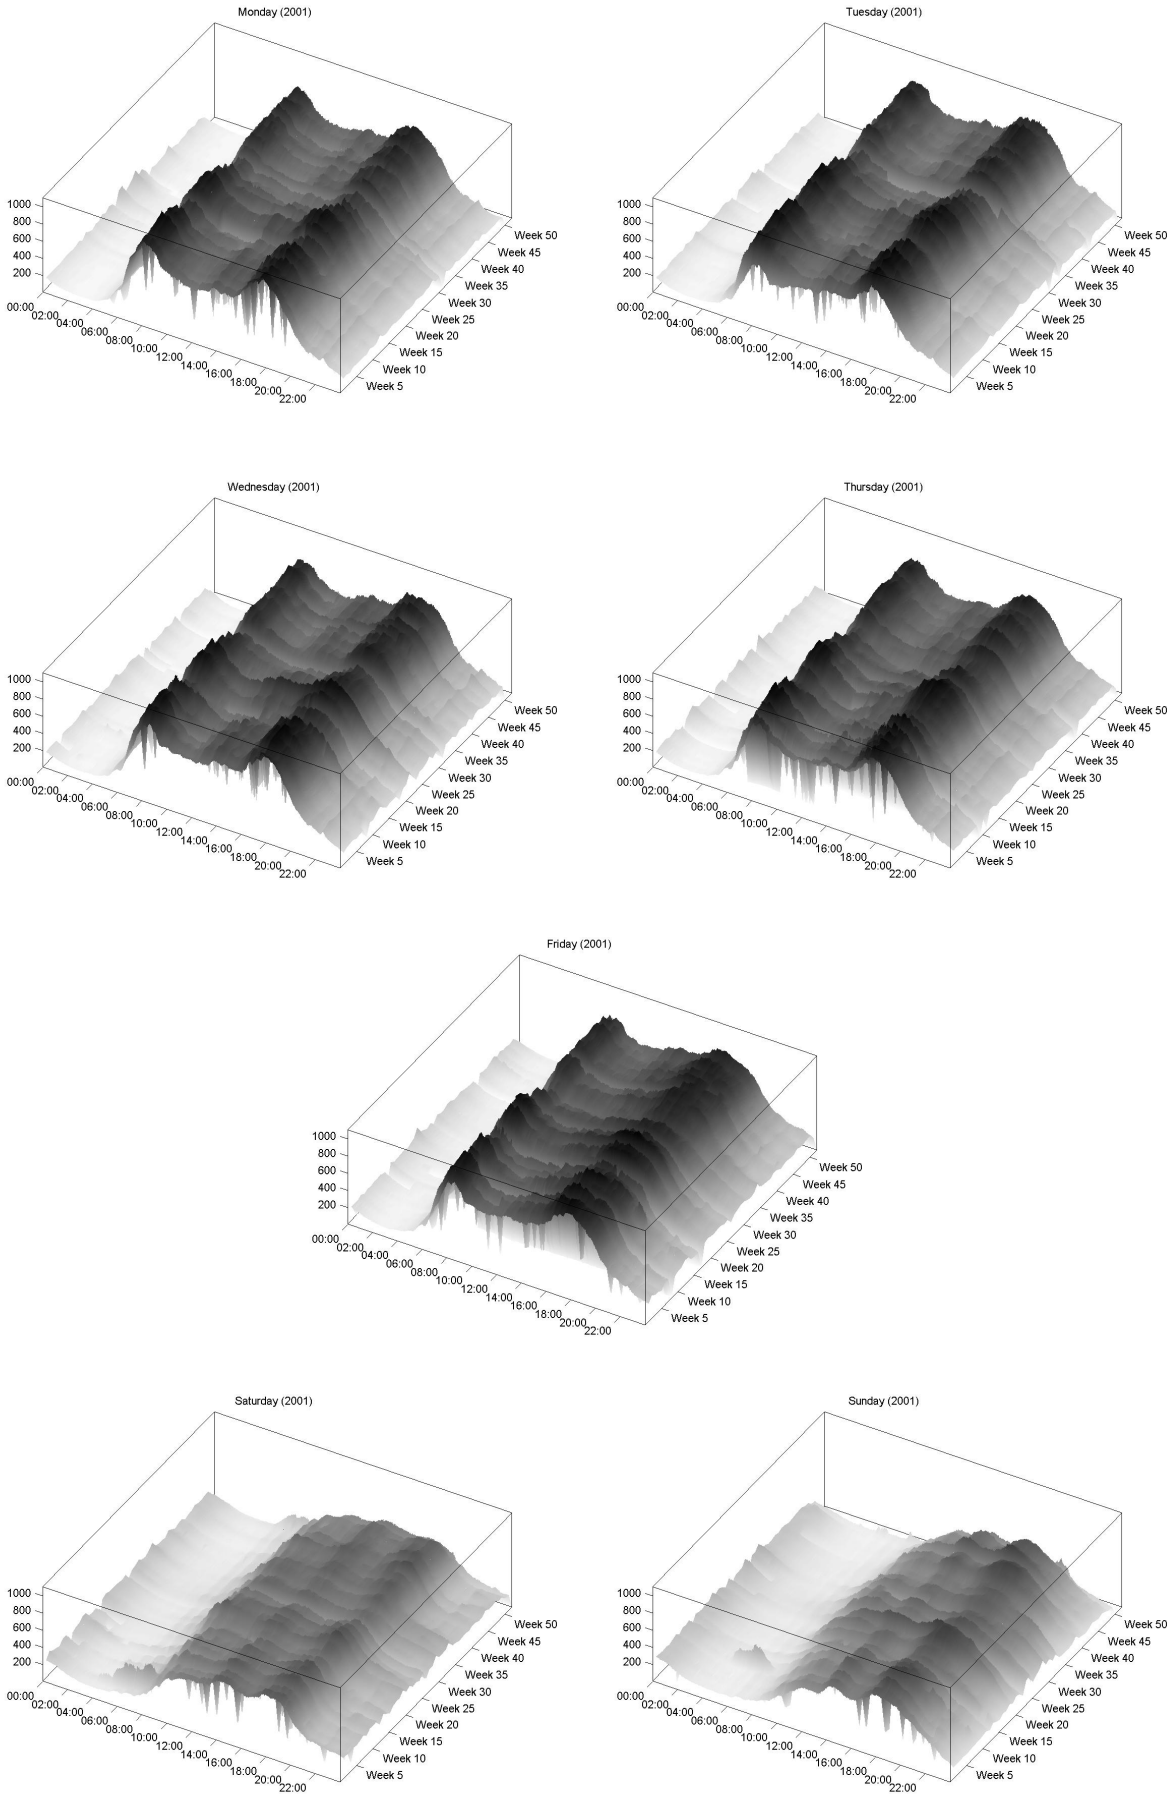

FIG. 6: Three-dimensional charts that show the total flows (cars plus trucks) for all days in 2001; each day of the week is represented separately, with holidays filtered from them. For each of these days, the shown flows are an average taken over all detectors in Flanders' motorway road network, represent an average travel behaviour over the complete network.

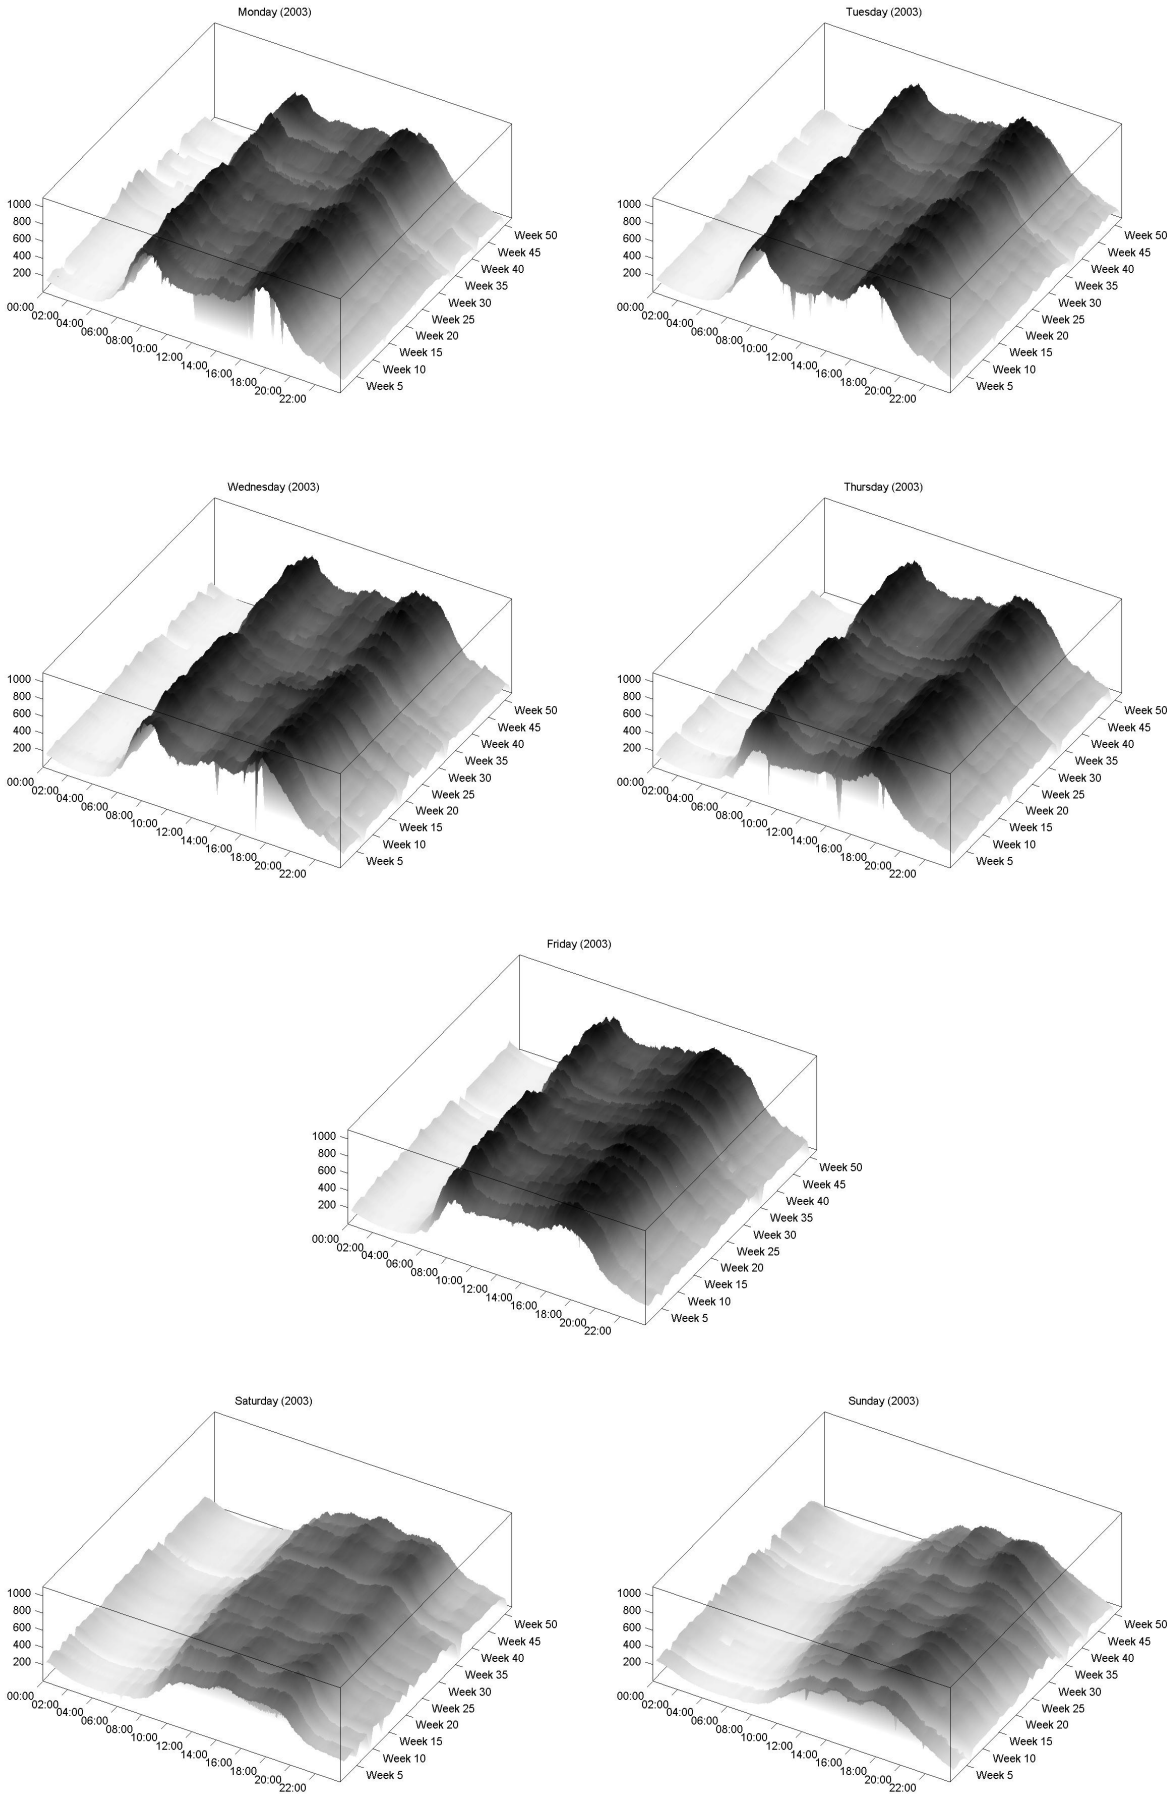

FIG. 7: Three-dimensional charts that show the total flows (cars plus trucks) for all days in 2003; each day of the week is represented separately, with holidays filtered from them. For each of these days, the shown flows are an average taken over all detectors in Flanders' motorway road network, represent an average travel behaviour over the complete network.

In order to more rigorously assess the differences between the different days of the week, and the evolution from 2001 to 2003, we have provided contourplots of the standard deviations of the flows as they are averaged over all existing detectors in Flanders' motorway road network. In Fig. 8 we show the results for 2001, whereas Fig. 9 shows the results for 2003; each time, we plot the standard deviations for Mondays, Fridays, Saturdays, and Sundays (because Tuesdays, Wednesdays, and Thursdays are quite similar to the Mondays). Note that the white 'streaks' in the images are due to missing values, as we are working with the raw data.

Looking at the information contained in the top left (Mondays) and right parts (Fridays) of Fig. 8 and Fig. 9, we can see that the standard deviation is the highest during the morning and evening rush hours; it can also be seen that the peak of the evening rush hour is more spreaded. Furthermore, from the darker regions it can be seen that the standard deviation for 2003 lies higher than the one for 2001, implying more diversity in the mean flows recorded by the different detectors.

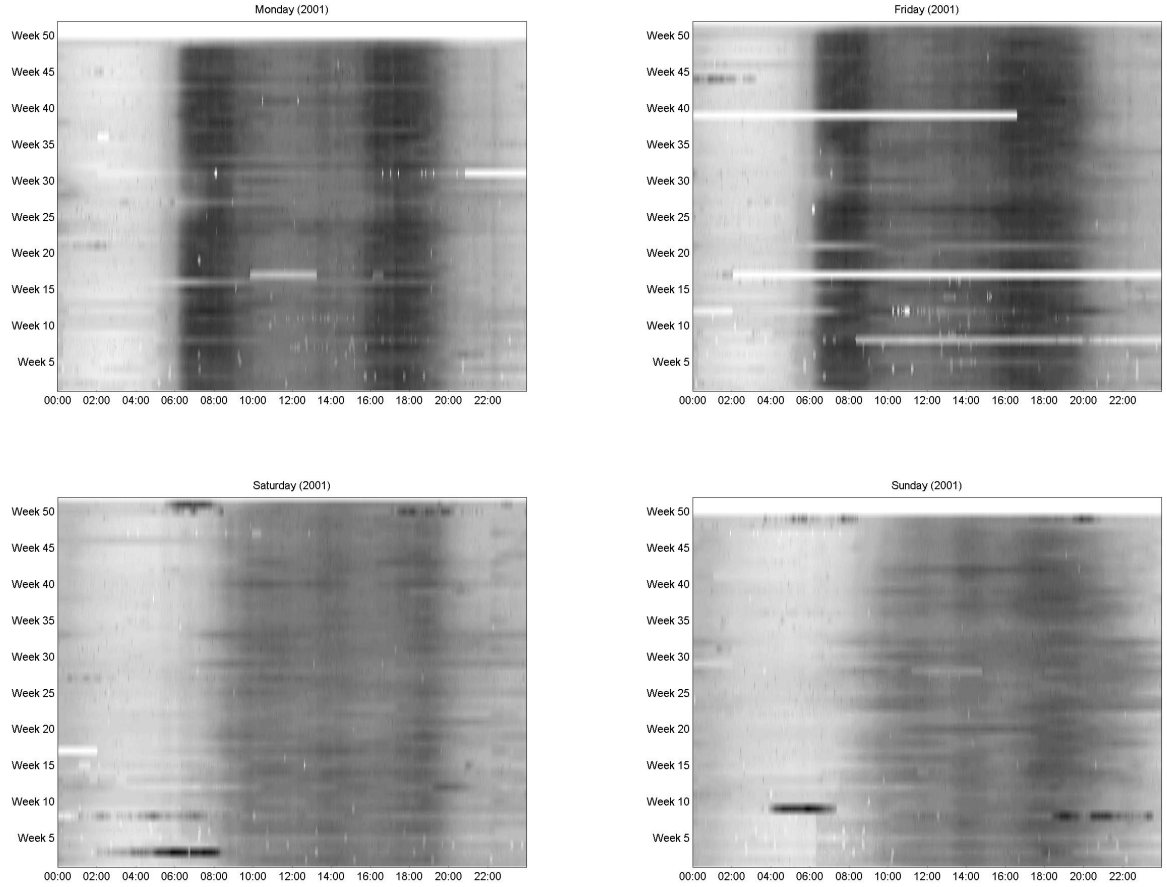

FIG. 8: Contourplots of the standard deviations of the flows for 2001 as they are averaged over all existing detectors in Flanders' motorway road network (see also Fig. 6). *Top-left:* Mondays. *Top-right:* Fridays. *Bottom-left:* Saturdays. *Bottom-right:* Sundays.

II. QUALITY ASSESSMENT OF THE MEASUREMENTS

With respect to quality of the measurements, it is well known that single inductive loop detectors are notorious for their errors. These errors are caused by factors such as external influences by a change in the environmental temperature, faulty calibrations, detector cross talk, chattering, transmission failures, ... Some of them cause high values for the flow to be reported, and in some cases a detector blanks completely resulting in no measurements at all. When transmission errors to the front-end computers occur (see section I B), typically the measurements of a whole detector station complex (grouping several SLDs) are lost, resulting in large gaps in the stored time series.

In this section, we first compare estimations of the mean speeds obtained from the algorithm explained in the previous section with those recorded by the detectors. We then take a look at the kinds of measurement errors that occur and the automatic detection of statistical outliers, after which we provide a methodology for quickly assessing area-wide detector malfunctioning.

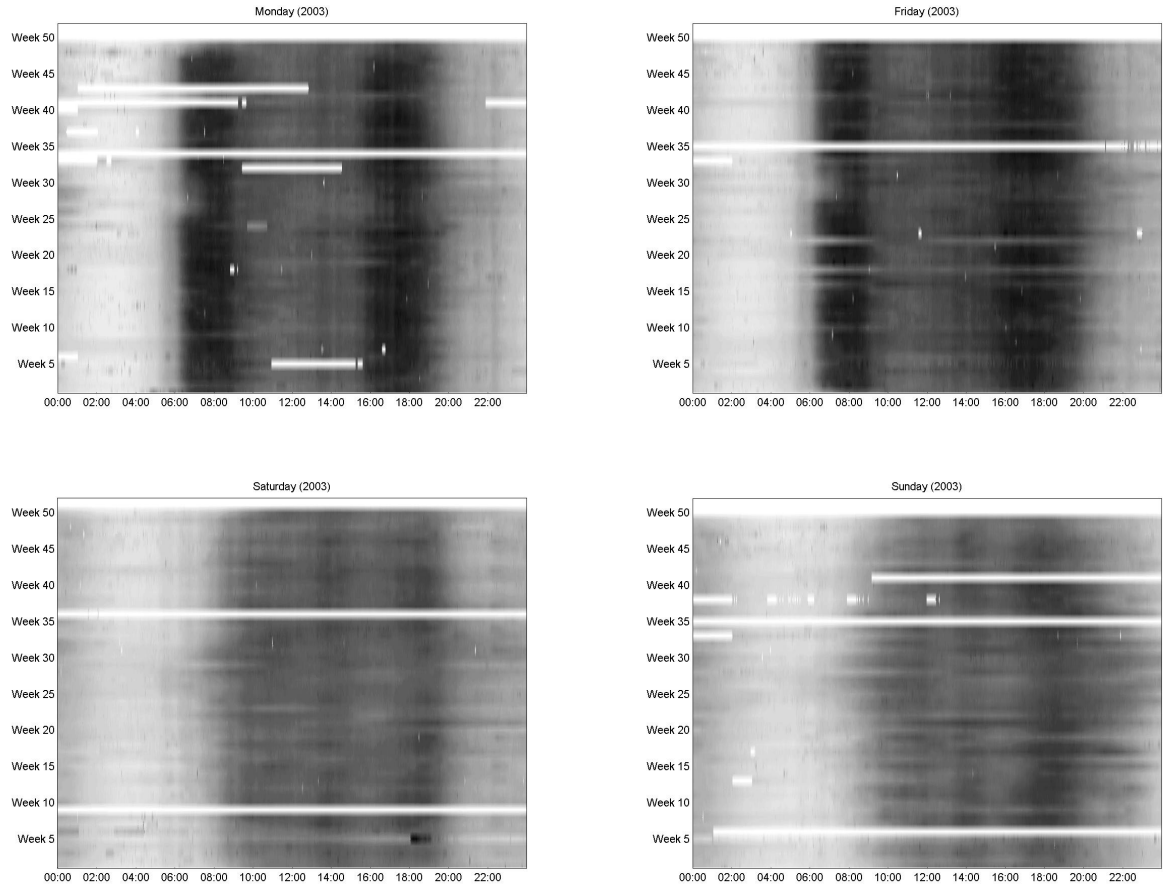

FIG. 9: Contourplots of the standard deviations of the flows for 2003 as they are averaged over all existing detectors in Flanders' motorway road network (see also Fig. 7). *Top-left:* Mondays. *Top-right:* Fridays. *Bottom-left:* Saturdays. *Bottom-right:* Sundays.

