# Hovering Data Clouds: A Decentralized and Self-organizing Information System*

Axel Wegener[1], Elad M. Schiller[2], Horst Hellbrück[1], Sándor P. Fekete[2], and Stefan Fischer[1]

[1] Institut für Telematik, Universität zu Lübeck
{wegener, hellbrueck, fischer}@itm.uni-luebeck.de
http://www.itm.uni-luebeck.de
[2] Institut für mathematische Optimierung, Technische Universität Braunschweig
{s.fekete, e.schiller}@tu-bs.de
http://www.math.tu-bs.de/mo

**Abstract.** With ever-increasing numbers of cars, traffic congestion on the roads is a very serious economic and environmental problem for our modern society. Existing technologies for traffic monitoring and management require stationary infrastructure. These approaches lack flexibility with respect to system deployment and unpredictable events (e.g., accidents). Moreover, the delivery of traffic reports from radio stations is imprecise and often outdated. In the project AutoNomos we aim at developing a decentralized system for traffic monitoring and managing, based on vehicular ad-hoc networks (VANETs). Our objective is to design a system for traffic forecasting that can deliver faster and more appropriate reactions to unpredictable events. In our design, cars collect traffic information, extract the relevant data, and generate traffic reports. A key concept are so-called *Hovering Data Clouds* (HDCs), which are based on the insight that many crucial structures in traffic (e.g., traffic jams) lead an existence that is independent of the individual cars they are composed of. The result is an elegant, robust and self-organizing distributed information system. In this paper we demonstrate first experimental results.

## 1 Introduction

Mobility and individuality are cornerstones of our modern society; this explains why car traffic has become so crucial for many aspects of our life. Unfortunately, this importance leads to ever-increasing numbers of cars; given the limited road capacities, frequent traffic hold-ups are common in urban regions and highways all over the world. Beyond the individual loss of freedom and mobility due to time spent in blocked traffic, these predicaments also have a very serious large-scale impact, due to increasing pollution and economic loss. As a consequence, ever-increasing efforts have been spent on optimization of road usage by means of traffic monitoring and management.

* Supported by DFG Focus Program “Organic Computing” (SPP 1183) project “AutoNomos” (Fe 407/11-1).

H. de Meer and J.P.G. Sterbenz (Eds.): IWSOS 2006, LNCS 4124, pp. 243–247, 2006.

Traditional online traffic forecasting systems require stationary infrastructure, such as magnetic loops, roadside relay stations, a massive central processing unit, and finally radio stations that broadcast the generated traffic reports (see [1]). Unfortunately, this centralistic approach suffers from a number of inherent design problems: a huge amount of data has to be gathered and communicated to a central server; the necessary central computations are enormous; and the results have to be distributed and implemented in a timely fashion.

Quite often, such a centralized approach is not just complicated, but also inappropriate. Many phenomena are local by nature, so they only affect a subset of traffic participants. Moreover, fast and flexible response is essential for defusing many local traffic situations. All this makes a distributed approach more flexible, faster and more accurate for dealing with dynamic, local traffic phenomena.

Recent advances in wireless communication technologies such as WLAN and GPS allow short-to-medium-range communication systems among vehicles on an ad-hoc basis. Vehicular ad-hoc networks (VANETs) aim at a number of commercial applications such as the improvement of car safety (see [2], [3]) or entertainment such as online gaming. The same equipment is also suited for decentralized traffic forecasting. In contrast to a fixed infrastructure, this approach provides precise results with cost-efficient equipment [4].

Existing centralized implementations require carefully planning with respect to the location of the sensing infrastructure. If the number of sensors is too small or their position is suboptimal, forecasting precision is compromised. On the other hand, just the situations that are critical (e.g., traffic jams) provide a sufficient number of cars to form a VANET and collect traffic statistics throughout the road, making a distributed approach simple and appealing.

Our approach provides fine granularity that improves the forecasting precision; it can also adapt more quickly to sudden, unpredictable events like accidents and traffic congestions, which defy even advanced forecasting models (see [5]).

This work is part of a project that aims at develop a system that provides the drivers with online reports of the current (and upcoming) traffic condition.

## 2 Design Criteria and System Concepts

Traditional system designs and ad-hoc unicast approaches such as [6] do not meet our requirements of scalability (i.e., the number of cars) and the unpredictable nature of car movements. Moreover, flooding of unrefined traffic information across the network does not scale.

In the project AutoNomos, we take a different, localized and distributed approach, in which cars collect information about traffic events, aggregate the relevant data, and generate dynamic traffic reports in a self-organized manner. A crucial concept for this challenge is what we call a *hovering data cloud* (HDC).