A. Comparing estimations of mean speeds

During our investigation of the measurements stored in the database, we uncovered to our surprise a significant discrepancy between the mean speeds as estimated by the single inductive loop detectors and those explicitly calculated by the algorithm elaborated in section I A 3; instead of what we expected, i.e., the same results, we obtained different estimations. In the remainder of this section, 'estimated mean speed' refers to the mean speed obtained by the SLD, whereas 'calculated mean speed' refers to the mean speed corresponding to the algorithm.

To illustrate this, we considered a sequence of SLDs (810 – 815), each of which belonged to a single measurement complex (at the E40 motorway near Ternat, two directions each consisting of three lanes). For each of these detectors, we calculated the mean speeds based on the algorithm from section I A 3, and compared them with those as estimated by the SLDs. The results are shown in the scatter plots in Fig. 10; all measurements were taken from the month november (30 days $\times$ 24 hours/days $\times$ 60 minutes/hour = 43,200 measurements), with black data points corresponding to the year 2001, and the gray data points to the year 2003.

Considering the results in these scatter plots, we can see that there is good agreement for low speeds, corresponding to either low flows or high occupancies (i.e., congested conditions). However, at moderately to high speeds (i.e., free-flow conditions), the discrepancy between estimated and calculated speeds starts to grow. Looking at the difference between the black and gray data points, it is clear that the estimations in 2003 differ significantly from those of 2001, indicating a possible recalibration; there is less scatter as the points are located in a more densely packed area. Furthermore, it would seem that the calculated speeds typically lie lower than the estimated speeds, especially in 2003.

In Fig. 11, we show the same type of scatter plot, but now for detector 668 (which actually is a camera called CLOF, i.e., an acronym for 'Camera Linkeroever'; F stands for the hexadecimal numbering scheme used, i.e., the 15th camera), located at the E17 Gent-Antwerpen near Kruibeke. The scatter plot in the left part of the figure exhibits the same type of behaviour for low and moderately high speeds as explained in the previous paragraph. There is however one more visible artefact: at a relatively low estimated mean speed of 40 km/h, there is cluster of black data points (highlighted by the thick black ellipse). For these points, the estimated mean speed is fixed whereas the calculated mean speed differs significantly.

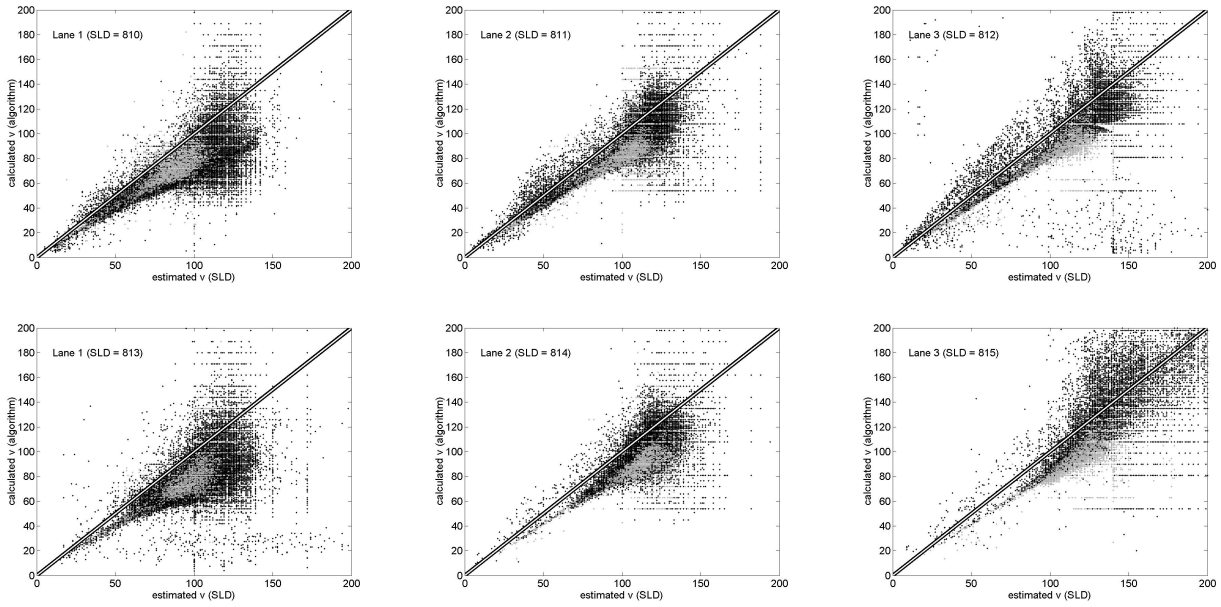

FIG. 10: Scatter plots showing the differences between mean speeds as estimated by the single inductive loop detectors and those calculated by the algoritm from section I A 3. The detectors belong to a single measurement complex (at the E40 motorway near Ternat, two directions each consisting of three lanes); all measurements were taken from the month november (30 days $\times$ 24 hours/days $\times$ 60 minutes/hour = 43,200 measurements), with black data points corresponding to the year 2001, and the gray data points to the year 2003 (the bissectrice is shown as the thick black/white line). As can be seen, there is good agreement for low speeds, but at moderately to high speeds the discrepancy between estimated and calculated speeds starts to grow. It is clear that the estimations in 2003 differ significantly from those of 2001, indicating a possible recalibration.

After explicit investigation of the time series corresponding to these data points, it seems that the estimated mean speed fluctuates smoothly around the fixed value of 40 km/h, whereas the algorithm of section I A 3 is better able to track the changes in occupancies (which carry more weight than the changes in car and truck flows).

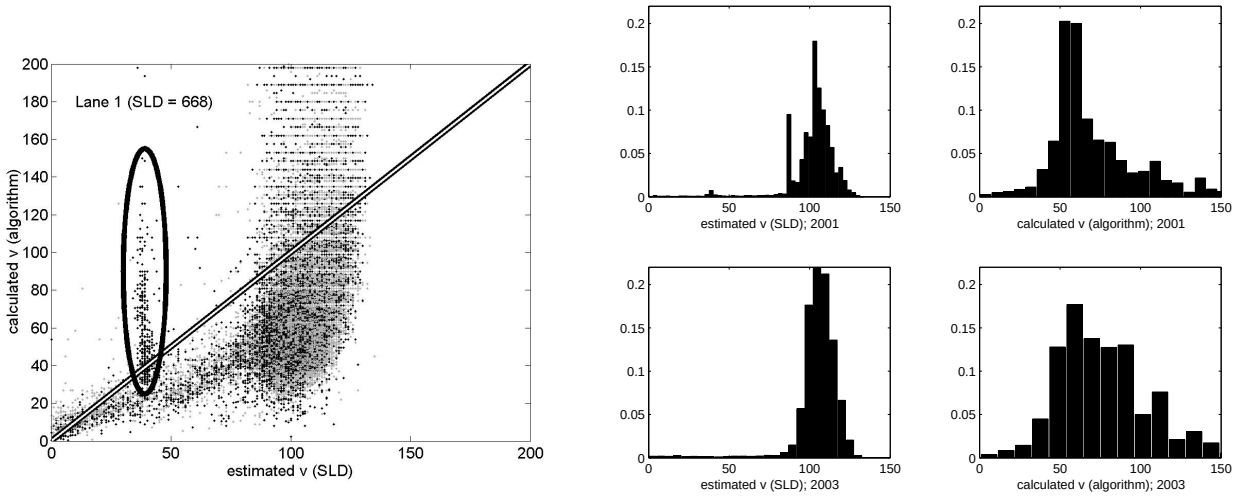

FIG. 11: *Left:* a scatter plot showing the differences between mean speeds as estimated by the single inductive loop detectors and those calculated by the algoritm from section I A 3. The shown detector is located at the E17 Gent-Antwerpen near Kruibeke; black data points correspond to the year 2001, the gray data points to the year 2003 (the bissectrice is shown as the thick black/white line). Note that at a relatively low estimated mean speed of 40 km/h, there is cluster of black data points (highlighted by the thick black ellipse); for these points, the estimated mean speed is fixed whereas the calculated mean speed differs significantly. *Right:* four histograms corresponding to the estimated and calculated mean speeds for 2001 and 2003. The former have a wider distribution than their latter counterparts.

For a more quantitative comparison, the right part of Fig. 11 provides histograms of the estimated and calculated mean speeds. It can immediately be seen that for 2001, the calculated mean speeds have a wider distribution than their estimated counterparts (see also Table II). As expected from the results of the scatter plots, the differences between all four

distributions are more pronounced for higher than for lower mean speeds.

	Est. $\overline{v}$ (2001)	*Calc. $\overline{v}$ (2001)*	*Est. $\overline{v}$ (2003)*	*Calc. $\overline{v}$ (2003)*
Mean	102	74	114	80
Std.dev	15	33	38	34

TABLE II: The means and standard deviations of the histograms from Fig. 11, corresponding to the estimated and calculated mean speeds. For 2001, the calculated mean speeds have a wider distribution than their estimated counterparts. Furthermore, the differences between all four distributions are more pronounced for higher than for lower mean speeds.

B. Measurement errors and outlier detection

Faulty measurements and the like are a plague for single inductive loop detectors; as such, we take a look at what causes these errors, giving an automatic detection of statistical outliers. We first describe what is meant by these kinds of outliers, after which we explain our methodology, discuss the results and provide some pointers for dealing with missing values in the data sets.

1. Outliers in a statistical sense

When considering faulty measurements from detectors, we can in general distinguish between structural failures versus occasional errors. The former can be due to a miscalibration, resulting in consistently faulty data (e.g., over- and underestimations of flows, detectors that get stuck in an on-/off-position, ...). Spotting and correcting these failures is not a difficult task (it requires e.g., a recalibration), in comparison with the latter class of occasional errors. These can have very different causes, such as detector cross talk, chattering, transmission failures, ... As a result, the detector logic can report incorrect data, for example, values that can easily be spotted are the sentinel values that get stored in the central database due to transmission failures.

Considering these 'strange values', we can look at them from a statistical perspective; as such, they are called *outliers*. From this point of view, *"outliers are observations that appear to be inconsistent with the remainder of the collected data"* according to Iglewicz and Hoaglin [29]. The phrase *"being inconsistent with the remainder"* can be given a more mathematical characterisation by taking into account the distributions of the measurements. Values that fall outside these distributions, or those that occur in the tails of them, can then be considered as outliers. Note that from a statistical point of view, outliers are not necessarily bad values, as it is possible that these data points might come from another population/distribution [65].

2. Explanation of the methodology

As mentioned at the end of section II B 1, we consider outliers to be values that are not conforming to the distribution of the measurements. In statistics, the process of automatically identifying outliers in univariate data is typically done based on the assumption that the measurements are *normally distributed*, with known mean and variance. The outliers are detected by comparing z-scores, which are measures that indicate how far a sample is located from the distribution's mean [52]:

$$z_i = \frac{x_i - \mu}{\sigma}, \tag{11}$$

with x_i a sample taken from a distribution with mean μ and standard deviation σ. Outliers are then samples for which the z-score (expressed in units of the standard devation) is greater than 3. Another method for assessing whether or not a sample is considered as an outlier, is by drawing a box-plot [52].

The above methods might seem fine, but they are insufficient when dealing multivariate data, consisting of n data points (observations) in p dimensions (variables): $\mathbf{x_i} = (x_{i1}, \ldots, x_{ip})$. Each of these observations can be stored as a row in a $n \times p$ matrix $X = (\mathbf{x_1}, \ldots, \mathbf{x_p})^{\mathrm{T}}$ with mean μ and covariance matrix Σ. So in order to tackle the problem of detecting outliers, we follow a similar methodology as with the z-score: for each point $\mathbf{x_i}$ in a multivariate data set, its so-called *Mahalanobis distance* (MD) is calculated [40, 52]:

$$\mathrm{MD}_i = \sqrt{(\mathbf{x_i} - \mu)^{\mathrm{T}}\, \Sigma^{-1}\, (\mathbf{x_i} - \mu)}. \tag{12}$$

However, outliers contaminating a data set can introduce a severe bias of the mean and variance. To take care of this problem, we use a robust estimator, called the *minimum covariance determinant* (MCD) estimator, for which a a computationally fast algorithm is available [50, 51]. Note that it is assumed that $n > 2p$, i.e., low-dimensional data. Although the Mahalanobis distance measure explicitly takes into account the correlations of the data set, it still exhibits the bias attributed to the classic mean and variance. To this end, we now replace the standard mean μ and covariance matrix Σ by their robustly-estimated counterparts $\hat{\mu}_{\mathrm{MCD}}$ and $\hat{\Sigma}_{\mathrm{MCD}}$. The resulting *robust distance* (RD) is thus written as follows:

$$\mathrm{RD}_i = \sqrt{(\mathbf{x_i} - \hat{\mu}_{\mathrm{MCD}})^{\mathrm{T}}\, \hat{\Sigma}_{\mathrm{MCD}}^{-1}\, (\mathbf{x_i} - \hat{\mu}_{\mathrm{MCD}})} \tag{13}$$

Detection of outliers is now based on comparing this distance against some specified treshold. Under the assumption that the data is normally distributed, the Mahalanobis distance is χ^2 distributed; thus,we say that an observation $\mathbf{x_i}$ is considered to be an outlier when its robust distance is greater than the specified treshold, i.e., $\mathrm{RD}_i \geq \sqrt{\chi^2_{p,0.975}}$ (corresponding to a significance level α =5%).

Applying this methodology to the traffic flow measurements, we consider a data set that is bivariate ($p = 2$) in nature: each data point consists of the occupancy and the total flow, i.e., $\mathbf{x_i} = (\rho_i, q_i)$ with $q_i = q_{c_i} + q_{t_i}$. Note that even though we are working with the raw data, no correction for the number of trucks is needed as the total count is the most reliable measurement an SLD can give (it is the classification that gives problems, as mentioned at the end of section I A 3). We used the following procedure to calculate the percentages of outliers in MATLAB:

1. For both years 2001 and 2003, we pooled together all *similar weekdays*, each time resulting in seven different data sets that contain all Mondays, Tuesdays, Wednesdays, Thursdays, Fridays, Saturdays, and Sundays, with holidays filtered from them [13].

2. Out of these seven data sets, we constructed seven bivariate data matrices containing the occupancies and total flows $\mathbf{x_i} = (\rho_i, q_i)$. In order to reduce the data size, we removed all duplicate data points. Knowing that the detectors' speed estimations are the least reliable measurements, we furthermore explicitly calculated the space-mean speeds based on the detector's recorded occupancy and flow measurements using equation (10) from section I A 4.

3. Because outliers in the free-flow traffic regime should not be compared to those in the congested traffic regime (due to the different distributions), we split all data sets into two non-overlapping parts. The criterion for discriminating between them was based on a combination of the free-flow speed and the critical occupancy; their respective tresholds were set at $\overline{v}_{\mathrm{ff}} = 85$ km/h and $\rho_c = 35\%$, respectively. Measurements below both tresholds were classified as being in the free-flow traffic regime, all other measurements are assumed to belong to the congested traffic regime. Note that some part of the synchronised flow traffic regime also belongs to our classification into a free-flow traffic regime, as this corresponds to a state of high flows at a relatively high speed [39].

4. For all these data sets, we now calculate a robust mean and covariance by means of the MCD estimator; to this end, we used the *Library for Robust Analysis* (LIBRA) [65]. Calculation of the MCD automatically gave us a classification for each data point as being either a regular observation or an outlier.

As an example of this methodology, we show the results for one detector in Fig. 12. The small dots denote measurements belonging to the free-flow traffic regime; the small crosses belong to the congested traffic regime. For this particular example, the data set consisted of 49 Mondays, with 44 and 451 unique points in the free-flow and congested traffic regimes, respectively. The thick solid and dashed ellipses denote the 97.5% tolerance boundaries for both regimes, based on the results of the MCD estimator. This means that for the selected significance level, there is an probability of 5% that any data points out of a large sample from a bivariate normal distribution, are misclassified as outliers outside the ellipse.

3. Discussion of the results

In order to interpret our results, we constructed illustrative gray-scale images. Because there are over 1500 detectors in each year present, and only seven weekdays, the resulting images are very thin. In order to increase the visual clarity of the images, we enlarged them vertically using a rescaling factor. The results for the year 2001 can be seen in Fig. 13, those for the year 2003 in Fig. 14; each time, the top part shows the percentages in the free-flow traffic regime, the middle part

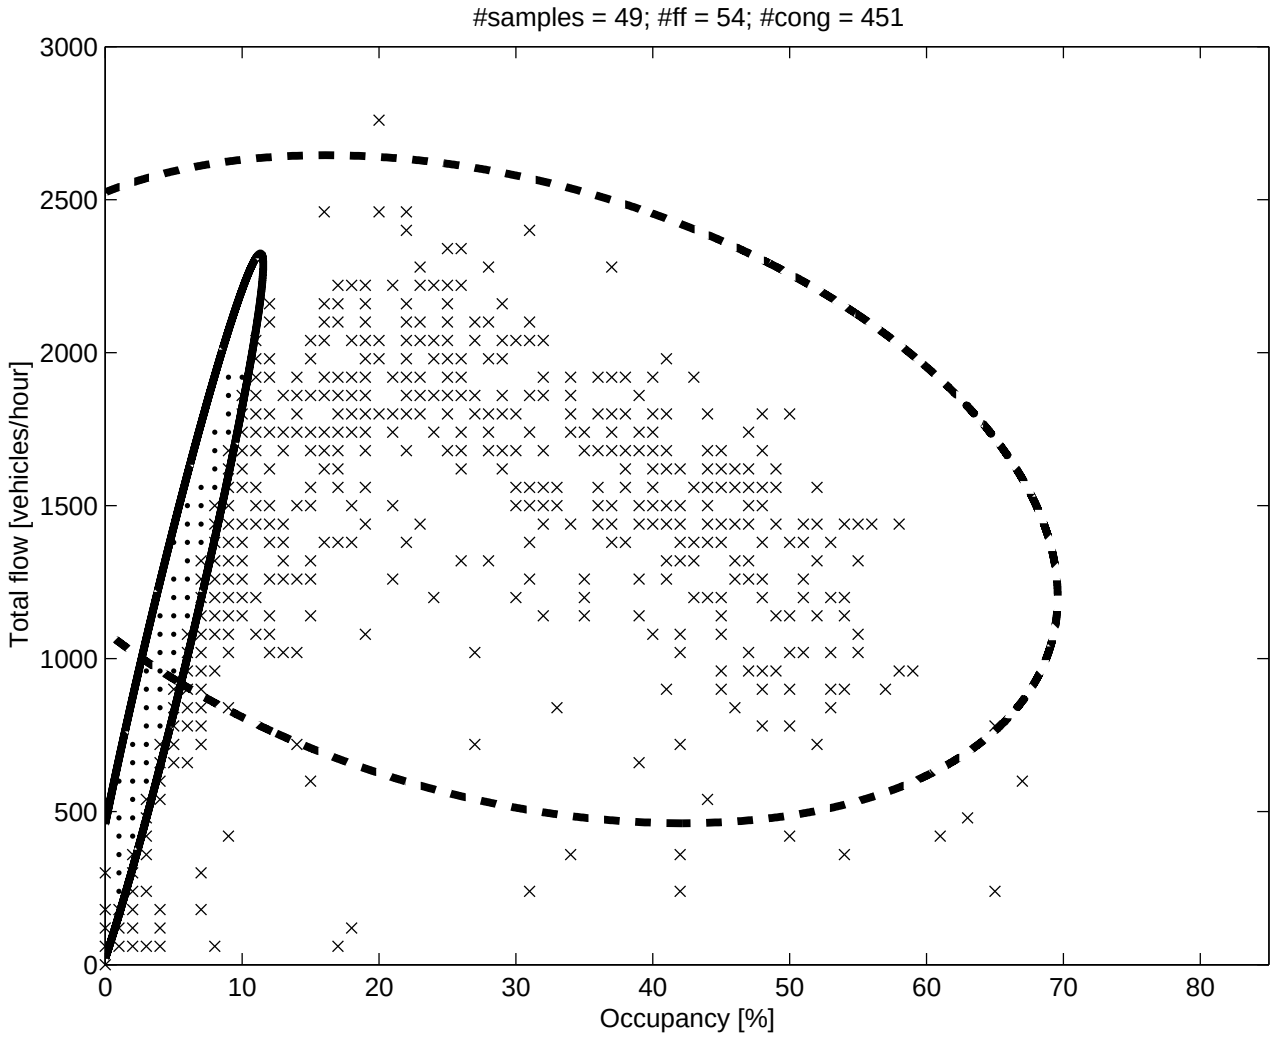

FIG. 12: Detecting outliers in the bivariate data from the (ρ,q) diagram. The small dots denote measurements belonging to the free-flow traffic regime; the small crosses belong to the congested traffic regime. For this particular example, the data set consisted of 49 Mondays, with 44 and 451 unique points in the free-flow and congested traffic regimes, respectively. The thick solid and dashed ellipses denote the 97.5% tolerance boundaries for both regimes, based on the results of the MCD estimator.

shows the percentages in the congested traffic regime, and the bottom part shows the average of both regimes. Lighter colours denote lower percentages, whereas a black colour denotes an upper bound of 30% outliers for that detector at that weekday. The detectors are arranged from left to right, with each 'column' in the image containing seven thin bars, one for each day of the week.

Looking at the resulting images, we can already spot errors that occurred when storing the measurements to the central database: in both Fig. 13 and Fig. 14 there exist 'white vertical gaps', denoting several detectors that malfunctioned at all weekdays, probably due to transmission failures. At these gaps, the percentage of outliers is zero, indicating that the corresponding detectors remained stuck in an on- or off-position during the entire measurement period. Furthermore, comparing the bottom parts of both figures, it seems that there were more outliers in 2003 than 2001; note that this can be an indication of an area-wide change in the calibration of the detectors. Another observation we can make is that for both years, the number of outliers in the congested traffic regime during Saturdays and Sundays is different from the other weekdays. This can be seen in the middle parts as the darker intensities in the two lower bands.