HDCs are motivated by traffic phenomena that exist and prevail independent of the individual vehicles they consist of. They are self-organizing entities that are not restricted to particular hosting nodes, nor to fixed regions. This makes them somewhat similar to the virtual mobile nodes described in [7], but the distributed concept allows for a wider range of structures. Practically speaking, HDCs are associated with phenomena such as traffic jams, they are responsible for capturing the phenomena's events and characteristics, and they arise with the onset of the phenomenon. At any time, an HDC has a distinct origin defined by a center and an expanse, i.e. the propagation range. Both can change over time, accounting for the represented event. In this first paper we focus on the description of a traffic density with motionless HDCs.

HDCs expand when nodes selectively forward *HDC messages* in the direction of interest, e.g. towards traffic that is directed to HDC's origin. The forwarding is bounded by the parameter expanse of the HDC, which depends on the underlying traffic event. To extend the propagation range with affordable bandwidth, HDCs are aggregated by their similarity of underlying events.

The desired behavior can only be achieved by self-organizing systems, as the set of hosting nodes changes over time. Locality is also required. The novelty of our approach is that we not only take the traditional storage-centric point of view that assigns the responsibility for data processing and data forwarding to a particular set of nodes.

## 3 A Distributed Algorithm for Implementing HDCs

A HDC includes a generic dataset that hosts traffic statistics. Cars that travel through an HDC region inform the other cars in the region about their present location and datasets. Periodically, a single car updates the HDC with its recent traffic statistics and initiates the propagation of HDC messages. The initiating car can be the one that is closest to the HDC origin.

We have taken an approach of random broadcasts for message propagation and neighborhood discovery, because of the elegant balance between communication loads and refresh rates (see the MILE project [8] for more details). The *update interval* for HDC messages creation and delivery set the trade-off between network loads and refresh rates. In order to reduce the communication overhead, we limit the lifetime of HDCs in absence of refreshing HDC messages. Their propagation expanses are restricted to a predefined "horizon" standing for area (and time) of interest of cars.

We have implemented an HDC that is fed by a simple periodic event: the measurement of car density in a stationary region. Possible extensions may include traffic information such as drive-through speed and lane change behavior. These parameters can reveal interesting irregularities, say, due to blocked lanes or road hazards that drivers are trying to avoid. The potential of the HDC concept to detect unpredictable events is far greater than that of existing implementations.

## 4 Evaluation

We present simulation results for stationary events that monitor car density on predefined locations. These events are set every 1 km (where the events' radius is set to 500 m) on a road with two lanes in one direction. The traffic simulator SUMO [9] generates appropriate movement traces of cars (see [10]). We extended the network simulator ns2 [11] with our HDCs algorithm and fed ns2 with the mobility traces. The communication model is based on the standard IEEE 802.11 with a communication range set to 250 m. We produce variations of the car density and compare it to the propagated and received HDC messages. We are mainly interested in the accuracy of the measurements (car density and time), propagation delay and network overhead produced by the HDCs.

The flow of cars starts with 1500 cars per hour, and increases sharply to 4000 during the simulation. The average speed is stable at around 26 m/s. Figure 4.a depicts the real number of cars passing the 6.5 km mark of the road.

Figure 4.b and Fig. 4.c show the projected traffic density at 6.5 km as received by cars 5.5 km behind (at position 1 km). Measurements are illustrated as points in the curves. In the first simulation run (Fig. 4.b) the update interval for HDC messages is set to 2.5 s, which consumes about 19 kbit/s of bandwidth per km of the road. In the second run the update interval is increased to 5 s (Fig. 4.c), which reduces bandwidth consumption to approximately 9 kbit/s. Note that the trade-off between the HDC message update interval and the bandwidth consumption affects the granularity of traffic reports, because every point on the curve in Fig. 4 represents a received HDC message. Moreover, the curves show that the HDC message update interval influences the propagation delay. Delay decreases from 66 s in Fig. 4.c to 22 s in Fig. 4.b. With advanced forwarding techniques we aim at further decreasing the delay without increasing bandwidth.