4. Dealing with missing values

To conclude this section about measurement errors, we provide some pointers with respect to dealing with missing values. When working with contaminated data, a frequently followed scheme is to first find all invalid data points (i.e., outlier detection), after which all these points are removed from the data set. As such, they are converted into missing values

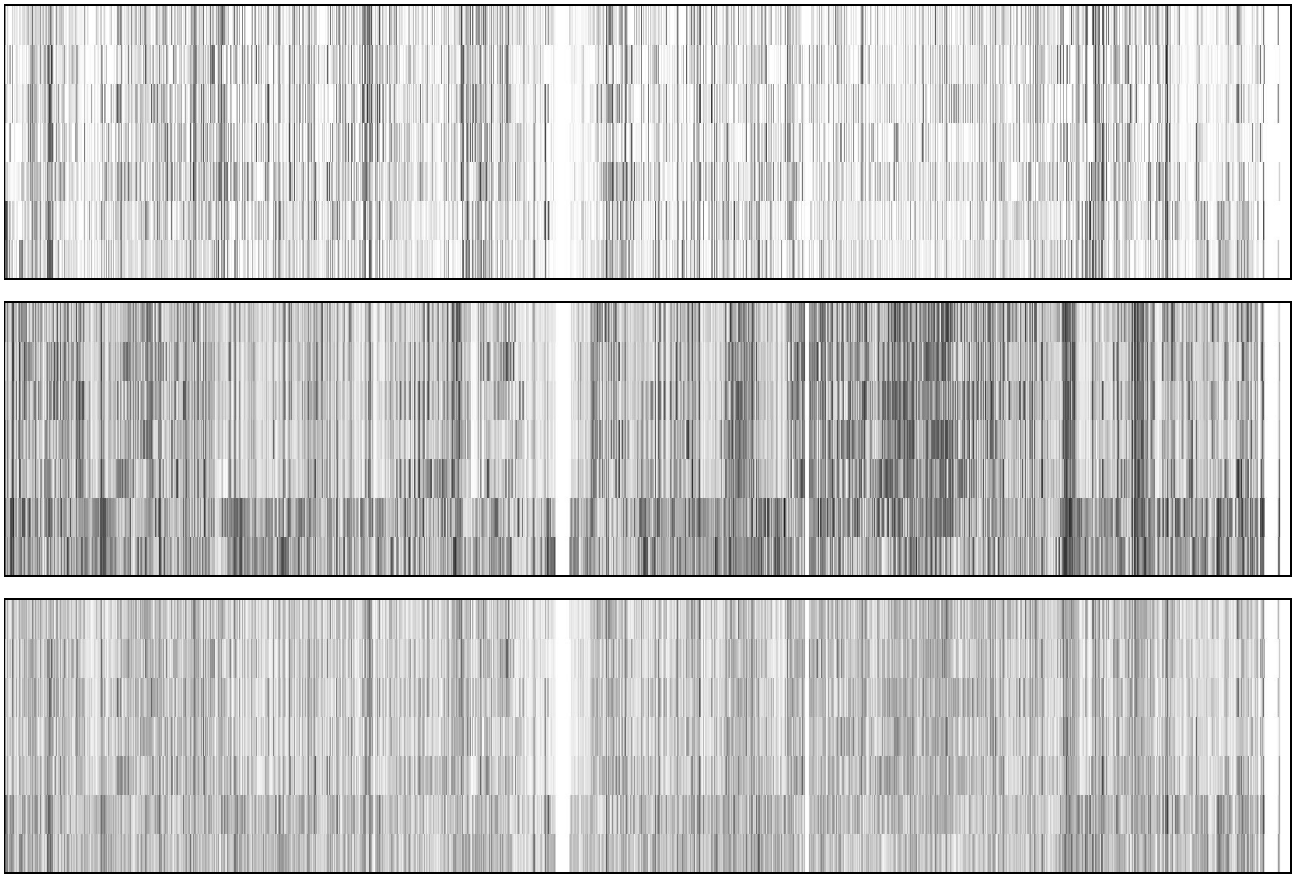

FIG. 13: The percentages of outliers in the free-flow traffic regime (*top*), the congested traffic regime (*middle*), and the average of both regimes (*bottom*); the small thin rectangles correspond to the 1654 detectors for the year 2001, whereas the seven rows correspond to the different days of the week. Lighter colours denote lower percentages, whereas a black colour denotes an upper bound of 30% outliers for that detector at that weekday.

and the preprocessing problem now becomes one of filling in all these missing values. We highlight a few of the many possible approaches:

- **Using reference days**
  As opposed to the use of classic interpolation schemes (based on e.g., linear or polynomial functions, splines, ...), Bellemans et al. proposed a method that is based on a *reference day*. In their work, they assumed the existence of an a priori known reference day that is representative of the day for which missing values have to be estimated. Based on the measurements $x(t-1)$ and $x_{\text{ref}}(t-1)$ at the previous time step, and the reference measurement $x_{\text{ref}}(t)$ at the current time step, the new measurement $x(t)$ is estimated as follows[77] [4–6]:

$$x(t) = \frac{x(t-1)}{x_{\text{ref}}(t-1)}\, x_{\text{ref}}(t). \tag{14}$$

  The fraction in the previous equation plays the role for scaling the reference measurement such that it corresponds to the traffic dynamics of the day under study.

- **Multiple imputation**
  One popular way for dealing with missing values is by means of *imputation*, i.e., 'filling them in' based from samples drawn from a probability distribution [37]. In principle, using Bayesian methods is a suited methodology for obtaining valid estimates for these missing values: once their distributions are known or estimated, the sought-after posterior probability can be calculated as the ratio between the likelihood times the prior and a normalising constant. In practice however, it is not always feasible to carry out such a full Bayesian analysis due complexity issues, normality assumptions, ...

  In short, *multiple imputation* (MI) can be summarised as follows: given an incomplete data set, the first step is to detect and fill in the missing values based on an imputation model that gives values drawn from a distribution. This

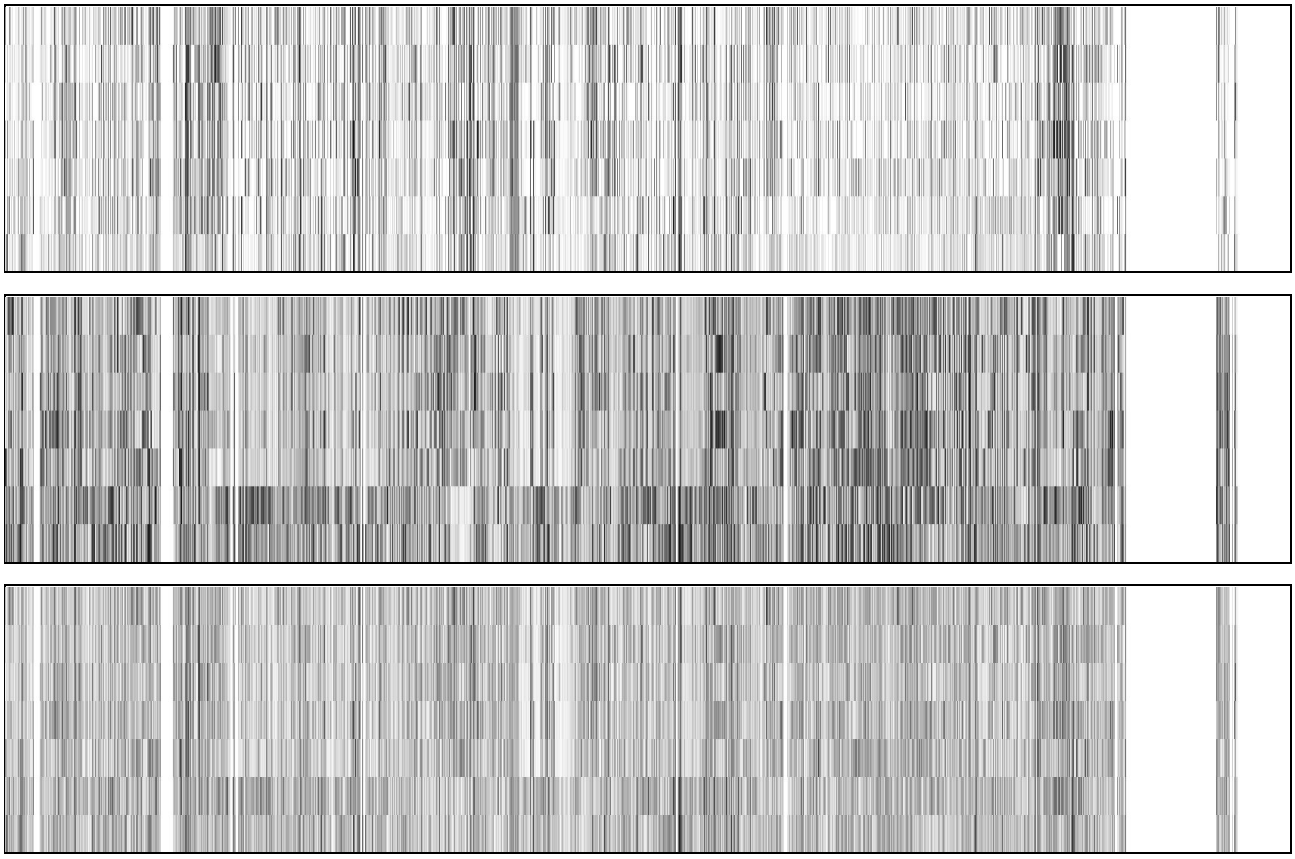

FIG. 14: The percentages of outliers in the free-flow traffic regime (*top*), the congested traffic regime (*middle*), and the average of both regimes (*bottom*); the small thin rectangles correspond to the 1800 detectors for the year 2003, whereas the seven rows correspond to the different days of the week. Lighter colours denote lower percentages, whereas a black colour denotes an upper bound of 30% outliers for that detector at that weekday.

is done not once but m times (hence the name 'multiple' imputation), resulting in m different complete data sets. Each data set is then analysed separately, after which the m results can easily be combined. A nice advantage of the MI method is that the value for m does not need to be large, e.g., $m = 10$ is typically sufficient [53]. In the first step, the small number of imputed values can be drawn from predictive distributions by e.g., a *Markov Chain Monte Carlo[78]* (MCMC) method [55].

Alternate approaches would be to use *maximum likelihood estimations* (MLE), which can be iteratively computed by a technique such as *expectation-maximisation* (EM) [22, 43]. Advantages of using MI compared to MLE, are that it can work better with smaller sample sizes, and the fact that model used for analysing the results can be different than the imputing model that was used to obtain values. The main difference between both approaches, is that missing values are dealt with implicitly in the MLE method, whereas they are dealt with prior to the analysis in the case of MI [56].

- **Time series analysis**
  Probably the most employed methodology in classic time series analysis (TSA) is the approach towards forecasting known as *Box-Jenkins analysis*. Simply put, the analysis is based on what is known as a *autoregressive integrated moving average* (ARIMA) model. Box et al. provided a complete method for removing trends and seasonal effects, by means of differencing the time series. As such, an ARIMA(p,d,q) model expresses a time series as a combination of current and past observations, with p, d, and q the orders of the autoregression, integration (the differencing), and the moving average, respectively. Autoregression determines the relevance of previous values with respect to the current value, integration takes care of detrending the time series to make it stationary (i.e., the mean and variance remain constant over time), and the moving average allows smoothing of the time series [9, 10].

  The previously outlined methodology works well for finite-dimensional linear models; however, when considering e.g., chaotic processes, the technique fails due to the inherent chaotic transitions and the presence of a continuous Fourier spectrum. To cope with this, we can look in another way at time series analysis, i.e., by means of *non-parametric models* that rely on the *state space* of the underlying dynamical system. Because a mathematical model of such a system is not always available, the state space can be constructed from a single time series by means of a process called *attractor reconstruction*. The principal method to this end is called *delay coordinate embedding*

(DCE); it was derived by Packard et al. and put into a rigid mathematical formulation by Takens [45, 60]. The idea behind DCE is that from the single time series, a set of new time series is constructed; each of these series is a *time-shifted version* of the original one[79]. If we assume that the time series is expressed as a sequence of observations $\mathbf{x}(t) = \{x_1(t), \ldots, x_n(t)\}$, belonging to an n-dimensional space, then the DCE method results in a vector $\mathbf{r}(t) = \{\mathbf{x}(t), \mathbf{x}(t-\tau), \ldots, \mathbf{x}(t-(m-1)\tau)\}$. In this derivation, τ is called the *delay* and m the *embedding dimension*. The powerful result of Takens proves that if both the embedding dimension m and the delay τ are selected in an optimal fashion, then the dynamics of both the reconstructed state space and the system's original state space are topologically identical [46, 54, 60]. The search for the optimal values for both parameters is guided by techniques such as *average mutual information* (AMI) for the delay, and the *false nearest neighbours* (FNN) algorithm for the embedding dimension [25, 33]. Practical implementations for this kind of analyses can be performed using e.g., the *TISEAN* package [28].

If the embedding dimension gets larger than 2 or 3, then a visualisation of this high-dimensional data becomes problematic. A way for dealing with this is by means of so-called *recurrence plots* (RP), invented by Eckmann et al. [23, 24]. In this kind of plot, all information from the trajectory of the time series that was constructed by applying the DCE method is converted into a two-dimensional image: each point (i,j) in such a plot is then shaded according to the distance between two corresponding trajectory points $\mathbf{r}(i)$ and $\mathbf{r}(j)$. If each point in such a recurrence plot is compared to a predefined treshold, then the resulting black-and-white image is called a *tresholded recurrence plot* (TRP). As an example, we show three RPs in Fig. 15. The left part is obtained from a time series that essentially is generated from uniformly distributed noise; there are no clearly delineated structures present. The middle part shows a TRP of a sinusoidal time series; as can be seen, the image exhibits a large degree of periodicity in its structures. The right part shows the results for a time series that contained a drift due to slowly varying parameters.

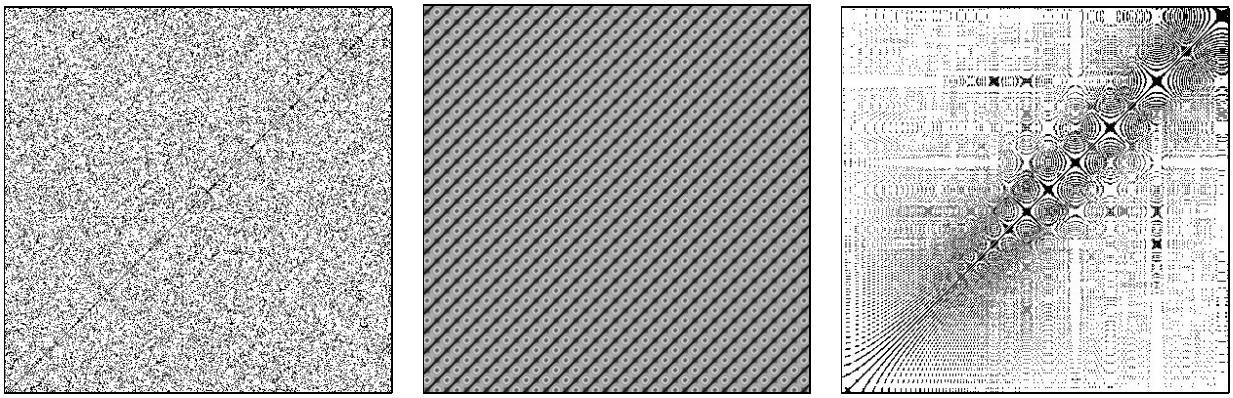

FIG. 15: *Left:* a recurrence plot obtained from a time series that essentially is generated from uniformly distributed noise; there are no clearly delineated structures present. *Middle:* a tresholded recurrence plot obtained from a sinusoidal time series; as can be seen, the image exhibits a large degree of periodicity in its structures. *Right:* a recurrence plot obtained from a time series that contained a drift due to slowly varying parameters.

Assessing the structures in these RPs remains a somewhat 'visual discipline'; to cope with this, Zbilut and Webber extended a techniques called *recurrence quantification analysis* (RQA) that allowed a more quantitative treatment of RPs. Their technique is based on five statistics that describe phenomena such as recurrence, determinism, entropy, trend, and the largest positive Lyapunov exponent (which is a measure for the chaoticity of a system) [69, 70]. Intuitively, these measures are related to visual features such as the percentage of lines to the main diagonal, the distribution of the lengths of diagonal lines, ... In light of the difficulties encountered by selecting the optimal embedding dimension, a promising result was obtained by Iwanski and Bradley, who state that it is possible to get the same RQA results without embedding [30].

With respect to the application of time series analysis to traffic flow data, we note the interesting result from Smith et al. In their work, they compared the use of classic ARIMA modelling to non-parametric modelling based on DCE. Their results indicate that the latter did not approach the performance of the former; this leads them to the belief that traffic data is rather stochastic as opposed to chaotic [59].

C. Assessing detector malfunctioning

As already hinted at in the introduction of this section, most single inductive loop detectors exhibit a large degree of errors, missing and/or incorrect values, ... In order to provide a more qualitative assessment of these errors, we adopt a screening

methodology that has been used in the PeMS project (see the end of section I B for more information); this allows us to provide clearly structured maps that allow a quick visual inspection of all detectors in Flanders' motorway network and their operations during the years 2001 and 2003. As a first part of this section, we explain the methodology behind the screening of the detector data, after which we provide and discuss the detector maps.

1. Explanation of the methodology

When a detector malfunctions (or even an entire measurement complex), its errors typically result in under- or overestimations of the flow, high occupancy values, blank data, ... Early methods for screening the measurements are based on acceptance and rejection regions in the scatter plots of a (k,q) diagram. For example, Payne et al. created tests on the bounds of minimum and maximum flows, occupancies, and mean speeds, in order to discriminate between good and bad 20-second and 5-minute samples of detector data [47]. Another algorithm was constructed by Jacobson et al. at the University of Washington; their *Washington Algorithm* provides an explicit acceptance region within the (k,q) diagram [31].

A more recent approach was followed by Chen et al., which resulted in the *Daily Statistics Algorithm* (DSA), currently used in the PeMS project [7, 12]. The idea behind this algorithm is to consider all measurement samples of a loop detector for one day, calculating four different scores based on these samples, and then, by comparison with some predefined tresholds, deciding whether or not the detector is considered to be malfunctioning. These scores check (1) for zero occupancy samples, (2) strictly positive occupancy samples with zero flow, (3) high occupancy samples, and (4) the entropy of these occupancy samples. The main strength of this algorithm is that it allows to test for detectors that continuously report faulty data, e.g., being stuck in the on/off position (although it is also possible that no vehicles crossed the detectors at all). For our study, we used the following scores:

$$S_1(i, \Delta T) = \text{number of samples in } \Delta T \text{ with } \rho_i = 0, \tag{15}$$
$$S_2(i, \Delta T) = \text{number of samples in } \Delta T \text{ with } \rho_i > \rho^*, \tag{16}$$
$$S_3(i, \Delta T) = \text{entropy of the occupancy samples in } \Delta T. \tag{17}$$

The sentinel values in the database (typically denoting transmission failures), are automatically catched by $S_2(i, \Delta T)$. The entropy $S_3(i, \Delta T)$ is calculated as follows:

$$S_3(i, \Delta T) = -\sum_{\rho_i \in p(\rho_i) > 0} p(\rho_i) \log(p(\rho_i)), \tag{18}$$

with $p(\rho_i)$ the estimated probability density function, defined as the histogram of the occupancies ρ_i (we selected 100 bins for the estimation). The entropy provides a measure for the randomness of a stochastic variable, i.e., constant values will result in a zero entropy. In their work, Chen et al. also discuss several shortcomings of the DSA approach, most importantly the lack exploiting spatial and temporal correlations between measurements stemming from neighbouring detectors [12]. As previously mentioned, in the original DSA, the daily decision on a detector being good or bad hinged on the scores that were compared to some predefined tresholds. Instead of adopting this methodology, we discard this binary classification and allow the complete range of results.

Note that we are working with the raw unprocessed vehicle counts from the detectors, instead of converting them to passenger car equivalents. The reason is that the latter introduces an incorrect percentage of trucks, due to the problems with misclassification as mentioned at the end of section I A 3.

2. Illustrative detector maps

Based on the scores S_1, S_2, and S_3 as explained in the previous section, we now provide charts of all detectors in the Flander's region (as already mentioned in section I B, the available data spans 1654 detectors for the year 2001, and 1800 detetectors for the year 2003). To this end, we calculate these scores for each hour in both years; they scores are stored in matrices that have 24 hours $\times$ 365 days = 8760 columns (i.e., ΔT = 60 minutes). All matrices are then normalised, after which they are converted to gray-scale images with each matrix element corresponding to one pixel in the image. As the predefined treshold for S_2, we chose $\rho^* = 35\%$, as was done in the work of Chen et al. [12].

In Fig. 16 we show the results after calculating the scores for all detectors during the entire year 2001, aggregated for each hour. The top, middle, and bottom row indicate the S_1, S_2, and S_3 scores, respectively. The darker a pixel is coloured, the higher the specific score is (black meaning that all samples during ΔT contribute to the score). Fig. 17 gives the same results, but for the year 2003.

Before we discuss both these detector maps, it is worthwhile to take a look at some general patterns that seem to occur. As these maps are highly detailed (i.e., spanning a width of some 8760 pixels), we provide two close-ups in Fig. 18. As can be seen in the close-up to the left, there seem to be some slanted 'streaks'; these may indicate detector malfunctions at successive detectors at successive time periods. The close-up to the right reveals another more frequently occurring phenonemon, namely vertical and horizontal lines: a darker horizontal line may indicate detector failure during a certain time period; a darker vertical line may indicate several (probably neighbouring) detectors (i.e., at a measurement post or complex) that are failing. The wider a vertical line, the more extended the time period of failure. A vertical line that runs completely from top to bottom on the map, typically indicates a problem during transmission or archival of measurements to the central database; as it is highly unlikely that an area-wide malfunctioning seems to occur, it is more logical to assume that an error occurred at the database level. Note that the regular grouping of short lines in the right close-up is related to the fact that at night time, the occupancy is very low as few vehicles cross the detectors; as a result, a high number of zero occupancies is reported and shown as darker segments.

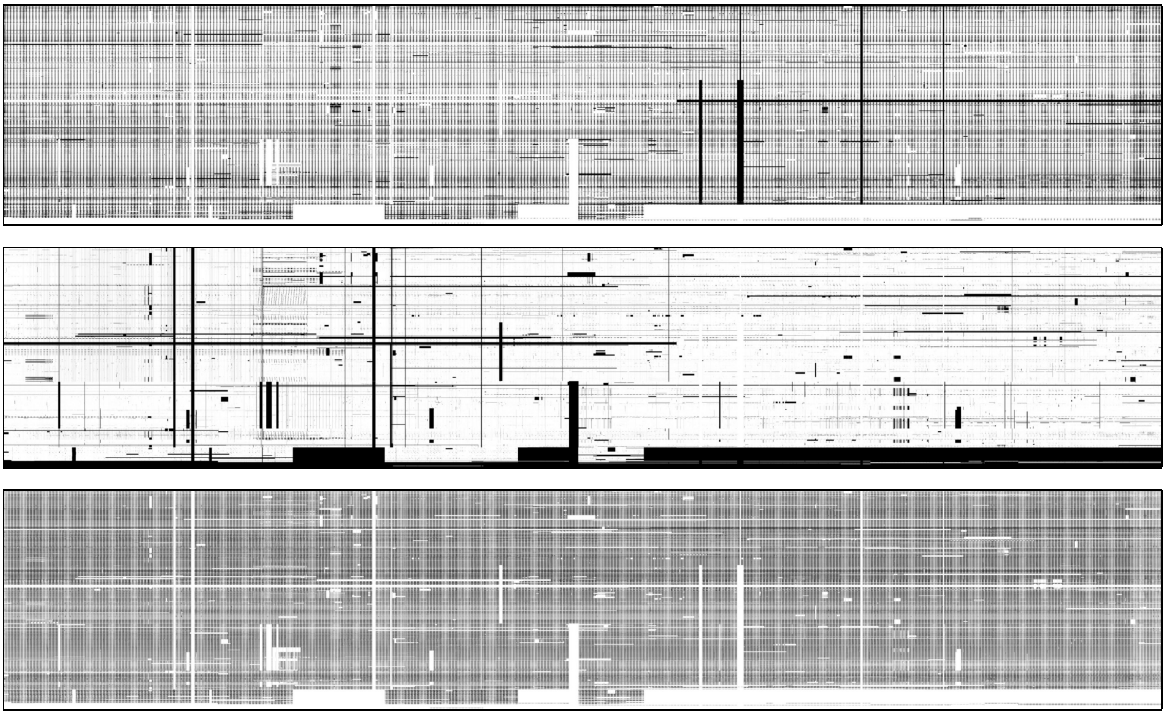

FIG. 16: Illustrative detector maps of the S_1 (*top*), S_2 (*middle*), and S_3 (*bottom*) scores, for all 1654 detectors in the year 2001. The scores were calculated for each hour; the darker a pixel is coloured, the higher the specific score is (black meaning that all samples during ΔT contribute to the score).

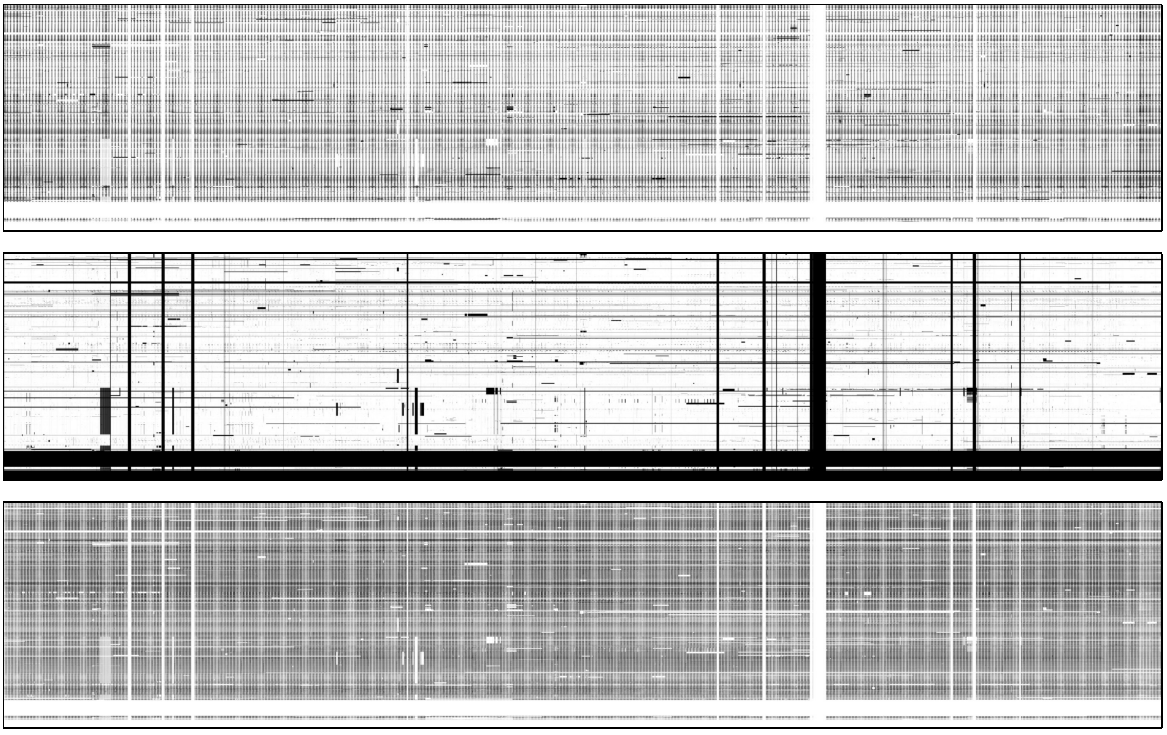

FIG. 17: Illustrative detector maps of the S_1 (*top*), S_2 (*middle*), and S_3 (*bottom*) scores, for all 1800 detectors in the year 2003. The scores were calculated for each hour; the darker a pixel is coloured, the higher the specific score is (black meaning that all samples during ΔT contribute to the score).

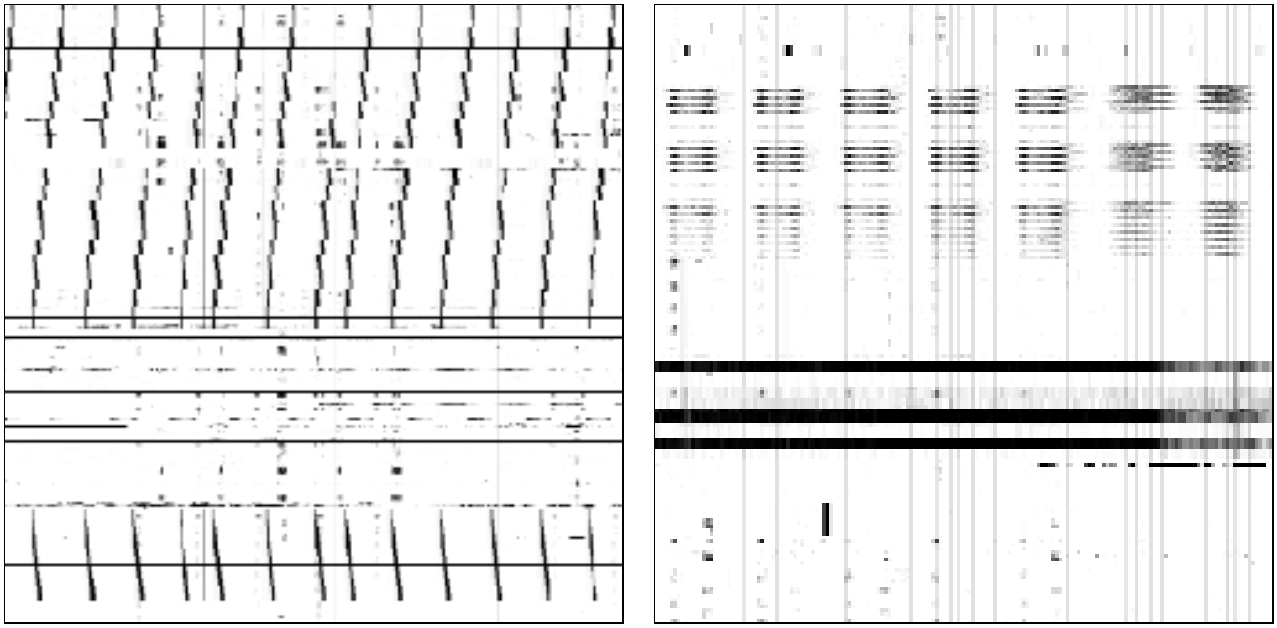

FIG. 18: Some close-up examples of general patterns that occur in the detector maps. *Left:* the presence of slanted 'streaks' may indicate detector malfunctions at successive detectors at successive time periods. *Right:* a more frequently occurring phenonemon, namely vertical and horizontal lines: a darker horizontal line may indicate detector failure during a certain time period; a darker vertical line may indicate several (probably neighbouring) detectors (i.e., at a measurement post or complex) that are failing. The wider a vertical line, the more extended the time period of failure.

Returning to the detailed detector maps provided in Fig. 16 and Fig. 17, we can see that from 2001 to 2003, the number of detector malfunctions seems to have decreased, based on the occurrence of darker regions in score S_2 (i.e., high occupancy values). Still, as can be seen from the middle part of Fig. 17, there are numerous detectors that seem to be malfunctioning during the entire year 2003, as is indicated by the frequent occurring of darker horizontal lines in the map. There were also several problems during transmission or archival to the central database, as is evidenced by the dark vertical lines. Finally, with respect to the entropy of the occupancy samples, we note that there seem to slightly less stuck detectors, as the white empty regions (indicating zero entropy) diminish from 2001 to 2003.

In order to more quantitatively consider these detector maps, Fig. 19 presents histograms showing the distributions of all scores. The top row displays the results for the year 2001, the bottom row for the year 2003; the left, middle, and right histograms correspond to scores S_1, S_2, and S_3, respectively. The distinct bars in both left histograms correspond to an increasing number of zero occupancy samples, representative of traffic at night time. Furthermore, as already highlighted in the previous paragraph, the number of high occupancy values has decreased (see both middle parts). Finally, the probability of a low entropy (around 0.5) seems to have diminished from 2001 to 2003 (see both right parts).

To conclude, we provide six more detector maps in Fig. 20 for 2001 and Fig. 21 for 2003. The difference between the previous maps, is that these ones show aggregated scores for whole days instead of every hour. As such, they are smaller in width, spanning only 365 pixels. In a sense, they convey the same information as presented in Fig. 16 and Fig. 17. Note the big black and white regions near the bottom of all six maps; they are most likely indicative of place holders in the central database for new detectors, resulting in default values for the flow and occupancies and correspondingly giving high scores.

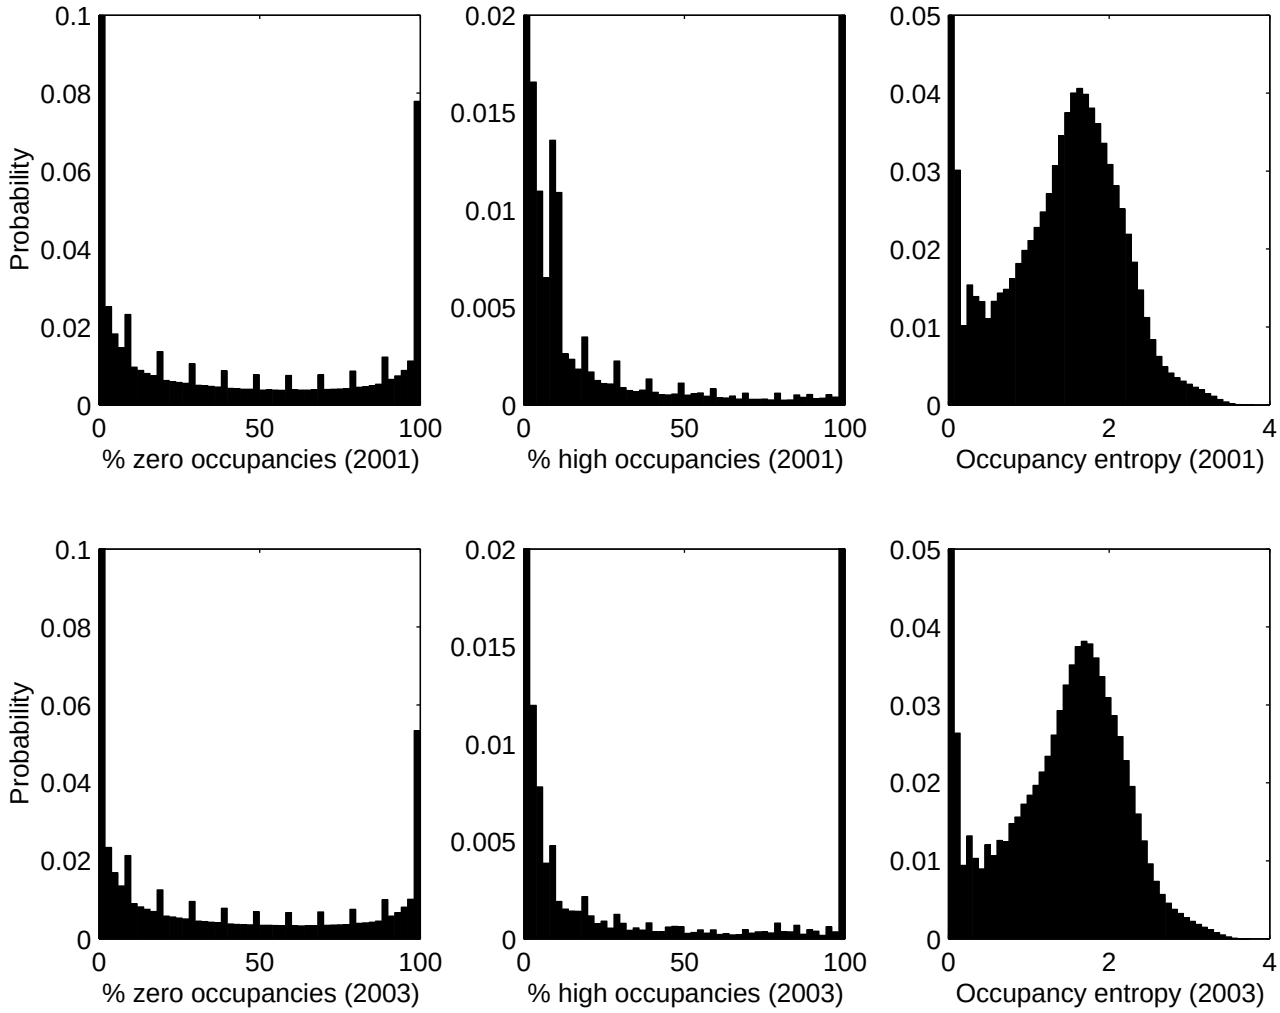

FIG. 19: Histograms showing the distributions of all scores for the years 2001 (*top row*) and 2003 (*bottom row*). Each time the left, middle, and right histograms correspond to scores S_1, S_2, and S_3, respectively.

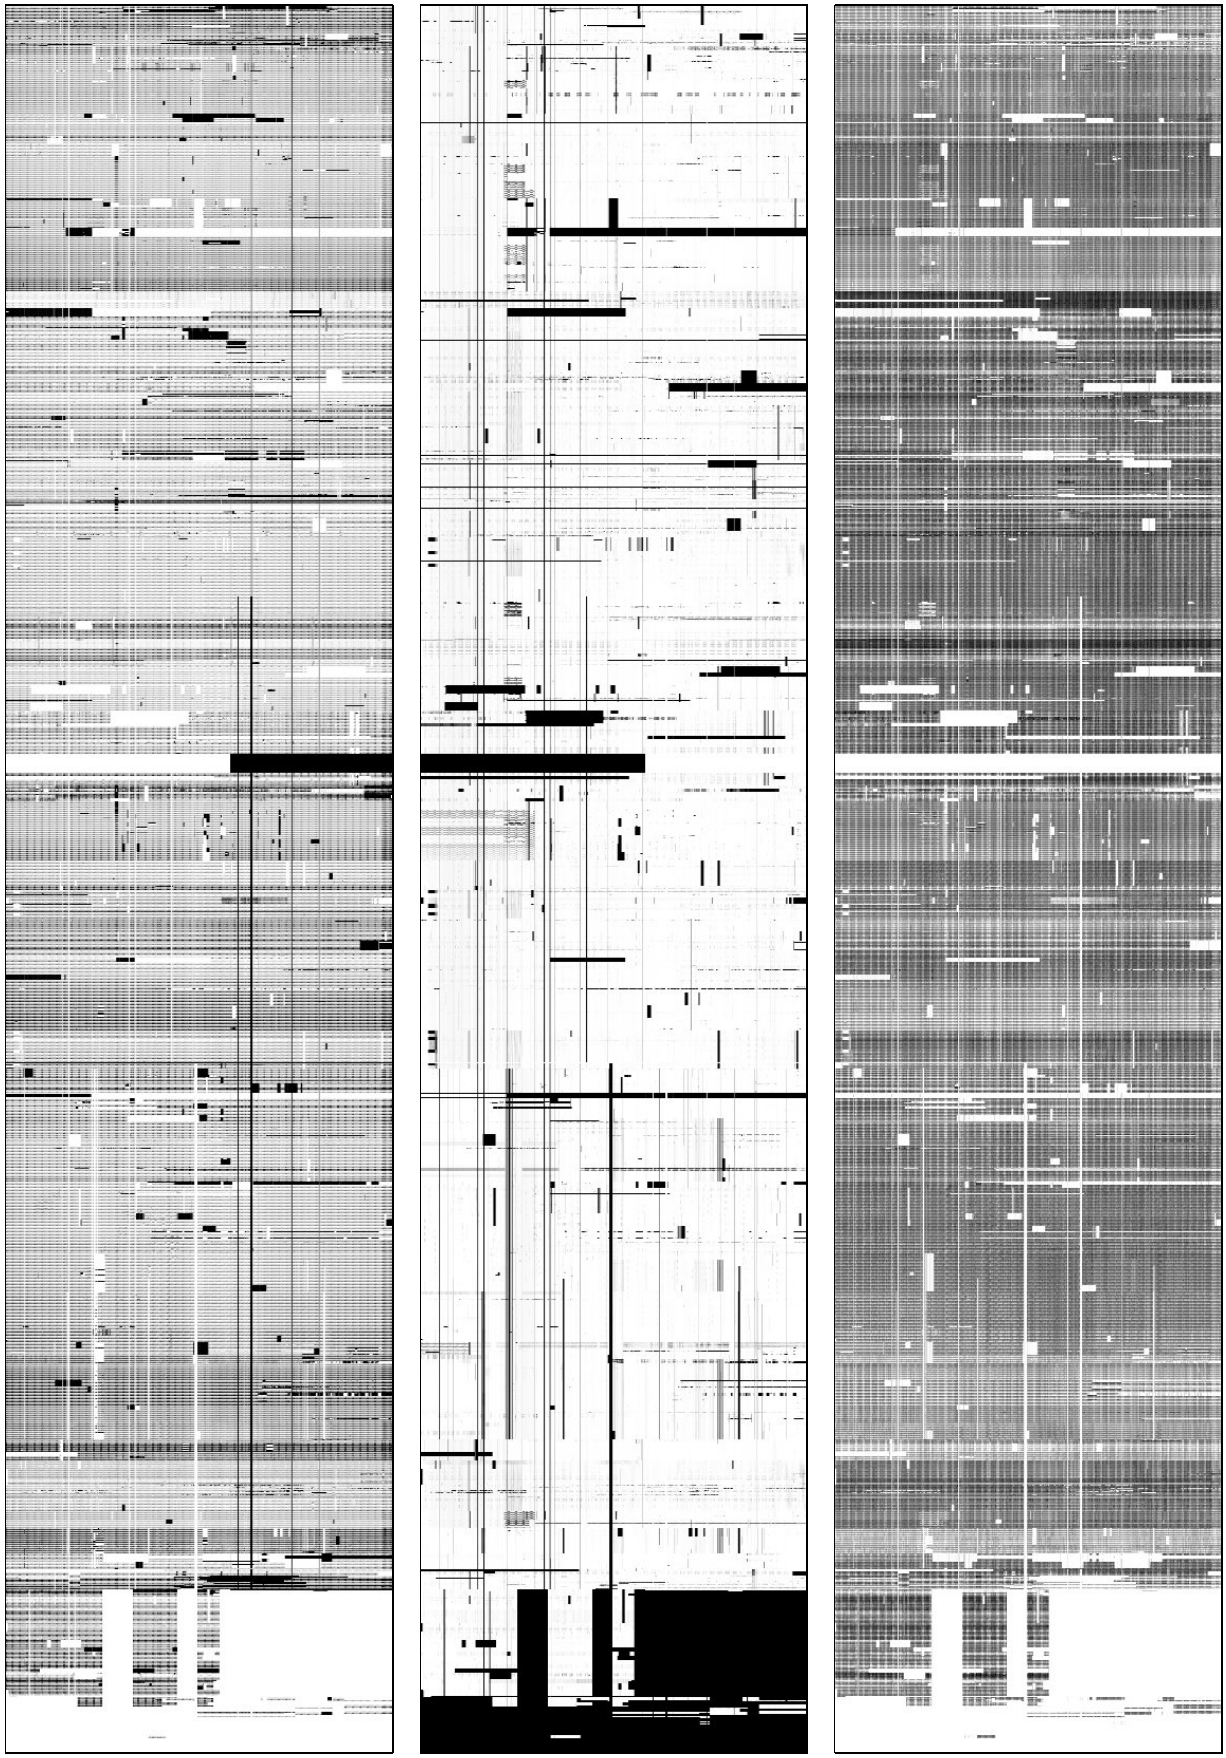

FIG. 20: Illustrative detector maps of the S_1 (*top*), S_2 (*middle*), and S_3 (*bottom*) scores, for all 1654 detectors in the year 2001. The scores were calculated for each day; the darker a pixel is coloured, the higher the specific score is (black meaning that all samples during ΔT contribute to the score). Note the big black and white regions near the bottom of all six maps; they are most likely indicative of place holders in the central database for new detectors.