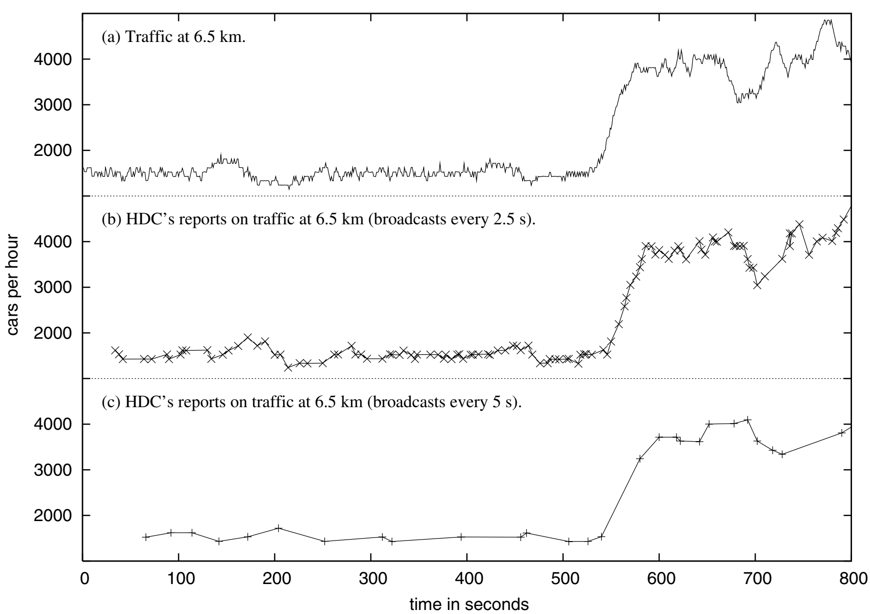


**Fig. 1.** Car density over time. Comparison of real car density and that reported by HDC messages.

## 5 Conclusion and Outlook

The AutoNomos project proposes a new approach for reporting traffic conditions. In contrast to centralized implementations, we facilitate a self-organizing car-to-car communication system. In this work, we introduce the key design concepts for sampling and reporting traffic information. The concept of HDCs allows information flow that is independent of the underlying car flow. Our preliminary tests demonstrate the feasibility of the proposed approach and the HDC concept.

A key topic in this project is to illuminate chain reaction scenarios that may occur when reporting on traffic. For example, the traffic congestion in a network of roads may oscillate, or the Braess paradox may appear (see [12]). Traffic forecasting services that includes additional traffic management aspects, e.g. route planning and road hazard notification, can increase drivers' confidence. We plan to investigate pattern recognition techniques for marking and mapping road hazards. Moreover, we will experiment with large-scale and complex traffic settings, and develop novel design concepts.

## References

1. Mazur, F., Chrobok, R., Hafstein, S., Pottmeier, A., Schreckenberg, M.: Future of traffic information - online-simulation of a large scale freeway network. IADIS International Conference WWW/Internet 2004 **1** (2004) 665–672
2. Franz, W., Hartenstein, H., Mauve, M.: Inter-Vehicle-Communications Based on Ad Hoc Networking Principles – The FleetNet Project. Universitätsverlag Karlsruhe (2005)
3. Car2car communication consortium. Web (2005) `http://www.car-to-car.org/`.
4. Varshney, U.: Vehicular mobile commerce. Computer **37**(12) (2004) 116–118
5. Chrobok, R., Pottmeier, A., Marinosson, S., Schreckenberg, M.: On-line simulation and traffic forecast: Applications and results. In: Proc. of the Internet and Multimedia Systems and Applications. (2002) 113–118
6. Nadeem, T., Dashtinezhad, S., Liao, C., Iftode, L.: Trafficview: A scalable traffic monitoring system. In: IEEE International Conference on Mobile Data Management. (2004) 13–26
7. Dolev, S., Gilbert, S., Schiller, E., Shvartsman, A., Welch, J.: Autonomous virtual mobile nodes. In: 3rd Workshop on Foundations of Mobile Computing (DIAL-M-POMC). (2005)
8. Hellbrück, H., Fischer, S.: Mine and mile: Improving connectivity in mobile ad-hoc networks. Mobile Computing and Communications Review MCCR **8**(4) (2004) 19–36
9. Krajzewicz, D., Hertkorn, G., Rössel, C., Wagner, P.: Sumo: An open-source traffic simulation. In: Proc. of the 4th Middle East Symposium on Simulation and Modeling. (2002) 183–187
10. Krauß, S., Wagner, P., Gawron, C.: Metastable states in a microscopic model of traffic flow. Physical Review E **55**(304) (1997) 5597–5602
11. University of Southern California: Ns-2: Network simulator-2. Web (2005) `http://www.isi.edu/nsnam/ns/`.
12. Murchland, J., Braess, D.: Braess's paradox of traffic flow. Transpn. Res. **4** (1970) 391–394