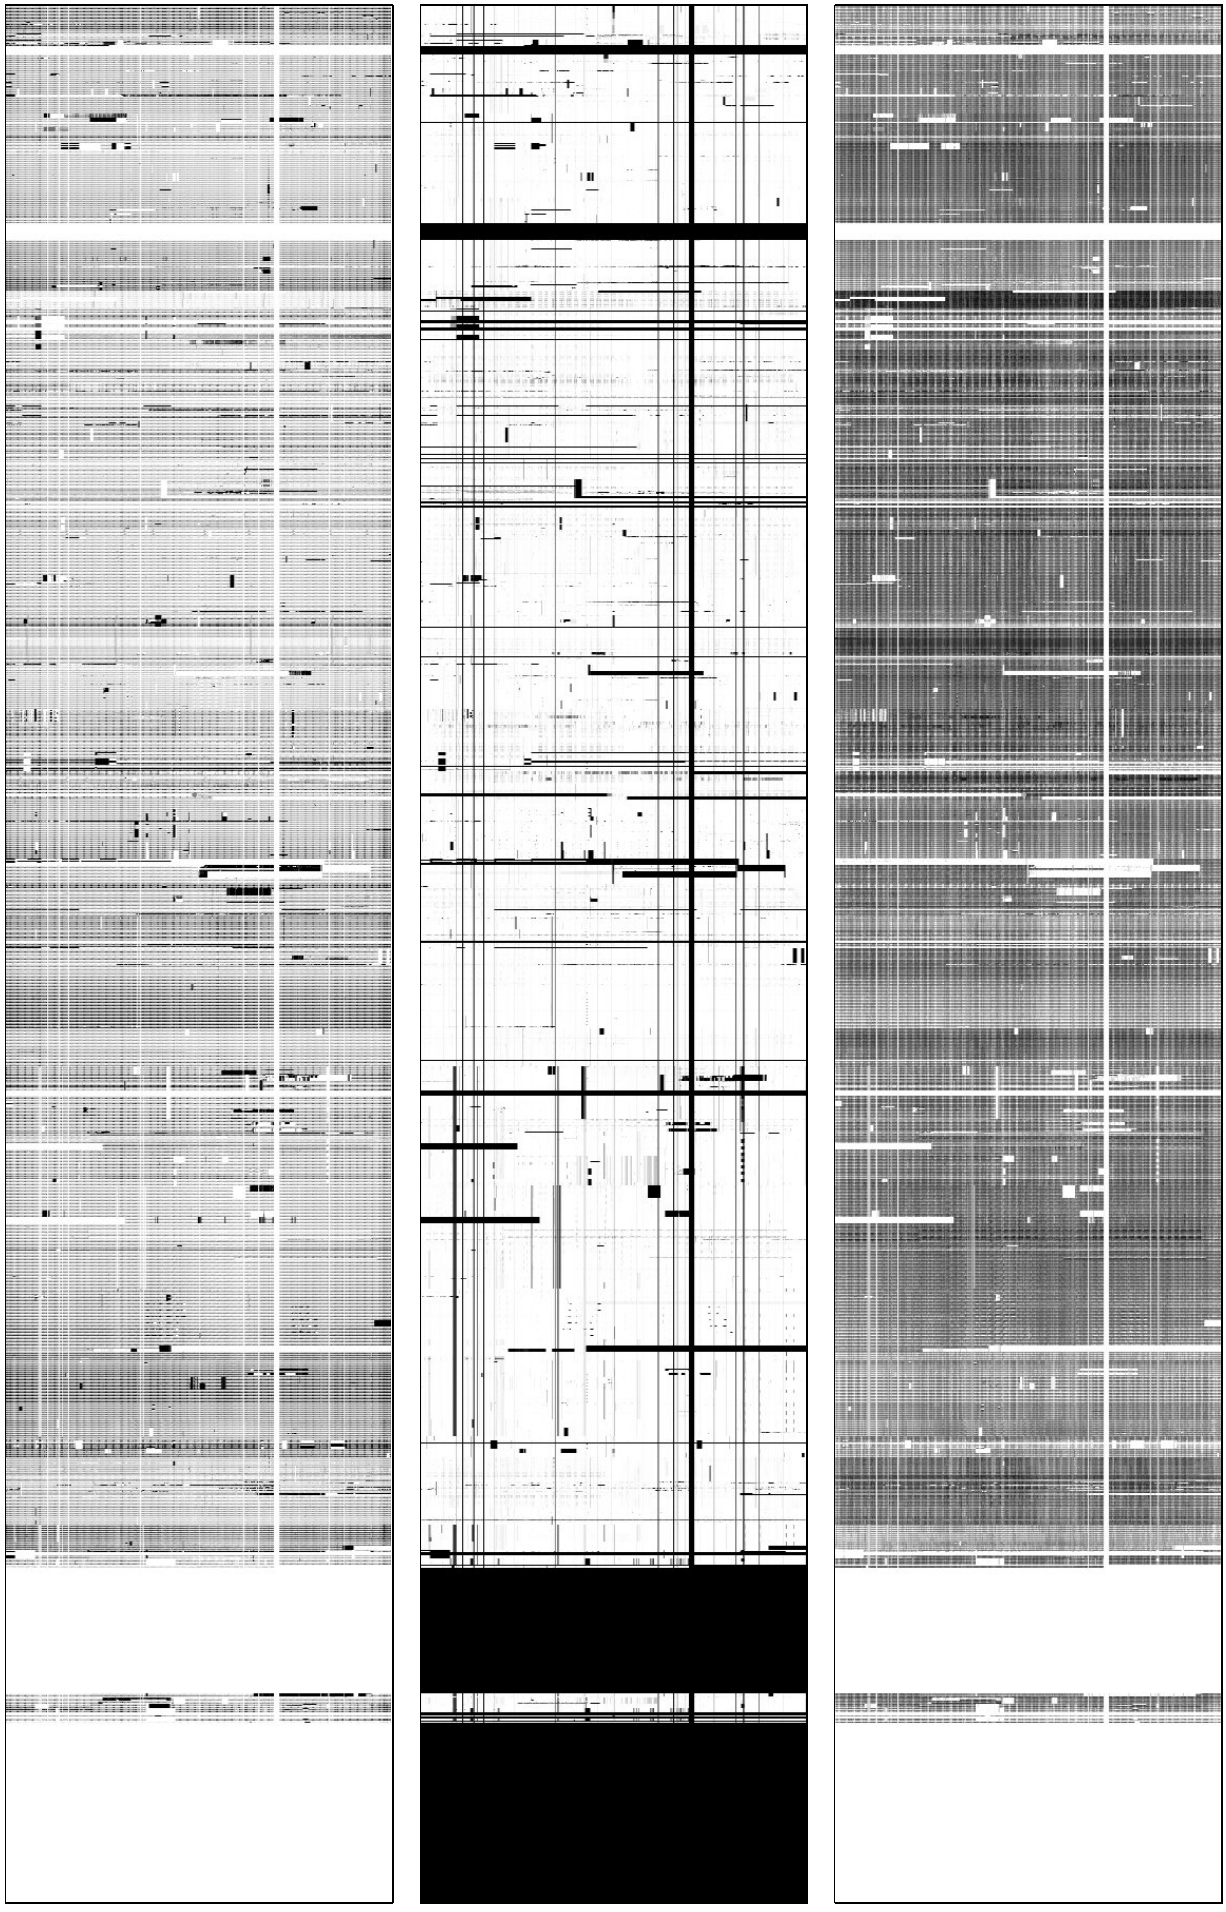

FIG. 21: Illustrative detector maps of the S_1 (*top*), S_2 (*middle*), and S_3 (*bottom*) scores, for all 1800 detectors in the year 2003. The scores were calculated for each day; the darker a pixel is coloured, the higher the specific score is (black meaning that all samples during ΔT contribute to the score). Note the big black and white regions near the bottom of all six maps; they are most likely indicative of place holders in the central database for new detectors.

III. OFF-LINE TRAVEL TIME ESTIMATION AND RELIABILITY INDICATORS

The final section of this paper is concerned with reliable *estimations* of the travel time in an off-line setting. As such, we are not predicting the travel times for unknown future traffic conditions. The methodology elaborated in this section can be classified according to Van Lint's taxonomy as a *data-driven modelling approach using a short-term prediction horizon; it constitutes indirect pre-trip estimation using a flow-based technique, having individual motorway sections as its spatial scope* [63].

The first part of this section discusses some of the techniques that are commonly applied when estimating travel times based on historical traffic flow measurements. In the subsequent part we present our methodology for estimating travel times between detector stations, after which we present an application to real-world measurements stemming from single inductive loop detectors. The final part discusses some reliability and robustness properties related to travel times and traffic flow dynamics.

Note that we exclude a large part of the literature dealing with predictions; for an example of this approach we refer the reader to e.g., the work done by Tampère et al. in the *PredicTime* project that mainly focusses on the use of Kalman filtering [34, 67], and the doctoral dissertation of Van Lint who employs state-space neural networks (SSNN) for reliable on-line predictions of travel times [63].

A. Common approaches towards travel time estimation

The most prominent technique employed, is by deriving the travel time based on the inverse of the measured mean speed[80]. We make the distinction between the *experienced dynamic travel time*, starting at a certain time t_0, over a road section of length K [8]:

$$T(t_0) = \int_0^K \frac{1}{v(t,x)} dx \qquad \forall\, t \geq t_0. \tag{19}$$

If not all local instantaneous vehicle speeds $v(t,x)$ are known at all points along the route, a simplification can be used, resulting in the *experienced instantaneous travel time*:

$$\widetilde{T}(t_0) = \int_0^K \frac{1}{v(t_0,x)} dx. \tag{20}$$

As an example, the following relation gives the experienced instantaneous travel time for a section, based on the mean speed measurements from an upstream and downstream detector post:

$$\widetilde{T}_{\mathrm{D_{up}}\rightarrow\mathrm{D_{down}}}(t_0) = \frac{K}{2}\left(\frac{1}{\overline{v}_{s_{\mathrm{D_{up}}}}(t_0)} + \frac{1}{\overline{v}_{s_{\mathrm{D_{down}}}}(t_0)}\right). \tag{21}$$

In the previous equation, it is assumed that the measured mean speeds at both detector posts correspond to the traffic situation within the section; during congested conditions, this assumption can lead to an underestimation of the travel time if both detectors are still registering free-flow traffic.

In practice, the travel times are based on the consideration of all vehicle trajectories; to this end, it is necessary to estimate the mean speed of each trajectory. When performing this step, a possible approach is to estimate the trajectories' speeds as piece-wise linear functions. This latter can be done by sampling them at discrete distances using e.g., single inductive loop detectors [8].

Some examples of this type of travel time estimation are the work of Dailey who uses volume and occupancy measurements from SLDs, incorporating the stochastic nature of these quantities in order to derive the travel time; a Kalman filter is included to address the variability of the observations [20], Petty et al. who adopt a similar methodology by using a stochastic model to accurately estimate the distribution of travel times, based on the measured upstream and downstream arrivals at SLDs [48], Coifman who provides a methodology for estimating vehicle trajectories, and hence also link travel times, from the measurements at DLDs, based on a triangular (k,q) fundamental diagram [15], and Coifman and Cassidy who adopt a vehicle reidentification algorithm that matches vehicle measurements made at a downstream detector to those made at an upstream detector; as such, it is possible to derive the link travel time [17, 18].

B. Estimating travel times based on flow measurements

As previously noted, many travel time estimation techniques are based on the inverse of the measured mean speed. In our work, we are dealing with single inductive loop detectors, of which the mean speed estimations are notoriously known to be unreliable. However, as mentioned at the end of section I A 4, the vehicle counts measured by these SLDs are quite reliable (for which we ignore the misclassification issues between cars and trucks, as only the total vehicle counts matter). As such, our methodology is based on the use of flows, i.e., the *cumulative curves*.

Briefly recapitulating, these cumulative curves represent the cumulative number of passing vehicles (denoted by N) with respect to time at different locations. Consider now a closed section of the road that conserves the number of vehicles (i.e., no on- or off-ramps); this section is demarcated by two detector posts which measure its inflow, respectively outflow of vehicles. The result is a pair of monotonically increasing functions $N_{\text{up}}(t)$ and $N_{\text{down}}(t)$, which increase each time a vehicles passes by [39].

The time needed to travel from one location to another can now easily be measured as the horizontal distance between the respective cumulative curves. Similarly, the vertical distance between these curves allows us to derive the accumulation of vehicles on the road section, which gives an excellent indication of growing and dissipating queues (i.e., congestion). Finally, at each time instant t the slope of this function corresponds to the flow $q(t)$.

In the remainder of this section we outline our methodology by first discussing how we extract the cumulative curves from the data, then how deal with systematical errors and synchronisation issues, and finally how we estimating the distribution of the travel time.

1. *Constructing the cumulative curves*

The first step in the estimation of travel times, is to create the cumulative curves, by extracting the total flow measurements $q_{\text{up}}(t)$ and $q_{\text{down}}(t)$ from both upstream and downstream detector posts, respectively. The underlying assumption we make is that all vehicles in the total inflow upstream to and outflow downstream from the section are accounted for by both detector posts. Note that because we are dealing with multi-lane traffic, vehicles are allowed to overtake each other, and as such the *first-in, first-out* (FIFO) condition no longer holds [19]. However, this is not a significant problem because we are measuring the total flows at both detector posts.

Once the total flows are extracted, the cumulative curves $N_{\text{up}}(t)$ and $N_{\text{down}}(t)$ are constructed as follows:

$$N(t) = \sum_{t'=t_0}^{t} q(t') = N(t-1) + q(t). \tag{22}$$

In this equation, $N(t)$ thus represents the number of the last vehicle that passed the detector station at time t. Because the original flow measurements were aggregated at periods of one minute, the resulting $N(t)$ increases each time from $t \longrightarrow t+1$ with the total number of vehicles that passed the detector post.

2. *Dealing with synchronisation issues and systematic errors*

When plotting both cumulative curves $N_{\text{up}}(t)$ and $N_{\text{down}}(t)$ in a single diagram, it is necessary to make them comparable. This entails the following two tasks:

- For starters, when both curves are reporting the same number of counted vehicles, they need to be *synchronised* with each other. To understand this, suppose a reference vehicle passes the upstream detector post at a certain time instant t_{up}; after a certain time period Δt, the vehicle reaches the downstream detector post at a later time instant t_{down}. The amount $\Delta t = t_{\text{down}} - t_{\text{up}}$ is the time it takes to cross the distance between both detector stations, allowing the synchronisation mechanism to shift the respective cumulative curves over this time period (i.e., initialising them with the passing of the reference vehicle). As a result, both curves are synchronised such that $N_{\text{down}}(t_{\text{down}}) = N_{\text{up}}(t_{\text{up}} + \Delta t)$ for the reference vehicle (with Δt being the horizontal separation of both curves at time instant t_{up}).
- Secondly, all lane detectors in both detector posts can be *differentially tuned*, or their measurements can be obtained by SLDs at one post and cameras at the other post. As a result, they can be systematically biased towards lower or higher vehicle counts (note that this is not related to the internal vehicle classification by the local controller).

Synchronising both cumulative curves

The first method we tried for synchronising both cumulative curves and obtaining Δt, was to estimate the experienced instantaneous travel time $\widetilde{T}_{\text{ff}}(t_{\text{up}})$ under free-flow conditions, according to equation (20):

$$\Delta t = \widetilde{T}_{\text{ff}}(t_{\text{up}}) = \frac{K}{\overline{v}_{\text{ff}}}, \tag{23}$$

with K the distance between both upstream and downstream detector posts and $\overline{v}_{\text{ff}}$ the mean speed of all vehicles in free-flow traffic. However, due to the nature of the original flow measurements that are aggregated at periods of one minute, it is necessary that $\widetilde{T}_{\text{ff}}(t_{\text{up}}) \geq 1$ minute. A further complication is the fact that this travel time should be an integral value because the flow measurements are equidistant in time. As a consequence, the synchronisation methodology can only be applied to detector posts for which the travel time under free-flow conditions is at least one minute; e.g., when travelling at $\overline{v}_{\text{ff}} = 100$ km/h, then the minimal distance between both detector posts should be (100 ÷ 3.6) m/s × 60 seconds ≈ 1.6 km. Further problems arise due to the rounding of the travel time towards the nearest whole minute. As a result, it is possible that both cumulative curves intersect each other, which is physically impossible as this would mean that the total number of vehicles contained in the section becomes negative, which can even lead to negative travel times.

In order to tackle the previous problem, and to be able to work with detector posts that are separated at arbitrary distances, we derived a new methodology. The key idea is to 'undo' the aggregation of the flow measurements. To this end, we convert the $N(t)$ curves to equivalent $t(N)$ curves. The conversion is done by constructing a new curve, such that all N vehicles are time instant t are distributed over the time interval $t \longrightarrow t + 1$. Note that the vehicles' detector passing times are now expressed in seconds (which are allowed to be fractional values). The original $N(t)$ curves were equidistant in time t, the corresponding $t(N)$ curves are now equidistant in the vehicle number N. Furthermore, whereas we assumed that $N_{\text{up}}(t) > N_{\text{down}}(t)$, we now assume that $t_{\text{up}}(N) < t_{\text{down}}(N)$.

At this stage, we have for each individual vehicle that passed the detector posts, it corresponding t_{up} and t_{down} time instants. With this knowledge, we can now again compute $\widetilde{T}_{\text{ff}}(t_{\text{up}})$. Instead of using equation (23) which is explicitly based on the mean speed $\overline{v}_{\text{ff}}$ of all vehicles in free-flow traffic, we adopt another strategy. Abandoning the mean speed in free-flow traffic, we directly work with the 'history of the traffic stream', i.e., its characteristic properties that travel downstream. These undulations in both $t_{\text{up}}(N)$ and $t_{\text{down}}(N)$ curves propagate freely and allow us to achieve a better and more stable synchronisation.

One way to achieve this, is by looking at the respective shapes of both curves during light traffic conditions (e.g., the early morning period when free-flow conditions are prevailing). *The idea now is to shift one curve such that the difference between the two curves' shapes is minimal.* This approach for finding the correct time lag $\widetilde{T}_{\text{ff}}(t_{\text{up}})$ loosely corresponds to the concept of *platoon matching* where the characteristic features of groups of vehicles are compared and matched, as described in the Travel Time Data Collection Handbook of Turner et al. [62]. Our method also corresponds to what Bovy and Thijs call *travel time estimation based on mass-balance*, where the inflow of a section is compared to its outflow [8]. And finally, as explained by Muñoz and Daganzo, finding the minimal difference in shapes between the two cumulative curves is also analogous to the maximisation of the *cross-correlation* of both curves [42].

In practice, we adjust the $t(N)$ curves, not by looking for the correct time difference Δt, but instead by finding the equivalent difference ΔN:

1. Choose a morning time t_{morning} with prevailing free-flow traffic.
2. Find the corresponding number of the first vehicle that passes the upstream detector post during this morning time period, i.e., $N_{\text{up,morningBegin}}$.
3. Choose the length $N_{\text{morningLength}}$ (expressed in vehicles) of the morning time period (e.g., 400 vehicles, which corresponds to some 10 minutes on a three-lane motorway); this allows us to retrieve the number of the last vehicle that passes the upstream detector, i.e., $N_{\text{up,morningEnd}} = N_{\text{up,morningBegin}} + N_{\text{morningLength}} - 1$.
4. Setup a region in which we will synchronise the cumulative curves:

$$N_{\text{syncRegionBegin}} = N_{\text{up,morningBegin}} - \frac{N_{\text{syncRegionLength}}}{2}, \tag{24}$$

$$N_{\text{syncRegionEnd}} = N_{\text{up,morningEnd}} + \frac{N_{\text{syncRegionLength}}}{2}. \tag{25}$$

If we set $N_{\text{morningLength}} = 400$ vehicles, then we choose the length of the synchronisation region slightly higher, e.g., 500 vehicles. Note that these values should not be taken to large, as this will introduce a systematic error due to the possibly differential tuning of both upstream and downstream detector posts.

5. Define the following objective function that represents the total squared difference of the travel time and the average time difference between both curves $t_{\text{up}}(N)$ and $t_{\text{down}}(N)$:

$$f(n) = \int_{i=1}^{N_{\text{morningLength}}} \left[t_{\text{down}}(n+i) - t_{\text{up}}(i) - \overline{\Delta t}(n)\right]^2 dN, \tag{26}$$

with the $\overline{\Delta t}(n)$ being the average time difference, defined as follows:

$$\overline{\Delta t}(n) = \frac{1}{N_{\text{morningLength}}} \sum_{i=1}^{N_{\text{morningLength}}} \left[t_{\text{down}}(n+i) - t_{\text{up}}(i)\right]. \tag{27}$$

Within the region $[N_{\text{syncRegionBegin}}, N_{\text{syncRegionEnd}}]$, we now search for the value of n_{min} at which the objective function is minimal (note that during the optimisation, we approximate the integral with a discrete sum over N):

$$\min_{n \in [N_{\text{syncRegionBegin}}, N_{\text{syncRegionEnd}}]} f(n). \tag{28}$$

In Fig. 22 we graphically illustrate how the synchronisation mechanism works: the $t_{\text{down}}(N)$ curve is shifted with respect to the $t_{\text{up}}(N)$ curve within the synchronisation region until the objective function becomes minimal.

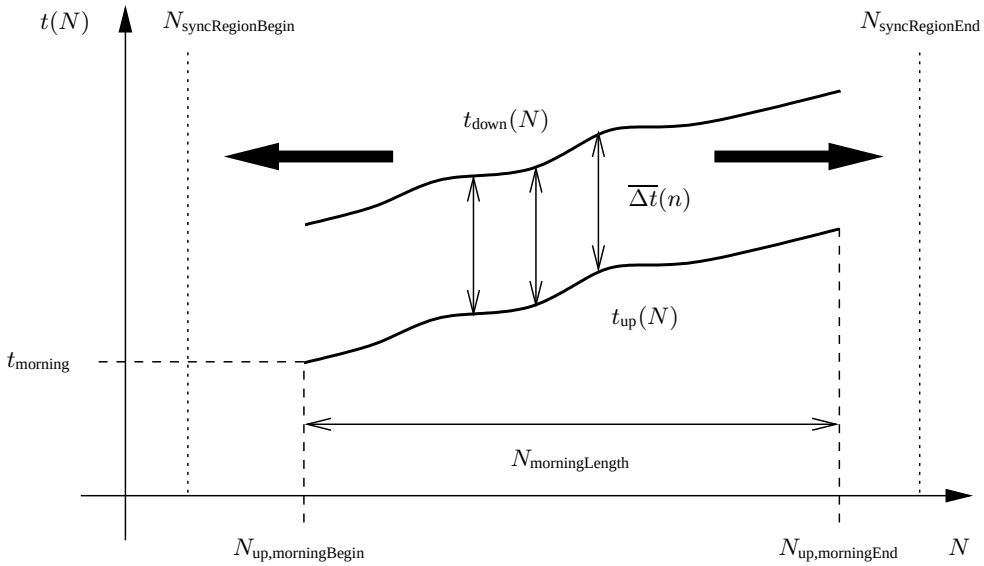

FIG. 22: A graphical illustration of the synchronisation mechanism: the $t_{\text{down}}(N)$ curve is shifted with respect to the $t_{\text{up}}(N)$ curve within the synchronisation region until the objective function becomes minimal; $\overline{\Delta t}(n)$ represents the average time difference between the two curves. After the minimisation, the shapes of both curves are qualitatively the same.

Finally, if we assume that all vehicles are driving with the same mean speed (i.e., the free-flow traffic regime), then the shapes of the curves are qualitatively the same. As such, the difference ΔN for which the two shapes differ the least is defined as follows:

$$\Delta N = N_{\text{up,morningBegin}} - (N_{\text{syncRegionBegin}} + n_{\text{min}} - 1). \tag{29}$$

The two curves are now synchronised, because $t_{\text{down}}(N) = t_{\text{up}}(N - \Delta N)$ holds for the selected morning period.

Correcting for systematic errors

When comparing cumulative curves from both upstream and downstream detector posts, it is possible that they exhibit a bias towards under- or overcounting of vehicles, e.g., the downstream post consistently counts more vehicles that the upstream post. In order to compensate for these systematic errors, we apply the same synchronisation mechanism, but this time for an evening period with free-flow traffic. After this step, we now have $\Delta N_{\text{morning}}$ and $\Delta N_{\text{evening}}$ available. In the case of bias, $\Delta N_{\text{morning}} \neq \Delta N_{\text{evening}}$, so the idea is to adjust the counts of $t_{\text{down}}(N)$ so that it counts the same number of vehicles as $t_{\text{up}}(N)$ does. There are two possibilities:

- $|t_{\text{down}}(N)| \geq |t_{\text{up}}(N)| \Longrightarrow$ we shorten $t_{\text{down}}(N)$ as follows:

$$\forall n \in |t_{\text{up}}(N)| : t_{\text{down}}(n) \leftarrow t_{\text{down}}(\left\lceil n \frac{|t_{\text{down}}(N)|}{|t_{\text{up}}(N)|} \right\rceil) \tag{30}$$

- $|t_{\text{down}}(N)| < |t_{\text{up}}(N)| \Longrightarrow$ we enlarge $t_{\text{down}}(N)$ by means of interpolation. Let us define the following index:

$$n_{\text{new}} = n \frac{|t_{\text{down}}(N)|}{|t_{\text{up}}(N)|}, \tag{31}$$

  then the interpolation assigns the new indices as follows:

$$\begin{aligned} \forall n \in |t_{\text{up}}(N)| : t_{\text{down}}(n) \leftarrow t_{\text{down}}(\lfloor n_{\text{new}} \rfloor) + \\ (t_{\text{down}}(\lceil n_{\text{new}} \rceil) - t_{\text{down}}(\lfloor n_{\text{new}} \rfloor)) \cdot |[n_{\text{new}}] - n_{\text{new}}|. \end{aligned} \tag{32}$$

An underlying assumption in the previous derivation is that the measurement error of the detectors occurs uniform in time, and is not related to the current traffic regime (e.g., free-flow versus congested traffic). Note that in similar spirit, Muñoz and Daganzo correct for bias by explicitly multiplying the cumulative curves of each detector post with a detector-specific correction factor, such that both stations count the same number of vehicles [42].

3. Estimating the distribution of the travel time

After synchronising both upstream and downstream cumulative curves and correcting for systematic errors, the travel time is now equal to:

$$T(N) = t_{\text{down}}(N) - t_{\text{up}}(N). \tag{33}$$

It is also possible to derive the space-mean speed, which is equal to the ratio of the distance K between both detector stations and the average travel time $\overline{T}(N)$ of all the vehicles [39]:

$$\overline{v}_s = \frac{K}{\overline{T}(N)}. \tag{34}$$

Once all travel times are known for all vehicles during the measurement period, we can estimate the empirical distribution of the travel time by creating histograms [39].

C. Indicators of reliability

Considering the methodology elaborated upon in the previous section, we now give some examples of estimating the travel time for a real-world example. To this end, we first give an overview of the area that encompasses our case study. We then estimate the travel time for a motorway within this case study, indicating what the expected travel time will be, as well as its deviation. Finally, we construct several reliability maps that show for a complete motorway the locations of recurrent congestion, as well as the encountered fluctuations which give indications of stop-and-go traffic.

1. Overview of the case study area

The area of our case study encompasses what we call *"Flanders' Triangle"* (see also Fig. 23); its corners are formed by three major cities, called Antwerpen, Gent, and Brussel (which is the capital of Belgium). Three bidirectional motorways are connecting these cities, i.e., the **E17** (having a length of approximately 2×54 km between Antwerpen and Gent), the **E40** (having a length of approximately 2×43 km between Gent and Brussel), and the **E19** (having a length of approximately 2×34 km between Brussel and Antwerpen). These motorways connect to the cities via ring roads; the **R0** around Brussel, the **R1** around Antwerpen, and the **R2** around Gent.

All motorways are composed of three lanes in each direction, with some exceptions such as the E19 which has a part where there are four lanes (separated in two) in each direction. The small grey filled circles represent the locations of the single loop detectors (a minority of the detectors are cameras). We collect traffic flow measurements for the year 2003, based on data stemming from 288 detectors (clockwise around the Triangle) and 264 detectors (counter-clockwise around the Triangle), which gives a grand total of 552 detectors.

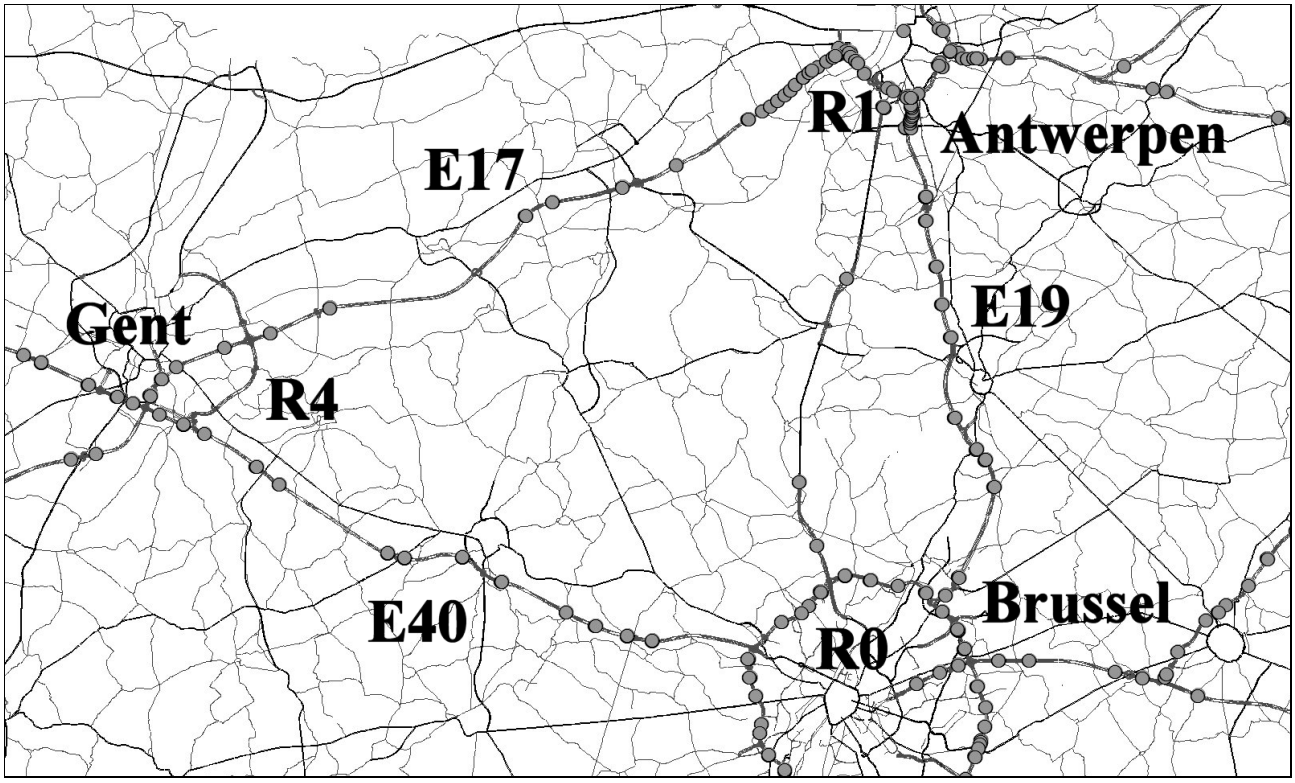

FIG. 23: The area of our case study, encompassing what we call *"Flanders' Triangle"*; its corners are formed by the three major cities Antwerpen, Gent, and Brussel, which are connected by the three bidirectional motorways, i.e., the E17 between Antwerpen and Gent, the E40 between Gent and Brussel, and the E19 between Brussel and Antwerpen. These motorways connect to the cities via ring roads; the R0 around Brussel, the R1 around Antwerpen, and the R2 around Gent. Within our case study, we collect traffic flow measurements for the year 2003, based on data stemming from 552 detectors in total (represented by the grey filled circles).

2. Travel time reliability

We now apply our methodology for off-line travel time estimation as explained in section III B, to a small real-world example. To this end, we take the following steps:

1. We chose two detector posts that demarcate a road section with no intermediate on- or off-ramps (hence we have conservation of the number of vehicles). For convenience, we set all sentinel values of the flows, occupancies, and mean speeds equal to zero (see section I B for an explanation of these sentinel values in the database).

2. Next, the flows, occupancies, and mean speeds of all individual lanes at all detector posts are aggregated together. Note that we normalise the occupancy to obtain the average occupancy across all lanes [39].

3. We then collect all groups of similar weekdays (e.g., all Mondays, Tuesdays, ...), as described in section I C.

4. For each individual day in each group of these weekdays, we consider all 60 $\times$ 24 = 1440 measurements for each lane detector. Using equation (22), we then construct the cumulative curves $N_{\text{up}}(t)$ and $N_{\text{down}}(t)$ for both detector posts, convert them to equivalent $t_{\text{up}}(N)$ and $t_{\text{down}}(N)$ curves while synchronising them using the methodology explained in section III B 2. Finally, we derive the travel times according to equation (33).

5. Based on these travel times, we construct histograms that represent the distributions of the travel times for the road section on the specified days, including the means and standard deviations of these travel times, as well as the 90% percentile. We also consider the medians and the median absolute deviations (MAD). Note that we include these latter robust estimators, in order to eliminate outliers such as special occasions, accidents and incidents, ...

6. Finally, all histograms belonging to the same day of the week (e.g., all Mondays) are averaged to obtain an estimate of the distribution of travel times on a typical day.

As a real-world example, we consider a three-lane section from the E40 motorway between Brussel and Gent (in the direction of Gent), located between the detector posts ate Erpe-Mere and Wetteren; the distance K between these two posts is 8.1 kilometres. For this section, we estimate the travel times on a typical Monday and a typical Friday, each time during the morning (08:00 – 10:00) and evening rush hours (17:00 – 20:00). The resulting histograms are shown in Fig. 24; Table III summarises the resulting statistics.

Whereas the mean and the median tell us what the expected travel time on the section will be, the standard deviation and the median absolute deviation give a clue about the expected fluctuations around this travel time [68]. The more reliable a road section is, the smaller the fluctuations will be. The histograms show that for a typical Friday, the distribution of the travel time is in general narrower than on a Monday. From the statistics, we can see that the mean and median values of the travel times lie close to each other; their deviations however show a difference. The reason for this is that our methodology is quite capable of estimating a qualitatively good travel time. Furthermore, there is no real difference between the expected travel time during the evening rush hours on a Monday and a Friday; there is however an increase in the deviation for a Monday.

With respect to the reliability of the travel time, we also note the evolution of the 90% and 95% quantiles; they indicate that a traveller can expect, with a probability of 90, respectively 95 percent, a travel time less than the percentiles. As can be seen in Table III and the histograms in Fig. 24, the section has more robustness against congestion (both recurrent and non-recurrent due to incidents) on a Friday than on a Monday. Furthermore, we can also spot an asymmetry between the morning and evening rush hours on a Friday. Further investigation of this phenomenon revealed that it was caused by some outliers (they could be cause by accidents, but this was not confirmed), which led to an increase in the calculated travel times.

Period	*Mean*	*Median*	*Std. Dev.*	*MAD*	*90%*	*95%*
Monday morning	5.1734	5.1661	0.5390	0.4295	7.2335	7.6579
Monday evening	4.9795	4.9712	0.5631	0.4546	7.4946	8.9535
Friday morning	4.7165	4.6630	0.4815	0.4029	6.1798	7.1738
Friday evening	4.9540	4.9743	0.4037	0.3281	5.7519	5.7519

TABLE III: The means and medians of the travel times, the corresponding standard deviations and median absolute deviations (MAD), as well as the 90% and 95% quantiles for the E40 motorway example from the case study.

3. Constructing reliability maps

Another aspect related to reliability of a transportation system, is expressed by on the one hand the expected mean speed a traveller will encounter, and on the other hand how the mean speed varies from one day of the week to another. To this end, we provide a methodology that allows us to construct time-space diagrams of all similar weekdays. As such, these containing 'typical days' clearly show the locations where *recurrent congestion* occurs. After explaining our technique, we apply it to some of our roads in the case study, i.e., the E19 between Antwerpen and Brussel and the R0 ring road around Brussel.

Our methodology encompasses the following steps:

1. For all motorways and ring roads in the case study, we consider all lane detectors at all detector posts. For convenience, we set all sentinel values of the flows, occupancies, and mean speeds equal to zero (see section I B for an explanation of these sentinel values in the database).

2. We then collect all groups of similar weekdays (e.g., all Mondays, Tuesdays, ...), as described in section I C.

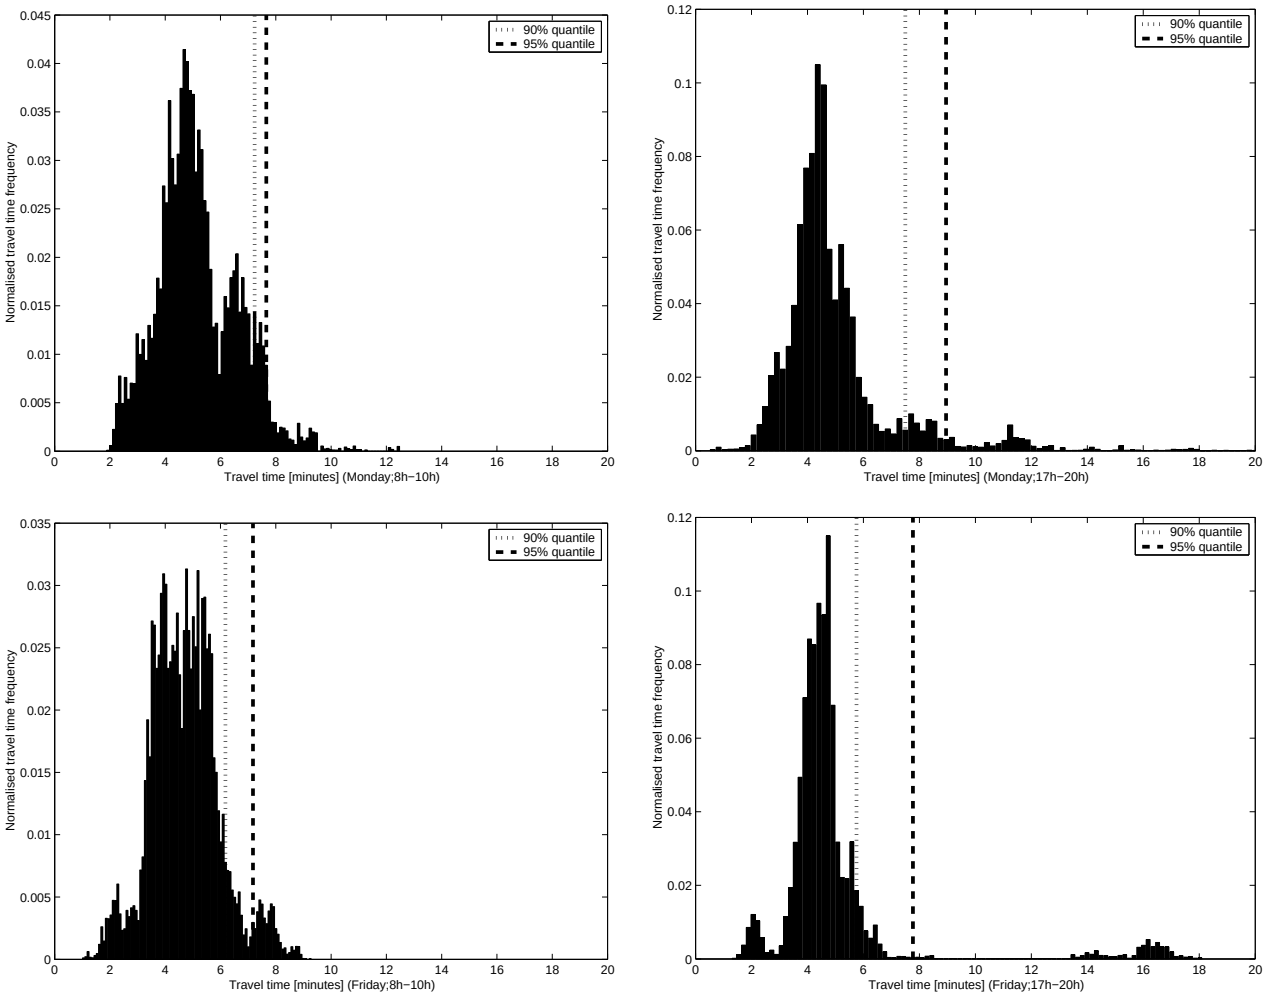

FIG. 24: Histograms of the estimated travel time for the E40 motorway example from the case study; the dotted line indicates the 90% quantiles, whereas the dashed line indicates the 95% quantile. Note how the distribution of the travel time is more narrow for a typical Friday then a typical Monday.

3. For each individual day in each group of these weekdays, we consider all $60 \times 24 = 1440$ measurements for each lane detector; we then calculate the median over all days for each measurement separately, as well as the median absolute deviation (MAD). See also Fig. 25 for a graphical depiction of this process. Note that we use these robust estimators instead of the classic mean, in order to eliminate outliers such as special occasions, accidents and incidents, ...

4. Next, the flows, occupancies, and mean speeds of all individual lanes at all detector posts are aggregated together. Note that we normalise the occupancy to obtain the average occupancy across all lanes [39].

5. Once all aggregated measurements are available, we collect them in time-space diagrams that each contain the medians and MADs of either the flows, the occupancies, or the speeds.

6. Finally, we apply Treiber and Helbing's tempo-spatial filter to the previously constructed time-space diagrams containing the medians; this filter reduces the noise in the measurements as it essentially is a low-pass filter (LPF), whilst interpolating locations and times between detectors in the time-space plane. The filter is also adaptive, in the sense that it keeps track of the current traffic conditions surrounding a measurement point: instead of just considering a points nearest Euclidean neighbours, the filter takes into account the direction of information flow (i.e., forward moving in free-flow traffic and backward moving in congested traffic) [61]. As a result, the traffic flow characteristics are more clearly pronounced, increasing the visibility of the resulting diagrams.

Applying the previously explained methodology to some of our roads in the case study, we obtain the filtered time-space diagrams in Fig. 26 and Fig. 28 for the E19 between Antwerpen and Brussel and the R0 ring road around Brussel (extended with a part of the E19 between Brussel and Mons), respectively. In all maps, the driving direction is upwards, while time advances to the right. Note that due to the spatial sparseness of the detectors in the Flanders' region (they are

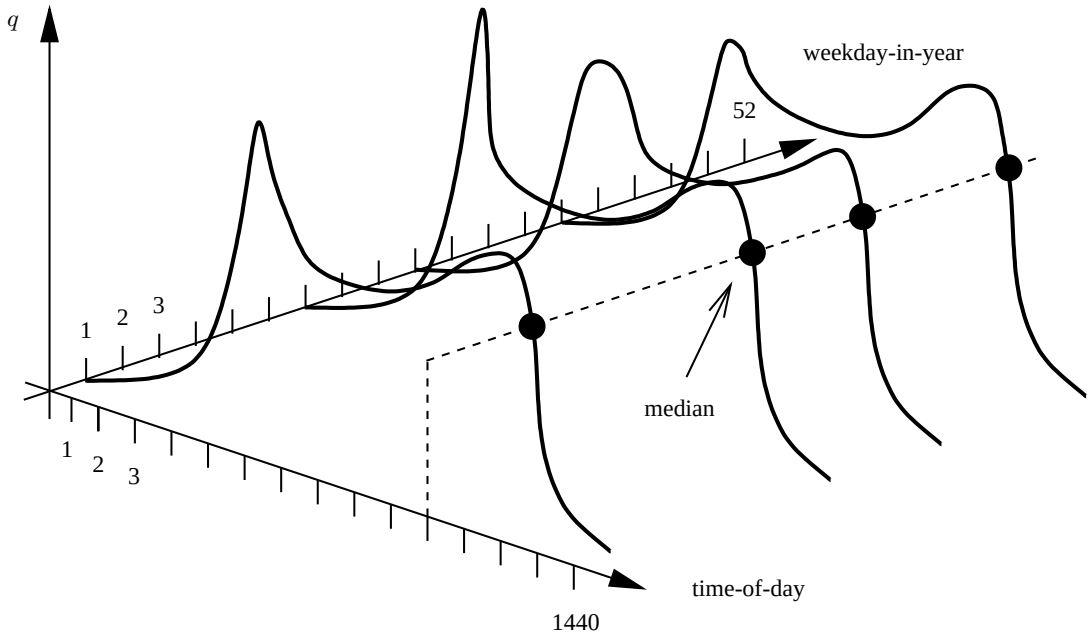

FIG. 25: For each individual day in each group of similar weekdays (e.g., all Mondays), we consider all 60 × 24 = 1440 measurements for each lane detector; we then calculate the median over all days for each measurement separately, as well as the median absolute deviation (MAD). Note that the graph shows the calculation for one detector; we for purposes of illustration, we plot the daily profiles of four distinct weekdays.

only located right before and right after an on- and off-ramp), it is possible that we miss certain jams that originate and dissolve in a motorway section without ever being recorded by an upstream or downstream detetector.

In practice, a good indicator for congestion can be found in the occupancy. However, because the mean speed is tied uniquely to the density, and hence also to the occupancy, we use its median and variance as the main macroscopic characteristics in our reliability maps.

In these maps, dark spots represent regions where the mean speed is rather low, which is indicative of congestion: jams start forming at the top part of these structures, they grow downwards and recede upwards. Because the mean speed is considered over all similar weekdays, the maps of these ‘typical days’ clearly show the locations where *recurrent congestion* occurs. The information contained in these charts thus gives an indication of the expected mean speed a traveller will encounter at a certain moment and location in the time-space diagram of a motorway. Anyone who wants to avoid traffic congestion, should therefore try to stay clear of the dark spots; this can be accomplished by not entering the motorway and thus choosing an alternative route, and/or by selecting another departure time.

Another interesting quantity to consider is the variance, as plotted in the maps in Fig. 27 and Fig. 29. These maps represent how the mean speed varies from one day of the week to another; dark spots now represent regions where a traveller can expect strong fluctuations in the traffic pattern (note that no filtering was applied to these maps).

We now discuss the results of our case study; in Figures 26, 27, 28, and 29, the left parts each time relate to a typical Monday, whereas the right parts relate to a typical Friday.

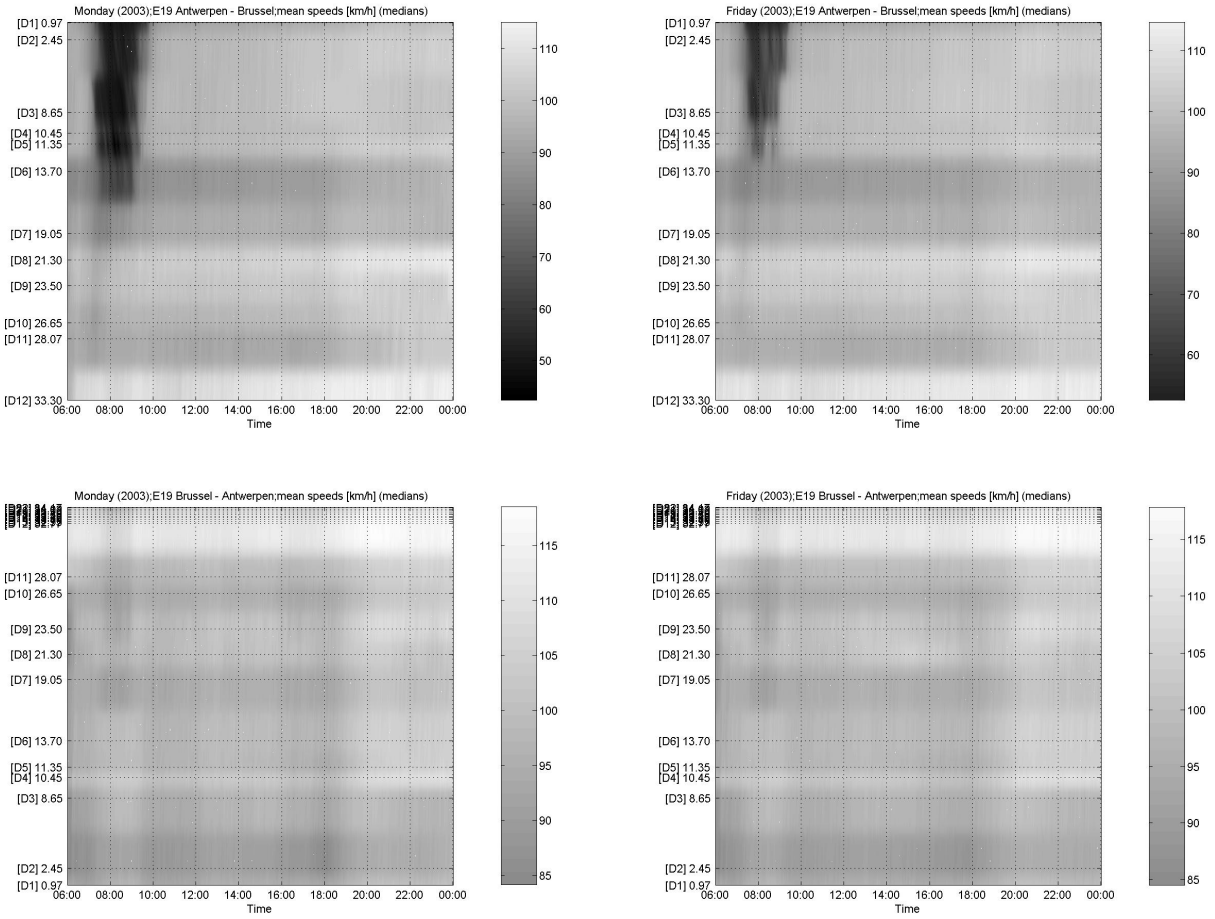

FIG. 26: Time-space diagrams showing the evolution of the mean speed and recurrent congestion on a typical Monday (*left* and a typical Friday (*right*), for the E19 motorway in the direction of Brussel (*top*) and the direction of Antwerpen (*bottom*). Note the region of severe congestion during the Monday morning rush hour near Brussel (top parts) between detectors D1 and D6; in contrast, the Friday morning rush hour is less pronounced. There is also a band of slower traffic between detectors D6 and D7, where the complex near Mechelen, a major city located right next to the motorway, is located.

E19 motorway between Antwerpen and Brussel

- The morning rush hour when entering Brussel on a Monday (top-left part of Fig. 26), exhibits severe congestion over a length of some 16 kilometres (ending between detectors D6 and D7, located at the complex near Mechelen-Zuid and Mechelen-Noord), during from 07:00 until 09:30. For a Friday (top-right part of Fig. 26), the length is reduced to some 13 kilometres (ending at detector D6, located at Mechelen-Zuid). In both cases, the congestion starts at the connection of the E19 with the R0 ring road near Vilvoorde (detector D1). Note that there is no evening rush hour for traffic in the direction of Brussel.

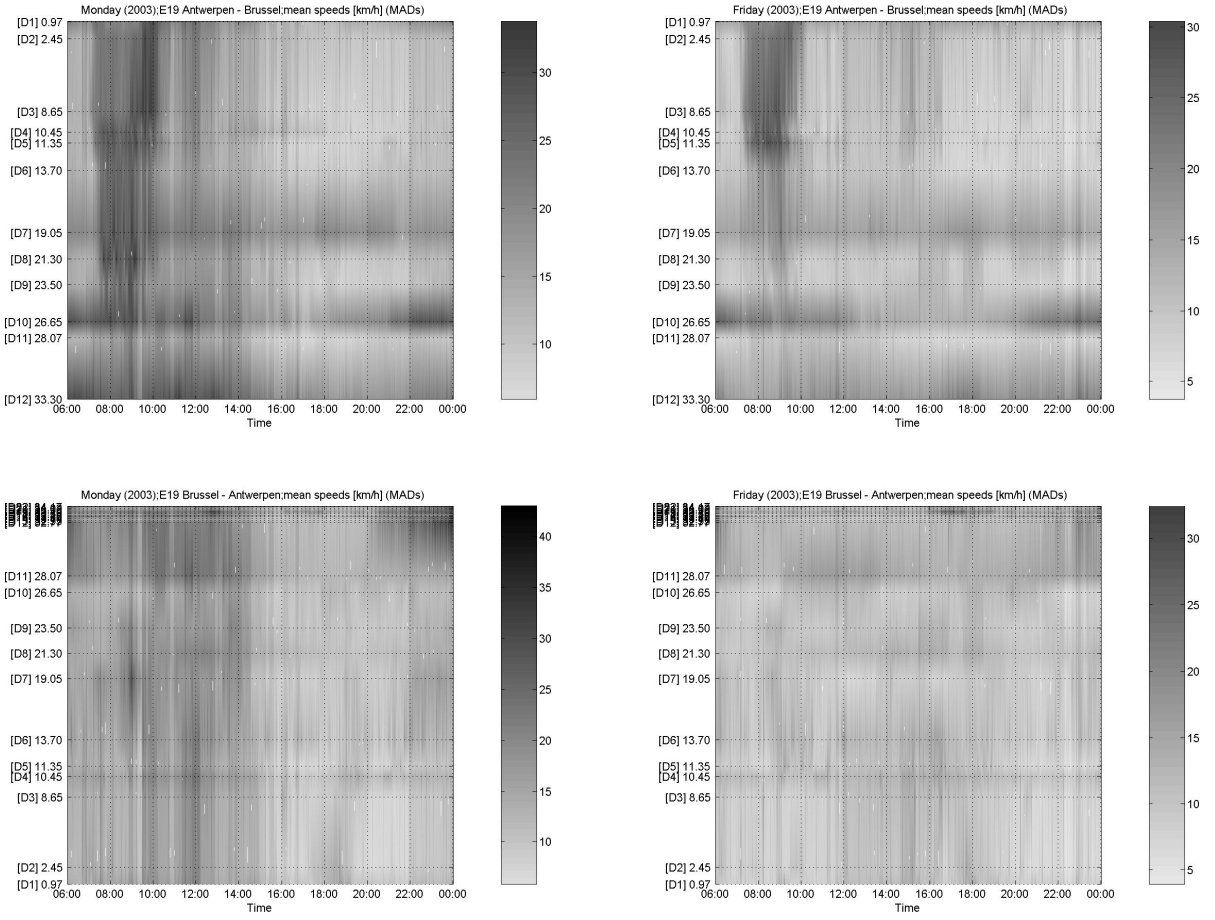

FIG. 27: Time-space diagrams showing the evolution of the median absolute deviation (MAD) of the mean speed on a typical Monday (*left* and a typical Friday (*right*), for the E19 motorway in the direction of Brussel (*top*) and the direction of Antwerpen (*bottom*). Note how the fluctuations in the top parts have a magnitude of approximately 30 to 35 km/h, probably indicating heavy stop-and-go traffic during the morning rush hour.

Within this congested region, the mean speed drops to some 50 km/h; the fluctuations in the top-left and top-right parts of Fig. 27 have a magnitude of approximately 30 to 35 km/h, which leads us to the conclusion that there probably is heavy stop-and-go traffic during the morning rush hour. Note the increased fluctuations near detector D10, located at Kontich: right before this point, the E19 was split up into 2×2 lanes, separated by a verge. Detector D10 lies right after an on- and off-ramp where the 2×2 lanes join together again.

Outside the congested region, the mean speed of traffic is approximately 90 km/h, with a small speed-up to some 100 km/h between detectors D7 (Mechelen-Noord) and D10 (Kontich). The darker region between detectors D6 and D7 represents traffic at the complex near Mechelen, which is a major city located right next to the motorway.

- Considering the E19 in the direction of Antwerpen (bottom-left and bottom-right parts of Fig. 26), we note that there is no associated morning or evening rush hour present. Instead, we observe that the mean speed approximately lies between 90 and 100 km/h. A mild slowdown towards some 95 km/h is spotted from detector D7 (Mechelen-Noord) on when entering Antwerpen.

  An interestingly observation is the light band in both parts: this band represents an increase in the mean speed to some 110 km/h. At this point, traffic enter the Craeybeckx tunnel which has four lanes; in this tunnel, traffic

is neatly split up for the three outgoing directions towards the city centre of Antwerpen, the R1 ring road around Antwerpen, and the direction towards The Netherlands.

Considering the bottom-left and bottom-right maps in Fig. 27, we note that there are no clearly visible large fluctuations, except for a small increase from detector D7 (Mechelen-Noord) on when entering Antwerpen, corresponding to the previously mentioned area where mild congestion occurs in the morning.

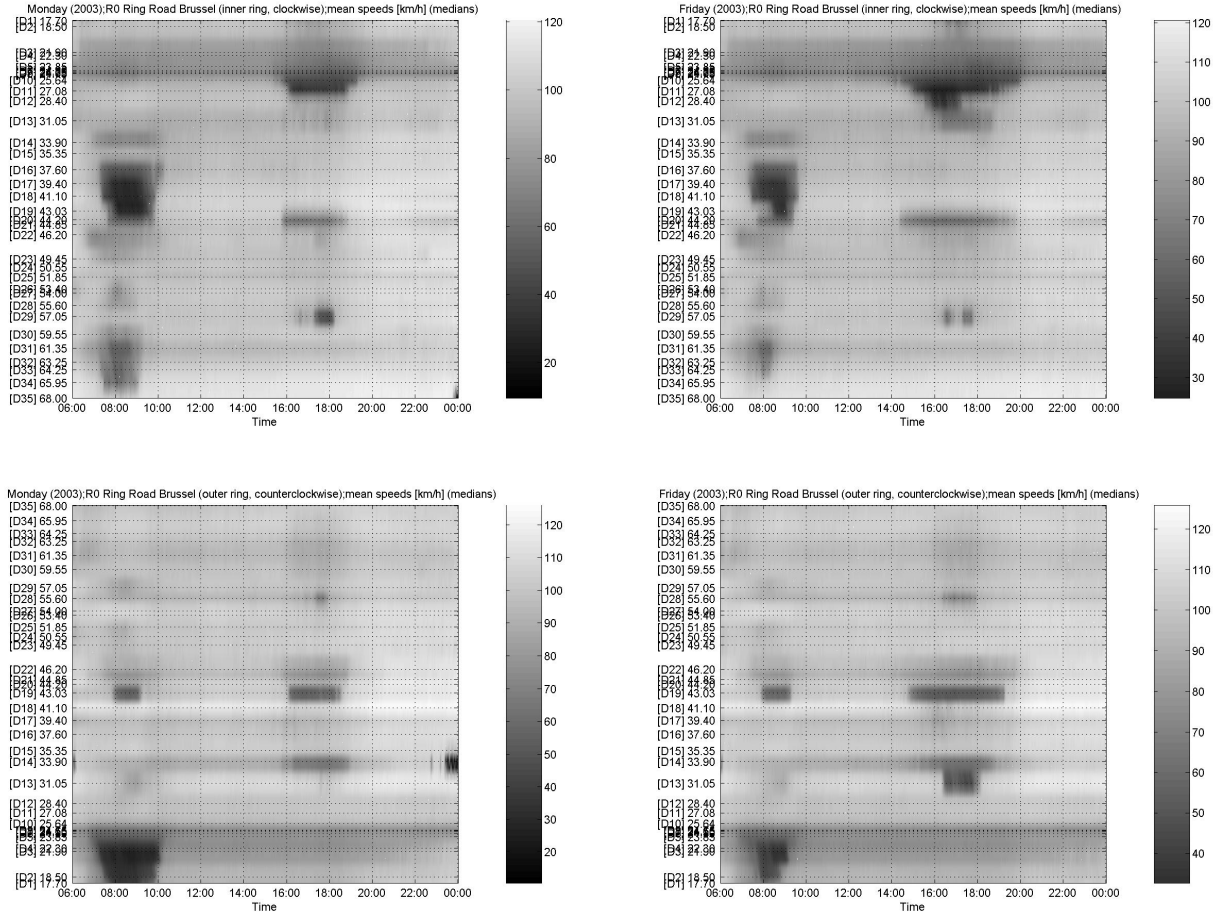

FIG. 28: Time-space diagrams showing the evolution of the mean speed and recurrent congestion on a typical Monday (*left* and a typical Friday (*right*), for the R0 clockwise inner ring road around Brussel (*top*) and the counter-clockwise outer ring road (*bottom*). Note the area of severe congestion during the morning rush hour between detectors D15 and D21 (i.e., the viaduct over Vilvoorde, spanning the R0 near Strombeek-Bever until the detectors at Wemmel and Merchtem where the A12 joins the R0), as well as the congestion during the evening rush hour between detectors D3 and D12 (i.e., Vierarmenkruispunt, Tervuren, Wezembeek-Oppem near the junction with the E40 motorway in the direction of Luik).

R0 ring road around Brussel

- During the morning rush hour, we can spot a region of severe congestion on the inner-ring road (i.e., the top-left and top-right parts of Fig. 28) between detectors D15 and D21; this corresponds to the region at the viaduct over Vilvoorde (see the satellite image in the bottom-left part of Fig. 30), spanning the R0 near Strombeek-Bever until the detectors at Wemmel and Merchtem where the A12[81] between Antwerpen and Brussel joins the R0 (see the satellite image in the top-left part of Fig. 30). The mean speed here lies around 30 km/h. The congested area covers some 8 kilometres, during a period between 07:00 and 10:00.

  Right above and below this congested location, there are two regions in which the mean speed is approximately 80 km/h; the upper region is located near detector D14 (Machelen) where the E19 motorway between Antwerpen and Brussel joins the R0. The lower region is located near detector D22 (Merchtem) where the E40 motorway between Gent and Brussel joins the R0 (see the satellite image in the top-right part of Fig. 30). There is also a small region of slower traffic (60 km/h) between detectors D31 (Beersel) and D34 (Huizingen), where the R0 blends into the E19 between Brussel and Mons in the Walloon region of Belgium.

  The evening rush hour is clearly visible between 15:30 and 19:00; within it, the mean speed drops to some 40 km/h between detectors D3 and D4 (Vierarmenkruispunt, see the satellite image in the bottom-right part of Fig. 30),

D5 and D10 (Tervuren), and D11 and D12 (Wezembeek-Oppem near the junction with the E40 motorway in the direction of Luik). Note that we can see a darker band in the maps between detectors D2 and D9 (viaduct Vilvoorde) during the entire day; at this location, there is a lower speed limit which results in a mean speed of some 80 km/h. Finally, there are also two small areas of congestion near detectors D29 (Anderlecht) and D20 (Wemmel), where the mean speed drops to some 60 km/h.

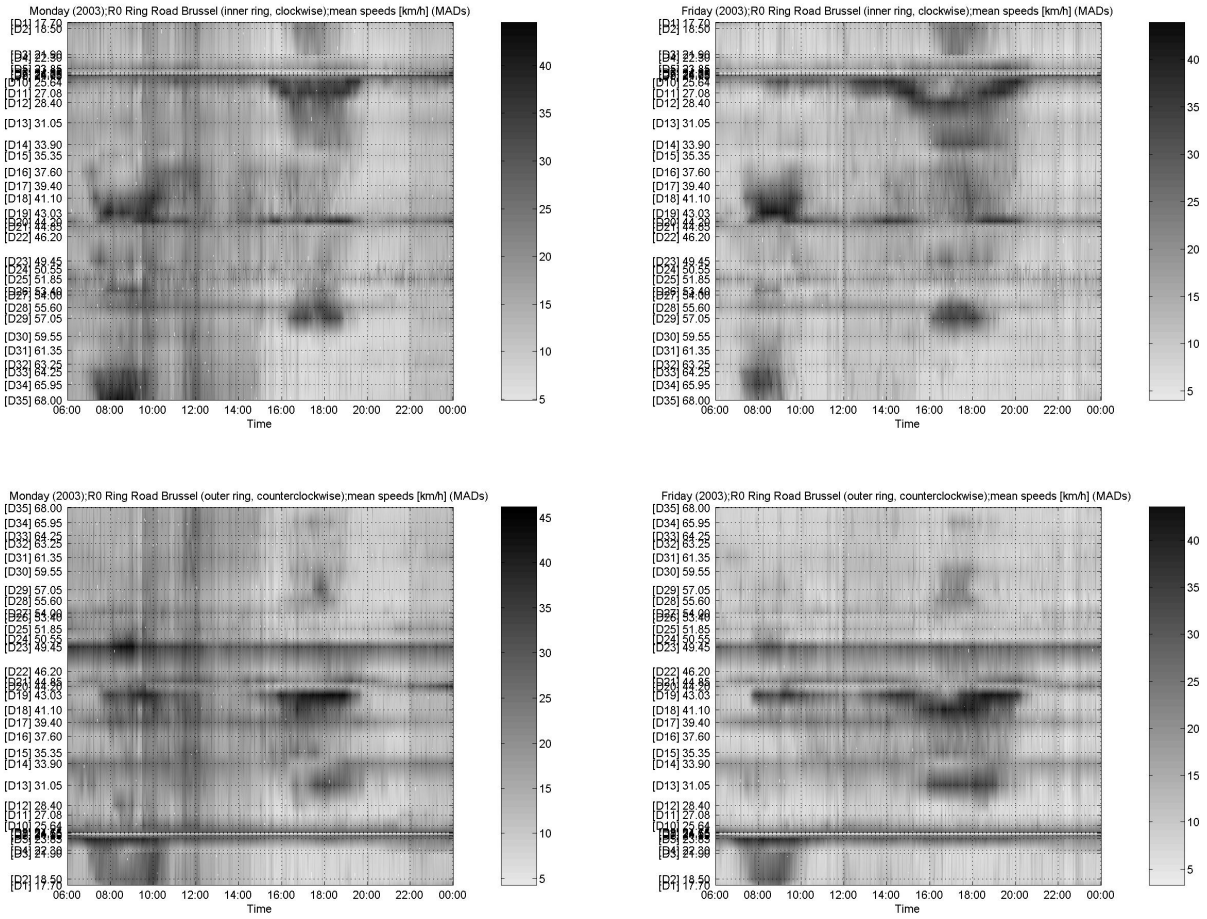

FIG. 29: Time-space diagrams showing the evolution of the median absolute deviation (MAD) of the mean speed on a typical Monday (*left* and a typical Friday (*right*), for the R0 clockwise inner ring road around Brussel (*top*) and the counter-clockwise outer ring road (*bottom*). Note the fluctuations in the top parts near detector D10; for Mondays, there appear to be more fluctuations in the morning period than for Fridays; conversely, the evening rush hour is more sensitive to disturbances on a Friday then on a Monday.

Comparing the left and right parts of the maps, we can see that in general the congestion is worse on a Friday then on a Monday; the evening rush hour previously discussed now already starts at the early time of 14:00, lasting until 20:00. Note that with respect to the mean speed, the congestion also results in low measurements of some 20 km/h. Furthermore, during the entire day, the darker band in the maps between detectors D2 and D9 (viaduct Vilvoorde) is more pronounced, indicative of slower traffic (60 km/h).

Looking at the top-left and top-right parts of Fig. 29, we note an important observation with respect to the reliability of the mean speed: on a Friday, there occur a lot of fluctuations near detector D10 (Tervuren), starting already at 12:00, lasting until 15:00. They reprise at 18:00, lasting until 20:30. Considering both parts, we can also see that for Mondays, there appear to be more fluctuations in the morning period than for Fridays; conversely, the evening rush hour is more sensitive to disturbances on a Friday then on a Monday.

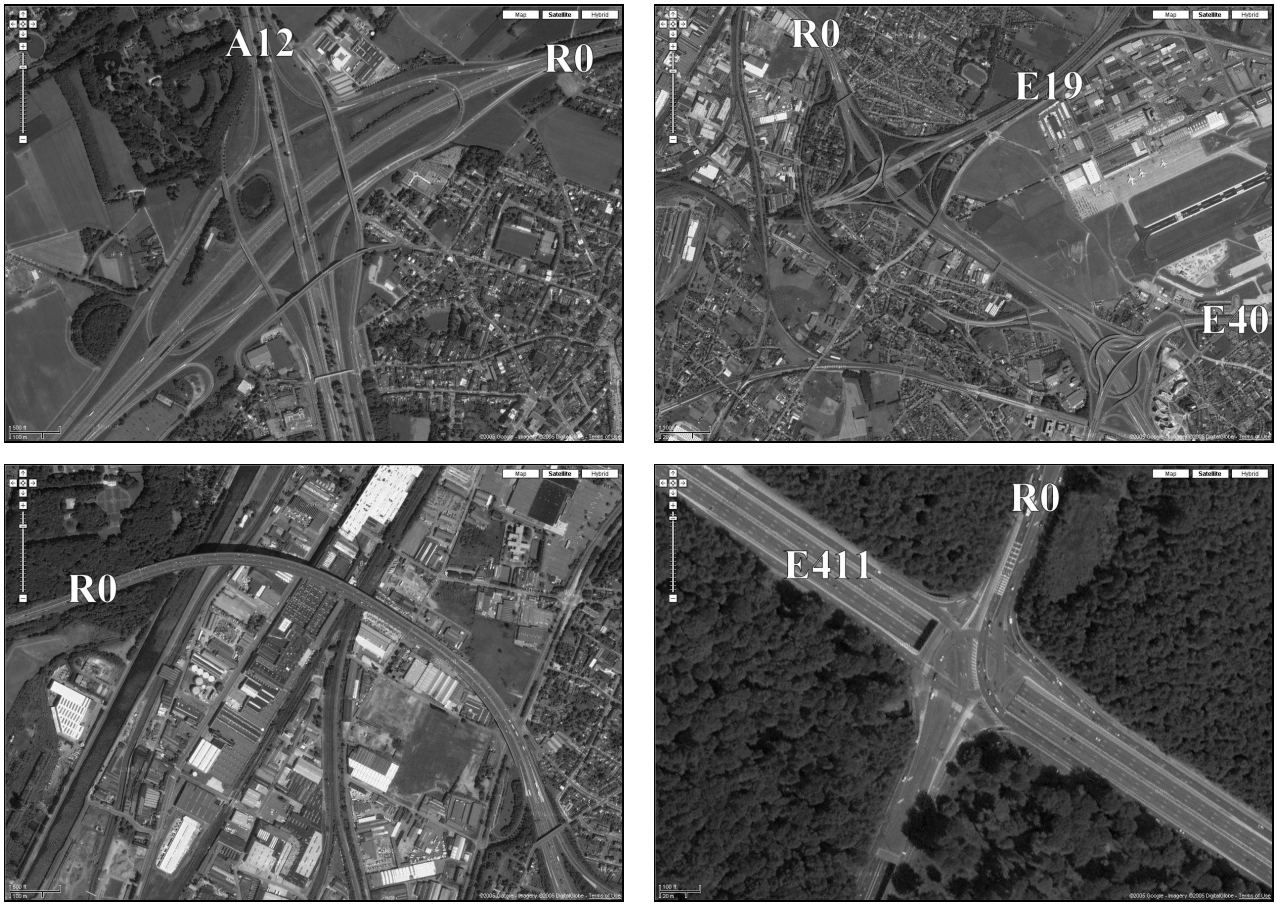

FIG. 30: Some satellite images of the more severely congested locations on the R0 ring road around Brussel *Top-left:* the junction between the A12 (between Antwerpen and Brussel) and the R0. *Top-Right:* the junctions between the E19 (between Antwerpen and Brussel), the E40 (between Brussel and Luik) and the R0. *Bottom-left:* the R0 on the large bridge over the city Vilvoorde and the Renault factory. *Bottom-Right:* the junction Vierarmenkruispunt ('4-arms crossing') between the E411 (between Brussel and Namen) and the R0 (all images reproduced after [26]).

- Considering the outer-ring road (i.e., the bottom-left and bottom-right parts of Fig. 28), we note that most congestion during the morning rush hour between 07:00 and 10:00 occurs between detectors D1 (Groenendaal) and D9 (Tervuren), where the junction of the R0 and the E411 (between Brussel and Namen) at the Vierarmenkruispunt ('4-arms crossing') is located. Here, the mean speed drops to some 30 km/h, indicating severe congestion. The outer ring also exhibits a slower speed of some 50 km/h near detector D19 (Strombeek-Bever). The same darker band of slower traffic is also visible between detectors D2 and D9 (viaduct Vilvoorde) during the entire day.

  With respect to the evening rush hour, we can see some congestion (60 km/h) at detector D19 (Strombeek-Bever), between 16:00 and 18:00 approximately. Also note the same level of congestion near detector D14 (Machelen) where the E19 motorway between Antwerpen and Brussel joins the R0. Looking at the difference between Monday and Fridays, we can see that on a Friday, the morning rush hours has a smaller duration than on a Monday. The opposite is true for the evening rush hour, which now lasts from 15:00 until 19:00.

  Looking at the fluctuations of the mean speed in the bottom-left and bottom-right parts of Fig. 29, we can make the same remarks as for the inner-ring road.

Note that with respect to the locations of congestion, we find a close match between those discussed in our research, and the ones reported in the *"Belgium's Congestion Top-25"* of Logghe and Vanhove [38].

IV. CONCLUSIONS

This paper provided several techniques which can assist in the analysis of traffic flow measurements gathered on Flanders' motorways. We describe how all these measurements are obtained and how they are stored in a central database. We then discuss the quality of the measurements, from a statistical point of view by giving a technique that tracks outliers. We also provided a methodology for quickly assessing structural and incidental detector malfunctioning, by means of

creating maps that give a clear visual indication of when and where the problems occurred. Subsequently, we provided a methodology for the off-line estimation of travel times, based on flow measurements (as opposed to the much used technique based on speed measurements). Finally, we gave some reliability and robustness properties related to travel times and traffic flow dynamics, which gives us an extra instrument for the analysis of recurrent congestion.

With respect to the tackling of congestion on Flanders' road network, the current policy is to keep the congestion on the motorways, avoiding its spreading towards the underlying secondary road network. At the moment, Flanders' government has a plan of action to investigate 25 'missing links' and bottlenecks in the road network (see Fig. 31 for an overview). At the current rate of progress, the government hopes by 2009 to have accomplished one third of its goal [3].

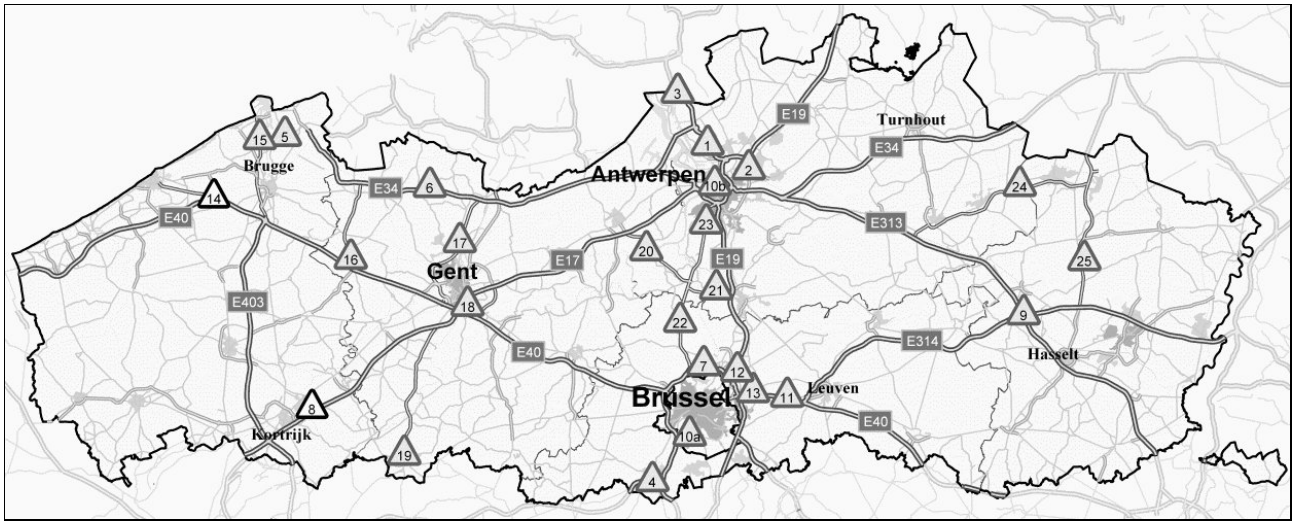

FIG. 31: An overview of the plan of action to investigate 25 'missing links' and bottlenecks in Flanders' road network; at the current rate of progress, the Flemish government hopes by 2009 to have accomplished one third of its goal [3].

In light of the *socio-economic use* of traffic data in Flanders, we note that at the moment, many information with respect to incidents is distributed by radio, mobile phone, Internet, ... to the individual road traveller. However, most of the information has a low quality, especially when it comes to structural congestion (everybody already *knows* where the congestion occurs. In this respect, there is a clear market available for supplying qualitative traffic information towards the road travellers; this is especially useful for individual people that have a high value-of-time.

Providing useful information to the road traveller is of paramount importance. With respect to important quantities such as travel times that get advertised on VMS/DRIPS systems, the agencies and corporations engaged in this business, should therefore not neglect the fact that the human interpretation of probability most likely fails. People tend to assign too high probabilities to extreme events (e.g., incidents). Furthermore, it is also important to mention the cause of a delay, as humans want to know *why* they are entering congestion; this fact should not be neglected as it is important for a driver's perception and peace of mind. Given a good enough reason, people get less annoyed by non-recurrent congestion; especially when they are supposed to take an alternative route in order to reach their destination.

Acknowledgements

Dr. Bart De Moor is a full professor at the Katholieke Universiteit Leuven, Belgium. Our research is supported by: **Research Council KUL**: GOA AMBioRICS, several PhD/postdoc & fellow grants, **Flemish Government**: **FWO**: PhD/postdoc grants, projects, G.0407.02 (support vector machines), G.0197.02 (power islands), G.0141.03 (identification and cryptography), G.0491.03 (control for intensive care glycemia), G.0120.03 (QIT), G.0452.04 (new quantum algorithms), G.0499.04 (statistics), G.0211.05 (Nonlinear), research communities (ICCoS, ANMMM, MLDM), **IWT**: PhD Grants, GBOU (McKnow), **Belgian Federal Science Policy Office**: IUAP P5/22 ('Dynamical Systems and Control: Computation, Identification and Modelling', 2002-2006), PODO-II (CP/40: TMS and Sustainability), **EU**: FP5-Quprodis,

ERNSI, **Contract Research/agreements**: ISMC/IPCOS, Data4s,TML, Elia, LMS, Mastercard.

---

[1] Jorge A. Acha-Daza and Fred L. Hall. The application of catastrophe theory to traffic flow variables. *Transportation Research B*, 28B:235–250, 1994.

[2] *Het wegenpatrimonium in cijfers*. Administratie Wegen en Verkeer, Departement Leefmilieu en Infrastructuur (LIN) van het Ministerie van de Vlaamse Gemeenschap, August 2004.
(URL: `http://wegen.vlaanderen.be/wegen/patrimonium`).

[3] *Knelpunten en 'missing links'*. Administratie Wegen en Verkeer, Departement Leefmilieu en Infrastructuur (LIN) van het Ministerie van de Vlaamse Gemeenschap, February 2006.
(URL: `http://wegen.vlaanderen.be/wegen/missinglinks`).

[4] Tom Bellemans. *Traffic Control on Motorways*. PhD thesis, Katholieke Universiteit Leuven, Department of Electrical Engineering ESAT-SCD (SISTA), May 2003.

[5] Tom Bellemans, Bart De Schutter, and Bart De Moor. On data acquisition, modelling and simulation of highway traffic. In *Proceedings of the 9th IFAC Symposium on Control in Transportation Systems*, pages 22–27, Braunschweig, Germany, June 1999.

[6] Tom Bellemans, Bart De Schutter, and Bart De Moor. Data acquisition, interfacing and pre-processing of highway traffic data. In *Proceedings of the Telematics Automotive 2000*, volume 1, pages 1–7, Birmingham, U.K., April 2000.

[7] Peter Bickel, Chao Chen, Jaimyoung Kwon, John Rice, Erik van Zwet, and Pravin Varaiya. *Measuring Traffic*, May 2004.

[8] Piet L. Bovy and Remmelt Thijs, editors. *Estimators of Travel Time for Road Networks – New developments, evaluation results, and applications*. Delft University Press, 2000. ISBN 90-407-2048-7.

[9] George E.P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. Cambridge University Press, Cambridge, 1976.

[10] George E.P. Box, Gwilym M. Jenkins, and G.C. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice Hall, Englewood Cliffs, New Jersey, 3rd edition, 1994.

[11] M.J.G. De Ceuster and L.H. Immers. *Verkeersindices voor het gebruik van het autosnelwegennet in België*. Technical report, Transport & Mobility Leuven, October 2001.

[12] Chao Chen, Jaimyoung Kwon, John Rice, Alexander Skabardonis, and Pravin Varaiya. Detecting errors and imputing missing data for single loop surveillance systems. In *Proceedings of the 82nd Annual Meeting of the Transportation Research Board*, Washington, D.C., 2003.

[13] Roland Chrobok, Oliver Kaumann, Joachim Wahle, and Michael Schreckenberg. Three categories of traffic data: Historical, current and predictive. In E. Schnieder and U. Becker, editors, *9th IFAC Symposium Control in Transportation Systems*, pages 250–255, IFAC, Braunschweig, 2000, 2000.

[14] Benjamin Coifman. Improved velocity estimation using single loop detectors. *Transportation Research A*, 35(10):863–880, 2001.

[15] Benjamin Coifman. Estimating travel times and vehicle trajectories on freeways using dual loop detectors. *Transportation Research Part A*, 36(4):351–364, 2002.

[16] Benjamin Coifman. Estimating median velocity instead of mean velocity at single loop detectors. *Transportation Research C*, 11(3-4):211–222, June 2003.

[17] Benjamin Coifman and Michael Cassidy. *Automated Travel Time Measurement Using Vehicle Lengths From Loop Detectors*. UCB-ITS-PRR-2000-12, Institute of Transportation Studies, University of California, Berkeley, August 2000.

[18] Benjamin Coifman and Michael Cassidy. Vehicle reidentification and travel time measurement on congested freeways. *Transportation Research Part A*, 36(10):899–917, 2002.

[19] Carlos F. Daganzo. The cell transmission model. part II: Network traffic. *Transportation Research B*, 29B(2):79–93, 1995.

[20] D.J. Dailey. Travel time estimates using a series of single loop volume and occupancy measurements. In *Proceedings of the 76th Annual Meeting of the Transportation Research Board*, Washington, D.C., January 1997.

[21] D.J. Dailey. A statistical algorithm for estimating speed from single loop volume and occupancy measurements. *Transportation Research B*, 33(5):313–322, 1999.

[22] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood estimation from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society, Series B*, 39:1–8, 1977.

[23] J.-P. Eckmann, S. Kamphorst, and D. Ruelle. Recurrence plots of dynamical systems. *Europhysics Letters*, 4:973–977, 1987.

[24] J.-P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics*, 57(3):617–656, 1985.

[25] A.M. Fraser and H.L. Swinney. Independent coordinates for strange attractors from mutual information. *Physical Review A*, 33(2):1134–1140, 1986.

[26] Google. *Google Maps*, February 2005.
(URL: `http://maps.google.com`).

[27] F. Hall and B.N. Persaud. Evaluation of speed estimates made with single-detector data from freeway traffic management systems. *Transportation Research Record*, 1232:9–16, 1989.

[28] Rainer Hegger, Holger Kantz, and Thomas Schreiber. Practical implementation of nonlinear time series methods: The TISEAN package. *Chaos*, (9):413–435, 1999.

[29] B. Iglewicz and D.C. Hoaglin. *How to Detect and Handle Outliers*. American Society for Quality Control, Milwaukee, Wisconsin, 1993.

[30] Joseph S. Iwanski and Elizabeth Bradley. Recurrence plot analysis: To embed or not to embed ? *Chaos*, 8:861–871, 1998.

[31] L. Jacobson, N. Nihan, and J. Bender. Detecting erroneous loop detector data in a freeway traffic management system. *Transportation Research Record*, 1287:151–166, 1990.

[32] Zhanfeng Jia, Chao Chen, Ben Coifman, and Pravin Varaiya. The pems algorithms for accurate, real-time estimates of g-factors and speeds from single-loop detectors. In *Proceedings of IEEE Intelligent Transportation Systems 2001*, pages 536–541, 2001.

[33] M.B. Kennel, R. Brown, and H.D.I. Abarbanel. Determining minimum embedding dimension using a geometrical construction. *Physical Review A*, 45:3403–3411, 1992.
[34] G. Klunder, H.H. Versteegt, and C.M.J. Tampère. *PredicTime: Model specifications and initial test results*. 037N-142-73196, TNO INRO, The Netherlands, July 2003.
[35] Wolfgang Knospe, Ludger Santen, Andreas Schadschneider, and Michael Schreckenberg. Single-vehicle data of highway traffic: microscopic description of traffic phases. *Physical Review E*, 65(05):6133–6147, May 2002.
[36] Jaimyoung Kwon, Pravin Varaiya, and Alexander Skabardonis. Estimation of truck traffic volume from single loop detector using lane-to-lane speed correlation. In *Proceedings of the 82nd Annual Meeting of the Transportation Research Board*, Washington, D.C., January 2003.
[37] R.J.A. Little and D.B. Rubin. *Statistical Analysis with Missing Data*. J. Wiley & Sons, New York, 1987.
[38] Steven Logghe and Filip Vanhove. *Het Belgische verkeer in cijfers*. 2004-01, Transport & Mobility Leuven, January 2004.
[39] Sven Maerivoet and Bart De Moor. *Traffic Flow Theory*. 05-154, Katholieke Universiteit Leuven, Department of Electrical Engineering ESAT-SCD (SISTA), July 2005.
[40] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Science of India*, 2(49), 1936.
[41] B. Mikhalkin, H. Payne, and L. Isaksen. Estimation of speed from presence detectors. *Highway Research Record*, 388:73–83, 1972.
[42] Juan Carlos Muñoz and Carlos F. Daganzo. Structure of the transition zone behind freeway queues. *Transportation Science*, 37(3):312–329, August 2003.
[43] In Jae Myung. Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology*, 47:90–100, 2003.
[44] L. Neubert, L. Santen, A. Schadschneider, and M. Schreckenberg. Single-vehicle data of highway traffic: A statistical analysis. *Physical Review E*, 60(6):6480–6490, December 1999.
[45] N.H. Packard, J.P. Crutchfeld, J.D. Farmer, and R.S. Shaw. Geometry from a time series. *Physical Review Letters*, 45:712–716, 1980.
[46] Ulrich Parlitz. Nonlinear time-series analysis. In Johan A.K. Suykens and Joos Vandewalle, editors, *Nonlinear Modeling: Advanced Black-Box Techniques*, chapter 8, pages 209–253. Kluwer Academic Publishers Boston, June 1998. ISBN 0-7923-8195-5.
[47] H.J. Payne, E.D. Helfenbein, and H.C. Knobel. *Development and testing of incident detection algorithms*. FHWA-RD-76-20, Federal Highway Administration, Washington, D.C., 1976.
[48] Karl F. Petty, Peter Bickel, Michael Ostland, John Rice, Frederic Schoenberg, Jiming Jiang, and Ya'acov Ritov. Accurate estimation of travel times from single-loop detectors. *Transportation Research A*, 32(1):1–17, 1998.
[49] A. Pushkar, F. Hall, and J.A. Acha-Daza. Estimation of speeds from single-loop freeway flow and occupancy data using cusp catastrophe theory model. *Transportation Research Record*, 1457:149–157, 1994.
[50] Peter J. Rousseeuw. Least median of squares regression. *Journal of the American Statistical Association*, 79:871–880, 1984.
[51] Peter J. Rousseeuw and Katrien Van Driessen. A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 41:212–223, August 1999.
[52] P.J. Rousseeuw and A.M. Leroy. *Robust Regression and Outlier Detection*. John Wiley & Sons, New York, 1987.
[53] D.B. Rubin. *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons, New York, 1987.
[54] Tim Sauer. Time series prediction by using delay coordinate embedding. In A.S. Weigend and N.A. Gershenfeld, editors, *Time Series Prediction: Forecasting the Future and Understanding the Past*, volume XV. Addison-Wesley, July 1993.
[55] J.L. Schafer. Multiple imputation: a primer. *Statistical Methods in Medical Research*, 8(1):3–15, March 1999.
[56] J.L. Schafer and R.M. Yucel. Missing data: our view of the state of the art. *Psychological Methods*, 7(2):147–177, 2002.
[57] Siemens. *Voertuigtellingen - Functionele Analyse*. v1.00 T643, Siemens N.V., August 1992.
[58] Siemens Nederland N.V. – Afdeling Traffic & Infra Solutions. *Siemens Lusdetector LD4*.
[59] Brian L. Smith, Billy M. Williams, and R. Keith Oswald. Comparison of parametric and nonparametric models for traffic flow forecasting. *Transportation Research Part C*, 10:303–321, 2002.
[60] Floris Takens. Detecting strange attractors in turbulence. In D.A. Rand and L.-S. Young, editors, *Dynamical Systems and Turbulence*, pages 366–381. Springer, Berlin, 1981.
[61] Martin Treiber and Dirk Helbing. Reconstructing the spatio-temporal traffic dynamics from stationary detector data. *Cooperative Transportation Dynamics*, 1(1):1–20, February 2002.
[62] Shawn M. Turner, William L. Eisele, Robert J. Benz, and Douglas J. Holdener. *Travel Time Data Collection Handbook*. Report FHWA-PL-98-035, Texas Transportation Institute, The Texas A&M University System, March 1998.
[63] J.W.C. van Lint. *Reliable travel time prediction for freeways*. PhD thesis, Technische Universiteit Delft, May 2004.
[64] Pravin Varaiya. Reducing highway congestion: An empirical approach. *European Journal of Control*, 11:301–309, 2005.
[65] Sabine Verboven and Mia Hubert. LIBRA: a MATLAB library for robust analysis. *Chemometrics and Intelligent Laboratory Systems*, 75:127–136, 2005.
[66] Verkeerscentrum Vlaanderen, Departement Leefmilieu en Infrastructuur, Administratie Wegen en Verkeer, Vuurkruisenplein 20, 2020 Antwerpen. *MINDAT – Databank ruwe verkeersdata Vlaams snelwegennet – 2001 en 2003*, February 2003.
[67] H.H. Versteegt and C.M.J. Tampère. *PredicTime: State of the art and functional architecture*. 03-7N-024-73196, TNO INRO, The Netherlands, February 2003.
[68] Pim Warffemius. Using the standard deviation of the travel time distribution as an indicator for valuing the reliability of travel time. In *ECTRI – European Conference of Transport Research Institutes, Young Researchers Seminar 2005*, The Hague, The Netherlands, May 2005. National Institute for Coastal and Marine Management.
[69] C.L. Webber and J.P. Zbilut. Dynamical assessment of physiological systems and states using recurrence plot strategies. *Journal of Applied Physiology*, 76:965–973, 1994.
[70] J.P. Zbilut and C.L. Webber. Embeddings and delays as derived from quantification of recurrence plots. *Physics Letters A*, 171:199–203, 1992.
[71] `http://www.macqel.be`

[72] `http://www.traficon.be`

[73] Note that in contrast to some other countries such as The Netherlands and Germany, our sensors do not measure microscopic variables such as time and space headways, . . . [35, 44].

[74] The estimation stems from the fact that the fundamental relation is based on the space-mean speed, whereas these single inductive loop detectors are only capable of dealing with time-mean speeds.

[75] The traffic centre in the Walloon region is located in Perex, Daussoulx.

[76] Note that we did not take into account special events (e.g., concerts, sports, . . .) or different weather conditions (e.g., summers versus winters).

[77] Smith et al. later used a similar technique, called a *naive forecast* which served as a worst-case approach for predicting traffic flows [59].

[78] In MCMC methods, the values are drawn from probability distributions based on Markov chains. These latter are discrete-time stochastic processes, in which the past is irrelevant for predicting the future.

[79] It is also possible to use *derivative coordinates* instead of delay coordinates, as was done in the original work of Packard et al. However, these derivatives have proven to be susceptible to noise and are therefore not generally useful [46].

[80] Note that we do not cover another popular approach, i.e., using floating car data (FCD) to directly measure travel times.

[81] Note that the A12 road is considered as an alternative that lies parallel to the E19 motorway between Antwerpen and Brussel. However, it does not constitute a full-fledged motorway, as it contains traffic lights and road crossings at the part located in Antwerpen.